

Dynamical and stationary properties of on-line learning from finite training sets

Peixun Luo* and K. Y. Michael Wong†

Department of Physics, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China

(Received 29 August 2002; published 21 January 2003)

The dynamical and stationary properties of on-line learning from finite training sets are analyzed by using the cavity method. For large input dimensions, we derive equations for the macroscopic parameters, namely, the student-teacher correlation, the student-student autocorrelation and the learning force fluctuation. This enables us to provide analytical solutions to Adaline learning as a benchmark. Theoretical predictions of training errors in transient and stationary states are obtained by a Monte Carlo sampling procedure. Generalization and training errors are found to agree with simulations. The physical origin of the critical learning rate is presented. Comparison with batch learning is discussed throughout the paper.

DOI: 10.1103/PhysRevE.67.011906

PACS number(s): 87.18.Sn, 75.10.Nr, 05.20.-y, 07.05.Mh

I. INTRODUCTION

In recent years, there have been many attempts to analyze the dynamics of learning from examples in classification and regression [1]. From a general perspective, this dynamics is typical of many other nonequilibrium complex many-body systems, in which individual components collectively contribute to the achievement of global objectives. Interesting examples can be drawn from the modeling of social behavior [2], economics and ecology [3], patterns of routing in road traffic [4], and telecommunications networks [5].

The dynamics of learning from examples in classification and regression refers to the dynamical process of minimizing the risk functions of the classifier or regressor, often via gradient descent, until a steady state is reached. Despite the progress in understanding the *steady-state* behavior of learning processes [6], the *dynamics* of learning was much less understood. This is probably due to the high complexity in its analysis, since it typically involves the evolution of many microscopic parameters, each strongly interacting with others in a convolutional way. Yet, a number of important issues in improving the learning efficiency depend on a better understanding of its dynamics, including the speed of convergence, the early stopping point for optimal generalization, the shortening of the plateau regime, and the avoidance of getting trapped in local minima [7–9]. Hence, it would be both useful and challenging to analyze the dynamics of learning.

On-line learning is a common mode of implementing learning, in which an independent example is presented at each learning step. Significant progress has been made in the case of on-line learning of *infinite* training sets [7,10,11]. Since statistical correlations among the examples can be ignored, the dynamics can be described by instantaneous dynamical variables, leading to great advances in our understanding of on-line learning. However, in reality, the same restricted set of examples is recycled during the learning process. This introduces temporal correlations of the weights in the learning history, rendering the analysis at best as an approximation to the reality.

There were some attempts to understand on-line learning with *recycled* examples. Early researchers used the approximate Fokker-Planck equation to describe the learning process [12,13]. The use of perturbative expansions of the master equation was shown to be insufficient for a precise calculation of global properties of on-line learning [14]. The difference between batch learning and on-line learning was investigated to the first order of the learning rate [15]. For general learning rates, the exact solution for the Hebbian rule was derived in Ref. [16]. Exact solutions were found for linear networks, and the generalization ability of on-line learning was found to outperform batch learning if a bias is present in the input [17]. The dynamics of on-line learning in multilayer neural networks was analyzed by using the dynamical replica method and solutions were found in the limit of large sizes of training sets [18].

A recent work based on the generating functional approach is a good step forward towards a general theory of describing the dynamical and stationary properties of on-line learning [19]. It illustrates the mean-field character of the dynamics in its description in terms of an effective single example. For random choices of the sequence of presented examples, the dynamics is characterized by the appearance of an example as a Poisson event in the learning sequence. Steady-state properties were discussed by neglecting fluctuations in the learning forces (referred to as the mean-force approximation hereafter).

In this paper, we propose an analysis of on-line learning with recycled examples using the cavity method. The cavity method is a mean-field analysis first used in magnetic systems [20]. It enables us to understand the properties of a system by focusing on the response of the system to a single element added to it. It was later generalized to study learning in neural networks with the advantages of a clear physical picture and microscopic insights to both their equilibrium and dynamical properties [8,9,21]. The cavity method was subsequently applied to analyze the dynamics of *batch* learning, in which the entire set of examples is provided for each learning step [8]. It provides dynamical equations and obtains important results on the overtraining, early stopping, noise effects, and average learning strategy.

To adapt the cavity method from batch learning to on-line learning in this paper, there is a need to account for the

*Electronic address: physlpx@ust.hk

†Electronic address: phkywong@ust.hk

following subtleties. (a) Averaging over the choice of sequencing the examples is now necessary. (b) The measurements of an example observed at an instant is now correlated with the instants when it was learned. This is due to the giant boost of that example at a learning step, which upsets the uniformity of the examples as in the case of batch learning.

We remark that the subtleties of sequence-dependent dynamics are common in many other complex dynamical systems. In the modeling of economics, for instance, different sequences of the market state appearing in the minority game may lead to different evolutions of the system [22]. System responses to a market state should also be correlated with the instants when the state appeared in the market history. Similar correlations should also exist in other ecological and traffic models. Hence, we hope that the present study of on-line learning can provide insights to other dynamical problems.

The purposes of this paper are: (a) to perform an exact analysis of the learning dynamics as far as the formulation allows, so that minimal approximations are made, and deeper physical insights can be extracted; (b) to illustrate the analytical approach using the simple example of a linear learning rule, which can act as a benchmark for verifying the validity of the theory, and a theoretical framework for more complicated systems, such as nonlinear learning rules and multilayer networks; (c) to explore efficient Monte Carlo procedures implied by the distribution of example activations predicted by the theory, which can be applied to the more complicated cases; (d) to study the difference between on-line learning and batch learning for general learning conditions.

The paper is organized as follows. In Sec. II, we describe the dynamics of on-line learning. In Sec. III, we introduce the cavity method and derive the dynamical equations for the macroscopic measurements: (a) $G(t,s)$ and $D(t,s)$, the Green's functions of weights and examples in response to stimuli; (b) $R(t)$, the correlation between the teacher and student weight vectors, and $C(t,s)$, the autocorrelation between student weight vectors at different times; (c) the fluctuation of the learning force $\langle F^2(t) \rangle$. The Monte Carlo procedure to calculate the training error is also presented. In Sec. IV, we compare theoretical predictions with simulation results. The average learning strategy in the long time limit is proposed and compared with the performance of batch learning. In Sec. V, we summarize our work and propose some future directions. In the Appendixes, we describe the mathematical details of the procedure of sequence averaging.

II. FORMULATION

We consider a training set of p examples generated by a teacher network with N weights B_j , $j=1, \dots, N$. For definiteness, we set $|\mathbf{B}|=1$. Each example $\mu=1, \dots, p$ consists of an N dimensional input vector ξ^μ , and a teacher generated output \tilde{y}_μ . It is convenient to introduce the parameter $\alpha \equiv p/N$. The inputs ξ_j^μ are Gaussian variables with zero mean and unit variance. The outputs are

$$\tilde{y}_\mu = \begin{cases} \text{sgn}(y_\mu + \varepsilon z_\mu) & \text{for classification} \\ y_\mu + \varepsilon z_\mu & \text{for regression,} \end{cases}$$

where $y_\mu \equiv \mathbf{B} \cdot \xi^\mu$ is the teacher activation, z_μ is a Gaussian variable with zero mean and unit variance, and ε is the noise amplitude.

The examples are learned by a student network with the same number of inputs and output. At each learning step t , one example in the training set is randomly drawn, and the sequence of example labels is denoted as $\sigma(t)$. If the example drawn out at time t is $\sigma(t)$ then the weights are modified according to

$$J_j \left(t + \frac{1}{N} \right) = J_j(t) + \frac{v}{N} \{ F(x_{\sigma(t)}(t), \tilde{y}_{\sigma(t)}) \xi_j^{\sigma(t)} - \lambda J_j(t) \} + \frac{1}{N} \eta_j(t), \quad (1)$$

where $x_{\sigma(t)}(t) \equiv \mathbf{J}(t) \cdot \xi^{\sigma(t)}$ is the student activation. v is the learning rate and λ is the weight decay (the factor v is excluded to facilitate comparisons with the steady-state results). The force $F(x, \tilde{y})$ describes the learning rule,

$$F(x, \tilde{y}) = \begin{cases} \tilde{y} & \text{for Hebbian rule} \\ \tilde{y} - x & \text{for Adaline rule} \\ \tilde{y}(\kappa - x\tilde{y})\Theta(\kappa - x\tilde{y}) & \text{for Adatron rule,} \end{cases} \quad (2)$$

where Θ is the step function and κ is a parameter often referred to as the stability. The last term in Eq. (1) is the dynamical noise term, often added to avoid the learning procedure being trapped in local minimum, with $\langle \eta_j(t) \eta_k(s) \rangle = 2T \delta_{ts} / N$ and T is the dynamical noise level.

In the limit of vanishing learning rate v , the on-line dynamics described by Eq. (1) is equivalent to the batch learning formulation in Ref. [8] when the time scale, weight decay and the dynamical noise in the latter are multiplied by factors of α/v , $1/\alpha$, and v^2/α^2 , respectively. However, for finite learning rate v , the randomness of the learning sequence adds noise to the dynamics.

III. THE CAVITY METHOD

A. The cavity activation and the Green's functions

Consider a new example 0 that is not included in the original training set. We define its activation at time t in the network trained without that example as its *cavity activation* $h_0(t)$, i. e., $h_0(t) \equiv \mathbf{J}(t) \cdot \xi^0$. It is a random variable since the network has not learned the information of this new example. When the size of the network N is very large, $h_0(t)$ is a Gaussian variable with mean $R(t)y_0$ and covariance $C(t,s) - R(t)R(s)$, where $R(t) \equiv \mathbf{B} \cdot \mathbf{J}(t)$ is the student-teacher correlation at time t , and $C(t,s) \equiv \mathbf{J}(t) \cdot \mathbf{J}(s)$ is the student-student autocorrelation at times t and s . Both $R(t)$ and $C(t,s)$ are self-averaging in the limit of large N .

Now we consider the evolution of another network $J_j^0(t)$, in which the example 0 is added to its training set. To ensure that the probability of occurrence of the new example 0 and the old ones remain identical, the new example sequence $\sigma^0(t)$ is obtained from the original example sequence $\sigma(t)$ according to

$$\sigma^0(t) = \begin{cases} \sigma(t) & \text{probability} = 1 - \frac{1}{1+p} \\ 0 & \text{probability} = \frac{1}{p+1}. \end{cases} \quad (3)$$

In the new example sequence $\sigma^0(t)$, at each learning step, the weight is modified according to

$$J_j^0\left(t + \frac{1}{N}\right) = J_j^0(t) + \frac{v}{N} \{F(x_{\sigma^0(t)}^0, \tilde{y}_{\sigma^0(t)}) \xi_j^{\sigma^0(t)} - \lambda J_j^0(t)\} + \frac{1}{N} \eta_j(t), \quad (4)$$

where $x_{\sigma^0(t)}^0 \equiv J^0(t) \cdot \xi^{\sigma^0(t)}$. Comparing the networks $J_j^0(t)$ and $J_j(t)$, we obtain from Eqs. (1) and (4),

$$\left(\hat{S} - 1 + \frac{v\lambda}{N}\right) [J_j^0(t) - J_j(t)] = \frac{v}{N} \{F(x_{\sigma^0(t)}^0, \tilde{y}_{\sigma^0(t)}) \xi_j^{\sigma^0(t)} - F(x_{\sigma(t)}, \tilde{y}_{\sigma(t)}) \xi_j^{\sigma(t)}\}, \quad (5)$$

where \hat{S} is the time shift operator. Let $G^{(0)}(t-t')$ be the bare Green's function,

$$G^{(0)}(t-t') = \Theta\left(t^+ - t' - \frac{1}{N}\right) \left(1 - \frac{v\lambda}{N}\right)^{N(t-t' - \frac{1}{N})}. \quad (6)$$

It satisfies

$$\left(\hat{S} - 1 + \frac{v\lambda}{N}\right) G^{(0)}(t-t') = \delta_{tt'}. \quad (7)$$

We assume that the adjustment from $J_j(t)$ to $J_j^0(t)$ is small so that linear response theory is applicable. Then on separating the contributions from the new example and the old ones, we have

$$\begin{aligned} J_j^0(t) - J_j(t) &= \frac{v}{N} \sum_{k, t' < t} G^{(0)}(t-t') \delta_{\sigma^0(t')0} \\ &\quad \times [F_0(t') \xi_j^0 - F_{\sigma(t')}(t') \xi_j^{\sigma(t')}] \\ &\quad + \frac{v}{N} \sum_{k, t' < t} G^{(0)}(t-t') \delta_{\sigma^0(t')\sigma(t')} \xi_j^{\sigma(t')} \\ &\quad \times F'_{\sigma(t')}(t') \xi_k^{\sigma(t')} [J_k^0(t') - J_k(t')], \end{aligned} \quad (8)$$

where $F_{\sigma(t)}(t)$ is the shorthand notation of the force due to example $\sigma(t)$ at time t , and F' represents the derivative of the force with respect to the activation x . We can now interpret this result from the viewpoint of the linear response theory. The first term on the right hand side describes the primary effects of adding example 0 to the training set and is the driving term for the difference between the two networks. This occurs at the discrete instants with $\sigma^0(t')=0$ by adding the force due to example 0 and removing that due to the original example $\sigma(t')$. The second term describes the many-body reactions due to the change of the original ex-

amples caused by the added example, and is referred to as the Onsager reaction term. Describing the response to the driving term by the Green's function, Eq. (8) reduces to

$$J_j^0(t) - J_j(t) = \frac{v}{N} \sum_{k, t' < t} G_{jk}(t, t') \delta_{\sigma^0(t')0} \times [F_0(t') \xi_k^0 - F_{\sigma(t')}(t') \xi_k^{\sigma(t')}], \quad (9)$$

where $G_{jk}(t, s)$ is the time-dependent Green's function with the iterative expression

$$G_{jk}(t, s) = \delta_{jk} G^{(0)}(t-s) + \frac{v}{N} \sum_{t' < t} \sum_l G^{(0)}(t-t') \delta_{\sigma^0(t')\sigma(t')} \times \xi_j^{\sigma(t')} F'_{\sigma(t')}(t') \xi_l^{\sigma(t')} G_{lk}(t', s). \quad (10)$$

The Green's function $G_{jk}(t, s)$ is the response of the weight J_j at time t due to a unit stimulus introduced at time s to the right hand side of Eq. (1) corresponding to weight J_k , in the limit of vanishing magnitude of the stimulus.

In the limit of large N , we can apply a diagrammatic analysis similar to the case of batch learning [8]. In contrast with batch learning, we need to first average Eqs. (9) and (10) over the distribution of example sequence using Eq. (3). This can then be followed by the usual averaging over the distribution of background examples, as in the case of batch learning. The result is that we can neglect the effect of removing the background example represented by the second term in the square bracket of the right hand side of Eq. (9). $G_{jk}(t, s)$ is self-averaging and diagonal in the large N limit, so that $G_{jk}(t, s) = G(t, s) \delta_{jk}$, where $G(t, s)$ satisfies the Dyson's equations,

$$G(t, s) = G^{(0)}(t-s) + v \int dt_1 \int dt_2 G^{(0)}(t-t_1) \times \langle D_{\sigma(t_2)}(t_1, t_2) F'_{\sigma(t_2)}(t_2) \rangle G(t_2, s), \quad (11)$$

$$D_{\sigma(s)}(t, s) = \delta(t-s) + \frac{v}{\alpha} \int dt_1 G(t, t_1) F'_{\sigma(s)}(t_1) D_{\sigma(s)}(t_1, s), \quad (12)$$

where $G^{(0)}(t-s) \equiv \Theta(t-s) \exp[-v\lambda(t-s)]$ is the bare Green's function. $D_{\sigma(s)}(t, s)$ is the *example Green's Function*, and $\langle \cdot \rangle$ represents average over distributions of both example sequences and examples.

We emphasize that the average $\langle D_{\sigma(t_2)}(t_1, t_2) F'_{\sigma(t_2)}(t_2) \rangle$ is different from the average $\langle D_{\mu}(t_1, t_2) F'_{\mu}(t_2) \rangle$. The former specifies that the function $F'(t_2)$ and $D(t_1, t_2)$ are due to the example that was picked from the example sequence for learning at the particular instant t_2 . During on-line learning, the activation of this example receives a giant boost at the learning instant, as mentioned later in the text discussion of Fig. 1. This makes its distribution different from that of a randomly drawn example μ , whose previous learning instant remains unspecified. Hence, the former average will be referred to as an *active* average, in contrast to the latter, which is referred to as a *passive* average.

Nevertheless, in the case of linear rules used for illustration later in this paper, $F'(t)$ is a constant independent of t .




FIG. 1. The evolution of the activations of a randomly selected example in a network with $N=1000, \alpha=0.3, \kappa=0.8, v=0.1$, and $\lambda=0.1$ using the Adatron rule. See text discussions for the explanations of the lines and symbols.

Hence, the active average in Eq. (11) becomes identical to the passive average. Thus, the Dyson's equations (11) and (12) becomes identical to those of batch learning [8], after rescaling the time and the weight decay.

In the case of Hebbian rule, $F'(x)=0$ and $D_\mu(t,s)=\delta(t-s)$. The Green's function becomes identical to the bare Green's function.

In the case of the Adaline rule, $F'(x)=-1$ and $D_\mu(t,s)=D(t,s)$ independent of example μ . The weight Green's function becomes invariant under translation of time, and can be written as

$$G(t,s)=G(t-s,0)=\int \rho(x)e^{-x(t-s)}dx, \quad (13)$$

where $\rho(x)$ is the density of state

$$\rho(x)=(1-\alpha)\Theta(1-\alpha)\delta(x-v\lambda)+\frac{\sqrt{(x_{\max}-x)(x-x_{\min})}}{2\pi\frac{v}{\alpha}(x-v\lambda)}, \quad (14)$$

with x_{\max} and x_{\min} as the edges of the spectrum given by $x_{\max}, x_{\min}=v\lambda+v(1\pm 1/\sqrt{\alpha})^2$, respectively.

The number of times m that the new example 0 appears in time t follows a Poisson distribution with mean t/α . If these appearances occur at times t_1, \dots, t_m ($t_m < \dots < t_1 < t$), Eq. (9) reduces to

$$J_j^0(t)=J_j(t)+\frac{v}{N}\sum_{r=1}^m G(t,t_r)F_0(t_r)\xi_j^0. \quad (15)$$

Multiplying both sides by ξ_j^0 and summing over j , one derives the relationship between the cavity activation and the generic activation of example 0,

$$x_0(t)=h_0(t)+v\sum_{r=1}^m G(t,t_r)F_0(t_r). \quad (16)$$

This relation enables us to express the cavity activation $h(t)$ of any example as a function of its generic activation $x(t_1), \dots, x(t_m), x(t)$ at the previous and current learning instants, and attributes physical meaning to the single effective example in Ref. [19]. Hereafter, we omit the example label if no confusion occurs.

The simulation results in Fig. 1 verify the relationship between the cavity activation and the generic activation for a randomly selected example. Up to $t=3$, the example is drawn from the learning sequence $\sigma(t)$ nine times, close to the Poisson average of $t/\alpha=10$. The solid line describes the evolution of $x(t)$, which exhibits giant boosts at the nine learning instants indicated by the vertical dashed lines. The dotted line describes the evolution of the cavity activation $h(t)$, which is obtained in a second network that uses the same learning sequence $\sigma(t)$, except that learning is paused when the example is drawn. Since the example and this network are uncorrelated, $h(t)$ evolves as a random walker with appropriate means and covariances. The filled circles indicate the values of the cavity activations predicted by Eq. (16), using the Green's functions measured by comparing learning with and without stimuli [23]. They show remarkable agreement with the simulated $h(t)$.

To derive the distribution of generic activations, we first consider the distribution of cavity activations, which is given in the Gaussian form at m learning steps and time $t_0(=t)$ by

$$P(h(t_0), \dots, h(t_m)|y)=\frac{\exp\left\{-\frac{1}{2}\sum_{i,j=0}^m [h(t_i)-R(t_i)y](C-RR^T)_{ij}^{-1}[h(t_j)-R(t_j)y]\right\}}{\sqrt{(2\pi)^{m+1}\det(C-RR^T)}}, \quad (17)$$

where $C-RR^T$ is a square matrix with size $m+1$ and $(C-RR^T)_{ij}=C(t_i,t_j)-R(t_i)R(t_j)$. The corresponding distribution of generic activations can be written as

$$P(x(t_0), \dots, x(t_m)|y, \tilde{y})=P(h(t_0), \dots, h(t_m)|y)\left|\frac{\partial[h(t_0), \dots, h(t_m)]}{\partial[x(t_0), \dots, x(t_m)]}\right|, \quad (18)$$

where $h(t_i)$ is a function of $x(t_i), \dots, x(t_m)$ defined by Eq. (16), and the dependence on \tilde{y} may arise from the learning forces. Since $\partial h(t_i)/\partial x(t_j)=0$ for $t_j > t_i$, and $\partial h(t_i)/\partial x(t_i)=1$, the Jacobian reduces to 1. Therefore, the distribution of generic activations can be expressed as

$$P(x(t_0), \dots, x(t_m)|y, \tilde{y}) = \frac{\exp\left\{-\frac{1}{2} \sum_{i,j=0}^m [h(t_i) - R(t_i)y](C - RR^T)^{-1}_{ij} [h(t_j) - R(t_j)y]\right\}}{\sqrt{(2\pi)^{m+1} \det(C - RR^T)}}, \quad (19)$$

where $h(t_i)$ are given by Eq. (16). In general, $h(t_i)$ can be a nonlinear function of $x(t_i), \dots, x(t_m)$. Hence, the generic activation distribution in Eq. (19) is no longer Gaussian, although the cavity activation distribution in Eq. (17) is. This characteristic of on-line learning is demonstrated in numerical simulations for the Adatron rule in Ref. [27].

We now illustrate how the above result can be applied to specific cases. For the Hebbian rule, Eq. (16) implies that $h(t)$ is not an explicit function of $x(t_r)$ at the previous learning instants t_r ,

$$h(t) = x(t) - \tilde{y} \sum_{r=1}^m \exp[-v\lambda(t-t_r)]. \quad (20)$$

This enables us to write down the instantaneous activation distribution, given the learning instants t_1, \dots, t_m of the example,

$$P(x, t_0|y, \tilde{y}; t_1, \dots, t_m) = \frac{\exp\left\{-\frac{1}{2} [C(t, t) - R^2(t)]^{-1} \left[x(t_0) - \tilde{y} \sum_{r=1}^m e^{-v\lambda(t-t_r)} - R(t)y\right]^2\right\}}{\sqrt{2\pi [C(t, t) - R^2(t)]}}. \quad (21)$$

The distribution is then averaged over the time distribution and the Poisson distribution of learning instants

$$P(x, t|y, \tilde{y}) = \langle P(x, t|y, \tilde{y}; t_1, \dots, t_m) \rangle_\sigma, \quad (22)$$

where $\langle \cdot \rangle_\sigma$ represents averaging over the distribution of sequences. The sequence average of an instantaneous quantity ψ at time t depending on the previous learning instants t_1, \dots, t_m is

$$\langle \psi(t|t_1, \dots, t_m) \rangle_\sigma = \sum_{m=0}^{\infty} \frac{e^{-t/\alpha}}{\alpha^m} \int_0^t dt_1 \dots \times \int_0^{t_{m-1}} dt_m \psi(t|t_1, \dots, t_m), \quad (23)$$

where the factor of $m!$ in the Poisson distribution is canceled by the number of permutations in ordering t_1, \dots, t_m . Using the Hubbard-Stratonovich identity, we can factorize the integrals over $t_r (1 \leq r \leq m)$. We arrive at the result

$$P(x, t|y, \tilde{y}) = \int \frac{d\hat{x}}{2\pi} \exp\left\{i\hat{x}[x - R(t)y] - \frac{1}{2} [C(t, t) - R^2(t)] \hat{x}^2 + \frac{1}{\alpha} \int_0^t ds [\exp(-i\hat{x}\tilde{y}e^{-v\lambda(t-s)} - 1)]\right\}, \quad (24)$$

which agrees with the rule-specific derivation in Ref. [16].

For the Adaline rule, substituting $F(x, \tilde{y}) = \tilde{y} - x$ into Eq. (16) yields a linear relation between the student activation and the cavity ones,

$$x(t_0) = \sum_{r=0}^m (1 + vG)_{0r}^{-1} \left[h(t_r) + v \sum_{s=r+1}^m G_{rs} \tilde{y} \right], \quad (25)$$

where G is a square matrix with size $m+1$ and $G_{rs} = G(t_r - t_s, 0)$ for $t_r > t_s$ and $G_{rs} = 0$ for $t_r \leq t_s$. Inserting the mean and variance of the cavity activation, we see that $x(t_0)$ is a Gaussian variable with mean and variance

$$\langle x(t_0) \rangle = \sum_{r=0}^m (1 + vG)_{0r}^{-1} \left[R(t_r)y + v \sum_{s=r+1}^m G_{rs} \tilde{y} \right],$$

$$\Delta^2(t_0) = \sum_{r,s=0}^m (1 + vG)_{0r}^{-1} (1 + vG)_{0s}^{-1} [C(t_r, t_s) - R(t_r)R(t_s)]. \quad (26)$$

To obtain the activation distribution in such an application as the average training error, we need to further average the Gaussian distribution in Eq. (26) over the learning sequences.

In general, for nonlinear learning rules, the linear inversion of Eq. (16) to obtain the student activation is not possible, and the activation distribution becomes non-Gaussian, even for a given sequence of learning instants. Nevertheless, a useful identity exists for the sequence average pertaining to an example, as derived in Appendix A,

$$\langle x(t) \rangle_\sigma = h(t) + \frac{v}{\alpha} \int_0^t dt' G(t, t') \langle F(t') \rangle_\sigma. \quad (27)$$

This equation of the sequence-averaged activation is the same as that of the self-consistent equation of the activation in batch learning, after rescaling the time and weight decays [8].

B. The student-teacher correlation

To analyze the student-teacher correlation, we multiply both sides of Eq. (1) by B_j and sum over j , yielding in the limit of large N ,

$$\left(\frac{d}{dt} + v\lambda\right)R(t) = v\langle\langle y_\mu F_\mu(t) \rangle\rangle_\sigma, \quad (28)$$

where $\langle\cdot\rangle_\mu$ represents averaging over the distribution of examples.

For the Adaline rule, the solution of $R(t)$ in Eq. (28) involves $\langle x_\mu(t) \rangle_\sigma$. By virtue of Eq. (27), and exploiting the example Green's function in Eq. (12), we obtain

$$\begin{aligned} \langle x_\mu(t) \rangle_\sigma &= \int_0^t dt_1 D(t-t_1) \\ &\times \left[h_\mu(t_1) + \frac{v}{\alpha} \int_0^{t_1} dt_2 G(t_1-t_2) \tilde{y}_\mu \right]. \end{aligned} \quad (29)$$

Applying Laplace transform to Eq. (29) and then Eq. (28),

$$\begin{aligned} \langle \tilde{x}_\mu(z) \rangle_\sigma &= \tilde{D}(z) \left[\tilde{h}_\mu(z) + \frac{v\tilde{y}_\mu}{\alpha z} \tilde{G}(z) \right], \\ (z+v\lambda)\tilde{R}(z) &= v \left\{ \frac{a}{z} - \tilde{D}(z) \left[\tilde{R}(z) + \frac{va}{\alpha z} \tilde{G}(z) \right] \right\}, \end{aligned} \quad (30)$$

where

$$a \equiv \langle y_\mu \tilde{y}_\mu \rangle_\mu = \begin{cases} \sqrt{\frac{2}{\pi(1+\varepsilon^2)}} & \text{for classification,} \\ \frac{1}{\sqrt{1+\varepsilon^2}} & \text{for regression.} \end{cases}$$

Here $\tilde{G}(z) \equiv \int_0^\infty dt e^{-zt} G(t,0)$ and $\tilde{D}(z) \equiv \int_0^\infty dt e^{-zt} D(t,0)$. Inverse Laplace transform yields

$$R(t) = a \int dx \rho(x) (x-v\lambda) \frac{1-e^{-xt}}{x}, \quad (31)$$

which is the same as that in batch learning after rescaling.

C. The student-student autocorrelation and the force fluctuation

To analyze the student-student autocorrelation, we multiply both sides of Eq. (1) by $J_j(s)$ and sum over j , thus obtaining in the limit of large N ,

$$\left(\frac{d}{dt} + v\lambda\right)C(t,s) = v\langle\langle x_\mu(s)F_\mu(t) \rangle\rangle_{\sigma(t)=\mu}. \quad (32)$$

We remark that $\langle x_\mu(s)F_\mu(t) \rangle_{\sigma(t)=\mu}$, is an active average, which is distinct from the *passive average* $\langle x_\mu(s)F_\mu(t) \rangle_\sigma$, where the example learned at time t is not necessarily μ .

For the Adaline learning rule, Appendix B shows that the average difference $W_\mu(s,t) \equiv \langle x_\mu(s)F_\mu(t) \rangle_{\sigma(t)=\mu} - \langle x_\mu(s)F_\mu(t) \rangle_\sigma$ can be expressed in terms of Green's functions and learning force,

$$W_\mu(s,t) = v \int_t^s dt' D(s,t') G(t',t) \langle F_\mu^2(t) \rangle_\sigma. \quad (33)$$

A similar equation for the passive average is also derived in Appendix B,

$$\begin{aligned} \langle x_\mu(s)F_\mu(t) \rangle_\sigma &= \int_0^s dt_1 D(s,t_1) \left[h_\mu(t_1) \langle F_\mu(t) \rangle_\sigma \right. \\ &\quad \left. + \frac{v}{\alpha} \int_0^{t_1} dt_2 G(t_1,t_2) \tilde{y}_\mu \langle F_\mu(t) \rangle_\sigma \right] \\ &\quad - \frac{v^2}{\alpha} \int_0^{\min(t,s)} dt_1 \int_{t_1}^s dt_2 \int_{t_1}^t dt_3 D(s,t_2) \\ &\quad \times G(t_2,t_1) D(t,t_3) G(t_3,t_1) \langle F_\mu^2(t_1) \rangle_\sigma. \end{aligned} \quad (34)$$

Therefore, one can perform the Laplace transforms

$$\tilde{C}(w,z) \equiv \int_0^\infty ds \int_0^\infty dt e^{-ws-zt} C(s,t),$$

$$\begin{aligned} &\langle\langle \tilde{F}_\mu(w) \tilde{F}_\mu(z) \rangle\rangle_\sigma \\ &\equiv \int_0^\infty ds \int_0^\infty dt e^{-ws-zt} \langle\langle F_\mu(s) F_\mu(t) \rangle\rangle_\sigma. \end{aligned}$$

After substituting Eqs. (33) and (34) into Eq. (32), and performing elaborate algebra, one obtains an equation of the weight correlation

$$\begin{aligned} \tilde{C}(w,z) &= v \frac{\tilde{D}(w) - \tilde{D}(z)}{w-z} \left\{ \frac{b}{wz} + \frac{\alpha a^2 [\tilde{D}(w) + \tilde{D}(z) - 1]}{wz} \right. \\ &\quad \left. + \alpha \langle\langle \tilde{F}_\mu^2(w+z) \rangle\rangle_\sigma + \frac{2\alpha T}{v^2(w+z)\tilde{D}(w)\tilde{D}(z)} \right\}, \end{aligned} \quad (35)$$

and an equation of the force autocorrelation

$$\begin{aligned} &\langle\langle \tilde{F}_\mu(w) \tilde{F}_\mu(z) \rangle\rangle_\sigma \\ &= \tilde{D}(w)\tilde{D}(z) \left\{ \frac{1}{wz} + \frac{\alpha a^2}{wz} [\tilde{D}(w) + \tilde{D}(z) - 2] \right. \\ &\quad \left. + \tilde{C}(w,z) + \frac{v^2}{\alpha} \tilde{G}(w)\tilde{G}(z) \langle\langle \tilde{F}_\mu^2(w+z) \rangle\rangle_\sigma \right\}. \end{aligned} \quad (36)$$

Here,

$$b \equiv \langle \tilde{y}_\mu^2 \rangle_\mu = \begin{cases} 1 & \text{for classification} \\ 1 + \varepsilon^2 & \text{for regression.} \end{cases}$$

We note the presence of the force fluctuation term $\langle \langle \tilde{F}_\mu^2(w+z) \rangle_\sigma \rangle_\mu$ in Eqs. (35) and (36). This term is absent in the corresponding equations in the case of batch learning. This can be seen by observing the scalings $z^{-1} \sim \tilde{G}(z) \sim \langle \langle \tilde{F}_\mu^2(w+z) \rangle_\sigma \rangle_\mu \sim v^{-1}$ in Eqs. (35) and (36), so that the coupling of weight correlation $\tilde{C}(w, z)$ and force autocorrelation $\langle \langle \tilde{F}_\mu(w) \tilde{F}_\mu(z) \rangle_\sigma \rangle_\mu$ via the force fluctuation $\langle \langle \tilde{F}_\mu^2(w+z) \rangle_\sigma \rangle_\mu$ will approach zero when v is vanishingly small, which indicates that this temporal correlation is unique to on-line learning. In contrast to batch learning, the presence of the force fluctuation term increases the weight correlation via Eq. (35), which in turn increases the force fluctuation itself via Eq. (36). As we shall see, this coupling leads to a collective relaxation mode.

Performing the inverse Laplace transform, one can obtain the autocorrelation

$$\begin{aligned} C(s, t) = & \int dx \rho(x) (x - v\lambda) \left[\frac{bv}{\alpha} + a^2 \left(x - v\lambda - \frac{v}{\alpha} \right) \right] \\ & \times \left(\frac{1 - e^{-xs}}{x} \right) \left(\frac{1 - e^{-xt}}{x} \right) \\ & + v \int dx \rho(x) (x - v\lambda) \int_0^{\min(t, s)} dt' e^{-x(t+s-2t')} \\ & \times \langle F^2(t') \rangle + T \int dx \rho(x) \frac{e^{-x|t-s|} - e^{-x(t+s)}}{x}, \end{aligned} \quad (37)$$

where $\langle F^2(t) \rangle \equiv \langle \langle F_\mu^2(t) \rangle_\sigma \rangle_\mu$. In Eq. (37), the first and third terms are similar to those in the autocorrelation function in batch learning. The second term represents the contribution induced by the force fluctuations arising from on-line learn-

ing. It vanishes when $v \rightarrow 0$. However, for finite learning rate v , one can see that this term plays a similar role as the dynamic noise. This can be seen by considering the asymptotic limit of the second term, where $\langle F^2(t') \rangle$ approaches the steady-state value of $\langle F^2 \rangle$, yielding

$$\frac{v}{2} \langle F^2 \rangle \int dx \rho(x) (x - v\lambda) \frac{e^{-x|t-s|} - e^{-x(t+s)}}{x}. \quad (38)$$

Comparing with the dynamical noise in the third term, we see that $\langle F^2 \rangle$ is a measure of an effective temperature, and the two noise contributions differ slightly in their spectrum of relaxation rates. Therefore, in practice, force fluctuations can also assist the learning dynamics to avoid being trapped in metastable states, playing the same role as dynamical noises to batch learning. Hereafter, we let $T=0$ in our final results.

To obtain the force autocorrelation $\langle \langle F_\mu(s) F_\mu(t) \rangle_\sigma \rangle_\mu$, we substitute Eq. (35) into Eq. (36) and perform the inverse Laplace transform, which yields,

$$\begin{aligned} \langle \langle F_\mu(t) F_\mu(s) \rangle_\sigma \rangle_\mu = & b + \frac{1}{v} \int dx \rho(x) \left[\frac{bv}{\alpha} + a^2 \left(x - v\lambda - \frac{v}{\alpha} \right) \right] \\ & \times \left[(x - v\lambda)^2 \left(\frac{1 - e^{-xt}}{x} \right) \times \left(\frac{1 - e^{-xs}}{x} \right) \right. \\ & - (x - v\lambda) \left(\frac{1 - e^{-xt}}{x} \right) - (x - v\lambda) \\ & \left. \times \left(\frac{1 - e^{-xs}}{x} \right) \right] + \int_0^{\min(t, s)} dt' \langle F^2(t') \rangle \\ & \times \int dx \rho(x) (x - v\lambda)^2 e^{-x(t+s-2t')}. \end{aligned} \quad (39)$$

Equating t and s , the force fluctuation is given by the inverse Laplace transform of

$$\langle \langle \tilde{F}_\mu^2(z) \rangle_\sigma \rangle_\mu = \frac{\frac{b}{z} + \frac{1}{v} \int dx \rho(x) \frac{x - v\lambda}{x^2} \left[\frac{bv}{\alpha} + a^2 \left(x - v\lambda - \frac{v}{\alpha} \right) \right] \left[\frac{x - v\lambda}{z + 2x} + \frac{2v\lambda}{z + x} - \frac{x + v\lambda}{z} \right]}{1 - \int dx \rho(x) \frac{(x - v\lambda)^2}{z + 2x}}. \quad (40)$$

The final expression consists of four contributions. First, the pole at $z=0$ gives the steady-state value of

$$\langle F^2 \rangle = \frac{b - \frac{1}{v} \int dx \rho(x) \left(1 - \frac{v^2 \lambda^2}{x^2} \right) \left[\frac{bv}{\alpha} + a^2 \left(x - v\lambda - \frac{v}{\alpha} \right) \right]}{1 - \int dx \rho(x) \frac{(x - v\lambda)^2}{2x}}. \quad (41)$$

The second and third contributions come from the relaxation spectrum of force fluctuations ranging through (x_{\min}, x_{\max}) and $(2x_{\min}, 2x_{\max})$, respectively. The fourth contribution is described by the existence of a collective relaxation mode arising from the force-weight coupling, which is a unique phenomenon of on-line learning. Its relaxation rate λ_σ is given by the pole

$$1 - \int dx \rho(x) \frac{(x - v\lambda)^2}{2x - \lambda_\sigma} = 0. \quad (42)$$

This is called slow mode in Ref. [17]. When λ_σ approaches zero, the steady-state force fluctuation and student weight will diverge. The critical learning rate at which the weight diverges is given by

$$v_c = \frac{4}{2 - \lambda[1 + \alpha + \alpha\lambda - \sqrt{(1 + \alpha + \alpha\lambda)^2 - 4\alpha}]}, \quad (43)$$

which is also derived through the spectral analysis [17].

D. The training and generalization errors

The performance of learning is measured by the training and generalization errors. Here, we provide their expressions for noiseless examples. Expressions for other cases can be derived similarly. In classification, the generalization error is defined as the probability that a *new* example presented to the network is misclassified, $E_g(t) = \langle \Theta[-(\mathbf{B} \cdot \boldsymbol{\xi})(\mathbf{J}(t) \cdot \boldsymbol{\xi})] \rangle_{\boldsymbol{\xi}}$. It is determined by the magnitude of the student vector $C(t, t)$ and its correlation with the teacher vector $R(t)$ [8], that is,

$$E_g(t) = \frac{1}{\pi} \cos^{-1} \frac{R(t)}{\sqrt{C(t, t)}}. \quad (44)$$

Analytical expressions of $R(t)$ and $C(t, t)$ are derived in preceding sections for the Adaline rule.

The training error is defined as the fraction of examples in the training set that are classified wrongly, i.e.,

$$A_{m-i, m-j} = \frac{C(t_{m-i}, t_{m-j}) - R(t_{m-i})R(t_{m-j}) - \sum_{k=1}^{j-1} A_{m-i, m-k} A_{m-j, m-k}}{A_{m-j, m-j}},$$

$$A_{m-i, m-i} = \left[C(t_{m-i}, t_{m-i}) - R^2(t_{m-i}) - \sum_{k=1}^{i-1} A_{m-i, m-k}^2 \right]^{1/2}. \quad (48)$$

(5) Compute $x(t_i)$ according to Eq. (16). This enables us to collect samples for the distribution $P(x(t_0)|y, \tilde{y})$, and hence estimate the training error.

(6) Steps (1) to (5) are repeated to yield sufficient amount of statistics.

For the Adaline learning, step (5) of the Monte Carlo sampling procedure can be further simplified by exploiting the Gaussian nature of the generic activation distribution. Using Eq. (26) to find the mean $\langle x(t_0) \rangle$ and variance $\sigma^2(t_0)$, the contribution of an example to the training error is given by $\text{erfc}[\langle x(t_0) \rangle \tilde{y} / \sqrt{2} \sigma(t_0)]$.

The above procedure assumes that the Green's function $G(t, s)$ and the correlations $R(t)$ and $C(t, s)$ are known *a*

$$E_t(t) = \int dy \int d\tilde{y} P(y, \tilde{y}) \sum_{m=0}^{\infty} \frac{e^{-t/\alpha}}{\alpha^m} \int_0^t dt_1 \dots \int_0^{t_m-1} dt_m \times \int dx(t) P(x(t)|y, \tilde{y}; t_1, \dots, t_m) \Theta[-\tilde{y}x(t)]. \quad (45)$$

This can be computed by a Monte Carlo sampling procedure, which has been shown to be free from finite size effects [24]. For general learning rules, we adopt the procedure with the following steps:

(1) For a given training example, generate the teacher activation y and \tilde{y} according to $P(y, \tilde{y})$.

(2) For time $t_0(=t)$, generate the number of times m the example appears in a training sequence from time 0 to t_0 according to a Poisson distribution with mean t_0/α .

(3) Generate the instants t_1, \dots, t_m that the example appears in the training sequence according to a uniform distribution between 0 and t_0 , with $0 < t_m < \dots < t_1 < t_0$.

(4) Generate the cavity activations $h(t_r)$, $r=0, \dots, m$ according to the Gaussian distribution with mean $R(t_r)y$ and covariance $C(t_r, t_s) - R(t_r)R(t_s)$. This can be carried out by generating the independent Gaussian variables z_{ik} ($k=0, \dots, m$) with mean 0 and variance 1, and transforming them to $h(t_i)$ via

$$h(t_i) = R(t)y + \sum_{k=i}^m A_{ik} z_k, \quad 0 \leq i \leq m \quad (46)$$

and the matrix elements A_{ik} are obtained from the recursion relations

$$A_{mm} = \sqrt{C(t_m, t_m) - R^2(t_m)} \quad (47)$$

and for $1 \leq j < i \leq m$,

priori. For general learning rules, $G(t, s)$ can be obtained by solving the Dyson's equations Eqs. (11) and (12). Since the equation involves an average over the distribution of learning sequence and examples, it can again be obtained by a Monte Carlo sampling procedure. Similarly, $R(t)$ and $C(t, s)$ can be obtained by solving Eqs. (28) and (32) by Monte Carlo sampling. These will be left for further studies. Here, for the exposition of the cavity method, we focus on the Adaline learning, where these functions can be obtained analytically as described in the preceding sections.

For the purpose of consistency check, one can also obtain these functions directly from simulations and plug into the Monte Carlo sampling procedure to check whether the gen-

erated activation distribution agrees with simulations.

The proposed Monte Carlo procedure is similar to the effective single pattern process in Ref. [19]. The difference lies in the generation of the learning instants of an example according to the Poisson distribution. In Ref. [19], an individual Poisson number with mean Δ/α is generated for every time increment Δ in the learning history. Here, the sampling efficiency is improved by a single Poisson number m with mean t_0/α and m learning instants with uniform distributions. For general learning rules, even if the Green's functions and correlation functions have to be generated from the Monte Carlo sampling procedure, it is possible to use similarly efficient samplings. This will be left for further studies.

While the Monte Carlo sampling procedure is useful in studying the transient behavior of learning, it can also be used to extract the stationary properties. At a very large observation time t , we look back at the learning history of an example by the network. Reasonably, only those learning events occur recently have detectable contributions to the distribution of the student activation $x(t)$ in Eq. (26). Therefore, we can calculate this distribution to any desired precision by adding the earlier learning events one by one, until certain stopping criteria are satisfied. Since the time intervals between successive learning events obey an exponential distribution, we have

$$P(x|y, \tilde{y}) = \lim_{m \rightarrow \infty} \lim_{t \rightarrow \infty} \prod_{r=1}^m \left[\int_0^\infty \frac{ds_r}{\alpha} e^{-(s_r/\alpha)} \right] \times P\left(x, t \left| y, \tilde{y}; t-s_1, \dots, t-\sum_{r=1}^m s_r \right.\right). \quad (49)$$

For the Adaline learning, the distribution $P(x, t|y, \tilde{y}; t-s_1, \dots, t-\sum_{r=1}^m s_r)$ is replaced by a Gaussian distribution with mean and variance given in Eqs. (26). We find that the contribution of earlier events approaches zero very quickly as m increases. Thus, we only need to invert small matrices in evaluating the training error at a steady state.

IV. RESULTS AND DISCUSSIONS

Figure 2(a) shows the transient behavior of the three macroscopic parameters, student-teacher correlation, student autocorrelation, and force fluctuation for typical learning parameters. The training error and generalization error are shown in Fig. 2(b). The theoretical predictions have an excellent agreement with simulations.

Figure 3 shows the generalization and training errors at the steady state. The theoretical predictions agree well with the simulation results. The learning dynamics diverges at the critical learning rate v_c . It is also observed that strong weight decays tend to restrain this divergence at large learning rates, pushing v_c to higher values. On the other hand, strong weight decays increase the generalization error when v is small.

In Fig. 3, we also present the results obtained by the mean-force approximation adopted in Ref. [19]. As we have shown above, in the steady state, the sequence noise due to

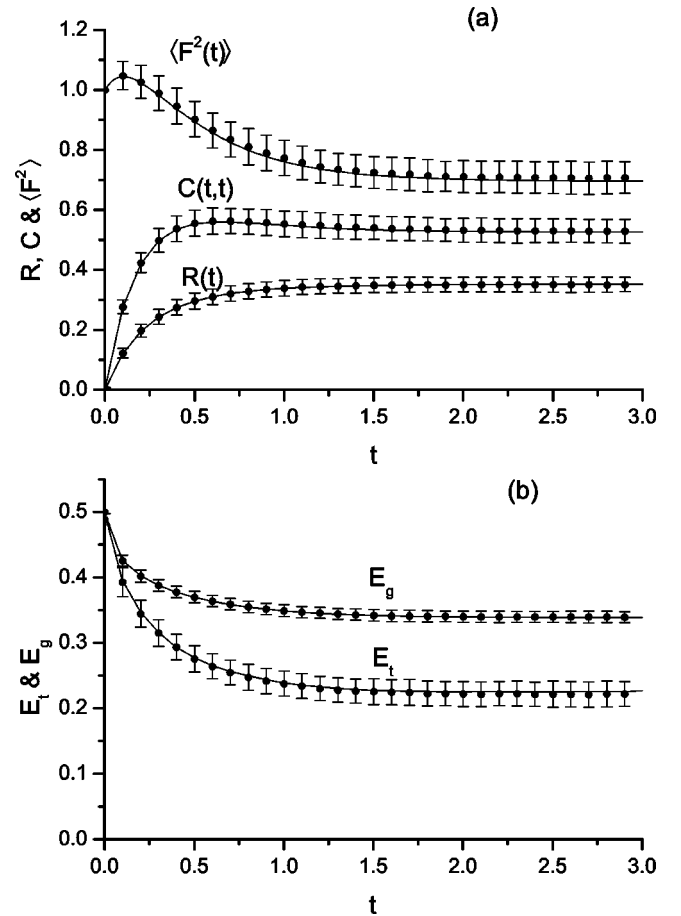


FIG. 2. The evolution of (a) teacher-student correlation $R(t)$, student autocorrelation $C(t,t)$, and force fluctuation $\langle F^2(t) \rangle$, (b) training error E_t and generalization error E_g for Adaline learning at $\alpha=1.2, v=1.9$, and $\lambda=0.8$. Solid lines: the cavity method, with analytical results for $R(t), C(t,t), \langle F^2(t) \rangle, E_g$ and Monte Carlo results for E_t , averaged over 500 000 samples. Symbols: simulations averaged over 100 samples with $N=500$.

the random drawing of examples at each learning step is equivalent to the external dynamical noise in batching learning. Thus, the dynamical variables will fluctuate around their temporal average even without other external noises. As a result, the mean-force approximation is only valid when the learning rate and the sequence noise is small. As shown in Fig. 3, it has an increasing discrepancy with simulations when v becomes large. The critical learning rate $v_c = 2(1 + \lambda)$ estimated by the mean-force approximation is larger than the simulation result. This discrepancy can be attributed to the omission of the force fluctuations therein.

An important question is whether on-line learning can perform as well as batch learning [25,26]. We have proposed an averaged strategy in the context of batch learning, predicting that dynamical noise can be averaged over to yield performances approaching noiseless learning [8]. Since sequence noise in on-line learning has similar effects as dynamical noise, we adopt the same strategy to improve performance of on-line learning.

In average learning, we first wait for the system to settle to the steady state, and then monitor the student weight vec-

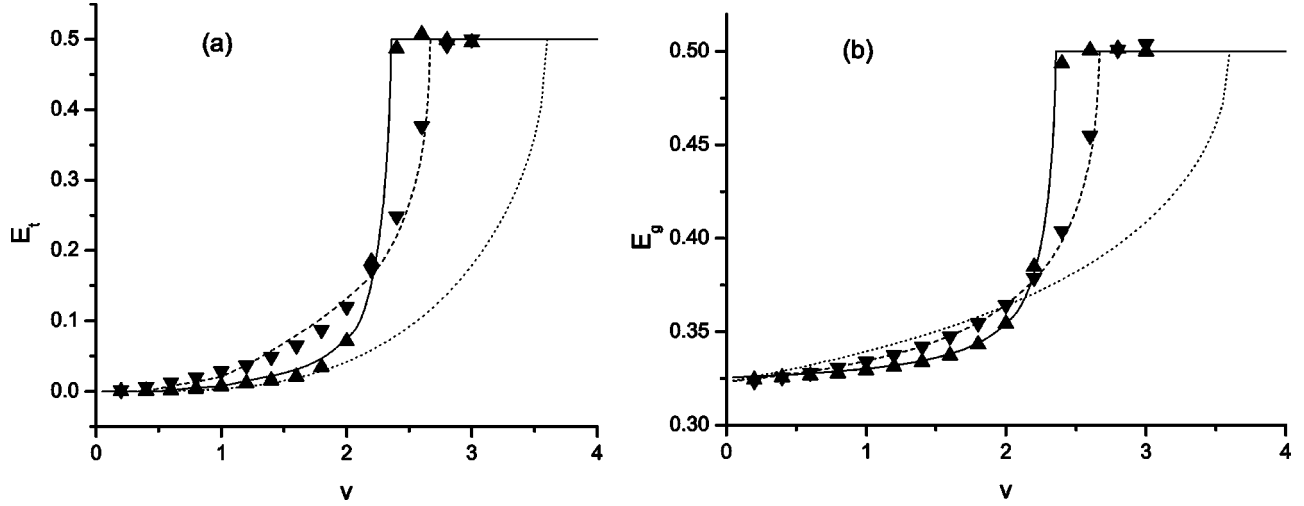


FIG. 3. The dependence of (a) the training error E_t and (b) generalization error E_g on the learning rate ν for $\alpha=0.5$. Solid lines: steady-state results of the cavity method for $\lambda=0.4$. Dashed lines: steady-state results of the cavity method for $\lambda=0.8$. Corresponding symbols: simulations averaged over 100 samples with $N=1000$ and $t=20$. Dotted lines: the mean-force approximation for $\lambda=0.8$.

tor for an extended period of time. We use the weight vector averaged over the monitoring period τ , $\bar{\mathbf{J}} = \lim_{t \rightarrow +\infty} 1/\tau \int_t^{t+\tau} \mathbf{J}(t') dt'$, as the estimated student vector. The average weight amplitude is given by

$$\bar{C}(\tau) = \lim_{t \rightarrow +\infty} \frac{1}{\tau} \int_t^{t+\tau} dt' C(t, t+t'). \quad (50)$$

Using Eq. (37), we obtain

$$\begin{aligned} \bar{C}(\tau) = & \int dx \rho(x) \frac{x - \nu\lambda}{x^2} \left[\frac{b\nu}{\alpha} + a^2 \left(x - \nu\lambda - \frac{\nu}{\alpha} \right) \right] \\ & + \frac{\nu}{2} \langle F^2 \rangle \int dx \rho(x) \left(\frac{x - \nu\lambda}{x} \right) \left(\frac{e^{-x\tau} - 1 + x\tau}{x^2 \tau^2 / 2} \right), \end{aligned} \quad (51)$$

where $\langle F^2 \rangle$ is given by Eq. (41). We note that when τ becomes very large, the contribution due to force fluctuations vanishes, making $\bar{C}(\tau)$ approach the result for batch learning. Hence, the *average learning strategy* yields a generalization performance as good as that of the batch mode, *independent of the learning rate, as long as it is below the critical value for divergent learning*. Figure 4 shows that the generalization error of on-line learning is equal to that of the batch mode at $\nu=0$, and gradually grows larger when ν increases. When the monitoring time τ increases, there is an impressive reduction of the generalization error compared with its instantaneous values. The longer the monitoring time τ , the smaller generalization error. In the limit of very long monitoring time, the generalization ability becomes equal to that of the batch mode for all values of ν below ν_c , and jumps discontinuously to divergence at ν_c .

V. CONCLUSION

We have analyzed the dynamics of on-line learning with restricted sets of examples, which are randomly recycled. Using the cavity approach, we can derive equations for the macroscopic parameters describing the learning dynamics. They are solvable for linear rules such as the Adaline rule, yielding results which agree with simulations. We also show that the student in on-line mode can learn as well as that in batch mode, after it is averaged over a monitoring period at the steady state.

Our work represents a step forward from two recent treatments [17,19]. Compared with Ref. [17], we have found that a functional relationship exists between the generic and cavity activations, and made explicit the distribution of activa-




FIG. 4. The steady-state generalization error E_g averaged over different monitoring times at $\alpha=2$ and $\lambda=0.2$. Symbols: simulations averaged over 10 samples with $N=1000$, $\tau=0$ (●), 2 (■), 5 (◆), 10 (▲), and 20 (▼). Corresponding lines: theory.

tions, which is a superposition of Gaussians for the Adaline rule and each weighted by a Poisson distribution. As a result, our framework can be used to analyze the training error E_t . More importantly, it has the potential to be extended to analyze nonlinear and multilayer networks. Compared with Ref. [19], we have based our analysis on a more physical picture. For example, we have made explicit the difference between the active and passive averages involving the activation of an example and its learning force at a previous instant. This enables us to analyze correctly the network behavior at large learning rates, where the mean-force approximation does not apply. The physical insights will be useful when analytical approximation schemes are devised for more complex networks.

However, our analysis shows that it is difficult to obtain a concise description of the steady-state behavior of on-line learning. The mean-force approximation in Ref. [19] represents a bold attempt towards this goal, but it works for small learning rates only. Here, we derive the steady-state activation distribution in Eq. (49), but it depends on the previous history of learning instants, and no instantaneous description of the distribution is available. As a result, it seems that a Monte Carlo sampling procedure is so far the best approach to find the steady-state distribution.

We have achieved our objective of benchmarking the cavity approach using the Adaline learning. The next step will be to extend the method to more complicated situations such as nonlinear and multilayer networks. We may devise efficient Monte Carlo sampling procedures to solve numerically the equations for the Green's functions, the teacher-student correlations, and the student autocorrelations, making use of the Poisson distribution of the learning events for a single example, the Gaussian distribution of the cavity activations, and their causal relations with the generic student activations.

The cavity analysis of linear networks is also the foundation of approximate descriptions of the stationary behavior of nonlinear and multilayer networks. One may describe the steady state by fluctuations about an averaged state. The fluctuations can be approximated by linear deviations which can be analyzed by the cavity approach analogous to the Adaline benchmark. The present work is moving along this direction.

Recently, various approximation schemes have been proposed in different learning regimes of complicated networks [17,25,27], yielding results with varying degree of success. For example, the conditionally Gaussian approximation cannot capture the sharp peaks in the activation distribution developed at the late learning stage of nonlinear rules (e.g., Adatron learning). On the other hand, from the perspectives of the cavity method, the nonlinear mapping from the Gaussian cavity activations to the generic activations can lead to sharp peaks in a natural way [9]. It is hoped that the cavity framework can provide useful insights to improved approximations in the future.

From a broader perspective, our work has developed concepts which will be useful in the dynamical studies of other on-line systems, say, in economics, ecology, and traffic models. We have demonstrated how sequence averages can provide precise descriptions of the typical behavior of these sys-

tems. Furthermore, we have illustrated the need to differentiate between active and passive averages when we analyze responses to system states appearing as sequences in on-line processes. Also, the cavity method provides a means to compute the difference between the active and passive averages. It also provides a Monte Carlo sampling procedure based on Poisson-generated events to study the dynamics. One possible application of these concepts may be found in the minority game, where the sequence of market states determines the evolution of the market behavior, analogous to the way the sequence of training examples determine the network behavior in on-line learning [22].

ACKNOWLEDGMENTS

We thank Ton Coolen for meaningful communications. This work is supported by the Research Grant Council of Hong Kong (Grants No. HKUST6157/99P and HKUST6153/01P).

APPENDIX A: SEQUENCE AVERAGE OF THE STUDENT ACTIVATION

According to the Eqs. (23) and (16), the sequence average of the student activation of an example is

$$\langle x(t) \rangle_\sigma = h(t) + v \sum_{m=1}^{\infty} \frac{e^{-t/\alpha}}{\alpha^m} \int_0^t dt_m \cdots \int_0^{t_2} dt_1 \sum_{r=1}^m G(t, t_r) F(t_r). \tag{A1}$$

Reordering the summations and the integrations before and after t_r ,

$$\langle x(t) \rangle_\sigma = h(t) + v \sum_{m=1}^{\infty} \frac{e^{-t/\alpha}}{\alpha^m} \sum_{r=1}^m \int_0^t dt_r \int_{t_r}^t dt_{r-1} \cdots \times \int_{t_2}^t dt_1 \left[\int_0^{t_r} dt_{r+1} \cdots \int_0^{t_{m-1}} dt_m G(t, t_r) F(t_r) \right], \tag{A2}$$

where we have made implicit the dependence of $F(t_r)$ on the previous learning instants t_{r+1}, \dots, t_m . Since the integrand in the square bracket does not depend on the values of t_1, \dots, t_{r-1} , the integration over these variables simply gives the factor of $(t-t_r)^{r-1}/(r-1)!$. We further note that

$$\sum_{m=1}^{\infty} \sum_{r=1}^m \frac{e^{-t/\alpha}}{\alpha^m} \equiv \frac{1}{\alpha} \sum_{m-r=0}^{\infty} \frac{e^{-t_r/\alpha}}{\alpha^{m-r}} \sum_{r-1=0}^{\infty} \frac{e^{-(t-t_r)/\alpha}}{\alpha^{r-1}} \tag{A3}$$

and then derive

$$\begin{aligned} \langle x(t) \rangle_\sigma &= h(t) + \frac{v}{\alpha} \int_0^t dt_r G(t, t_r) \\ &\times \left[\sum_{m-r=0}^{\infty} \frac{e^{-t_r/\alpha}}{\alpha^{m-r}} \int_0^{t_r} dt_{r+1} \cdots \int_0^{t_{m-1}} dt_m F(t_r) \right]. \end{aligned} \quad (\text{A4})$$

From Eq. (23), the summation in the square bracket is just the sequence average of the activation at time t_r , yielding Eq. (27).

APPENDIX B: SEQUENCE AVERAGE OF ACTIVE AND PASSIVE CORRELATIONS

From Eqs. (16) and (23), the active sequence average at time s of the activation of example μ that is learned at time t is

$$\begin{aligned} \langle x_\mu(s) \rangle_{\sigma|\sigma(t)=\mu} &= h_\mu(s) + v \sum_{m=0}^{\infty} \frac{e^{-t/\alpha}}{\alpha^m} \int_0^t dt_1 \cdots \\ &\times \int_0^{t_{m-1}} dt_m \sum_{n=0}^{\infty} \frac{e^{-(s-t)/\alpha}}{\alpha^n} \int_t^s dt_{m+1} \cdots \\ &\times \int_t^{t_{m+n-1}} dt_{m+n} \left[\sum_{r=1}^{m+n} G(s, t_r) F_\mu(t_r) \right. \\ &\left. + G(s, t) F_\mu(t) \right]. \end{aligned} \quad (\text{B1})$$

Using similar arguments as in Appendix A, the average for the Adaline rule can be written as a self-consistent equation for $s > t$,

$$\begin{aligned} \langle x_\mu(s) \rangle_{\sigma|\sigma(t)=\mu} &= h_\mu(s) + \frac{v}{\alpha} \int_0^t dt' G(s, t') \langle F_\mu(t') \rangle_\sigma \\ &+ v G(s, t) \langle F_\mu(t) \rangle_\sigma + \frac{v}{\alpha} \int_t^s dt' G(s, t') [\tilde{y}_\mu \\ &- \langle x_\mu(t') \rangle_{\sigma|\sigma(t)=\mu}]. \end{aligned} \quad (\text{B2})$$

Note that $\langle x_\mu(s) \rangle_{\sigma|\sigma(t)=\mu} = \langle x_\mu(s) \rangle_\sigma$ when $s \leq t$. Subtracting Eq. (27) from the above equation, one can derive

$$\begin{aligned} \langle x_\mu(s) \rangle_{\sigma|\sigma(t)=\mu} &- \langle x_\mu(s) \rangle_\sigma \\ &= v G(s, t) \langle F_\mu(t) \rangle_\sigma - \frac{v}{\alpha} \int_0^s dt' G(s, t') \\ &\times [\langle x_\mu(t') \rangle_{\sigma|\sigma(t)=\mu} - \langle x_\mu(t') \rangle_\sigma]. \end{aligned} \quad (\text{B3})$$

If we multiply $F_\mu(t)$ to both sides of Eq. (16) at time s , and perform sequence averages analogous to Eqs. (B1) and (B2),

we obtain equations for the active and passive averages of the activation-force correlation,

$$\begin{aligned} \langle x_\mu(s) F_\mu(t) \rangle_{\sigma|\sigma(t)=\mu} &= h_\mu(s) \langle F_\mu(t) \rangle_\sigma + \frac{v}{\alpha} \int_0^t dt' G(s, t') \\ &\times \langle F_\mu(t') F_\mu(t) \rangle_{\sigma|\sigma(t')=\mu} + v G(s, t) \langle F_\mu^2(t) \rangle_\sigma \\ &+ \frac{v}{\alpha} \int_t^s dt' G(s, t') \langle F_\mu(t') F_\mu(t) \rangle_{\sigma|\sigma(t)=\mu}, \end{aligned} \quad (\text{B4})$$

$$\begin{aligned} \langle x_\mu(s) F_\mu(t) \rangle_\sigma &= h_\mu(s) \langle F_\mu(t) \rangle_\sigma + \frac{v}{\alpha} \int_0^t dt' G(s, t') \\ &\times \langle F_\mu(t') F_\mu(t) \rangle_{\sigma|\sigma(t')=\mu} \\ &+ \frac{v}{\alpha} \int_t^s dt' G(s, t') \langle F_\mu(t') F_\mu(t) \rangle_\sigma. \end{aligned} \quad (\text{B5})$$

Subtracting Eq. (B4) with Eq. (B5) leads to

$$W_\mu(s, t) = v G(s, t) \langle F_\mu^2(t) \rangle_\sigma - \frac{v}{\alpha} \int_t^s dt' G(s, t') W_\mu(t', t). \quad (\text{B6})$$

Multiplying the example Green's function $D(r, s)$ to both sides of Eq. (B6), integrating over s , and applying Eq. (12), one obtains Eq. (33).

To obtain Eq. (34), one replaces $\langle F_\mu(t') F_\mu(t) \rangle_{\sigma|\sigma(t')=\mu}$ in the right hand side of Eq. (B5) with $\langle F_\mu(t') F_\mu(t) \rangle_\sigma - v \int_t^{t'} dt_1 D(t, t_1) G(t_1, t') \langle F_\mu^2(t') \rangle_\sigma$, and arrives at

$$\begin{aligned} \langle x_\mu(s) F_\mu(t) \rangle_\sigma &= h_\mu(s) \langle F_\mu(t) \rangle_\sigma + \frac{v}{\alpha} \int_0^s dt' G(s, t') \\ &\times \langle \tilde{y}_\mu F_\mu(t) \rangle_\sigma - \frac{v}{\alpha} \int_0^s dt' G(s, t') \\ &\times \langle x_\mu(t') F_\mu(t) \rangle_\sigma - \frac{v^2}{\alpha} \int_0^t dt_1 \\ &\times \int_{t_1}^t dt_2 G(s, t_1) D(t, t_2) G(t_2, t_1) \\ &\times \langle F_\mu^2(t_1) \rangle_\sigma. \end{aligned} \quad (\text{B7})$$

Multiplying both sides with $D_\mu(r, s)$, integrating over s and applying Eq. (12), one finally reaches to Eq. (34).

- [1] C.W.H. Mace and A.C.C. Coolen, *Stat. Comput.* **8**, 55 (1998).
- [2] Robert Axelrod, *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* (Princeton University Press, Princeton, NJ, 1997).
- [3] *The Economics as an Evolving Complex System*, edited by P. W. Anderson, K. Arrow, and D. Pines, (Addison-Wesley, Redwood City, CA, 1988).
- [4] B. Huberman and R. Lukose, *Science* **277**, 535 (1997).
- [5] A. Tanenbaum, *Computer Networks*, 4th ed. (Prentice-Hall, Englewood Cliffs, NJ, 2003).
- [6] C. M. Bishop, *Neural Networks for Pattern Recognition* (Clarendon Press, Oxford, 1995).
- [7] *On-line Learning in Neural Networks*, edited by D. Saad (Cambridge University Press, Cambridge, London, 1998).
- [8] K.Y.M. Wong, S. Li, and Y.W. Tong, *Phys. Rev. E* **62**, 4036 (2000).
- [9] K.Y.M. Wong, *Europhys. Lett.* **30**, 245 (1995).
- [10] M. Biehl and H. Schwarze, *J. Phys. A* **28**, 643 (1995).
- [11] D. Saad and S.A. Solla, *Phys. Rev. Lett.* **74**, 4337 (1995).
- [12] G. Radons, H. G. Schuster, and D. Werner, in *Parallel Processing in Neural Systems and Computers*, edited by R. Eckmiller, *et al.* (North-Holland, Amsterdam, 1990), p. 261.
- [13] L.K. Hansen, R. Pathria, and P. Salamon, *J. Phys. A: Math Gen.* **26**, 63 (1993).
- [14] T. Heskes, *J. Phys. A: Math Gen.* **27**, 5145 (1994).
- [15] T.M. Heskes and B. Kappen, *Phys. Rev. A* **44**, 2718 (1991).
- [16] H.C. Rae, P. Sollich, and A.C.C. Coolen, *J. Phys. A: Math Gen.* **32**, 3321 (1999).
- [17] D. Barber and P. Sollich, in *On-line learning in Neural Networks* (Ref. [7]), p. 279.
- [18] A.C.C. Coolen, D. Saad, and Y.S. Xiong, *Europhys. Lett.* **51**, 691 (2000).
- [19] J.A.F. Heimeel and A.C.C. Coolen, *J. Phys. A: Math Gen.* **34**, 9009 (2001).
- [20] M. Mézard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [21] P. Luo and K.Y. Michael Wong, *Phys. Rev. E* **64**, 061912 (2001).
- [22] D. Challet, M. Marsili, and R. Zecchina, *Phys. Rev. Lett.* **84**, 1824 (2000).
- [23] K. Y. M. Wong and F. Li, *Advances in Neural Information Processing Systems*, edited by Z. Ghahramani (MIT Press, Cambridge, 2002), Vol. 14, p. 535.
- [24] H. Eissfeller and M. Opper, *Phys. Rev. Lett.* **68**, 2094 (1992).
- [25] S.I. Amari, *Neural Comput.* **10**, 251 (1998).
- [26] M. Opper, *Phys. Rev. Lett.* **77**, 4671 (1996).
- [27] A.C.C. Coolen and D. Saad, *Phys. Rev. E* **62**, 5444 (2000).