Versatility and adaptive performance

Preben Alstrøm and Dimitris Stassinopoulos

Center for Chaos and Turbulence Studies, The Niels Bohr Institute, Copenhagen 2100 Ø, Denmark

We construct a class of adaptive networks whose purpose is to function in a volitile environment. The characteristic structure is a sparsely connected network with a low-activity internal control mechanism that functionally keeps the network adaptive and with a local-activity-dependent synaptic response to rewarded and penalized actions. The resulting dynamics, called *adaptive performance*, adjusts easily to changes in parameters induced by input or feedback, and it shows fluctuations over a wide range of time scales. Our specific studies concern conditioning and control. We speculate that adaptive performance may be an underlying mechanism in brain function, as well as provide a basis for new technology.

One of the most long-sought goals in neuroscience is to understand the mechanisms that enable the brain to undertake a diversity of tasks [1]. On the neural level, the underlying network dynamics is not task specific, and any realistic attempts to explain the network's ability of function must take this into account [2].

In this paper we introduce and analyze adaptive networks that by construction are versatile. We shall see that the network performance is intimately connected with an adaptive nature that allows information to flow through a variety of ever-changing paths. We call this *adaptive performance* to emphasize the conceptual differences from *adapted performance*, where an (artificial) neural network in a "learning mode" is adapted to a given task [3]. In this process the weights (corresponding to the synaptic efficacies) are set. The weights are then used in an "associative memory mode" in which only the state vector is changed. Conceptually, training and performance are separated as two different "modes." The usual argument for this separation in modeling brain functions is that the synaptic efficacies in general change only slightly on the basic time scale for the neural dynamics, i.e., the firing time. One might therefore naively expect that the dynamics for the synaptic efficacies can be separated from the dynamics for the neurons. This is not true: Even small changes in synaptic efficacies may change the activity substantially. Such highly adjustable states we find in turn to be an inevitable result of versatile goal-directed behavior. Thus, in our picture of adaptive networks the neural state n and the synaptic efficacies J_{ij} are necessarily interwoven dynamical variables.

A fundamental result that brings the above view into further perspective is the wide range of time scales that are observed in the brain, e.g., in EEG "brain wave" recordings of collective neural activity [1,4]. Spectra obtained from these recordings can often be broad with a long tail that resembles $1/f$ noise. This intriguing observation unveils an important characteristic of brain function, namely, a wide distribution of temporal fluctuations. Thus, one cannot divide brain function into a given number of distinct processes taking place on different time scales. The variety of time scales is a result of the dynamics itself. In addition, the spectra may show peaks (α, β, γ, δ, and θ waves) associated with specific types of brain activity. Below, we show numerically that adaptive performance gives rise to an adaptive complex dynamics with similar wide distributions of time scales.

In terms of brain function, we stress from the beginning that our goal is not to explain "intelligent" behavior from a microscopic description of the electrochemical processes important on the neuronal level. Our purpose is here to point out mechanisms that, we believe, underlie the cooperative neuronal activity, and, most importantly, to show that these mechanisms result in successful behavior in a changeable environment. We find that the sufficient and necessary ingredient is a sparsely connected network with a low-activity internal control mechanism that functionally keeps the network adaptive, and with a locally activity-dependent synaptic response to rewarded and penalized actions.

To be concrete, consider a network with the architecture depicted in Fig. 1. Periodic boundary conditions are assumed for the layers. To avoid the firing pattern becoming dominated by loops of firing neurons, and for computational simplicity, a feed-forward network is chosen with input at the first row and output at the bottom row. The network units are McCulloch-Pitts [5] "neurons" that can be in only two states, firing ($n=1$) or not firing ($n=0$) [6]. For biological implementation sparseness is required. By sparseness we mean that the number of synaptic junctions from one neuron is extremely small compared to the number of neurons. Here, spareness is introduced by assuming that each neuron has only three incoming and three outgoing "synaptic junctions" [Fig. 1(b)]. The number of neurons and/or synapses may be chosen to be larger. In the human brain, there are billions of neurons, and each neuron is associated with thousands of synaptic junctions, giving the brain faster performance and the ability to solve more complex problems.

The input h_i (postsynaptic potential) received by the ith neuron is the linear sum of the synaptic efficacies J_{ij} (from the jth neurons to th ith neuron) of those connections that were activated,

$$h_i = \sum_j J_{ij} n_j \; . \tag{1a}$$

If this input exceeds a threshold field T (here taken to be homogeneous), then the ith neuron fires,

$$n_i \to n_i' = \Theta(h_i - T) \; . \tag{1b}$$

The modification of a synapse, i.e., the change in efficacy, called reinforcement, is allowed to depend only locally on the coherence in activity. The goal of reinforcement is to strengthen firing paths that lead to a positive feedback (synaptic facilitation), and weaken the paths resulting in a negative feedback (synaptic depression). In our simple network, this is done using the following reinforcement rule [7]:

$$J_{ij} \to J_{ij}' = \tilde{J}_{ij} \Big/ \sum_i \tilde{J}_{ij} \; , \tag{2a}$$

where

$$\tilde{J}_{ij} = J_{ij} + f(J_{ij}) n_i' n_j \; , \tag{2b}$$

and where the reinforcement function $f(J)$ has the logistic form

$$f(J) = rJ(1 - J) + \eta \; , \tag{2c}$$

with $|r| \ll 1$. This specific form of the reinforcement function seems unimportant [$J(1-J)$ may be replaced by other positive functions with saturation points (here 0 and 1)]. In our "minimal" model, all neurons are taken

(a)

INPUT

OUTPUT

FEEDBACK EXTERNAL SYSTEM

(b)

FIG. 1. (a) Schematic diagram of an adaptive network in an environment. (b) Illustration of incoming and outgoing junctions.

to be "excitatory" [8], $0 < J_{ij} < 1$. Possible noise is represented by the term η, chosen randomly in space and time from a uniform distribution of values between $-\eta_0$ and $+\eta_0$. From an energetic point of view, noise is often helpful for finding relevant paths through the network.

The (global) "reaction" field r is associated with a feedback from the environment. Again, we cannot allow this feedback to be task specific without loosing versatility — error back propagation or similar instructive mechanisms applied to speed up training in artificial neutral networks are thus precluded. To be acceptable, the feedback must be only evaluative. We here consider the simplest case of a yes or no (right or wrong) signal [9], and r assumes only two values,

$$r(\text{yes}) = r_+ > 0, \quad r(\text{no}) = r_- < 0 \; . \tag{3}$$

An additional element crucial to the above type of network is an internal control mechanism for the activity, i.e., the number of firing neurons. We find that it is essential for the performance that this number is kept small [10]. In the brain, the interplay between excitatory and inhibitory neurons is partly responsible for such a control, but restrictions on the transport of chemicals probably play an important role too. Here, we implement the internal control on the threshold field T. Alternatively, one an implement the control on the synaptic efficacies. The total activity $A = \sum_{i \in C} n_i$ from a region C, taken to be the output region, is detected. If this exceeds the value A_0, the threshold T is slightly increased, while T is lowered if $A < A_0$,

$$T \to T' = T + \delta \, \text{sgn}(A - A_0) \; , \tag{4}$$

where $0 < \delta \ll 1$, and $\text{sgn}(0) = 0$. A_0 must be small compared to $A_{\max} = |C|$. Here we use $A_0 = 1$, and $\delta = 0.01/N$, where N is the number of layers.

Threshold control has also been used to regulate activity in more traditional neural network studies [11]. However, in these studies the threshold control is (as well as other rules) implemented to "tune" the network towards selected activity patterns, and the search for rules are motivated by a desire for more efficient learning and recall. This requirement naturally leads to different rules for the learning mode and the associative memory mode [11]. In the conceptually different one-mode network presented here, "most efficient" is ill-defined as no tasks are given a priori. The rules must be task independent. In choosing these rules, the demand for versatility and capability replaces the demand for speed and efficiency.

Our adaptive dynamics is applied to two tasks: A, conditioning, where a certain active behavior is rewarded, and B, control, where the task is to track a moving target. The motion is here restricted to be on a line. Under A we show that conditioning leads to learning through an adaptive process with temporal fluctuations that resemble $1/f$ noise. Task B is chosen to clearly demonstrate the difference between adaptive and adapted performance. While task A is a basic test for applications involving classification and analysis; task B is a paradigm for control systems and robot kinematics. In order to show a performance that is solely a result of the networks adap-
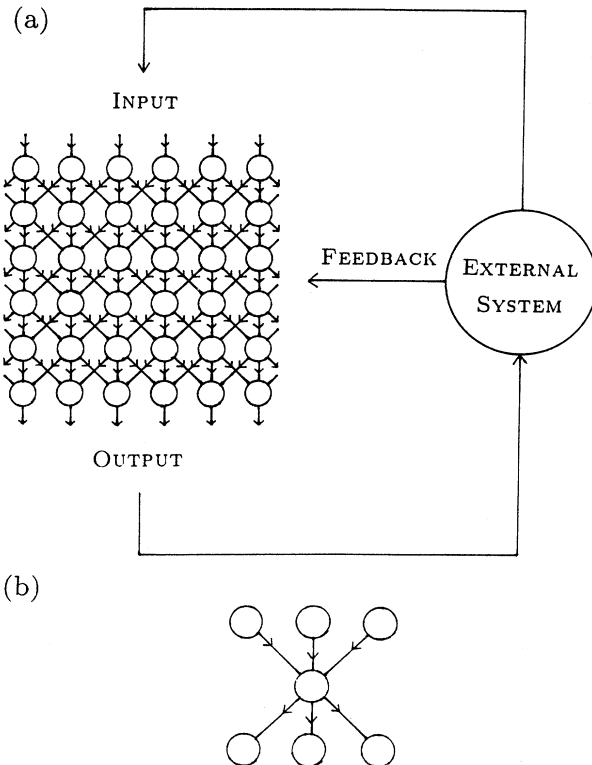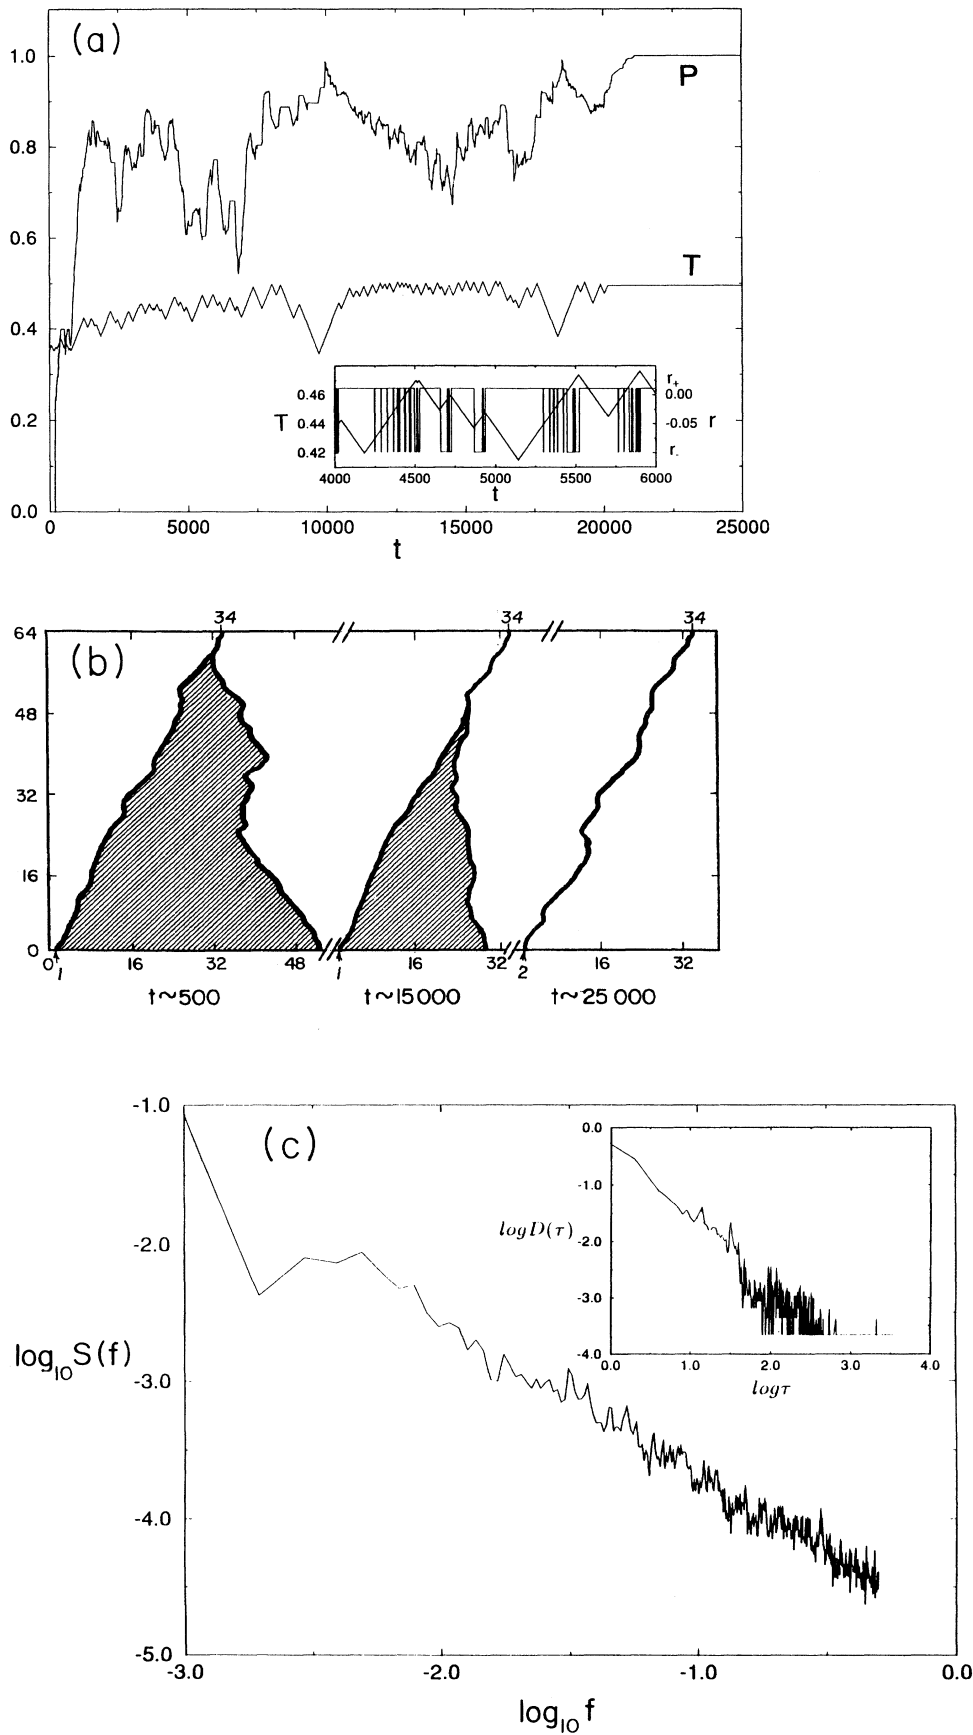
FIG. 2. (a) Performance $P(t)$ (averaged over the interval $[\max(0, t - 1000); t]$), and threshold $T(t)$ for task A. Inset: Sample of temporal fluctuations in the reaction field r and threshold $T(t)$. (b) The activity prerecorded over 1000 time steps is concentrated within a region that slowly narrows down to a single path. (c) Power spectrum $S(f)$ for the fluctuating reaction field. Inset: Distribution $D(\tau)$ of time intervals over which r is constant (9198 intervals total, log to the base 10).

tive nature, the input for both tasks is taken to be a single neuron that fires constantly. The initial condition is always taken as a random uniformly distributed set of J_{ij}'s $(0 < J_{ij} < 1)$, and for all neurons, $n_i = 0$ (except for the input).

Task A: A 64×64 lattice is used, and at every time step activation of an output neuron is associated with a certain action (out of 64). Behavior is the combination of actions, and is considered active if at least one action is performed. One of the actions (randomly selected) is considered desirable. The behavior is instantaneously rewarded if it includes the desired action. The feedback is taken to be positive $(r = r_+ = 0.01)$ if the behavior is re-

warded, or if no actions are performed at all. Otherwise, the feedback is negative $(r = r_- = -0.1)$. Noise is neglected, i.e., $\eta_0 = 0$. We define the *performance* P as the fraction of active behaviors that were rewarded. Figure 2(a) shows an example of how the performance P and the threshold T change with time. A threshold variation is observed that is bounded between $T_{\min} \simeq \frac{1}{3}$ and $T_{\max} \simeq \frac{1}{2}$. The performance P is found to rise quickly from 0 to about 0.8, around which value P fluctuates until it reaches unity.

In our approach, performance is taking place on the neural time scale (ms). From the point of neuroscience, we do not act and react on the neural time scale. Our
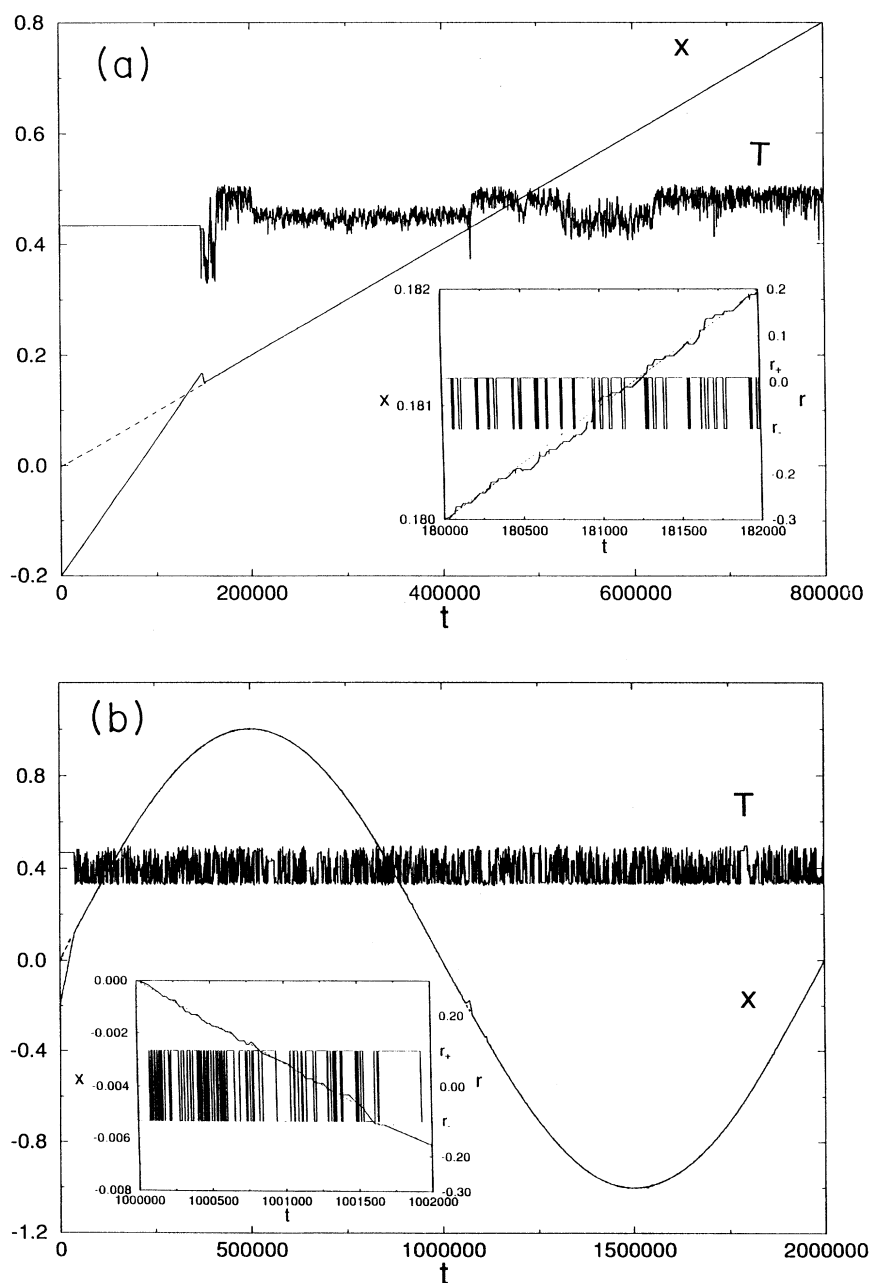

FIG. 3. Simulated motion of a tracking point $x(t)$ that tracks (a) a target moving with a constant speed; (b) an oscillating target. Also shown is that threshold variation $T(t)$. Insets: Local magnification of the motion $x(t)$ and the associated reaction field $r(t)$. In (a) $r_- = -0.1$, $r_+ = 0.01$, and $\eta_0 = 0$. In (b) $r_- = -0.1$, $r_+ - 0.1$, and $\eta_0 = 0.05$. (c) Power spectra $S(f)$ for the fluctuating reaction fields. The thin curve is for (a); the thick curve is for (b). Insets: Corresponding distributions $D(\tau)$ of time intervals over which r is constant [43 604 (respectively, 27 936) intervals total] the log in (c) is to the base 10.

neural dynamics may be extended, taking into account that the feedback from the external environment gives rise to a reinforcement signal on a time scale much longer than the neutral time scale considered here. Such an extension can reduce the number of the behavioral responses that are needed.

When the network performs perfectly, the activity is concentrated along a single path connecting the input site with the output site associated with the desired action. Through the preceding process the average activity is slowly narrowed down, as shown in Fig. 2(b). The observed behavior pattern is, however, very complex. Shown in the inset of Fig. 2(a) is a sample of the temporal fluctuations in the reaction field r and threshold T. From the signals it is clear that many time scales are represented. The power spectrum [12] $S(f)$ for the fluctuating reaction field is shown in Fig. 2(c). It displays a $1/f$ power law over a wide range of time scales,

$$S(f) \sim f^{-\alpha} , \tag{5a}$$

with $\alpha \simeq 1.1$. In addition, we have, on the basis of five independent runs, determined the distribution $D(\tau)$ of time intervals τ over which r is constant ($r = r_+$ or $r = r_-$) [inset of Fig. 2(c)]. The distribution is approximated by a power law,

$$D(\tau) \sim \tau^{-\beta} , \tag{5b}$$

with $\beta \simeq 1.2$. The results suggest that adaptive performance in the process of learning leads to behavioral changes on a great variety of time scales.

Task B: A 16×16 lattice is used, and activation of the ith output site results in motion of the tracking point by a distance d_i. The values of d_i are chosen to be within a range $-d < d_i < d$ of sufficient size, randomly distributed

along the output line. Whenever there is a firing at the output line, the tracking point is move accordingly. If more than one fires the average of the corresponding distances is chosen. In case of no motion (no output firing), or when motion brings the tracking point closer to the target, the feedback is positive ($r = r_+$); when the tracking point is moved further away from the target, the feedback is negative ($r = r_-$). Two cases are considered: (a) The target moves with a constant speed; (b) the target oscillates. In case (a), no noise is present ($\eta_0 = 0$); in case (b), noise is present ($\eta_0 = 0.05$). The results are shown in Fig. 3. It is seen that the target is found and followed.

A closer inspection [insets of Figs. 3(a) and 3(b)] shows that the network dynamics again is complex. Adaptive performance gives rise to motion fluctuations on all time scale [Fig. 3(c)]. However, peaks may now be presented in the power spectrum [13] $S(f)$ and in the distribution $D(\tau)$, indicating that some time scales are selected for the performing process. In case (a), we find two peaks in $D(\tau)$, a narrow peak at $\tau_1 = 15$, and a broad peak around $\tau_2 = 80$. A separation of the distribution $D(\tau)$ into two distributions $D_-(\tau)$ and $D_+(\tau)$ associated with $r = r_-$ and $r = r_+$ shows that τ_1 is related to the negative feedback, while τ_2 is related to the positive feedback. The result suggests that peaks may arise when learning has taken place. We emphasize, however, that the activity is not restricted to a few output sites. On the contrary, the activity keeps fluctuating over all output sites. In case (b), no peaks are present, may be as a result of noise. The spectrum and time distribution are approximated by the power laws (5a) and (5b) with $\alpha \simeq 1.5$ and $\beta \simeq 1.4$.

The conditioning test (Skinner box) and the one-dimensional control problem presented by tasks A and B are the most simple representatives of their kind. Further complications may be added: Pavlov's dog tests,
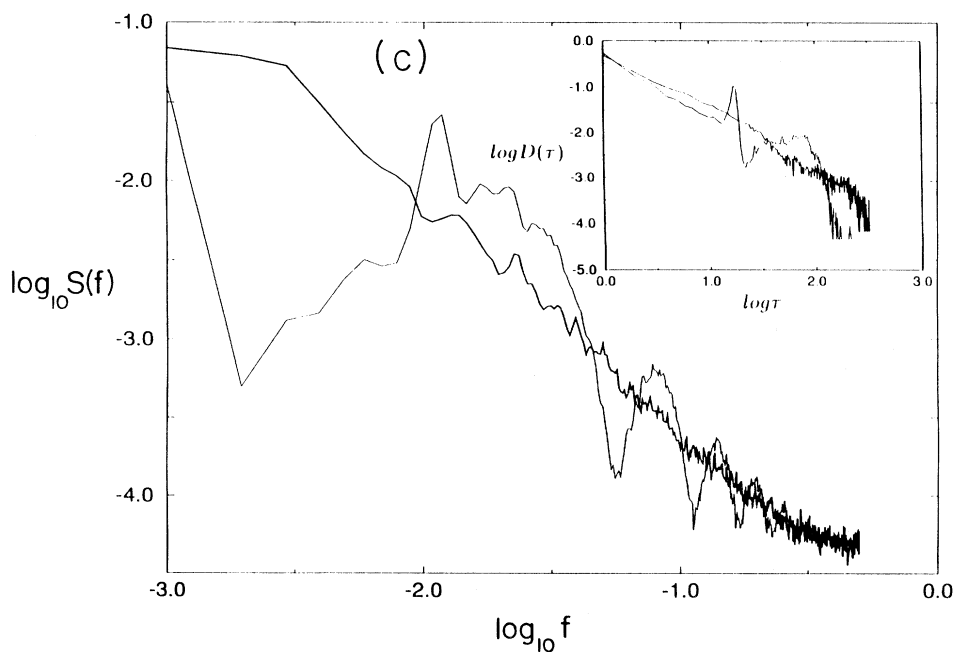

FIG. 3. (Continued).

recognition tasks, inverted pendulum problems, higher-dimensional control problems, etc. The challenges seem endless in number.

The need for adaptive networks of the type introduced above seems particularly clear where the external response possibilities are not known or cannot be defined at a sufficient level, or where the goals progressively change. In the real world in which we function, this is the rule. In industry it is often the case, which suggests adaptive performance as a basis for new technology. Unlike many other artificial neural networks, the networks introduced here do not need a supplementary agent, computing by standard means how to set the synaptic weights. Adaptive networks locally and dynamically set the weights themselves, according to the tasks selected by the user. Furthermore, the networks have by construction a black-box design, which means that they are interface independent—essentially, no manual is needed.

In conclusion, we believe that the concept of adaptive performance may be underlying mechanism in brain functions, as well as provide a basis for new technology. To this end, we have introduced a class of versatile adaptive networks whose main purpose is to function in a volatile environment. The essential ingredient is a sparsely connected network with an internal control mechanism for the activity, and with a local-coherence-dependent synaptic response to rewarded and penalized actions. On this basis, we find a behavior pattern with temporal fluctuations over a wide range of time scales, indicating that the neural network functionally is highly adaptive to environmental changes. We have shown that the resulting power spectra resembles $1/f$ noise, in some cases with peaks associated with certain types of activity.

[1] For a general introduction to neuroscience, see, e.g., E. R. Kandel, H. H. Schwartz, and T. M. Jessell, *Principles of Neural Science*, 3rd ed. (Elsevier, New York, 1991).

[2] This point and the epistemological consequences have recently been emphasized and discussed by W. J. Freeman and co-workers; see, e.g., C. A. Skarda and W. J. Freeman, Behav. Brain Sci. **10**, 161 (1987), and the following Open Peer Commentary; see also W. J. Freeman, Sci. Am. **264**, 78 (1991).

[3] For an introduction to artificial neural networks and computing, see, e.g., D. J. Amit, *Modeling Brain Function: The World of Attractor Neural Networks* (Cambridge University, Cambridge, 1989); J. A. Hertz, A. S. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computation* (Addison-Wesley, Redwood, 1991); J. M. Zurada, *Introduction to Artificial Neural Systems* (West, St. Paul, 1992).

[4] W. J. Freeman, *Mass Action in the Nervous System* (Academy, New York, 1975); Biol. Cybern. **56**, 139 (1987); see also Ref. [2].

[5] W. S. McCulloch and W. Pitts, Bull. Math. Biophys. **5**, 115 (1943).

[6] Physiologically, one may consider the neuronal unit as a group of neurons. See the discussion on this point in Ref. [2].

[7] Note that the definition $\sum_i J_{ij} = 1$. This strict condition may, however, be slightly weakened. For the brain it represents a constraint on the chemical power available for neurotransmission.

[8] The model can easily be modified to include inhibitory neurons by changing (2) slightly. Fluctuations can thereby be reduced, allowing for a faster response time to environmental changes.

[9] In the context of neural computation, this is called reinforcement learning (see Hertz *et al.* in Ref. [3]). The global implementation of the reaction field used here is, however, fundamentally different from standard implementations involving stochastic processes in artifical neural network.

[10] In the context of associative memories, the important and biological relevance of a low activity has been discussed and studied using the so-called Willshaw model with local inhibition [D. Golomb, N. Rubin, and H. Sompolinsky, Phys. Rev. A **41**, 1843 (1994)].

[11] See, e.g., D. J. Willshaw and J. T. Buckingham, Philos. Trans. R. Soc. London Sect. B **329**, 205 (1990); Network **4** 441 (1993).

[12] The signal over $\sim 25\,000$ time steps was divided in 50 segments of 1024 time steps (50% overlap), and a standard Parcen window was used to obtain the power spectrum from these segments.

[13] Obtained from 195 segments of size 1024.