

Phase transitions in the generalization behaviour of multilayer perceptrons: II. The influence of noise

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

1997 J. Phys. A: Math. Gen. 30 8541

(<http://iopscience.iop.org/0305-4470/30/24/017>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 171.66.16.110

The article was downloaded on 02/06/2010 at 06:07

Please note that [terms and conditions apply](#).

Phase transitions in the generalization behaviour of multilayer perceptrons: II. The influence of noise*

B Schottky† and U Krey‡

† Department of Computer Science and Applied Mathematics, Aston University, Birmingham B4 7ET, UK

‡ Institut für Physik II der Universität Regensburg, Universitätsstrasse 31, D-93040, Germany

Received 24 April 1997, in final form 16 September 1997

Abstract. We extend our study of phase transitions in the generalization behaviour of multilayer perceptrons with non-overlapping receptive fields to the problem of the *influence of noise* concerning, for example, the input units and/or the couplings between the input units and the hidden units of the second layer (= ‘input noise’) or the final output unit (= ‘output noise’). Without output noise, the output itself is given by a general, permutation-invariant Boolean function of the outputs of the hidden units. As a result we find that the phase transitions which we found in the deterministic case, mostly *persist* in the presence of noise. The influence of the noise on the position of the phase transition, as well as on the behaviour in other regimes of the loading parameter α , can often be described by a simple rescaling of α depending on the strength and type of the noise. We then consider the problem of the optimal noise level for Gibbsian and Bayesian learning, looking at replica symmetry breaking as well. Finally, we consider the question of why learning with errors is useful at all.

1. Introduction and overview

1.1. Introduction and basic definitions

Recently [1], one of us treated the problem of phase transitions in the generalization behaviour of two-layer neural networks with non-overlapping receptive fields. The architecture of the systems considered is shown in figure 1. It corresponds to a tree of totally N input units, which are grouped into K vectors ξ_1, \dots, ξ_K of $M := N/K$ binary components $\xi_{k,m} = \pm 1$, with $k = 1, \dots, K$ and $m = 1, \dots, M$.

Each one of these vectors ξ_k determines the binary output σ_k of a so-called ‘hidden unit’ according to the perceptron-rule

$$\sigma_k = \text{sgn} \left(\frac{1}{\sqrt{M}} \mathbf{w}_k \cdot \xi_k \right) \equiv \text{sgn} \left(\frac{1}{\sqrt{M}} \sum_{m=1}^M w_{k,m} \cdot \xi_{k,m} \right) \quad (1)$$

where the so-called coupling vectors \mathbf{w}_k have M arbitrary real components $w_{k,m}$, which are only constrained by the normalization $\mathbf{w}_k^2 = M$.

The final output σ (= ‘classification’, ‘answer’) of the machine for a given input (= ‘question’) results from a fixed Boolean function

$$\sigma = B(\sigma_1, \dots, \sigma_k) \equiv B(\{\sigma_k\}) \quad (2)$$

* Based on the PhD thesis of B Schottky, Regensburg 1996.




Figure 1. The architecture of the class of networks is considered. There are N binary input units separated into K different groups leading each to a ‘hidden unit’. The binary outputs of the ‘hidden units’ are fed into a final Boolean output function. Only the weights w from the inputs to the hidden units can be modified by learning processes. Two kinds of noise will be considered below: ‘input noise’ and ‘output noise’.

of the outputs σ_k of the hidden units. This Boolean function is arbitrary apart from the postulate that it should be invariant against a permutation of the arguments.

Now the task of this classification machine is to learn a certain ‘rule’ by *modification of the coupling vectors* through learning the correct classification of a set of input examples $\xi^\mu = (\xi_1^\mu, \dots, \xi_K^\mu)$, with $\mu = 1, \dots, p$. Here it is assumed that the so-called loading parameter $\alpha := p/N$ is finite, while the thermodynamic limit $N \rightarrow \infty$ is taken.

In the following it is also assumed that the ‘rule’, by which the correct answers follow from the questions, is implemented by a ‘teacher perceptron’ of the same architecture as given above, with fixed ‘teacher couplings’ w^t . In particular we assume that the Boolean function of the student machine is the same as that of the teacher. However, the noise levels can be different, unless otherwise stated (see below).

We consider the *generalization ability* $g(\alpha)$, see [2, 3], of the system after a training process with $p = \alpha \cdot N$ examples; $g(\alpha)$ is defined as the probability that after the training an additional random question is answered correctly, i.e. in the same way as the teacher would answer in the absence of noise. It should be stressed that after the training we switch off *any* noise, both for the teacher and student machine. In contrast, during the training, noise of various kinds will corrupt both the student and teacher behaviour (see below).

Of course $g(\alpha)$ generally does not only depend on α , but also on the architectures considered, i.e. on the Boolean function $B(\{\sigma_k\})$, and noise. Only in the limit $\alpha \rightarrow \infty$, as already shown in [1], in the absence of noise does the architecture not matter, and one obtains for $\alpha \rightarrow \infty$ asymptotically the universal result

$$g(\alpha) \rightarrow 1 - \frac{0.625}{\alpha} \quad (3)$$

for the so-called Gibbsian learning (see below), where a student is drawn randomly from an ensemble which consists in the deterministic case just of those students classifying the training set correctly.

This asymptotic result is independent of the choice of the Boolean function and the number K of hidden units. Although, as already mentioned, our Boolean functions are quite general, apart from the constraint of permutation invariance, and although the behaviour depends essentially only on a small set of characteristic numbers (see below), we mention for the following that the most important machines considered are:

- the *committee machine*: this machine classifies by a majority vote of the $K := 2n + 1$ hidden units, i.e. with

$$\sigma = \operatorname{sgn}\left(\sum_{k=1}^K \sigma_k\right) \quad (4)$$

- whereas the *parity machine* is defined for general $K \geq 2$ by

$$\sigma = \operatorname{sgn}\prod_{k=1}^K \sigma_k \quad (5)$$

- and finally the AND-machine by

$$\sigma = \operatorname{sgn}\left(\sum_{k=1}^K \sigma_k - K + 1\right) \quad (6)$$

i.e. a positive classification is only given, if all hidden units agree.

The main result of this paper concerns the possible existence of phase transitions in the generalization behaviour as a function of α . For example for the *parity* machine, in contrast to the *committee*, generalization only starts if α is larger than a critical value, [4–6] (‘Aha effect’). In the preceding paper, [1], this was discussed for the deterministic case, whereas in the present second part we discuss the influence of noise. Generally, we find that the phase transitions mostly persist, although with changed critical values, and we also find certain scaling laws combining the critical loading parameter α and the ‘noise strength’. Furthermore, the performance of the system is found to be optimal, when the noise of the ‘student machine’ adapts to that of the teacher in a certain way.

1.2. Overview

As there are a lot of categories considered in this paper, it is easy to lose track. So we give a brief overview for better orientation. The categories considered are as follows.

- Two types of noise, *input* and *output* noise, see sections 2.5 and 2.6 below.
- We consider: (i) the case where noise levels of the teacher and student are assumed to be *the same* (section 3); case (ii) where the student noise level can be *chosen to optimize* the learning behaviour (section 4).
- Although most results are discussed for general values of the reduced size $\alpha := p/N$ of the training set, the two limits $\alpha \rightarrow 0$ (more precise would be: ‘ α small’, see section 3.2.1) and $\alpha \rightarrow \infty$ are of special interest.
- In view of the noise strength, particular emphasis is put on the two limits of *small* and *large* noise levels, respectively.
- Our main results are for general two-layer perceptrons with K ‘hidden units’, where $K > 1$; however, emphasis is sometimes put on the case $K = 1$, i.e. the single-layer perceptron.
- There are different learning rules (section 2.2), and for $K > 1$ one has to distinguish the different Boolean output functions (section 2.3).
- Our main results are obtained with the *replica-symmetric* approach (see below), but we also discuss some results obtained with *broken replica symmetry* (see section 4).

Since in principle all these categories can be arbitrarily combined there is a very large number of combinations, but not all of them are considered in this paper. We also mention that some considerations are *only* made for the simple perceptron ($K = 1$).

The paper is now organized as follows.

- Section 2 outlines the theoretical framework used in this paper, describes the *learning algorithms* considered and introduces the two *types of noise* treated in this paper.

- Section 3 deals with case (i) mentioned above, i.e. there it is assumed that the *noise level* for the student is chosen to be *the same* as that of the teacher (which is not a bad choice). Both types of noise are considered, with special emphasis on the limiting cases $\alpha \rightarrow 0$ and $\alpha \rightarrow \infty$ of the loading parameter.

- In section 4 we investigate the impact of a *varying* student noise level at fixed teacher noise, i.e. case (ii), aiming at the optimal choice. This is only done for the case of output noise. Furthermore, here we concentrate mainly on the single-layer perceptron, taking ‘replica symmetry breaking’ (RSB) into account. The multilayer case is only treated for the limiting cases of large and small training sets.

- Section 5 deals with the question, why a finite student noise level (which means that some of the training patterns are not learned correctly) can be useful at all.

- Finally, section 6 presents our conclusions.

2. Basic theory

The answering behaviour of a student and teacher network for given weights is defined by the two functions

$$\phi_{s/t}(\sigma^\mu | \mathbf{w}^{s/t}, \boldsymbol{\xi}^\mu) \quad (7)$$

determining the probability of getting the final answer σ^μ to a question $\boldsymbol{\xi}^\mu$ if the weights \mathbf{w} are given. The sub-/superscripts ‘s’ and ‘t’ stand for student and teacher, respectively. So ϕ encodes both the underlying architecture and the noise process corrupting the answer.

2.1. The version space and free energy

From ϕ_t one derives the probability P_t of getting answers $\{\sigma^\mu\}$ for patterns $\{\boldsymbol{\xi}^\mu\}$, with $\mu = 1, \dots, p$, by the teacher rule:

$$\begin{aligned} P_t(\tilde{\sigma}^p) &= \int d\mathbf{w}^t P_w(\mathbf{w}^t) \cdot p(\tilde{\sigma}^p | \mathbf{w}^t) \\ &= \int d\mathbf{w}^t P_w(\mathbf{w}^t) \prod_{\mu=1}^p \phi_t(\sigma^\mu | \mathbf{w}^t, \boldsymbol{\xi}^\mu) \end{aligned} \quad (8)$$

where the so-called *prior* P_w takes care of the normalization constraints. Hereafter we use $\tilde{\sigma}^p$ as the notion for the set $\{\sigma^\mu\}$ of answers, with $\mu = 1, \dots, p$.

Using the Bayes theorem, the probability that a specific weight vector \mathbf{w}^s is the correct one given the training patterns and the answers by the teacher, is determined through ϕ_s by

$$p(\mathbf{w}^s | \tilde{\sigma}^p) = \frac{P_w(\mathbf{w}^s) p(\tilde{\sigma}^p | \mathbf{w}^s)}{\int d\mathbf{w}^s P_w(\mathbf{w}^s) p(\tilde{\sigma}^p | \mathbf{w}^s)} \quad (9)$$

with

$$p(\tilde{\sigma}^p | \mathbf{w}^s) = \prod_{\mu=1}^p \phi_s(\sigma^\mu | \mathbf{w}^s, \boldsymbol{\xi}^\mu). \quad (10)$$

This defines the *degree of membership* of a specific coupling vector w^s to the so-called *version space*. The version space contains all student coupling vectors with a weight proportional to the probability that these couplings agree with those of the actual teacher. So one defines the corresponding partition function

$$Z(\tilde{\sigma}^p) = \int dw^s P_w(w^s) \prod_{\mu=1}^p \Phi_s(w^s | \sigma^\mu, \xi^\mu) \tag{11}$$

where

$$\Phi_s(w^s | \sigma, \xi) = c_E \phi_s(\sigma | w^s, \xi) \tag{12}$$

with c_E being a positive free constant, which takes into account the fact that the degree of membership does not have to be normalized.

In our case we restrict the possible functions $\phi_{s/t}$ of the student resp. the teacher to depend only on the corresponding *local field values* $h_k^{s/t}$ of the hidden nodes,

$$h_k^{s/t} := \frac{1}{\sqrt{M}} w_k^{s/t} \cdot \xi_k \tag{13}$$

so

$$\phi_{s/t}(\sigma | w^{s/t}, \xi) \equiv \phi_{s/t}(\sigma | \{h_k^{s/t}\}). \tag{14}$$

We can now perform the Gardner analysis, [7], by calculating

$$\begin{aligned} \mathcal{F} &:= \frac{1}{N} \langle \langle \ln Z \rangle_{w^t} \rangle_{\{\xi^\mu\}_{\mu=1\dots p}} \\ &= \frac{1}{N} \text{Tr}_{\{\tilde{\sigma}^p\}} \langle P_t(\tilde{\sigma}^p) \ln Z(\tilde{\sigma}^p) \rangle_{\{\xi^\mu\}_{\mu=1\dots p}} \end{aligned} \tag{15}$$

which we will call ‘free energy’ although this is physically not precise.

To describe the structure of the version space in the thermodynamic limit $N \rightarrow \infty$ we introduce the order parameters

$$q_k := \frac{1}{M} w_k^1 \cdot w_k^2 \tag{16}$$

$$r_k := \frac{1}{M} w_k^t \cdot w_k^s. \tag{17}$$

So q_k is the overlap between the k th subperceptron of two students chosen randomly from the version space, and r_k is the corresponding overlap of a random student vector with the teacher couplings. Nevertheless, since B is restricted to be permutation symmetric, these quantities cannot depend on the node number k , and thus it is sufficient to use

$$q = q_k \quad r = r_k \tag{18}$$

as permutation-symmetric order parameters. In the replica-symmetric approximation, straightforward calculations, see [8], lead to the final result

$$\mathcal{F} = \text{extr}_{(q,r)} \left\{ K^{-1} \sum_{k=1}^K \left[\frac{1}{2} \ln(1-q) + \frac{1}{2} \frac{q-r^2}{1-q} \right] - \alpha W(q,r) \right\} \tag{19}$$

with the so-called *energy term*

$$W(q,r) = \int \prod_{k=1}^K \mathcal{D}t_k \sum_{\sigma=\pm 1} F_t(\sigma, \{q, r, t_k\}) \cdot \ln F_s(\sigma, \{q, t_k\}) \tag{20}$$

and the *architecture specification* for the teacher machine

$$F_t(\sigma, \{q, r, t_k\}) = \int \prod_{k=1}^K Ds_k \phi_t \left(\sigma, \left\{ s_k \sqrt{1 - \frac{r^2}{q}} - t_k \frac{r}{\sqrt{q}} \right\} \right) \quad (21)$$

and for the student machine

$$F_s(\sigma, \{q, t_k\}) = \int \prod_{k=1}^K Ds_k \phi_s \left(\sigma, \left\{ s_k \sqrt{1 - q} - t_k \sqrt{q} \right\} \right). \quad (22)$$

Here $Dx = (2\pi)^{-1/2} dx \exp(-x^2/2)$ is the Gauss measure, and the values of the order parameters q, r are fixed by the saddle-point conditions

$$\frac{\partial \mathcal{F}}{\partial q} = \frac{\partial \mathcal{F}}{\partial r} = 0. \quad (23)$$

These are very general formulae which allow us to calculate the order parameters for classification machines with tree architecture. Nevertheless, for *non*-permutation-symmetric Boolean functions B one would have to distinguish between the q_k, r_k for different nodes.

2.2. Gibbs and Bayes algorithms; Gardner analysis

As training algorithms we discuss: (i) the Gibbs algorithm and (ii) the Bayes algorithm. They are distinguished by the way the version space is utilized. For the Gibbs algorithm, a ‘typical student machine’ is drawn at random out of the version space according to the weight factor (9), and then an average is performed as usual. In contrast, for the Bayes algorithm one takes into account *all* members of the version space and gives that answer, which corresponds to their weighted majority vote. In this way, the *a posteriori* error probability is minimized. Therefore, the answers given by this so-called Bayes procedure usually cannot be obtained from only one machine of the kind considered.

2.3. Notation to encode the Boolean function

For a specific Boolean function B (with a given number K of hidden units) we define the following expression:

$$\Delta^\sigma(\{\sigma_l\}) = \begin{cases} 1 & B(\{\sigma_l\}) = \sigma \\ 0 & \text{constant.} \end{cases} \quad (24)$$

Thus, Δ^σ is 1 just for those internal representations which are mapped to ‘ σ ’ by the Boolean function B . We recall that only learnable problems are considered, so the same Boolean function specifies the architecture for the teacher and student networks, and thus we need no superscript to distinguish between them.

We also need a short code to denote a special architecture. We use the same convention as already introduced in [1]: a Boolean function is characterized by its number of nodes K and a special ‘mcode’ q with $q = \sum_{v=0}^{K-1} n_v 2^v$. Here $n_v = 0$ or $= 1$, respectively, if a positive vote of exactly v hidden units leads to a negative resp. positive final output σ of the Boolean function. (By convention, $v = K$ shall always imply a positive output.) The name is then combined to ‘ KK_mcodeq ’, so for example ‘ $K4_mcode2$ ’ is a network with four hidden units and positive output if exactly four or one hidden unit(s) have positive vote.

2.4. Error probability

The order parameters describe rather well the learning success (or failure) and we will often present just these values. Nevertheless the quantity, which is so to say of final interest, is the generalization ability. In the presence of noise, there exist of course different possibilities to define this quantity; our choice is, as stated above, to assume that *after training* the noise is switched-off completely, both for the teacher and student networks. Moreover, we refer rather to the generalization *error* $\epsilon(\alpha) := 1 - g(\alpha)$ measuring the probability that student and teacher *disagree* on a new question.

For a given architecture this generalization error is determined uniquely by the values of the order parameters $q(\alpha)$ and $r(\alpha)$, no matter which noise processes have influenced the learning.

For the Gibbs algorithm ϵ only depends on the typical student–teacher overlap r and is given by (see also equation (21) in [1]):

$$\epsilon(r) = \left(\frac{1}{2}\right)^K \sum_{\sigma=\pm 1} \left[\text{Tr}_{\{\sigma_k^t\}\{\sigma_k^s\}} \Delta^\sigma(\{\sigma_k^t\}) \Delta^{-\sigma}(\{\sigma_k^s\}) \prod_k \left(1 - \frac{1}{\pi} \arccos(\sigma_k^t \sigma_k^s r)\right) \right]. \quad (25)$$

For the Bayes algorithm we obtain the generalization error by

$$\epsilon_{\text{Bayes}} = \int \prod_{k=1}^K \text{D}t_k \min \left[\text{Tr}_{\{\sigma_k\}} \Delta^1(\{\sigma_k\}) \prod_k H(\sigma_k \gamma t_k), \text{Tr}_{\{\sigma_k\}} \Delta^{-1}(\{\sigma_k\}) \prod_k H(\sigma_k \gamma t_k) \right] \quad (26)$$

where the ‘min’ means that the error probability corresponds to the smaller fraction of the version space, which belongs to the minority votes. Note that in (26) the values $q(\alpha)$ and $r(\alpha)$ of both order parameters are important.

For the simple perceptron, and for the parity machine in general, there is a close relationship between (25) and (26); we will return to this point later.

2.5. Output noise

We have investigated the influence of two types of noise. The first one, called ‘output noise’, flips the final output with the probability

$$p_f = \frac{e^{-\beta}}{1 + e^{-\beta}}. \quad (27)$$

This corresponds to

$$\phi(\sigma, \{h_k\}) = \frac{\exp(-\beta \cdot \delta[-\sigma, B(\{\text{sgn}(h_k)\})])}{1 + \exp(-\beta)} \quad (28)$$

as the probability of obtaining the output σ for given fields $\{h_k\}$ at the hidden units. In equation (28), we have written $\delta[i, k]$ for the Kronecker symbol ($= 1$, if $i = k$, $= 0$ for $i \neq k$), and β^{-1} is the parameter characterizing the noise strength.

2.6. Input noise

The second type of noise, called ‘input noise’, causes a noise perturbation of the local fields at the hidden units, $h_k \rightarrow h_k + \eta$, where η is a random variable with the Gaussian distribution $\rho(\eta) = (2\pi\gamma)^{-1/2} \exp(-\eta^2/(2\gamma))$. Similarly to the case of output noise, it is natural to define the noise strength $\beta^{-1} := \gamma$.

The origin of this noise can be that the input pattern itself is subjected to corruption by noise or that there is weight noise in the couplings of the teacher.

The flip probability $p_f(\gamma)$ depends here on the architecture, namely it is

$$p_f(\gamma) = \epsilon \left(\sqrt{(1 + \gamma)^{-1}} \right) \quad (29)$$

with the generalization error $\epsilon(r)$ already known from equation (25).

The probability ϕ , to answer on $\{h_k\}$ with σ , is

$$\phi(\sigma, \{h_k\}) = \text{Tr}_{\{\sigma_k\}} \Delta^\sigma(\{\sigma_k\}) \prod_{k=1}^K H(-\sigma_k h_k \sqrt{\beta}) \quad (30)$$

where $H(x) := \int_x^\infty du \exp(-u^2/2)/\sqrt{2\pi}$.

Equation (29) can be seen as follows. The generalization error $\epsilon(r)$ defines the probability that student and teacher machine, which have overlap r , give a different answer on a question. The local fields at a node k of the respective machines can be written as $h^t = t$, $h^s = t \cdot r + v \cdot \sqrt{1 - r^2}$, with two independent, normally distributed random variables t and v . The flip probability, for comparison, defines the probability that the local field h^0 , by adding noise with average 0 and variance γ , is changed to h^1 , where $h^0 = t$ and $h^1 = t + v\sqrt{\gamma}$, such that the final answers differ. Thus, the problems are completely analogous, with the exact correspondences

$$\frac{\sqrt{1 - r^2}}{r} \hat{=} \sqrt{\gamma} \quad \text{or} \quad r \hat{=} \sqrt{\frac{1}{1 + \gamma}}. \quad (31)$$

With (29), one can easily discuss the small-noise limit in this case, since for $r \rightarrow 1$, according to equation (33) in [1], one gets

$$\epsilon(r) \rightarrow n_c \pi^{-1} \sqrt{2(1 - r)} \quad (32)$$

which follows, for example from equation (25) for $r \rightarrow 1$; thus one obtains

$$p_f(\gamma \rightarrow 0) \rightarrow n_c \cdot \pi^{-1} \sqrt{\gamma}. \quad (33)$$

Here n_c is, in the limit considered, the only architecture-dependent value determining the asymptotics $\epsilon(\gamma \rightarrow 0)$. It characterizes the ‘border-regime’ of the Boolean function $B(\{\sigma_k\})$, namely by

$$n_c = \left(\frac{1}{2}\right)^K N_c \quad (34)$$

where N_c is the number of all those possible $K \cdot 2^K$ bit-flips of the outputs of the hidden units, which would lead to a change in the final output, see equation (35) in [1].

3. Teacher and student machine have identical noise levels

In the following we assume at first $\phi_s \equiv \phi_t$, which means that the student machine uses the known noise level of the teacher machine. This assumption, which implies $r = q$, is natural, since in this way overfitting will be avoided; moreover, as we will see later, it is not too far from the optimal choice.

3.1. Free energy

After the preparations in the last section we can calculate the free energy for both noise types. One gets from equations (19) and (20)

$$\mathcal{F} = \text{extr}_{(q)} \left\{ \frac{1}{2} \ln(1 - q) + \frac{q}{2} - \alpha \cdot W(q) \right\} \quad (35)$$

where for the case of output noise

$$W(q) = - \int \prod_k Dt_k \sum_{\sigma} \frac{1}{1 + \exp(-\beta)} \text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \times \prod_k H(\sigma_k \gamma t_k) \ln[\text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \gamma t_k)] \quad (36)$$

with

$$\tilde{\Delta}^{\sigma}(\{\sigma_k\}) := \Delta^{\sigma}(\{\sigma_k\}) + \exp(-\beta) \Delta^{-\sigma}(\{\sigma_k\}) \quad (37)$$

and for input noise

$$W(q) = - \int \prod_k Dt_k \sum_{\sigma} \text{Tr}_{\{\sigma_k\}} \Delta^{\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \bar{\gamma} t_k) \ln \left[\text{Tr}_{\{\sigma_k\}} \Delta^{\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \bar{\gamma} t_k) \right] \quad (38)$$

with

$$\bar{\gamma} = \sqrt{\frac{q}{1 - q + (1/\beta)}}. \quad (39)$$

These formulae will be generalized below for the cases of different noise levels of the student and teacher machine, and for 1-step replica symmetry breaking (RSB1).

3.2. Identical noise levels: results for the case of output noise

3.2.1. *Small training set: $q \rightarrow 0$.* The limit of a ‘small’ training set is best characterized by the corresponding limit $q \rightarrow 0$; in some cases this implies $\alpha \rightarrow 0$, but sometimes the corresponding limiting α can have a finite values as well (see below). The opposite limit $q \rightarrow 1$, which implies $\alpha \rightarrow \infty$, will be considered in section 3.2.2.

For $q \rightarrow 0$ a redefinition of the *correlation moments* a_m^{σ} of equation (47) in [1] suffices to capture the influence of the noise considered. With

$$a_m^{\sigma} = \left(\frac{1}{2}\right)^K \text{Tr}_{\{\sigma_k\}} \Delta^{\sigma}(\{\sigma_k\}) \prod_{i=1}^m \sigma_i \quad (40)$$

one defines

$$b_m^{\sigma} := \frac{a_m^{\sigma} + e^{-\beta} a_m^{-\sigma}}{1 + e^{-\beta}} = \begin{cases} a_m^{\sigma} \tanh\left(\frac{\beta}{2}\right) & \text{for } m \geq 1 \\ a_0^{\sigma} \tanh\left(\frac{\beta}{2}\right) + \frac{e^{-\beta}}{1 + e^{-\beta}} & \text{for } m = 0. \end{cases} \quad (41)$$

Therefore, the b_m^{σ} fulfil the same algebraic relations, equation (51) in [1], as the a_m^{σ} , namely

$$b_m^{\sigma} = -b_m^{-\sigma} \quad \text{and} \quad b_0^{\sigma} + b_0^{-\sigma} = 1. \quad (42)$$

In particular, the so-called *order-index* n , see below, is unchanged; n is defined by

$$b_n^{\sigma} \neq 0 \quad b_m^{\sigma} = 0 \quad \text{for } 1 \leq m < n. \quad (43)$$

Moreover, for $W(q)$ one gets to lowest order in q

$$W(q) = -b_0^1 \ln b_0^1 - b_0^{-1} \ln b_0^{-1} - \ln(1 + e^{-\beta}) - \frac{q^n}{2} \left(\frac{2}{\pi}\right)^2 \binom{K}{n} \frac{(b_n^1)^2}{b_0^1 b_0^{-1}} + \dots \\ =: w_0 - q^n w_1 + \dots \quad (44)$$

which agrees completely with equation (54) in [1] apart from the replacement of a_m^{σ} by b_m^{σ} and by the non-essential additional term $\ln(1 + e^{-\beta})$. So the results for $q(\alpha)$ in [1] can be

simply generalized by these replacements. Only with the generalization error $\epsilon(\alpha)$ we have to keep in mind that after training the noise is switched off, such that for $\epsilon(\alpha)$ the a_m^σ must be kept.

In view of the order-index n , we thus can state as in [1] that:

- for $n = 1$, for example for the committee machine, the overlap $q(\alpha)$ increases for $\alpha \ll 1$ proportional to α , i.e. there is generalization right from the beginning, and one obtains

$$q(\alpha) \rightarrow 2\alpha \frac{K}{\pi} \frac{(b_1^1)^2}{b_0^1 b_0^{-1}}. \quad (45)$$

This happens for six of the nine cases with $K = 4$ hidden units given in figure 1 of [1].

- For $n = 2$, phase transitions of second order or of first order (or both) are possible, as discussed in detail in section 7.2 of [1]. Generally, for $n = 2$ the network is purely guessing, i.e. the error probability is $\frac{1}{2}$, as long as α is between 0 and the first critical value, whereas an increase of α beyond this critical value leads to a continuous (resp. discontinuous) increase of the generalization ability in the case of a second-order (resp. first-order) transition. These transitions with $q(\alpha) \equiv 0$ for $\alpha < \alpha_c$ are called by us ‘Aha-effect transitions’. As just stated, they appear only for $n \geq 2$, whereas for $n = 1$ only so-called ‘interim transitions’ (if at all) can occur: at an ‘interim transition’ $q(\alpha)$ is finite already below α_c .

If for $n = 2$ the second-order transition is not preceded by a first-order one, the critical loading is

$$\bar{\alpha}_c = \frac{\pi^2 (b_0^1 b_0^{-1})}{4K(K-1)(b_2^1)^2}. \quad (46)$$

The case $n = 2$ occurs twice in figure 1 of [1].

- For $n \geq 3$, as in the noise-free case, one always gets a first-order transition at a critical $\alpha_c > 0$. Nevertheless this α_c has to be obtained numerically since the behaviour around a $q = q_c$ with $q_c > 0$ is relevant.

This case happens for example for the parity machine with $K \geq 3$, which has $n = K$.

Thus, as long as there is no transition of first order, the behaviour can be described analytically by looking at how the noise strength β^{-1} changes the correlation moments b_m^σ . Some of the following statements are based on this fact; they are not exact as far as the *locations* of first-order transitions are concerned, but we do not always state this limitation explicitly.

If one considers only machines, which have the same probability for the two possible outputs $\sigma = \pm 1$, these results can be simply condensed into a rescaling

$$\alpha \rightarrow \alpha_{\text{eff}} := \alpha \cdot \tanh^2 \left(\frac{\beta}{2} \right) \equiv \alpha \cdot (1 - 2p_f)^2 \quad (47)$$

where p_f is the flip probability defined in equation (27). This can be intuitively understood as follows.

- $p \cdot (1 - 2p_f)$ is the ‘uncorrupted fraction’ of the training set.
- The results are affected by noise from both the teacher and the student machine, which explains the power of 2 in equation (47).

3.2.2. Identical noise levels; output noise; large training set: $q \rightarrow 1$. In this limit, the teacher network is approximated with arbitrary accuracy, for every noise strength β^{-1} . That this is possible, is not at all self-evident: the training set contains mistakes since the teacher makes errors, the student learns this set making errors as well, but despite these facts the

teacher machine is approximated perfectly. As we will see below, a bad training strategy could empede generalization; so the fact that the student accepts the noise level of the teacher, as assumed at present, is already a good strategy, although it is not yet optimal, as we will see later.

From equation (36) one obtains for $q \rightarrow 1$ the information gain

$$W(q) \rightarrow p_f[\ln(1 - p_f) - \ln p_f] + n_c w_1(\beta)\sqrt{1 - q} \tag{48}$$

with

$$w_1(\beta) = -\sqrt{\frac{2}{\pi}} \int_0^\infty du \left\{ \frac{1 - e^{-\beta}}{1 + e^{-\beta}} H(u) \ln \left[\frac{H(u) + e^{-\beta} H(-u)}{H(-u) + e^{-\beta} H(u)} \right] + \frac{e^{-\beta}}{1 + e^{-\beta}} \ln[e^\beta H(u) + H(-u)] + \frac{1}{1 + e^{-\beta}} \ln[e^{-\beta} H(u) + H(-u)] \right\}. \tag{49}$$

Determining $q(\alpha)$ from $W(q)$ and inserting the result again into equation (32), one obtains

$$\epsilon(\alpha) \rightarrow \frac{\sqrt{2}}{w_1(\beta)\pi\alpha}. \tag{50}$$

So again, as in the deterministic case, one has an $1/\alpha$ -asymptotics, and the prefactor does not depend on the Boolean function $B(\{\sigma_k\})$, but only on the noise level β . Therefore again, one is led to a rescaling

$$\alpha \rightarrow \alpha_{\text{eff}} := r(\beta) \cdot \alpha \equiv \frac{w_1(\beta)}{w_1(\infty)} \cdot \alpha \equiv \frac{w_1(\beta)}{0.720\,647} \cdot \alpha. \tag{51}$$

In figure 2, the scaling parameter $r(\beta) \equiv w_1(\beta)/0.720\,647$, which applies to the regime $q \rightarrow 1$, and the ‘intuitive’ scaling parameter $r^{\text{it}}(\beta) := [1 - 2p_f(\beta)]^2$, which applies to the limit $q \rightarrow 0$, are presented as a function of the ‘flip parameter’ $2p_f$, which corresponds to the ‘corrupted fraction of the training set’. Obviously $r(\beta)$ is $< r^{\text{it}}(\beta)$, which means that




Figure 2. The scaling factor $r(\beta)$ of equation (51), which applies to the case $q \rightarrow 1$ of high loading, and the ‘intuitive scaling factor’ $r^{\text{it}} = (1 - 2p_f)^2$, which applies to the limit $q \rightarrow 0$, are presented as a function of the ‘corrupted fraction’ $2p_f = 2e^{-\beta}/(1 + e^{-\beta})$ of the training set with respect to the output. $\beta := \beta_s = \beta_t = 1/T_s = 1/T_t$ characterizes the output noise levels of the student and the teacher machine.

for large loading the effect of noise is stronger than for low loading. But the difference is not large, so that the ‘intuitive reduction factor’ $r^{II}(\beta)$ will always give a good estimate of the effect. But it should be noted that in the limit $q \rightarrow 1$ already a slight corruption of the training set leads to a significant reduction of the generalization ability because of the infinite slope of $r(2p_f)$ in the limit $p_f \rightarrow 0$.

3.2.3. Output noise: Data collapsing in the large-noise limit. Strong noise is described by $\beta \rightarrow 0$. Calculating the behaviour in this limit for both, $q \rightarrow 0$ and $q \rightarrow 1$, one sees that $q(\alpha, \beta)$ is $\propto \beta^2 \alpha$. We have looked for (approximate) data-collapsing in the whole parameter region $0.2 < \beta < 1$, by plotting the curves $q(\alpha, \beta)$ not only as a function of α , with β as curve parameter, but using, instead, also the product $\beta^2 \alpha$ as a scaling variable. The results are compared in figure 3 and show that with the variable $\beta^2 \alpha$, for $\beta \leq 1$, a good, although still approximate, data-collapsing is obtained in the whole region $0 < \alpha \beta^2 < \infty$. This data-collapsing becomes asymptotically exact in the above-mentioned limits $q \rightarrow 0$ and $q \rightarrow 1$.

3.3. Identical noise levels: Results for the case of input noise

3.3.1. Small training set: $q \rightarrow 0$. Again we consider at first the limit $q \rightarrow 0$. In this case one gets from equation (39)

$$\bar{\gamma} \rightarrow \sqrt{q \cdot \frac{1}{1+\gamma}} =: \sqrt{q \cdot \zeta} \quad (52)$$

and for the free energy

$$-\mathcal{F} \rightarrow \text{extr}_{(q)} \left\{ -\frac{q^2}{4} - \alpha \cdot (w_0 - q^n \zeta^n w_1) \right\} \quad (53)$$

where w_0 and w_1 are defined as

$$w_0 = -a_0^1 \ln a_0^1 - a_0^{-1} \ln a_0^{-1} \quad w_1 = \frac{1}{2} \left(\frac{2}{\pi} \right)^n \binom{K}{n} \frac{(a_n^1)^2}{a_0^1 a_0^{-1}}. \quad (54)$$

Again, the order-index n , and thus the *qualitative* behaviour, remains unchanged by the noise. In view of the three cases of n , we now have the following.

- For $n = 1$, the limit $q(\alpha \rightarrow 0)$ is given by

$$q(\alpha) \rightarrow 2\alpha \zeta \frac{K}{\pi} \frac{(a_1^1)^2}{a_0^1 a_0^{-1}}. \quad (55)$$

Since $\zeta < 1$, the noise diminishes the overlap.

- For $n = 2$, the critical value $\bar{\alpha}_c$ of the second-order phase transition (if it is not preceded by a first-order one, see above) shifts to a higher value:

$$\bar{\alpha}_c(\zeta) = \frac{1}{\zeta^2} \bar{\alpha}_c(0). \quad (56)$$

- For $n \geq 3$, there is again a first-order phase transition, but the resulting critical loading must be determined numerically, since the behaviour around a $q = q_c$ with $q_c > 0$ is relevant.

Taking all three cases together, we find a rescaling, which also applies for $n \geq 3$, namely

$$\alpha \rightarrow \alpha_{\text{eff}} = \zeta^n \cdot \alpha. \quad (57)$$

We recall again the caution which has to be taken as far as first-order transitions are concerned.

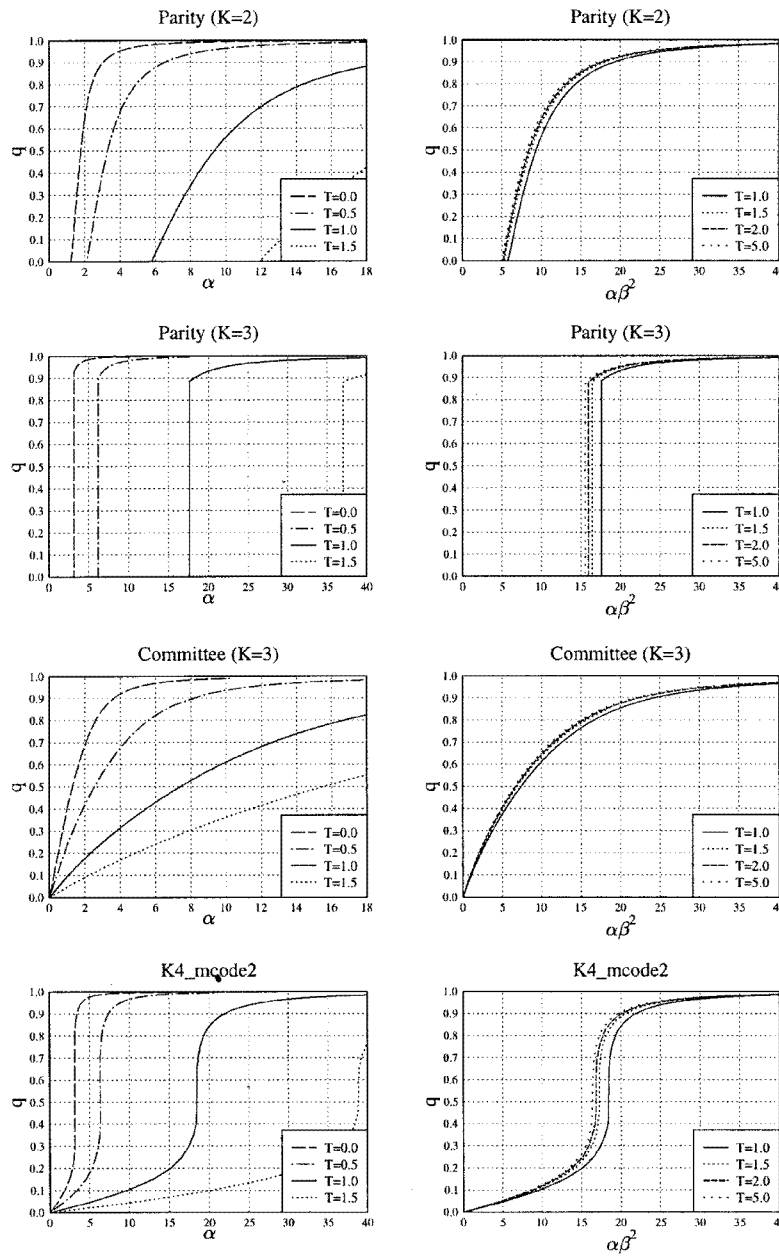


Figure 3. Data collapsing for output noise: On the left-hand side, the overlap $q(\alpha)$ of the couplings of the student and teacher machines is presented as a function of the loading $\alpha = p/N$, where p is the number of examples of the training set, for four different values of the common output noise level $T = 1/\beta$ of the student and the teacher ($T = 0; 0.5; 1.0$ and 1.5), whereas on the right-hand side the results (for $T = 1.0, 1.5, 2.0$ and 5.0) are presented as a function of $\alpha\beta^2$. Note that for the parity machine with $K = 2$ and 3 , respectively, one has an ‘Aha-effect’ phase-transition of second order ($n = 2$) and first order ($n = 3$) respectively, whereas for the committee machine and the K4_mcode2 machine, where the order index $n = 1$, the machine generalizes right from the beginning. For the committee machine, there is no phase transition at all, whereas for the K4_mcode2 machine, there is an ‘interim transition’ around $\alpha\beta^2 \approx 16$, a situation, which is also compatible with $n = 1$, see [1].

3.3.2. *Identical noise levels; input noise; large training sets: $q \rightarrow 1$.* Considering the limit $q \rightarrow 1$, i.e. $\alpha \rightarrow \infty$, one sees from equation (39) that now $\bar{\gamma}$ does not converge to ∞ , in contrast to the behaviour *without* noise, or with *output* noise. Instead

$$\bar{\gamma} \rightarrow \sqrt{\beta} [1 - \frac{1}{2}(1 + \beta)(1 - q)]. \quad (58)$$

Expanding equation (36) for $1 - q \rightarrow 0$, one obtains

$$W(q) \rightarrow w_0(\beta) + (1 - q) \cdot w_1(\beta). \quad (59)$$

If one determines $q(\alpha)$ from equation (23) and inserts it into equation (32), one obtains

$$\epsilon(\alpha) \rightarrow n_c \frac{1}{\pi \sqrt{w_1(\beta)\alpha}} \quad (60)$$

with

$$w_1(\beta) = \frac{(1 + \beta)\sqrt{\beta}}{2\sqrt{2\pi}} \int \prod_k \text{Dt}_k \sum_{\sigma=\pm 1} \text{Tr}_{\{\sigma_k\}} \Delta^\sigma(\{\sigma_k\}) \left[\sum_m \sigma_m t_m e^{-\beta t_m^2/2} \prod_{k(\neq m)} H(\sigma_k t_k \sqrt{\beta}) \right] \\ \times \ln \left[\text{Tr}_{\{\sigma_k\}} \prod_{k(\neq m)} H(\sigma_k t_k \sqrt{\beta}) \right]. \quad (61)$$

From equation (60) one can see that *input noise*, in contrast to output noise, leads to a drastic deterioration of the generalization ability, namely from an asymptotics as $\epsilon(\alpha) \rightarrow c/\alpha$ to the slower decrease $\epsilon(\alpha) \rightarrow \tilde{c}/\sqrt{\alpha}$. Additionally we find that the prefactor \tilde{c} of this behaviour—in contrast to c —depends on the architecture. In figure 4, for the $K = 2$ - and $K = 3$ -parity machines and for the $K = 3$ -committee, we present this *prefactor* \tilde{c} of the asymptotic behaviour $\epsilon(\alpha \rightarrow \infty) \rightarrow \tilde{c}/\sqrt{\alpha}$ as a function: (i) of γ and (ii) of the ‘flip probability’ $p_f = n_c \pi^{-1} \sqrt{\gamma}$. Interestingly, with the last-mentioned representation, the data almost collapse to a single curve, although the results look quite different when presented against γ .

For the simple perceptron a corresponding result is given in [3].

3.3.3. *Input noise: Data collapsing in the large-noise limit.* For the case of input noise it can further be shown that in the high-temperature limit $\beta \rightarrow 0$ and for $\alpha \rightarrow \infty$

$$\epsilon(\alpha) \rightarrow n_c \frac{\sqrt{2}}{\pi} \frac{1}{\sqrt{V_0 \beta^n \alpha}} \quad (62)$$

with

$$V_0 = \left(\frac{2}{\pi}\right)^n \sum_{\sigma=\pm 1} \frac{(a_n^\sigma)^n}{a_0^\sigma} \binom{K}{n}. \quad (63)$$

The corresponding behaviour for $\alpha \rightarrow 0$ can be seen from (57). Summarized, data collapsing in the limit $\beta \rightarrow 0$ both for $q \rightarrow 0$ and $q \rightarrow 1$, but with n -dependent rescaling $\alpha \rightarrow \alpha_{\text{eff}} = \beta^n \alpha$.

In figure 5, for two examples, we check whether for the input noise temperatures $\gamma = 1/\beta = 1.0, 2.0, \dots, 5.0$ one gets data collapsing. Thereby we use $\zeta = 1/(1 + \gamma)$ rather than β for rescaling which is a better choice if β is not close to 0. We compare results for q plotted as a function of α with results, where q is plotted against α/ζ^2 for the $K = 2$ -parity machine, resp. against α/ζ for the K4_mcode2 machine.

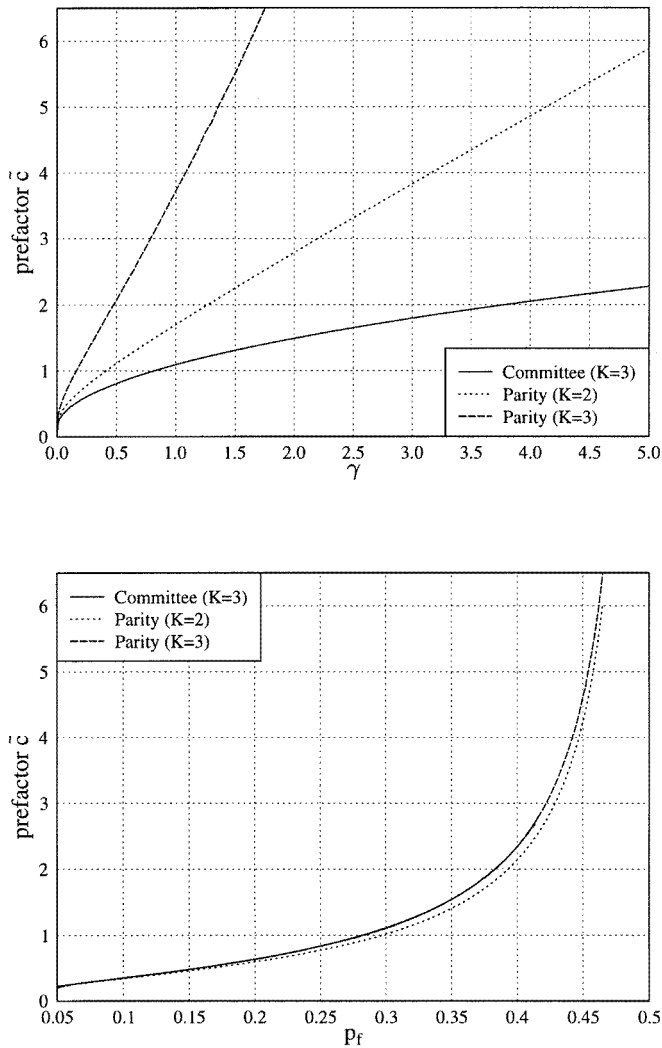


Figure 4. For the $K = 2$ and $K = 3$ parity machines, and for the $K = 3$ committee, the prefactor \tilde{c} of the $\tilde{c}/\sqrt{\alpha}$ -asymptotics of the error-probability $\epsilon(\alpha)$ is plotted against γ and p_f , respectively, for the case of input noise; $\gamma \equiv \beta^{-1}$ is the common noise level of teacher and student machines, and p_f is defined in (33) and (34). In the lower plot, the curves for the $K = 3$ parity and committee machines overlap to the accuracy of the drawing.

3.4. Identical noise levels: The different impact of output and input noise

As already seen there are some significant differences of the impact caused by output and input noise, respectively. Let us first have a closer look at the influence of a small disturbance by noise.

3.4.1. Small-noise limit: $\beta \rightarrow \infty$. To compare the impacts of noise in this limit we have to distinguish between the cases of small and large reduced size $\alpha = p/N$ of the training set, respectively. For small α the effect of input noise is for $\gamma \rightarrow 0$ (or $\beta \rightarrow \infty$) according

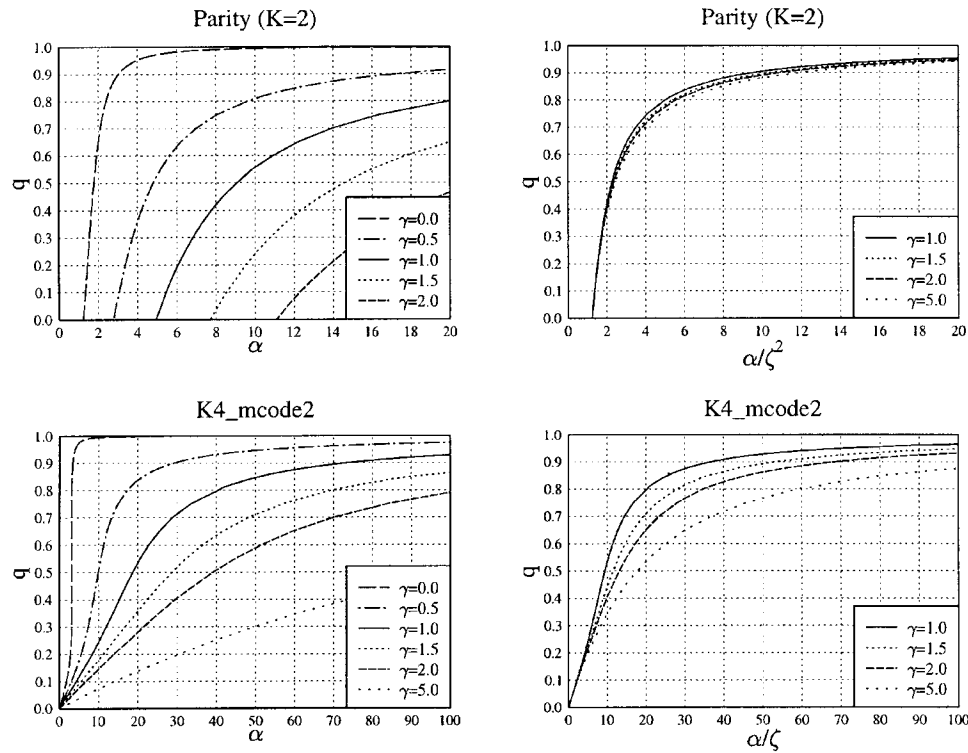


Figure 5. Data collapsing for input noise. For the $K = 2$ parity machine and for the K4_mcode2 machine, with common noise temperatures $\gamma = 1/\beta$ of the teacher and student machine, ranging from $\gamma = 0$ via 1, 1.5, 2 to 5, the overlap q between teacher and student couplings is plotted against α and α/ζ^n , respectively, where n is the order-index of the system and $\zeta := (1 + \gamma)^{-1}$. For the parity machine, the data collapsing is almost perfect, whereas for the second machine it applies only to the limits of small and high values of α/ζ . Note that here, in contrast to figure 3, the ‘interim transition’ of the second machine is destroyed by the input noise.

to equation (57)

$$\alpha_{\text{eff}} = \zeta^n \cdot \alpha = (1 + \gamma)^{-n/2} \cdot \alpha \rightarrow \left(1 - \frac{n}{2}\gamma\right) \alpha. \quad (64)$$

With the flip rate p_f given by equation (33) in this limit we have for *input noise*

$$\alpha_{\text{eff}} = \alpha \cdot \left[1 - \frac{n}{2} \left(\frac{\pi p_f}{n_c}\right)^2\right]. \quad (65)$$

The training set is therefore only reduced by a small amount $\propto p_f^2$. In contrast, for *output noise* this amount is $\propto p_f^1$. For machines with equal probability for final output $\sigma = \pm 1$ this can directly be seen from equation (47). This means that a small amount of input noise does hardly matter for the case of $\alpha \rightarrow 0$, in contrast to a small amount of output noise, which—so to say—instantly deteriorates the behaviour.

On the other hand, for $\alpha \rightarrow \infty$ a small amount of input noise induces a *qualitative change* in the asymptotics, since then the behaviour is shifted from the $1/\alpha$ -asymptotics to the slower $1/\sqrt{\alpha}$ -decrease, if α is beyond the corresponding (non-universal) crossover value. In contrast, for output noise the $1/\alpha$ -behaviour is qualitatively unchanged, only the prefactor is increased.

3.4.2. *Input noise: Disappearance of phase transitions.* A further difference concerns the phase transition of the K4.mcode2 machine: Here the *intermediate* phase transition as existent in the deterministic case (see [1]) disappears for the case of input noise. This does not happen for output noise. Whether other intermediate phase transitions are affected similarly has to be checked.

Nevertheless all other types of phase transitions (see section 3.2.1 above or [1]) persist for *both cases* of noise, although the shift of the critical parameters depends of course both on the strength and type of the noise.

4. Optimization of the noise level of the student machine

In this section we focus exclusively on *output noise*, but now different noise levels for student and teacher networks are allowed. The following two sections (4.1 and 4.2) give the theory for the general multilayer case. The formulae are evaluated mainly for the simple perceptron; for the general case just asymptotic results within the replica-symmetric formalism are given.

4.1. Different noise levels; output noise: Replica-symmetric formalism

In the replica-symmetric formalism of the preceding section, one has two different noise strengths β_t and β_s of the teacher resp. student machine, and additionally now $q \neq r$. Therefore, instead of equations (35) and (36) one has

$$\mathcal{F} = \text{extr}_{(q,r)} \left\{ \frac{1}{2} \ln(1-q) + \frac{q-r^2}{2(1-q)} - \alpha \cdot W(q,r) \right\} \quad (66)$$

where for the present case of output noise it is

$$W(q,r) = - \int \prod_k Dt_k \sum_{\sigma} \frac{1}{1 + \exp(-\beta_t)} \text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \gamma_r t_k) \\ \times \ln \left[\text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \gamma t_k) \right] \quad (67)$$

with

$$\gamma = \sqrt{\frac{q}{1-q}} \quad \text{and} \quad \gamma_r = \frac{r}{\sqrt{q-r^2}} \quad (68)$$

and with $\tilde{\Delta}$ defined by equation (37) with $\beta \equiv \beta_t$ and β_s , respectively. The functions $q(\alpha)$ and $r(\alpha)$ follow again from the saddle-point conditions (23). An additional quantity of interest is the *relative training-error* $\epsilon_{tr} := E_{tr}/p = -\alpha^{-1} \partial \mathcal{F} / \partial \beta_s$; ϵ_{tr} is thus the fraction of the training set, which is misclassified by the student machine. A straightforward calculation, see [8], yields

$$\epsilon_{tr} = \frac{1}{1 + e^{-\beta_t}} \int \prod_{k=1}^K Dt_k \sum_{\sigma=\pm 1} \frac{\text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \gamma_r t_k)}{\text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \gamma t_k)} \\ \times e^{-\beta_s} \text{Tr}_{\{\sigma_k\}} \Delta^{-\sigma}(\{\sigma_k\}) \prod_k H(\sigma_k \gamma t_k). \quad (69)$$

4.2. Different noise levels; output noise: RSB

Within the usual RSB1 scheme, see [9], one gets with $q^{a,b} \equiv 1$ for $a = b$, $q^{a,b} \equiv q_1$, if a and b are different, but belong to the same subgroup of m of the n replicas, and $q^{a,b} \equiv q_0$ else, and with $\Delta q := q_1 - q_0$:

$$\mathcal{F} = \text{extr}_{(r, q_0, q_1, m)} \left\{ \frac{q_0 - r^2}{2(1 - q_1 + m\Delta q)} + \frac{1}{2} \ln(1 - q_1) + \frac{1}{2m} \ln \left(1 + \frac{m\Delta q}{1 - q_1} \right) - \frac{\alpha}{m} \cdot W(r, q_0, q_1, m) \right\} \quad (70)$$

with

$$W(r, q_0, q_1, m) = - \int \prod_k D t_k \sum_{\sigma} \frac{1}{1 + \exp(-\beta_t)} \text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \prod_k H \left(\sigma_k \frac{r}{\sqrt{q_0 - r^2}} t_k \right) \times \ln \left\{ \int \prod_k D v_k \left(\text{Tr}_{\{\sigma_k\}} \tilde{\Delta}^{\sigma}(\{\sigma_k\}) \prod_k H \left[\sigma_k \frac{t_k \sqrt{q_0} + v_k \sqrt{\Delta q}}{\sqrt{1 - q_1}} \right] \right)^m \right\}. \quad (71)$$

4.3. Replica symmetry and RSB results for the simple perceptron

In the following we use the notation ‘perfect student’ or ‘perfect learning’ to describe the fact that the given training set \mathcal{TS} is learnt *without* any error, so ‘ \mathcal{TS} -perfect’ would be a more precise terminology. ‘Non-perfect learning’ (or better ‘ \mathcal{TS} -nonperfect learning’) means that errors with respect to \mathcal{TS} are made. We stress here that \mathcal{TS} may already be corrupted with respect to the original *rule*. (We also mention that ‘perfect learning’ is sometimes also used to describe coinciding architecture of student and teacher which is the case here anyway. To keep our special definition in mind, we always use primes in the terminology ‘perfect’.)

4.3.1. Replica symmetry results: ‘Non-perfect’ teacher, but ‘perfect’ student. In the following, for simple perceptrons ($K = 1$) we consider at first the case of a ‘perfect’ student machine (i.e. necessarily $\beta_s = \infty$) but allow for a ‘non-perfect’ teacher ($\beta_t < \infty$), which means that the training set itself is partially corrupted, since the answers given by the teacher on the questions ξ^{μ} ($\mu = 1, \dots, p = \alpha N$) do not always follow the rule, but are partially random.

Of course, ‘perfect learning’ of the *corrupted* training set is then possible only up to a specific α_c depending on the noise; for example if for *all* input patterns ξ^{μ} the outputs, prescribed by \mathcal{TS} , would be randomized with respect to the original rule, then one would get the famous result $\alpha_c = 2$ of Gardner, [7].

In figure 6(a), both for the case of output noise and input noise, α_c is presented as a function of the non-corrupted fraction $(1 - 2p_f)$ of the training set, with p_f taken from equations (27) and (29), respectively. Additionally, in figure 6(b), for output noise with $T_t = 1$, the curve $r(\alpha)$ for *maximal stability learning* (MSL) is shown (full curve) and compared with the corresponding result for Gibbs learning (broken curve). For the MSL case, the student is not a random member of the version space but has those couplings, which lead to *maximal stability* of the classification of the whole training set of patterns mapped to +1 and -1, respectively. This specific member can be obtained by the well-known AdaTron algorithm, [10].

Two overfitting effects can be seen: considering Gibbs learning, the overlap *decreases* at the end of the curve. Compared with MSL, Gibbs learning is worse for small α (which




Figure 6. (a) Storage capacity α_c of a deterministic student perceptron (i.e. $K = 1, T_s = 0$) in the presence of a noisy training set ($T_t \neq 0$): a fraction p_f of the binary answers of the teacher perceptron are misclassified. Both, input and output noise are considered. (b) Overlap r of a deterministic student perceptron (i.e. $K = 1, T_s = 0$) with the teacher vector in the presence of a noisy training set with output noise strength $T_t = 1$, plotted as a function of the reduced size $\alpha := p/N$ of the training set. The full curve is for maximal stability learning, i.e. the AdaTron algorithm, [10], while the broken curve is for Gibbs learning. The two overfitting effects appearing here are explained in the text.

is expected) but becomes better for $\alpha \rightarrow \alpha_c$, although MSL chooses a specific vector for the student which is supposed to perform the classification task very well.

These are hints that training with noise might avoid overfitting effects of the student and therefore could lead to an enhanced performance.

4.3.2. Replica symmetry results: ‘Non-perfect’ teacher and ‘non-perfect’ student. Here we consider again the perceptron ($K = 1$) and assume a given output noise level $\beta_t = 1$ of the

teacher machine.

In figure 7 we present the dependence of various quantities on the noise strength $T_s := 1/\beta_s$ of the student machine. These are as follows.

- The overlap $r(T_s)$ of the two coupling sets. This quantity determines the generalization ability of the Gibbs training algorithm.
- The typical overlap $q(T_s)$ of two student perceptrons.
- The overlap r_{cp} of the so-called ‘central-point network’ with the teacher machine: The coupling vector of the ‘central-point network’ is obtained as a (weighted) average of the coupling vectors of all student perceptrons forming the Gibbs ensemble. Analogously to equations (75) and (76) in [1] one can show that $r_{\text{cp}}(T_s) = r(T_s)/\sqrt{q(T_s)}$. The corresponding generalization error is obtained by plugging r_{cp} into (25).

Moreover, for $\beta_s = \beta_t$ the corresponding network turns out to have the same generalization ability as an exploitation of the version space performed by the Bayes algorithm, see [1].

From the results of figure 7, the following points should be noted.

(1) For all values of α , $r(T_s)$ and $q(T_s)$ cross at $T_s = T_t$. In fact, for this case the teacher has the same properties as a typical student of the Gibbs ensemble.

(2) The curves $r_{\text{cp}}(T_s)$ have a flat maximum at $T_s = T_t$. This is also obvious: since in this case the central-point network reaches the generalization ability of the Bayes algorithm, which is maximal, according to information theory. Nevertheless, since the curve $r_{\text{cp}}(T_s)$ is very flat around the maximum, the detailed value, $T_s = T_t$, is non-essential.

(3) In contrast, the overlap $r(T_s)$ for Gibbsian learning shows a pronounced maximum for a finite noise level T_s only for $\alpha = 2$ and $\alpha = 5$, but not for $\alpha = 1$. Obviously for Gibbs learning, training with output noise (i.e. $T_s > 0$) is only advantageous beyond a finite α , which of course depends on T_t . For a similar model this was reported in [11].

Figure 8 presents this optimal value of the student noise level T_s as a function of α for fixed $T_t = 1$. Beyond $\alpha = \tilde{\alpha}_c$ (≈ 0.6 in figure 8) training with noise leads to better generalization. We determined the value of $\tilde{\alpha}_c$ only numerically from the appearance of a maximum. For large α , the optimal T_s for $T_t = 1$ converges to 0.605 24.

In figure 9, the training error ϵ_{tr} is presented as a function of T_s for $T_t = 1$ and for different α -values, ranging from $\alpha = 0$ (lowest curve for small T_s) to $\alpha \rightarrow \infty$. For $T_s \rightarrow \infty$ all curves converge to $\epsilon_{tr} = 0.5$, as they should. For $\alpha \rightarrow 0$ the result is only determined by β_s , and for $\alpha \rightarrow \infty$ only by β_t , namely

$$\epsilon_{tr}(T_s)|_{\alpha \rightarrow 0} = \frac{e^{-\beta_s}}{1 + e^{-\beta_s}} \quad \epsilon_{tr}(T_s)|_{\alpha \rightarrow \infty} = \frac{e^{-\beta_t}}{1 + e^{-\beta_t}}. \quad (72)$$

4.3.3. Different noise levels; simple perceptron: RSB results. The non-monotoneous behaviour of $r(T_s)$ for $\alpha = 2$ in figure 7 is not an artefact of replica symmetry (RS): for $\alpha = 2$ we determined an optimal $T_s > 0$, and since the problem is *learnable* for this α and every T_t , RS is correct.

However, when α becomes $> \alpha_c$ (which is always larger than 2 according to figure 6(a)), one expects RSB. In figure 10, for $\alpha = 5$ and $T_t = 1$, the order parameters $r(\alpha)$ and $q(\alpha)$ are presented, as obtained in RS (heavy curve) and RSB1 theories. Obviously, RS is correct around the optimal student temperature $T_s \approx 0.35$ and beyond, but for smaller values ($T_s \lesssim 0.27$) RS is broken. In fact, in the RS approach, a paper of Györgyi, [11], predicts (for a similar model) a rather large and useful effect of ‘training with noise’ already for the case of small T_s . However, due to RSB corrections, see figure 11, the benefit of noise is weaker than expected in the RS calculation of [11] since for the case of small noise of the




Figure 7. For the perceptron ($K = 1$) and the case $T_t = 1$ of the teacher's output noise level and the three cases of $\alpha = 1, = 2,$ and $= 5$, the overlaps q and r for Gibbs learning, and r_{cp} for the central-point network, i.e. Bayesian learning, as explained in the text, are plotted over the student machine's output noise level T_s . Note that from r an *optimal* noise temperature $T_s = T_{opt}$ can be defined.




Figure 8. Here the optimal student machine's output noise temperature T_{opt} , as determined in figure 7, is presented over α for $T_t = 1$ (Gibbs learning, $K = 1$). The dashed line is the asymptotic limit for $\alpha \rightarrow \infty$.




Figure 9. For the perceptron ($K = 1$) with Gibbs learning, for various values of α and fixed teacher machine's output noise temperature $T_t = 1$, the training error ϵ_{tr} of equation (69) is plotted over the output noise temperature T_s of the student machine.

student machine the overlap r^{RSB} calculated in RSB1 is higher than calculated with RS, and the value obtained in RSB1 for $\alpha \lesssim 0.27$ is almost as high as that obtained at the optimal value $T_s \approx 0.35$, where RS is correct. An extension of RSB, for example to RSB2 (see [9]) may even enhance the value of $r(T_s)$ calculated in RSB1.

Here we mention that Uezu, in a recent preprint, [17], has independently treated problems studied in section 4.3, with largely overlapping results.




Figure 10. For Gibbs learning with $\alpha = 5$, $T_I = 1$ and $K = 1$ the order parameters r^{RS} and q (for RS), and r^{RSB} , q_0 and q_1 (in RSB1) are presented over the student perceptron's output noise temperature T_s . For values of T_s which are slightly smaller than the optimal value $T_s \approx 0.35$, RS is broken and the overlap r^{RSB} is somewhat enhanced with respect to the RS case, but still smaller than the optimum.




Figure 11. For $T_I = 1$ and $K = 1$, the optimal overlap $r(T_s = T_{opt}(\alpha); \alpha)$ is presented. For comparison, also the overlap obtained for $T_s = 0.1$, where RS is broken, is plotted, both in RS approximation and in RSB1, where the result is only slightly suboptimal.

4.3.4. Why RSB occurs: An intuitive explanation. As stated above, RSB can occur if α becomes $> \alpha_c$ (which is always larger than 2 according to figure 6(a)). Nevertheless RS




Figure 12. This figure suggests an analogy, making plausible that RS is broken for small student machine's noise level (left scenario), but is *restored* beyond a critical value of T_s , for fixed teacher machine's noise level T_t (right scenario).

is restored if the student temperature T_s increases above a critical point as demonstrated in the previous paragraph.

To understand this, let us consider the case $T_s \rightarrow 0$, which means that only student machines with minimal training error are allowed. The available phase space will then separate into disjunct parts: to every disjunct part student machines belong, which misclassify a certain (different) minimal set of patterns of the training examples. The student machines 'in between the disjunct components' make more errors. If T_s is increased, they also become 'more and more allowed' until finally the allowed region of phase space melts together to a single component, i.e. RS is restored. This scenario is presented qualitatively in figure 12.

4.4. Different noise levels; output noise: RS and RSB results for multilayer networks

Since the numerical effort increases drastically for multilayer networks we consider exclusively output noise and restrict ourselves to give just asymptotic results for this case. For $T_t = 0.2, 0.5, 1.0$ and 2.0 the ratio T_s/T_t is varied.

4.4.1. Large training sets: $q \rightarrow 1$. We ask, how the student temperature T_s should be chosen in the case of large training sets. The aim is to obtain an optimal prefactor for the asymptotic behaviour of the generalization error, $\epsilon(\alpha) \rightarrow c_0/\alpha$. This can be calculated analytically for given T_t , results are presented in figure 13. There, for four different values of T_t we present results for the ratio c_0/c_0^{opt} of the coefficients c_0 of the asymptotic behaviour $\epsilon(\alpha) \rightarrow c_0/\alpha$. Here c_0^{opt} refers to the optimal choice of T_s for given T_t . The results have been plotted as a function of T_s/T_t . Again we find that the optimal choice is $T_s/T_t \approx 0.6$, but with a 'flat behaviour'. Note that these results do not depend on the architecture, in contrast to corresponding results for input noise.

4.4.2. Small training sets: $q \rightarrow 0$. The crucial question in this limit is, in which way the above-mentioned phase-transitions shift when a noise strength $T_s \neq T_t$ is used, in particular whether there is an optimal $T_s^{\text{opt}} \neq T_t$ for which the 'Aha-effect' happens earlier, i.e. for smaller α .

For the *parity machine* this question can be immediately answered. There the choice $T_s = T_t$ means that the already-mentioned 'central-point network' with overlap r/\sqrt{q} reaches the Bayes generalization ability, which is optimal, i.e. $\epsilon_{\text{Bayes}}(q) = \epsilon_{\text{Gibbs}}(\sqrt{q})$. Therefore, for the parity machine the 'Aha-effect' phase transition cannot occur earlier than calculated for $T_s = T_t$.




Figure 13. For a general Boolean function B and the case of output noise, the ratio of the prefactors c_0/c_0^{opt} for the asymptotic behaviour $\epsilon(\alpha \rightarrow \infty) \rightarrow c_0/\alpha$ is plotted against T_s/T_t for the values of $T_t = 0.2, 0.5, 1.0$ and 2.0 . Note that the optimal value $T_s \approx 0.6T_t$ is more or less universal, and the behaviour in the vicinity of this point is rather flat.

For other networks, a slight addition to the argument is in order, since now the generalization ability of the ‘central-point network’ is smaller than that one obtained with the Bayes prescription. But for $T_s = T_t$ the typical student is also a typical teacher, i.e. we have an ensemble of student coupling vectors, which corresponds to the *a posteriori* probability that a certain coupling vector is that of the teacher machine. The Bayesian generalization error can be calculated from the expectation value of q for this ensemble by means of (26). This implies that the Bayes generalization ability becomes trivial (i.e. the error probability is $\frac{1}{2}$) below the critical α calculated for $T_s = T_t$ (since the Bayes generalization is optimal, a variation of T_s cannot lead to further improvement). So a transition to non-zero q cannot occur earlier for whatever T_s one chooses, and thus for all networks considered the phase transition cannot appear for smaller α .

5. Why is training with noise useful?

That training with noise can be useful, is already ‘folklore’, see e.g. [12, 13], and in the present context it is also known that in this way one can avoid overfitting (see [14, 2]). Here we want to go somewhat more into the details and look into the phase space structure.

5.1. Survey of the phase space of a small system

Let us first define the so called ‘genuine error’ E_0 of the student. This quantity counts the number of patterns from the training set where the (deterministic) answer of the student disagrees *with the original rule* and not with the partially corrupted answer of the teacher network.

Now we perform a simple survey of the phase space of a small perceptron with $K = 1$ and $N = 10$, with normalized randomly chosen coupling vector w^f with N components, and with a given set of $p = 20$ questions ξ^μ , i.e. we have $\alpha = p/N = 2$. By flipping




Figure 14. Comparison of the overlap $r(\alpha = 2)$ of our theory for $K = 1$ as a function of T_s for $K = 1$ and $T_t = 0.91$ (broken curve, calculated as in figure 7) with results from a simulation of the states of a small system as described in the text (full curve). In both cases there is an optimal output noise strength T_s . For more details see the text.

five of the 20 answers given by w^t , we have our teacher perceptron endowed with an output noise level of $p_f = \frac{5}{20} = 0.25$ and with a training set consisting of the 20 pairs of questions and the partially corrupted answers. Then a large number of *student vectors* w^s are drawn randomly from a uniform prior over all normalized real vectors with $N = 10$, where the random components are sampled from a Gaussian distribution with zero average and variance $1/N$. Finally, for each w^s , we evaluate: (i) the error E_{tr} with respect to the actual training set, i.e. the corrupted one; (ii) the error E_0 with respect to the *uncorrupted* answers on the training patterns, and (iii) the actual overlap r with w^t . The results are condensed into a two-parameter table: for each combination of the values of E_{tr} and E_0 the *number* of vectors (yielding these errors) as well as the corresponding averaged overlap r is stored.

This table can be used to calculate the $r(T_s)$ curve for the specific values of $\alpha = 2$ and $p_f = 0.25$. This is shown in figure 14 (full curve) and compared with the theoretical result with $\alpha = 2$ and $T_t = 0.91$ (corresponding to $p_f = 0.25$) which has been evaluated analogously to figure 7. We see that there is actually—despite the smallness of the simulated system—a nice qualitative similarity (more would be unexpected) in the behaviour of $r(\alpha)$ as a function of T_s ; in particular we find that it is useful to increase the noise until $T_s \approx 0.4$, whereas for larger T_s noise is more and more detrimental.

Looking more directly at the performance of this small system, we note the following numbers from a typical realization of the survey through its phase space. In this survey the combination ($E_0 = 5$, $E_{tr} = 0$) was realized 215 times. These 215 students made no error with respect to the actual (i.e. corrupted) training set, and thus they made $E_0 = 5$ errors with respect to the original rule. Possible values for $E_{tr} = 1$ are then: (i) $E_0 = 4$, i.e. one of the five errors of the student with respect to the original rule has been corrected, or (ii) $E_0 = 6$, i.e. one additional error has been made. In the simulation we found that (i) occurred in 1746 cases, whereas (ii) was less frequent, namely 1146 times, although there are many




Figure 15. For a teacher output noise level of $T_t = 1$ and $\alpha = 2$ the curves for r , ϵ_0 and ϵ_{tr} are shown for $K = 1$ as a function of the noise level T_s of the student machine. The increase in r is related to a decrease of ϵ_0 showing the error correcting behaviour of the non-perfect student.

more combinatorial possibilities for case (ii). The corresponding averaged overlaps are of course increased for the case $E_0 = 4$ and decreased for $E_0 = 6$. So obviously the system is able to correct errors previously made with respect to the original rule, and to increase the typical overlap r in this way.

5.2. The error ϵ_0 with respect to the uncorrupt training set

This can as well be confirmed theoretically. Defining $\epsilon_0 = E_0/p$ as the fraction of errors of the student with respect to the original training set, one derives the following equation:

$$\epsilon_0 = \frac{2}{1 + e^{-\beta_t}} \int Dt \frac{H(\gamma_r t)H(-\gamma t)e^{-\beta_s} + H(-\gamma_r t)H(\gamma t)e^{-\beta_t}}{H(\gamma t) + e^{-\beta_s}H(-\gamma t)} \quad (73)$$

with γ and γ_r defined in (68). In figure 15, ϵ_0 is plotted as a function of T_s for $\alpha = 2$ and $T_t = 1$. So with increasing T_s , ϵ_0 decreases at first, which means that at first the system mainly corrects the mistakes contained in the corrupted training set with respect to the original rule. But after a minimum, ϵ_0 increases again for larger T_s and approaches the ‘training error’ ϵ_{tr} (i.e. with respect to the corrupted training set) for $T_s \rightarrow \infty$.

The reason for this ability to produce the ‘right’ errors can be explained in terms of a sort of energy–entropy competition.

Allowing a given student to make errors there are more possibilities to increase the number of errors than to reduce this number; nevertheless, there are many more students in the version space classifying the training set with a decreased error number than with an increased one.

5.3. Multifractal phase space analysis

A more thorough analysis of the phase space structure is possible with the recent ‘multifractal technique’ of Monasson and co-workers [15, 16]. Extending this analysis to problems with noise, the distribution of the phase space volumina V_τ weighted with the corresponding degree of membership corresponding to a set $\tau = \{\sigma^\mu\}_{\mu=1,\dots,p}$ of answers on the training questions ξ^μ is given by the quantity

$$g(m) = \frac{1}{mN} \left\langle \ln \sum_\tau (e^{-\beta_s E_{rr}(\tau)} V_\tau)^m \right\rangle. \quad (74)$$

Here m is the inverse of a formal temperature and controls, how strong V_τ is weighted. Formally $(-g(m))$ is a free energy derived from the partition function

$$Z = \sum_\tau e^{-mE_\tau} \quad (75)$$

with the energy $E_\tau := -\ln[e^{-\beta_s E_{rr}(\tau)} V_\tau]$. Then the quantities

$$-k(m) := -\frac{\partial(mg(m))}{\partial m} \quad \text{and} \quad c(k) := \frac{\partial g(m)}{\partial(1/m)} \quad (76)$$

correspond formally to internal energies and entropies with respect to $T_m := 1/m$, and measure, how the phase-space volume V_τ and the corresponding number of realizations scale in the thermodynamic limit, namely as e^{Nk} and $e^{Nc(k)}$, respectively. Here $k < 0$, but $c(k) > 0$. These quantities can be calculated by a formalism, which resembles a RSB1 calculation, see [15, 16, 8]. The dominating behaviour, which is already calculated in the usual RS calculation, $m = 1$, is noted by the asterisks in figure 16 and follows from the identity $d[c(k) + k]/dk = 0$. In this figure 16, the results are for the simple perceptron. The overlap $r(k)$ is presented as a function of $-k$, for $\alpha = 5$ and the teacher output noise temperature $T_t = 1$, for different values of T_s , namely $T_s = 1.0, 0.8, 0.6, 0.4, 0.3, 0.2, 0.16, 0.14, 0.12$, and 0.1 , from the left. In this way, a differentiated picture of changes in the phase-space distribution induced by changes of the noise temperature T_s can be given. So there are two effects: with increasing temperature, regions with high student–teacher overlap become more and more active. At the same time, the volume determining the typical (dominant) overlap, i.e. the position of the asterisk, moves towards the maximum of these overlap curves. After having reached the optimal temperature, the curve decreases again, and regions with high training error and small overlap begin to dominate the phase space.

5.4. Why noise is useful: An intuitive picture

Summarizing the calculation and results of this section one can give an intuitive picture for the observed behaviour: In figure 16 we plot a scenario corresponding to a slice through a lake with a flat bank on the left, but a steep shore on the right.

To get the basic point let us think of the simple perceptron in a regime where RS is preserved even for $T_s = 0$, thus the (corrupted) training set is learnable. The deepest point in the lake corresponds to the error-free solution(s) with respect to this (corrupted) training set.

An increase of the student temperature T_s corresponds to an increase of the water level starting from this deepest point. Of course in the case of figure 16 the direction of the flat bank is favoured compared with the steep shore, so the centre-of-mass shifts to the left for increasing water level.




Figure 16. Phase-space analysis according to the formalism of [15, 16] of a perceptron (i.e. $K = 1$) with $T_t = 1$ and $\alpha = 5$, i.e. the overlap $r(k)$ has been presented as a function of the typical volume measure $(-k)$ as explained in the text. The student output noise temperatures of the different curves correspond to $T_s = 1.0, 0.8, 0.6, 0.4, 0.3, 0.2, 0.16, 0.14, 0.12$ and 0.1 , from the left. The asterisk denotes the ‘typical result’ as obtained according to the usual RS calculation.




Figure 17. This sketch shall make plausible, why training with noise can lead to an enhancement of the overlap of the couplings of the student and the teacher machine’s couplings. For details see the text.

This can be transferred to our model as follows.

- The direction of the flat bank (on the left) corresponds to noise-induced flips in the answers which *correct wrong answers* in the corrupt training set. The corresponding phase space of vectors is large (1746 students with a genuine error level, i.e. with respect to the original rule, reduced from $E_0 = 5$ to $E_0 = 4$ in the example of section 5.1).

- On the other hand, the steep shore on the right means that the phase space corresponding to detrimental changes of answers, which are correct in the training set, should be significantly smaller (1146 students with an enhanced error level, $E_0 = 5 \rightarrow E_0 = 6$).

So, corresponding to the ‘water dynamics’ scenario of figure 17, if the noise level is enhanced, the additional accessible phase space is dominantly situated on the left. The corresponding flat bank corresponds to just those solutions which correct errors instead of adding new ones. Since the former are also those having (intuitively) an increased overlap

with the teacher, ‘training with noise’ can be useful.

But what happens if T_s is increased too much, corresponding to flooding the lake? First of all, we have to notice that the ‘optimal’ solution is placed somewhere in the neighbourhood of the deepest point towards the flat shore. So flooding means that many solutions far from the optimum are included, and the performance of a random choice (corresponding to the Gibbs algorithm) should decrease. Nevertheless, looking at the *centre of mass* (corresponding to the Bayes algorithm), this quantity should be less sensitive to the deluge, especially if the shapes of the shores become more similar in the remote areas of the lake. So the Bayes algorithm should be less affected by a too high student temperature. This lower sensitivity of the Bayes algorithm to the negative effect of a very high noise, in contrast to the case of Gibbsian learning, can be nicely observed in figure 7.

6. Conclusions

We have studied the influence of input or output noise on the existence of *phase transitions* in the generalization behaviour of two-layer neural networks with non-overlapping receptive fields. Generally we find for Gibbs learning as a function of the reduced size $\alpha := p/N$ of the training set that the ‘Aha-effect’ phase transitions, where the system performs a simple guess for $\alpha < \alpha_c$ and only generalizes if α exceeds a critical value, *persist* in the presence of noise. However, the critical parameters scale with the strength of the noise, see for example figures 3 and 4. In particular, the order-index n of the system, which characterizes the behaviour at the transition, is unchanged: For $n = 1$ (e.g. the committee machine) the system starts to generalize already with arbitrarily small α , while for $n = 2$ (e.g. the $K = 2$ -parity machine) there is a continuous ‘Aha-effect’ transition (i.e. a second-order phase transition), whereas for $n \geq 3$ (e.g. the parity machine with $K \geq 3$, which has $n = K$) a discontinuous ‘Aha-effect’ transition happens. Only the so-called ‘interim phase transitions’, which appear in some cases for $n = 1$, for example the K4_mcode2 machine, get lost by input noise, but not by output noise. (At the ‘interim transitions’ only the strength of the generalization ability is enhanced from an already finite value for $\alpha < \alpha_c$ to a larger value above α_c .) Therefore, the ‘Aha-effect’ transitions are—so to say—more generic than ‘interim transitions’.

Concentrating on the simple perceptron and output noise we also studied the problem of an *optimal choice* for the noise level T_s of the student machine, given that of the teacher. Looking at *Gibbsian learning* there is a critical $\tilde{\alpha}_c$ above which a student trained with a finite noise level has a better performance than a ‘perfect student’, i.e. a student who has learnt the (partially corrupted) training set without errors. For $\alpha \rightarrow \infty$ this optimal noise level approaches $\approx 0.6T_t$ (this is true for all network types considered here). The case is slightly different for *Bayesian learning*. Here the choice $T_s = T_t$ is optimal for all α . Nevertheless the learning curve is quite flat around this optimal choice, so small deviations from the optimal noise have no large impact.

In view of RSB for the problem considered we found that for given T_t RS is typically conserved down to the optimal values of T_s and somewhat below. But if α is larger than a critical noise-dependent value, RSB occurs for low noise temperatures. Due to this fact, the suboptimal overlap r calculated in a RS theory for this region of small T_s is corrected to a higher value, which in a RSB1 calculation almost reaches the optimal number. This means that the ‘gain’ achieved by training with noise is less pronounced than predicted in for example the RS theory of [11]. RSB of higher order may lead to additional (slight) corrections in the region, where RS is broken. Nevertheless the effect of an optimal finite student noise rate T_s in some cases was shown to be not artificial since the values of α

resp. T_s , below resp. above which replica symmetry remains preserved, and also the optimal value of T_s in this RS region, as calculated in this paper, will remain unchanged.

Finally we took a quick look on mechanisms allowing improved learning for imperfect students. We showed that error correction is possible, due to a sort of energy–entropy competition, leading to an increased overlap compared with the minimal-error solutions.

Acknowledgments

The authors would like to thank G Dirscherl, F Gerl and M Probst for stimulating discussions. We would also like to thank one of the referees, who was so kind to send us a recent preprint of [17], in which the problem of the optimal noise level of the student is treated for single-layer perceptrons with different types of noise in an RS and RSB1 approach. Uezu's results strongly overlap with parts of our section 4.3.

References

- [1] Schottky B 1995 Phase transitions in the generalization behaviour of multilayer neural networks *J. Phys. A: Math. Gen.* **21** 4515–31
- [2] Watkin T H L, Rau A and Biehl M 1993 The statistical mechanics of learning a rule *Rev. Mod. Phys.* **65** 499–556
- [3] Oppen M and Kinzel W 1995 Statistical mechanics of generalization *Physics of Neural Networks* vol III, ed J L van Hemmen, E Domany and K Schulten (Berlin: Springer)
- [4] Hansel D, Mato G and Meunier C 1992 Memorization without generalization in a multilayered neural network *Europhys. Lett.* **20** 471–6
- [5] Schwarze H 1993 Learning a rule in a multilayer neural network *J. Phys. A: Math. Gen.* **26** 5781–94
- [6] Oppen M 1994 Learning and generalization in a two-layer neural network: The role of the Vapnik–Chervonenkis dimension *Phys. Rev. Lett.* **72** 2113–21
- [7] Gardner E 1987 Maximum storage capacity in neural networks *Europhys. Lett.* **4** 481–5
- [8] Schottky B 1996 Generalization behaviour of multilayer neural networks *PhD Thesis* University of Regensburg (in German)
- [9] Mezard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
- [10] Anlauf J K and Biehl M 1989 The AdaTron: An adaptive perceptron algorithm *Europhys. Lett.* **10** 687–92
- [11] Györgyi G 1990 Inference of a rule by a neural network with thermal noise *Phys. Rev. Lett.* **64** 1957–60
- [12] Pöppel G and Krey U 1987 A dynamical learning process for the recognition of correlated patterns in symmetric spin glass models *Europhys. Lett.* **4** 979–85
- [13] Gardner E, Stroud N and Wallace D J 1989 Training with noise *J. Phys. A: Math. Gen.* **22** 2019–30
- [14] Oppen M, Kinzel W, Kleinz J and Nehl R 1990 On the ability of the optimal perceptron to generalize *J. Phys. A: Math. Gen.* **23** 581–6
- [15] Monasson R and O’Kane D 1994 Domains of solutions and replica symmetry breaking in multilayer neural networks *Europhys. Lett.* **27** 85–90
- [16] Monasson R and Zecchina R 1995 Weight space structure and internal representations: A direct approach to learning and generalization in multilayer neural networks *Phys. Rev. Lett.* **75** 2432–5
- [17] Uezu T 1997 Learning from stochastic rules under finite temperature—optimal temperature and asymptotic learning curve *J. Phys. A: Math. Gen.* **30** L777–84