Springer Series *Physics of Neural Networks*

E. Domany
J.L. van Hemmen
K. Schulten (Eds.)

Models of Neural Networks III

Association, Generalization and Representation

Springer-Verlag New York, Berlin, Heidelberg 1996. ISBN 0-387-94368-4

Note: Due to a change of the style file, layout and length differ from the printed version.

1

Global Analysis of Recurrent Neural Networks

Andreas V.M. Herz[1]

ABSTRACT The article reviews recurrent neural networks whose retrieval dynamics have been analyzed on a global level using Lyapunov functions. Discrete-time and continuous-time descriptions are discussed. Special attention is given to distributed network dynamics, models with signal delays, and systems with integrate-and-fire neurons. The examples demonstrate that Lyapunov's approach provides powerful tools to study the retrieval of fixed-point memories, the recall of temporal associations, and the synchronization of action potentials in networks with spiking neurons.

1.1 Global Analysis — Why?

Information processing may be defined as the systematic manipulation of external data through the internal dynamics of some biological system or artificial device. In general, such a manipulation requires a highly nontrivial mapping between input data and output states. Important parts of this task can be accomplished with recurrent neural networks characterized by massive nonlinear feedback: triggered by an appropriate external stimulus, such systems relax towards attractors that encode some *a priori* knowledge or previously stored memories.

Within this approach to associative information processing, understanding the computational capabilities of a neural network is equivalent to knowing its complete attractor structure, that is, knowing what kind of input drives the network to which of its possibly time-dependent attractors. Understanding the computational properties of a recurrent neural network thus requires at least three levels of analysis. (i) What can be said about the existence and stability of fixed-point solutions? (ii) Are there static attractors only or are there also periodic limit cycles and aperiodic attractors, as expected for generic nonlinear systems? (iii) What is the structure of the basins of attraction?

Questions about the precise time evolution between the initial network state and the final output define a forth level of analysis. Though less im-

[1]Department of Zoology, University of Oxford, Oxford, OX1 3PS, England

portant within the framework of attractor neural networks, these questions are highly relevant for systems that extract information 'en route' without waiting for the arrival at some attractor [1]. At a fifth level of analysis, one might finally be interested in questions concerning the structural stability of a given network, that is, its robustness under small changes of the evolution equations.

With regard to the computational capabilities of a neural network, questions about the type of attractor and the structure of basins of attraction are of paramount importance. These questions deal with *global* properties of the network dynamics. Accordingly, they cannot be answered using local techniques only: a linear stability analysis of fixed-point solutions, the first level of analysis, may reveal helpful knowledge about the network behavior very close to equilibria, but it can never be used to rule out the existence of additional time-dependent attractors that may dominate large parts of the network's state space. Due to computational constraints, numerical simulations can offer limited additional information only.

Highly simplified network models provide a partial solution in that they often permit the application of global mathematical tools. However, such formal networks are characterized by bold approximations of biological structures. In the manner of good caricatures, they may nevertheless capture features that are also essential for more detailed descriptions.

One of the global mathematical tools is Lyapunov's "direct" or "second method" [2]. In the present context, it may be described as follows. Let the vector $x = (x_1, \ldots, x_N)$ denote the state variables of a neural network. These variables change in time according to some evolution equation, for example a set of coupled differential equations $\frac{d}{dt}x_i = f_i(x)$ if time is modeled as a continuous variable t. A solution will be denoted by $x(t)$. If there exists an auxiliary scalar state function $L(x)$ that is bounded below and nonincreasing along all trajectories, then the network has to approach a solution for which $L(t) \equiv L(x(t))$ does *not* vary in time.[2] The global dynamics can then be visualized as a down-hill march on an "energy landscape" generated by L. In this picture, every solution approaches the bottom of the valley in which it was initialized.

The asymptotic expression for $L(t)$ and the equation $\frac{d}{dt}L(t) = 0$ contain valuable information about the very nature of the attractors — the first and second level of analysis. Notice in particular that a solution that corresponds to a local minimum of the Lyapunov function has to be asymp-

[2]Special care has to be taken with respect to unbounded solutions and continuous families of solutions with equal L. Let me remark at this point that in the present article, formal rigor will often be sacrificed for transparency of presentation. A mathematically rigorous introduction to Lyapunov functions can be found in the monograph of Rouche, Habets and Laloy [3]. It contains — apart from a large number of interesting theorems and proofs — also some fascinating examples that illuminate possible pitfalls due to imprecise definitions.

totically stable, that is, it attracts every solution sufficiently close to it.

As an example, consider a gradient system

$$\frac{dx_i}{dt} = -\frac{\partial L(x)}{\partial x_i} \ . \tag{1.1}$$

Using the chain rule, the time derivative of L is given by

$$\frac{d}{dt}L(t) = \sum_{i=1}^{N} \frac{\partial L}{\partial x_i}\frac{dx_i}{dt} = -\sum_{i=1}^{N}\Big(\frac{dx_i}{dt}\Big)^2 \ . \tag{1.2}$$

The last expression is negative unless $x(t)$ is a fixed-point solution. It follows that if $L(x)$ is bounded below, the system has to relax to equilibria.

The most important feature of Lyapunov's direct method cannot be overemphasized: the method does not require any knowledge about the precise time evolution of the network; the mere existence of a bounded function that is nonincreasing along every solution suffices to characterize the system's long-time behavior. As a consequence, one can analyze the long-time dynamics of a feedback network without actually solving its equations of motion. Furthermore, most Lyapunov functions studied in this article play a role similar to that of the Hamiltonian of a conservative system: for certain stochastic extensions of the deterministic time evolution, the network dynamics approach a Gibbsian equilibrium distribution generated by the Lyapunov function of the noiseless dynamics. This has allowed the application of powerful techniques from statistical mechanics and has led to quantitative results about the performance of recurrent neural networks far beyond the limits of a local stability analysis. The existence of a Lyapunov function is thus of great conceptual as well as technical importance.

Lyapunov's method suffers, however, from one serious flaw: no systematic technique is known to decide whether a dynamical system admits a Lyapunov function or not. Finding Lyapunov functions requires experience, intuition, and luck. Fortunately, a wealth of knowledge on both practical and theoretical issues has been accumulated over the years.

The present article is intended as an overview of neural network architectures and dynamics where Lyapunov's method has been successfully employed to study the global network behavior. A general framework for modeling the dynamics of biological neural networks is developed in Section 1.2. This framework allows for a classification of various dynamical schemes found in the literature and facilitates the formal analysis presented in later sections.

Recurrent networks that relax to fixed-point attractors only have been used as auto-associative memories for static patterns. Section 1.3 reviews convergence criteria for a number of prototypical networks; the Hopfield model [4], the Little model [5], systems with graded-response neurons [6, 7], iterated-map networks [8] and networks with distributed dynamics [9, 10].

A statistical mechanical analysis of networks with block-sequential dynamics and results about the convergence to fixed points in networks with signal delays conclude the section.

Neural networks with signal delays can be trained to learn pattern sequences. Such systems are analyzed in Section 1.4. It is shown that with a discrete time evolution, these networks can be mapped onto "equivalent networks" with block-sequential updating and no time delays. This connection allows for a quantitative analysis of the storage of temporal associations in time-delay networks. Next, the time evolution of a single neuron with delayed feedback and continuous-time dynamics is discussed. Two different Lyapunov functions are presented. The first shows that under certain conditions, all solutions approach special periodic attractors; the second demonstrates that under less restrictive conditions, the system relaxes to time-varying solutions that need not be periodic.

The pulse-like nature of neural activity has frequently been modeled using (coupled) threshold elements that discharge rapidly when they reach a trigger threshold. With uniform positive couplings, some networks composed of such "integrate-and-fire neurons" approach globally synchronized solutions where all neurons fire in unison. With more general coupling schemes, the systems approach phase-locked solutions where neurons only exhibit locally synchronized pulse activity. Section 1.5 presents Lyapunov functions for such a class of integrate-and-fire models. An additional proof shows that the phase-locked solutions are reached in minimal time.

1.2 A Framework for Neural Dynamics

Starting with a brief description of the anatomy and physiology of single neurons, this section introduces a general framework for modeling neural dynamics.

1.2.1 Description of Single Neurons

Neurons consist of three distinct structures: dendrites, cell body, and axon. Dendrites are thin nerve fibers and form highly branched structures called dendritic trees. They extend from the central part of a neuron, the cell body or soma, which contains the cell nucleus. The axon, a single long fiber, projects from the soma and eventually branches into strands and substrands. Located along the axon and at its endings are synapses that connect one ("presynaptic") neuron to the dendrites and/or cell bodies of other ("postsynaptic") neurons [11].

Neurons communicate via an exchange of electrochemical signals. At rest, a cell is held at a negative potential relative to the exterior through selective ion pumps in the cell membrane. If the potential at the soma exceeds a firing threshold due to incoming signals, a strong electrical pulse is generated.

This excitation is called an action potential or spike. It is propagated along the axon by an active transport process that results in a soliton-like pulse of almost constant size and duration [12]. Following the generation of a spike, the membrane potential quickly drops to a subthreshold value. After the event, the neuron has to recover for a short time of a few milliseconds before it can become active again. This time interval is called the refractory period.

At synapses, action potentials trigger the release of neurotransmitters, chemical substances that diffuse to the postsynaptic cell where they bind to receptors. This process leads to changes of the local membrane properties of the postsynaptic neuron, causing either an increase or decrease of the local potential. In the first case, the synapse is called an excitatory synapse; in the second case, an inhibitory synapse. Through (diffusive) transport processes along the dendritic tree, an incoming signal finally arrives at the soma of the postsynaptic neuron where it makes a, usually minute, contribution to the membrane potential.

How can one construct a mathematical framework for neural dynamics that may be used to analyze large networks of interconnected neurons?

Let me begin with the description of neural output activity. A spike is an all-or-none event and may thus be modeled by a binary variable as was pointed out by McCulloch and Pitts [13]. It will be denoted by $S_i = \pm 1$ where i enumerates the neurons. This specific representation emphasizes the resemblance between McCulloch-Pitts neurons and Ising spins.[3] Following the conventional notation, $S_i = 1$ means that cell i is firing an action potential, and $S_i = -1$ means that the cell is quiescent.

In an alternative formulation, a quiescent cell is denoted by $S_i = 0$. Both representations are equivalent if the network parameters are transformed appropriately. In the integrate-and-fire models to be discussed in the article, the duration of action potentials is set to zero for simplicity. To obtain a nonvanishing pulse integral, a spike is modeled by a Dirac δ-function, so that formally speaking, one is dealing with a $0/\infty$-representation of action potentials.

An action potential is generated if the membrane potential u_i exceeds a firing threshold u_{thresh}. Since the trigger process operates without significant time lags, spike generation (in the ± 1-representation) may be written

$$S_i(t) = \mathrm{sgn}[u_i(t) - u_{\mathrm{thresh}}] \tag{1.3}$$

where $\mathrm{sgn}(x)$ denotes the signum function.

---

[3] The Ising model [14] provides an extremely simple and elegant description of ferromagnets and has become one of the most thoroughly studied models in solid-state physics. The formal similarity between certain extensions of this model, namely spin glasses, and neural networks such as the Hopfield model has stimulated the application of statistical mechanics to neural information processing, see also Section 1.3.6.

In most models to be analyzed in this article, the membrane potential u_i is not reset after the emission of an action potential. An important exception are networks with integrate-and-fire neurons whose precise reset mechanism is discussed in Section 1.2.3.

Some cortical areas exhibit pronounced coherent activity of many neurons on the time scale of interspike intervals, that is, $10-100$ ms [15, 16, 17]. Modeling this phenomenon requires a description of output activity in terms of single spikes, for example by using integrate-and-fire neurons.[4] In other cases, the exact timing of individual action potentials does not seem to carry any relevant information. One may then switch to a description in terms of a coarse-grained variable, the short-time averaged firing rate V. Unlike the binary outputs of McCulloch-Pitts neurons, the firing rate is a continuous variable. The firing rate varies between zero and a maximal rate $V_{\max}$ which is determined by the refractory period. Within a firing-rate description, model neurons are called "analog neurons" or "graded-response neurons."

In such a real-valued representation of output activity, the threshold operation (1.3) is replaced by an s-shaped ("sigmoid") transfer function to describe the graded response of the firing rate to changes of the membrane potential,

$$V_i(t) = g_i[u_i(t)] \tag{1.4}$$

with $g_i : \mathbb{R} \to [0, V_{\max}]$. The functions g_i can be obtained from neurophysiological measurements of the response characteristic of a cell under quasi-stationary conditions.

Once generated by a neuron, say neuron j, an action potential travels as a sharp pulse along the axon and arrives at a synapse with neuron i after some time lag τ_{ij}. The delay depends on the distance traveled by the signal and its propagation speed and may be as long as $10-50$ ms. It follows that the release of neurotransmitter at time t does not depend on the present presynaptic activity but that it should be modeled by some function whose argument is the earlier activity $S_j(t-\tau_{ij})$. Diffusion across the synaptic cleft adds a distributed delay which is usually modeled by an integral kernel with a single hump.

What remains in the modeling process is the formalization of the dendritic and somatic signal processing. The force driving the membrane potential u_i up or down will be called the local field and denoted by h_i. Formally, the local field can always be written as a power series of the synaptic input currents. The exact form of the coefficients depends on the microscopic cell properties.

Dendrites and cell bodies are complex extended objects with intricate

[4] Alternative approaches are discussed in the contribution of Gerstner and van Hemmen in this volume [18].

internal dynamics. This implies that within any accurate microscopic description, even the dendrites and soma of a *single* cell have to be represented by a large number of parameters and dynamical variables [19, 20].[5] However, such a detailed approach cannot be pursued to analyze the time evolution of large networks of highly interconnected neurons as they are found in the cerebral cortex where a neuron may be connected with up to $10,000$ other cells [21].

The theory of formal neural networks offers a radical solution to this fundamental problem. Following a long tradition in statistical physics, the theory is built upon the premise that detailed properties of single cells are not essential for an understanding of the *collective* behavior of large systems of interacting neurons: "Beyond a certain level complex function must be a result of the interaction of large numbers of simple elements, each chosen from a small variety." [22]. This point of view invites a long and controversial debate about modeling the brain and, more general, modeling complex biological systems. Such a discussion is beyond the scope and intention of the present article. Instead, I will cautiously adopt this position as a powerful working hypothesis whose neurobiological foundations require further investigation.[6] The advantage is obvious: under the assumption that the function of large neural networks does not depend on microscopic details of single cells, and knowing that in general, *many* incoming signals are necessary to trigger an action potential, it is sufficient to consider just the first terms of the power series defining the local field h_i. For the rest of this article, I will use the simplest approach and take only linear terms into account. The local field may then be written

$$h_i(t) = \sum_{j=1}^{N} \int_0^{\tau_{\max}} J_{ij}(\tau) V_j(t-\tau) d\tau + I_i^{\text{ext}}(t) \ . \tag{1.5}$$

For two-state neurons, the term $V_j(t-\tau)$ is replaced by $S_j(t-\tau)$. The weight $J_{ij}(\tau)$ describes the influence of presynaptic activity of neuron j at time $t-\tau$ on the local field of neuron i at time t. Input currents due to external stimuli are denoted by $I_i^{\text{ext}}(t)$.

The temporal details of signal transmission are reflected in the functional dependence of $J_{ij}(\tau)$ upon the delay time τ. Axonal signal propagation corresponds to a discrete time lag, diffusion processes across the synapses and along the dendrites result in delay distributions with single peaks. Distributed time lags with multiple peaks may be used to include pathways

---

[5] The argument applies to axons as well, but due to the emergent simplicity of axonal signal transport — action potentials are characterized by a dynamically stabilized, fixed pulse shape — a macroscopic description in terms of all-or-none events is justified.

[6] Unexpected support for this viewpoint comes from elaborate computer simulations of the dynamics of single cerebellar Purkinje cells [23].

via interneurons that are not explicitly represented in the model. A synapse is excitatory if $J_{ij}(\tau) > 0$ and inhibitory if $J_{ij}(\tau) < 0$. Self couplings $J_{ii}(\tau)$ that are strongly negative for small delays may be used to model refractoriness [24, 25].[7] In network models without synaptic and dendritic delays, the local field h_i is identical with the total synaptic input current to neuron i, often denoted by I_i in the neural network literature.

As shown in this section, there are three main variables to describe the activity of single neurons — the membrane potential u_i, the output activity V_i or S_i, and the local field h_i. These three variables correspond to the three main parts of a neuron — soma, axon, and dendritic tree. The strongly nonlinear dependence of V_i or S_i upon u_i captures the "decision process" of a neuron — to fire or not to fire. This decision is based on some evaluation of the weighted average h_i of incoming signals. To close the last gap in the general framework, one has to specify the dynamical relation between the membrane potential u_i and the local field h_i.

If there are no transmission delays, equations (1.3) — (1.5) contain only a single time argument and no time derivatives, that is, they do not describe any *dynamical* law. It follows that the relation between u_i and h_i has to be formulated as an evolution equation. If one opts for a description where time is treated as a discrete variable, the evolution equation will be a difference equation, otherwise a differential equation. As a first approximation, both types of dynamical description may be linear since the main source for nonlinear behavior, namely spike generation, is already described by equation (1.3) or (1.4).

1.2.2 Discrete-Time Dynamics

Within a discrete-time approach, time advances in steps of fixed length, usually taken to be unity. To obtain a consistent description, all signal delays should then be nonnegative integers. Accordingly, the temporal integral $\int_0^{\tau_{\max}} J_{ij}(\tau) S_j(t-\tau) d\tau$ in (1.5) is replaced by a sum $\sum_{\tau=0}^{\tau_{\max}} J_{ij}(\tau) S_j(t-\tau)$.

In a discrete-time model, the most straightforward dynamical relation between u_i and h_i is the shift operation

$$u_i(t+1) = h_i(t) \; . \tag{1.6}$$

At a first glance, this dynamical relation neglects any inertia of the membrane potential caused by a nonzero transmembrane capacitance. According to (1.6), the membrane potentials are just time-shifted copies of the local fields. Inertia could be included on the single-neuron level by an additive term $\alpha u_i(t)$ on the right-hand side of (1.6), however, a similar effect can be obtained through a proper choice of the update rule for the overall

[7] In some sense, the same is achieved in integrate-and-fire models where the membrane potential is explicitly reset after spike generation.

network as will be discussed at the end of this section.

For two-state neurons, equations (1.3), (1.5) and (1.6) may be combined to yield the single-neuron dynamics

$$S_i(t+1) = \mathrm{sgn}[h_i(t)] \tag{1.7}$$

where

$$h_i(t) = \sum_{j=1}^{N} \sum_{\tau=0}^{\tau_{\max}} J_{ij}(\tau) S_j(t-\tau) + I_i^{\mathrm{ext}}(t) \ . \tag{1.8}$$

The term u_{thresh} has been absorbed in I_i^{ext} without loss of generality. In passing, let me remark that in the exceptional case $h_i(t) = 0$, it is advisable to supplement (1.7) by the convention $S_i(t+1) = S_i(t)$ for (purely technical) reasons that will become apparent in Section 1.3.1.

For analog neurons, equations (1.7) and (1.8) are replaced by

$$V_i(t+1) = g_i[h_i(t)] \tag{1.9}$$

and

$$h_i(t) = \sum_{j=1}^{N} \sum_{\tau=0}^{\tau_{\max}} J_{ij}(\tau) V_j(t-\tau) + I_i^{\mathrm{ext}}(t) \ . \tag{1.10}$$

The membrane potential u_i does not appear in equations (1.7) - (1.10) anymore as the single-neuron description has been reduced from three to two variables — output activity and local field. Either one might be used as a state variable.

Neurotransmitters are released in small packages by a stochastic mechanism that includes spontaneous release at times when no spikes arrive at a synapse [26, 27]. This phenomenon, known as synaptic noise, is the most important source for stochasticity in neural signal transmission.

If one takes synaptic noise into account, the local field becomes a fluctuating quantity $h_i + \nu_i$ where ν_i denotes the stochastic contributions. The probability of spike generation is then equal to the probability that the local field exceeds the firing threshold. For two-state neurons, this probability may be written as

$$\mathrm{Prob}\big[S_i(t+1) = +1\big] = f[h_i(t)] \tag{1.11}$$

where Prob denotes probability and $f : \mathbb{R} \to [0,1]$ is a monotone increasing function.

A careful analysis of synaptic transmission reveals that under the assumption of linear dendritic processing, the stochastic variable ν_i is distributed according to a Gaussian probability distribution [22, 28]. In that case equation (1.11) can be approximated by

$$\mathrm{Prob}\big[S_i(t+1) = \pm 1\big] = \frac{1}{2}\big\{1 \pm \tanh\big[\beta h_i(t)\big]\big\} \tag{1.12}$$

where $T \equiv \beta^{-1}$ is a measure of the noise level. In the limit as $T \to 0$, one recovers the deterministic threshold dynamics (1.7). In the physics literature, the update rule (1.12) is known as Glauber dynamics [29]. It was invented as a heat-bath algorithm for the Ising model [14] and has become an important tool to analyze the collective properties of many-particle systems.

Equations (1.7) – (1.10) describe the time evolution of *individual neurons*. This leaves a number of options for the updating process at the level of the *overall network* [10].

First, there is the question of how many neurons may change their state at a time. Theoretical investigations of recurrent networks with discrete-time dynamics have almost exclusively focused on two cases: parallel dynamics (PD) and sequential dynamics (SD). In the former case, all neurons are updated in perfect synchrony which has led to the name "synchronous dynamics." In the latter case, only one neuron is picked at each time to evaluate its new state — "one-at-a-time updating" — while the activities of all other neurons remain constant. Parallel updating and sequential updating are two extreme realizations of discrete-time dynamics. Intermediate schemes will be called distributed dynamics (DD) and include block-sequential iterations where the network is partitioned into fixed clusters of simultaneously updated neurons.

Next, there is the question how groups (of one or more neurons) are selected at each time step. One may have a fixed partition of the network or one may choose random samples at each time step. Alternatively one may study selective mechanisms such as a maximum-field or greedy dynamics [30]. Here, the neuron with the largest local field opposite to its own activity is updated.[8]

Network dynamics are said to be *fair sampling* if on an intermediate time scale no neuron is skipped for the updating process on average. The terminology emphasizes the similarity with the idea of "fairness" used by the computer science community [31]. On a conceptual level, fair sampling assures that all neurons have a chance to explore the part of phase space accessible to them through their single-neuron dynamics. Most computationally useful iteration schemes are of this type. All updating schemes with a fixed partition or a random selection process are fair sampling. Exceptions may only occur in pathological situations within selective algorithms.

Finally, there is the question whether signal delays may or may not overlap, as is illustrated in Figure 1. The latter case is of utmost importance for the storage and retrieval of pattern sequences as will be discussed in Section 1.4.

---

[8] The network dynamics of integrate-and-fire neurons may also be viewed as a selective update process: only those neurons whose local fields are larger then the threshold are active for the duration of an action potential. After that time, both output S_i and membrane potential u_i are reset to their rest values.

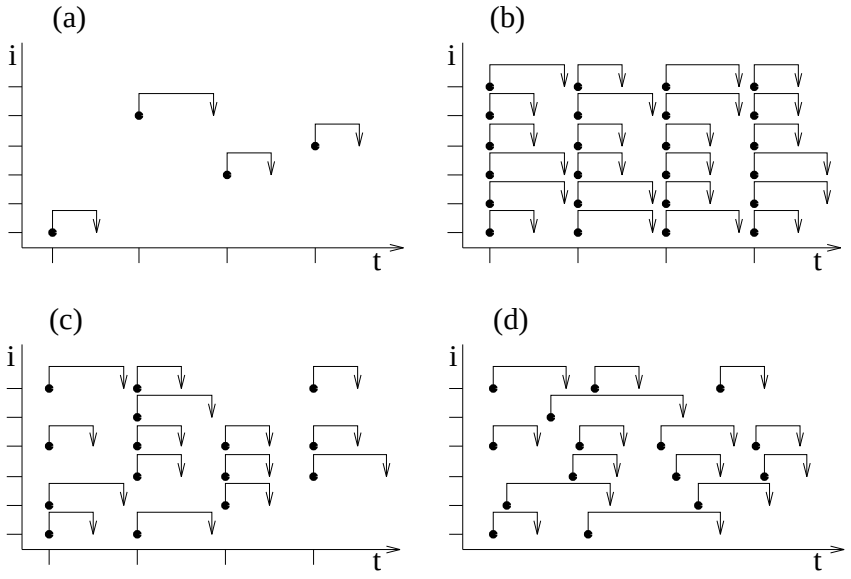

FIGURE 1.1. Schematic representation of discrete-time updating schemes. Horizontal axes represent time, ticks on the vertical axes label the neurons. Delays due to transmission and computation times are indicated by the finite duration of the updating "event" for a given neuron. Clocked networks have ticks on the time axis. (a) One-at-a-time or sequential dynamics (SD); (b) Synchronous or parallel dynamics (PD); (c) Distributed dynamics (DD): still clocked, but with arbitrary update groups at each time step; (d) Fully asynchronous dynamics, including overlapping delays.

Summarizing the above discussion, updating rules for networks with distributed discrete-time dynamics may be categorized according to the following five criteria:

1. Description of output activity: (a) discrete; (b) continuous.
2. Single-neuron dynamics: (a) deterministic; (b) stochastic.
3. Size of group to be updated at each time step:
   (a) all neurons — parallel dynamics (PD);
   (b) some neurons — distributed dynamics (DD);
   (c) one neuron — sequential dynamics (SD).
4. Selection of the update group at each time step: (a) fixed partition; (b) random sample; (c) selective choice.
5. Handling of delays: (a) overlapping not allowed; (b) overlapping allowed.

Most discrete-time descriptions appearing in the literature can be classified by these five criteria. For instance, Caianiello's model [32] uses McCulloch-Pitts neurons (rule 1a) and includes a broad distribution of transmission delays (rule 5b). All neurons are updated at the same time (rule 3a and 4a) according to a deterministic threshold operation (rule 2a). The Little model [5] differs from Caianiello's approach in that it describes single neurons as stochastic elements (rule 2b) with instantaneous interactions only (rule 5a). In the Hopfield model [4], neurons are updated one at a time (rule 3c), again without signal delays (rule 5a).

If neurons are picked in a random order, there is a nonzero chance that a neuron is skipped during an elementary cycle of the network dynamics. On the level of macroscopic order parameters, this leads to an effective inertia comparable to that generated by an additive term $\alpha u_i(t)$ in (1.6).[9]

In closing this section, let me introduce some helpful notation: networks with deterministic parallel dynamics, continuous neurons and no transmission delays (rule 1b, 2a, 3a, 4a, 5a) will be called iterated-map networks (IM), those with (a broad distribution of) transmission delays and a deterministic parallel dynamics (rule 2a, 3a, 4a, 5b) will be referred to as time-delay networks (TD).

---

[9] For a derivation of the evolution equations of macroscopic order parameters, see for example reference [33].

1.2.3 Continuous-Time Dynamics

The step size in a discrete-time description is usually identified with the duration of an action potential. This implies on the one hand that such a description cannot accomodate the time resolution required to study the synchronization of action potentials.[10] On the other hand, the feedback delay implicitly built into any discrete-time description may lead to dynamical artefacts such as spurious oscillations. To avoid both problems, one may alternatively study networks with continuous-time dynamics.

Graded-Response Neurons

Membrane potentials of real neurons are subject to leakage currents due to the finite resistivity of biological membranes. Once charged by a short input current modeled by the local field $h_i(t)$, the membrane potential $u_i(t)$ of cell i relaxes to some rest value which is set to zero for simplicity.

The physics of charging and leakage is best captured by the linear first-order differential equation

$$C\frac{d}{dt}u_i(t) = -R^{-1}u_i(t) + h_i(t) \; . \tag{1.13}$$

Here C denotes the input capacitance of a neuron and R is its transmembrane resistance. Model neurons whose membrane potential changes according to the differential equation (1.13) will be called graded-response neurons (GR).

Inserting equation (1.5) into (1.13), the time evolution of graded-response neurons may be written

$$C\frac{d}{dt}u_i(t) = -R^{-1}u_i(t) + \sum_{j=1}^{N}\int_0^{\tau_{\max}} J_{ij}(\tau)V_j(t-\tau)d\tau + I_i^{\rm ext}(t) \tag{1.14}$$

where, as in Section 1.2.1, the output activity V_j depends on the membrane potential u_j through the nonlinear response characteristic (1.4).

Similar to the discrete-time dynamics considered in Section 1.2.2, one of the original three variables to describe neural activity has become superfluous. In Section 1.2.2, the membrane potential $u_i(t)$ was expressed through the (time-shifted) local field $h_i(t-1)$, now the local field $h_i(t)$ has been replaced by the membrane potential $u_i(t)$ and its time derivative $\dot{u}_i(t)$.

---

[10] Decreasing the step size leads to a complication in the mathematical formulation because one is forced to introduce effective delayed interactions if one wants to assure that action potentials last for multiple elementary time steps.

Integrate-and-Fire Neurons

Below the firing threshold, (leaky) integrate-and-fire neurons operate in the same way as graded-response neurons (1.13). However, when the membrane potential of a cell reaches the threshold $u_{\rm thresh}$, the cell produces an action potential and resets its potential to $u_{\rm reset}$. For convenience, units can be chosen such that $u_{\rm thresh} = 1$ and $u_{\rm reset} = 0$.

Assuming vanishing signal delays and action potentials of negligible duration, the local field $h_i(t)$ of neuron i is then given by

$$h_i(t) = \sum_j J_{ij} f_j(t) + I_i^{\rm ext}(t) \tag{1.15}$$

where the instantaneous firing rate $f_j(t)$ is a sum of Dirac δ-functions

$$f_j(t) = \sum_n \delta(t - t_j^n) \tag{1.16}$$

and the t_j^n are the times at which neuron j generates an action potential. Throughout the remaining sections on integrate-and-fire neurons, the external input $I_i^{\rm ext}(t)$ is assumed to be constant in time, $I_i^{\rm ext}(t) = I_i^{\rm ext}$.

The general behavior of the system is now as follows. While none of the neurons is producing an action potential, equation (1.13) can be integrated and yields

$$u_i(t) = [u_i(t_0^+) - RI_i^{\rm ext}]{\rm e}^{-\frac{(t-t_0)}{RC}} + RI_i^{\rm ext} \qquad \text{for } t \geq t_0 \tag{1.17}$$

where t_0 denotes the last firing time. When the potential u_j of neuron j reaches 1 (the threshold) it drops instantaneously to 0. At the same time, the potential u_i of each neuron i to which j makes a synapse is increased by J_{ij}.

Because the duration of action potentials and of synaptic currents have been set equal to zero, the description given so far contains an ambiguity. To which value should neuron i be reset if at time t an action potential is produced by cell j, if the synapse from j to i is excitatory, $J_{ij} > 0$, and if $u_i(t^-) > 1 - J_{ij}$? For in this case, the action potential will raise u_i above 1, and cell i should generate its action potential during the flow of synaptic current produced by the synapse J_{ij}. When synaptic (and dendritic) time constants of the nerve cells to be modeled are longer than the duration of action potentials, what should actually happen in the model is that cell j should fire when its potential reaches $u_{\rm thresh} = 1$, and the synaptic current from synapse J_{ij} which arrives after i fires should be integrated to yield a positive potential (relative to $u_{\rm reset}$) afterward. Thus, if cell j fires first and at time t, and that event evokes a firing of neuron i, then after both action potentials have been generated, the two membrane potentials should be

$$u_j(t^+) = J_{ji} \tag{1.18}$$

and

$$u_i(t^+) = u_i(t^-) + J_{ij} - 1 \ . \tag{1.19}$$

The first equation represents the fact that j fired first when $u_j = 1$, was reset to 0, and when subsequently neuron i generated its action potential, this changed the potential of j to J_{ji}. The second equation represents the fact that i fired second, reduced its potential by 1 when it did so, but received the synaptic current J_{ij} when neuron j fired.

The updating rule can be generalized to a large network of neurons by the following algorithm. As the potentials all increase with time, a first neuron j reaches $u_j = 1$. Reset that potential to zero. Then change the potential of each neuron i by J_{ij}. If, following this procedure, some of the potentials become greater than 1, pick the neuron with the largest potential, say neuron k, and decrease its potential by 1.[11] Then change the potential of each neuron l by J_{lk}. Continue the procedure until no membrane potential is greater than 1. Then "resume the flow of time," and again let each potential u_i increase according to equation (1.17).

This deterministic algorithm preserves the essence of the idea that firing an action potential carries a neuron from u_{thresh} to u_{reset}, and effectively apportions the synaptic current into a part which is necessary to reach threshold, and a part which raises the potential again afterward. Because the firing of one neuron can set off the instantaneous firing of others, this model can generate events in which many neurons are active simultaneously.

When synaptic (and dendritic) time constants are shorter than the duration of an action potential, all contributions from the synaptic current that arrive during spike generation are lost, and equation (1.19) should be replaced by $u_i(t^+) = J_{ij}$. Generalizing from these two extreme cases, (1.19) becomes

$$u_i(t^+) = u_i(t^-) + \gamma(J_{ij} - 1) \tag{1.20}$$

with $0 \leq \gamma \leq 1$.

For models with $\gamma = 1$, the order in which the neurons are updated in an event in which several neurons fire at once does not matter as long as $J_{ij} \geq 0$. For these cases any procedure for choosing the updating sequence of the neurons at or above threshold will yield the same result because the reset is by a fixed negative amount (here: -1) regardless of whether immediately prior to reset $u_i = 1$ or $u_i > 1$.

If in addition to choosing $\gamma = 1$, the limit $R \to \infty$ is considered, one is dealing with perfectly integrating cells. For a network of such neurons, the cumulative effects of action potentials and slow membrane dynamics commute if $J_{ij} \geq 0$. This makes the model formally equivalent to a class of

[11] If several neurons exhibit the same, maximum potential, one may use some fixed, random, or selective update order to pick one of them.

"Abelian avalanche" models [34, 35]. Closely related earthquake models and (discrete-time) "sandpile models" relax to a critical state with fluctuations on all length-scales, a phenomenon known as "self-organized criticality" [36].

The similarity between the microscopic dynamics of such model systems and networks of integrate-and-fire neurons has led to speculations about a possible biological role of the stationary self-organized critical state [37, 38, 39]. However, whereas for earthquakes, avalanches, and sandpiles, the main interest is in the properties of the stationary state, for neural computation, it is the convergence process itself which does the computation and is thus of particular interest. Furthermore, computational decisions must be taken rapidly, and in any event the assumption of constant input from other cortical areas implicit in all models breaks down at longer times [40, 41].

1.2.4 Hebbian Learning

The previous sections focused on the dynamics of neural activity. Synaptic efficacies were treated as time-independent parameters. Real synapses, however, are often modifiable. As postulated by D.O. Hebb [42], their strengths may change in response to correlated pre- and postsynaptic activity: "When an axon of cell A is near enough to excite cell B and *repeatedly* or *persistently* takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B is increased."

Hebbian plasticity has long been recognized as a key element for associative learning [43].[12] How should it be implemented in a formal neural network that might include transmission delays?

Hebbian learning is local in both space and time: changes in synaptic efficacies depend only on the activity of the presynaptic neuron and the evoked postsynaptic response. Within the present framework, presynaptic activity is described by the axonal output V_j or S_j. Which neural variable should be chosen to model the postsynaptic response?

Neurophysiological experiments demonstrate that postsynaptic spiking is *not* required to induce long-term potentiation (LTP) of synaptic efficacies — "a critical amount of postsynaptic depolarization is normally required to induce LTP in active synapses, but sodium spikes do not play an essential role in the LTP mechanism" [45]. This result implies that the postsynaptic response is best described by the local field h_i — it represents the dendritic potential and is not influenced by the detailed dynamics of the cell body (u_i) or the spike generating mechanism (V_i or S_i).

Let us now study a discrete-time system where delays arise due to the

[12] Various hypotheses about the microscopic mechanisms of synaptic plasticity are the subject of an ongoing discussion [44].

finite propagation speed of axonal signals, and focus on a connection with delay τ between neurons j and i. Originally, Hebb's postulate was formulated for excitatory synapses only, but for simplicity, it will be applied to all synapses of the model network.

A presynaptic action potential that arrives at the synapse at time t was generated at time $t-\tau$. Following the above reasoning, $J_{ij}(\tau)$ should therefore be altered by an amount that depends on $V_j(t-\tau)$ and $h_i(t)$, most simply their product

$$\Delta J_{ij}(\tau) \propto h_i(t) V_j(t-\tau) \Delta t \ . \tag{1.21}$$

The bilinear expression (1.21) does not cover saturation effects. They could be modeled by an additional decay term $-\alpha J_{ij}(\tau)\Delta t$ on the right-hand side of (1.21).

The combined equations (1.3) – (1.5) and (1.21) describe a "double dynamics" where both neurons and synapses change in time. In general, such a system of coupled nonlinear evolution equations cannot be analyzed using Lyapunovs' direct method although there are some interesting counterexamples [46]. To simplify the analysis, one usually splits the network operation in two phases — learning and retrieval. For the learning phase, one frequently considers a "clamped" scheme where neurons evolve according to external inputs only, $h_i(t) = I_i^{\rm ext}(t)$. Once the learning sessions are over the $J_{ij}(\tau)$ are kept fixed.

In the following, I focus on deterministic discrete-time McCulloch-Pitts neurons in a clamped scheme with $I_i^{\rm ext}(t) = \pm 1$. This simplification implies that $S_j(t+1) = I_j^{\rm ext}(t)$. Starting with a *tabula rasa*, $J_{ij}(\tau) = 0$, one obtains after P learning sessions, labeled by μ and each of duration D_μ,

$$J_{ij}(\tau) = \varepsilon(\tau) N^{-1} \sum_{\mu=1}^{P} \sum_{t_\mu=1}^{D_\mu} I_i^{\rm ext}(t_\mu) I_j^{\rm ext}(t_\mu - 1 - \tau) \equiv \varepsilon(\tau) \tilde{J}_{ij}(\tau) \ . \tag{1.22}$$

The parameters $\varepsilon(\tau)$ model morphological characteristics of the axonal delay lines; N^{-1} is a scaling factor useful for the theoretical analysis. Let me mention that an input sequence should be offered already $\tau_{\rm max}$ time steps before the learning session starts so that all variables in (1.22) are well defined. According to (1.22) synapses act as microscopic feature detectors during the learning sessions: they measure and store correlations of the taught sequences in both space (i,j) and time (τ). This leads to a resonance phenomenon where connections with delays that approximately match the time course of the external input receive maximum strength. Note that these connections are also the ones that would support a stable sequence of the same duration. Thus, due to a subtle interplay between external stimulus and internal architecture (distribution of τ's), the Hebb rule (1.22), which *prima facie* appears to be instructive in character, exhibits in fact pronounced selective characteristics [47].

An external stimulus encoded in a network with a *broad* distribution of transmission delays enjoys a rather multifaceted representation. Synaptic couplings with delays that are short compared to the typical time scale of single patterns within the taught sequence are almost symmetric in the sense that $J_{ij}(\tau) \approx J_{ji}(\tau)$. These synapses encode the individual patterns of the sequence as *unrelated static objects*. On the other hand, synapses with transmission delays of the order of the duration of single patterns of the sequence are able to detect the transitions between patterns. The corresponding synaptic efficacies are asymmetric and establish various temporal relations between the patterns, thereby representing the complete sequence as *one dynamic object*.

Let me remark that the interplay between neural and synaptic dynamics, and in particular the role of transmission delays, has been a subject of intensive research [32, 42, 48, 49]. The full consequences for the learning and retrieval of temporal associations have, however, been explored only recently.

As a special case of (1.22), consider the Hebbian learning of static patterns, $I_i^{\rm ext}(t_\mu) = \xi_i^\mu$, offered during learning sessions of equal duration $D_\mu = D$ to a network with a uniform delay distribution. For mathematical convenience, the distribution is taken to be $\varepsilon(\tau) = D^{-1}$. In this case, (1.22) yields synaptic strengths that are independent of the delay τ,

$$J_{ij}(\tau) = J_{ij} = N^{-1} \sum_{\mu=1}^{P} \xi_i^\mu \xi_j^\mu \ , \tag{1.23}$$

and symmetric,

$$J_{ij} = J_{ji} \ . \tag{1.24}$$

The synaptic symmetry (1.24) plays a key role for the construction of Lyapunov functions as will be shown in the following sections.

Another kind of symmetry arises if all input sequences $I_i^{\rm ext}(t_\mu)$ are cyclic with equal periods $D_\mu = D$. If one defines patterns ξ_{ia}^μ by $\xi_{ia}^\mu = I_i^{\rm ext}(t_\mu = a)$ for $0 \le a < D$, one obtains from (1.22)

$$\tilde{J}_{ij}(\tau) = N^{-1} \sum_{\mu=1}^{P} \sum_{a=0}^{D-1} \xi_{ia}^\mu \xi_{i,a-1-\tau}^\mu \ . \tag{1.25}$$

Note that the synaptic strengths are now in general asymmetric. They do, however, obey the symmetry $\tilde{J}_{ij}(\tau) = \tilde{J}_{ij}(D-(2+\tau))$. For all networks whose *a priori* weights $\varepsilon(\tau)$ satisfy $\varepsilon(\tau) = \varepsilon(D-(2+\tau))$, this leads to an "extended synaptic symmetry" [50, 51],

$$J_{ij}(\tau) = J_{ij}(D-(2+\tau)) \ , \tag{1.26}$$

extending the previous symmetry (1.24) in a natural way to the temporal domain. This type of synaptic symmetry allows the construction of a

Lyapunov function for time-delay networks as will be explained in Section 1.4.1.

1.3 Fixed Points

This section focuses on the storage of static patterns in networks with instantaneous interactions. It will be shown that under certain conditions for the model parameters, various network dynamics exhibit the same long-time behavior: they relax to fixed points only.

Feedback networks with fixed-point attractors can be made potentially useful devices for associative computation as soon as on knows how to embed desired activity patterns as attractors of the dynamics. In such circumstances, an initial state or "stimulus" lying in the basin of attraction of a stored "memory" will spontaneously evolve towards this attractor. Within a biological context, the arrival at the fixed point may be interpreted as a cognitive event, namely the "recognition of the stimulus."

The hypothesis that the brain utilizes fixed-point attractors to perform associative information processing has led to quantitative predictions [52] which are in good agreement with neurophysiological measurements [53]. However, even if the hypothesis was refuted in its literal sense, it would nevertheless continue to provide an important conceptual tool to think about neural information processing.

1.3.1 Sequential Dynamics: Hopfield Model

Hopfield's original approach [4] is based on McCulloch-Pitts neurons with discrete-time dynamics, instantaneous interactions and constant external stimuli. Neurons are updated one at a time, either according to a deterministic threshold operation (1.7) or probabilistic Glauber dynamics (1.12). In the original model, neurons are chosen in a random sequential manner but in simulations, the update order is often fixed in advance, corresponding to a quenched random selection. Within the classification scheme of Section 1.2.2, the Hopfield model is thus characterized by rules 1a, 3c, and 5a.

If the single-neuron dynamics are deterministic, the time evolution of the network is a special realization of (1.7),(1.8) and may be written

$$S_k(t+1) = \mathrm{sgn}[h_k(t)] \tag{1.27}$$

where k is the index of the neuron updated at time t and

$$h_k(t) = \sum_j J_{kj} S_j(t) + I_k^{\mathrm{ext}} \ . \tag{1.28}$$

All other neurons remain unchanged, $S_j(t+1) = S_j(t)$ for $j \neq k$.

What can be said about the global dynamics generated by equations (1.27) and (1.28)? Consider the quantity

$$L_{\rm SD} = -\frac{1}{2}\sum_{i,j=1}^{N} J_{ij}S_iS_j - \sum_{i=1}^{N} I_i^{\rm ext}S_i \ . \tag{1.29}$$

The change of $L_{\rm SD}$ in a single time step, $\Delta L_{\rm SD}(t) \equiv L_{\rm SD}(t+1) - L_{\rm SD}(t)$, is

$$\begin{aligned}\Delta L_{\rm SD}(t) = \ &- \ \frac{1}{2}\sum_{i,j=1}^{N} J_{ij}[S_i(t+1)S_j(t+1) - S_i(t)S_j(t)] \\ &- \ \sum_{i=1}^{N} I_i^{\rm ext}[S_i(t+1) - S_i(t)] \ . \end{aligned}\tag{1.30}$$

Assume again that neuron k is updated at time t. The difference $\Delta S_j(t) \equiv S_j(t+1) - S_j(t)$ equals zero or ± 2 if $j = k$ and vanishes otherwise. For the special case where the synaptic efficacies satisfy the symmetry condition (1.24), one obtains

$$\begin{aligned}\Delta L_{\rm SD}(t) &= \Delta S_k(t)J_{kk}S_k(t) - \Delta S_k(t)[\sum_{j=1}^{N} J_{kj}S_j(t) + I_k^{\rm ext}] \\ &= -\frac{1}{2}J_{kk}[\Delta S_k(t)]^2 - \Delta S_k(t)h_k(t). \end{aligned}\tag{1.31}$$

According to (1.27) and the remark following (1.8), neuron k does not change its state if $h_k(t)S_k(t) \geq 0$. If this condition is not fulfilled, the neuron flips and $\Delta S_k(t) = 2S_k(t+1)$. The change of $L_{\rm SD}$ may then be written

$$\begin{aligned}\Delta L_{\rm SD}(t) &= -2[J_{kk} + S_k(t+1)h_k(t)] \\ &= -2[J_{kk} + |h_k(t)|] \ . \end{aligned}\tag{1.32}$$

The last line follows from the evolution equation (1.27) and the identity $|a| = a\ \mathrm{sgn}(a)$. Equation (1.32) proves that $L_{\rm SD}$ is nonincreasing along every solution if the self couplings J_{ii} are nonnegative.[13] As a finite sum of finite terms, $L_{\rm SD}$ is bounded. If $J_{ii} \geq 0$ for all neurons, $L_{\rm SD}(t)$ has to approach a limit as $t \to \infty$. Furthermore, $\Delta L_{\rm SD}(t)$ vanishes only if the

[13] This condition is satisfied in Hopfield's original model where all self couplings are set to zero.

neuron updated at time t does not change its state.[14] This proves that the Hopfield network relaxes to fixed-point solutions only. According to (1.27) and (1.28) these equilibria satisfy

$$S_i = \mathrm{sgn}[\sum_j J_{ij} S_j + I_i^{\mathrm{ext}}] \quad \text{for all } i \ . \tag{1.33}$$

The results obtained may be summarized as follows:

If the synaptic efficacies J_{ij} satisfy the symmetry condition (1.24) and if the self interactions J_{ii} are nonnegative, then the dynamics of the Hopfield model (1.27),(1.28) admit the Lyapunov function (1.29) and converge to fixed points (1.33) only.

Let me clarify a potentially confusing point. For neural networks with McCulloch-Pitts neurons, the state space consists of the corners of an N-dimensional hypercube $\{-1,+1\}^N$, also known as Hamming space. In this *discrete* space, the smallest state change possible is a single spin flip, $S_i \to -S_i$. As a consequence, the system may converge to fixed points that are *not* stable with respect to activity changes of single neurons, in the sense that a single spin flip made to a fixed-point solution could actually lower L. For instance, consider a network where for some neuron i, the self interaction J_{ii} dominates possible contributions from other neurons, $J_{ii} > \sum_{j\neq i} |J_{ij}|$. In such a case, the initial value of S_i will never be changed, independent of its sign. The earlier results about network convergence continue to hold, that is, the system evolves towards fixed-point solutions only, but those are not necessarily local minima of L in the discrete-space sense.

1.3.2 Parallel Dynamics: Little Model

The Little model [5] uses the most simple discrete-time dynamics conceivable: it is a network of McCulloch-Pitts neurons, updated in parallel using instantaneous interactions only (rule 1a, 3a, 4a, and 5a). Within a deterministic description of single neurons (rule 2a), the time evolution of the network is given by

$$S_i(t+1) = \mathrm{sgn}[h_i(t)] \quad \text{for all } i \tag{1.34}$$

where

$$h_i(t) = \sum_j J_{ij} S_j(t) + I_i^{\mathrm{ext}} \ . \tag{1.35}$$

Except for the update order, equations (1.34) and (1.35) are identical to (1.27) and (1.28). Accordingly, the fixed-point solutions of the Little model

[14] For zero self coupling J_{kk}, and in the exceptional case $h_k(t) = 0$, $\Delta L_{\mathrm{SD}}(t)$ vanishes for any update rule, even if one chooses $S_k(t+1) = -S_k(t)$ if $h_k(t) = 0$. However, if one sets $S_k(t+1) = S_k(t)$ as mentioned in Section 1.2.1, $\Delta L_{\mathrm{SD}}(t) = 0$ implies $\Delta S_k(t) = 0$ as desired.

are the same as those of the Hopfield model, given by (1.33). Are there additional time-dependent attractors?

For simplicity, only the case $I_i^{\rm ext} = 0$ will be analyzed in this section. Nonzero inputs will be treated in Sections 1.3.4 and 1.3.5. As in Section 1.3.1, let me focus on networks with symmetric couplings and study the time evolution of a suitable auxiliary function

$$L_{\rm PD} = -\sum_{i=1}^{N} |h_i| = -\sum_{i=1}^{N} h_i \ {\rm sgn}(h_i) \ . \tag{1.36}$$

If one evaluates this expression along a solution generated by the network dynamics (1.34) and (1.35), one obtains

$$\begin{aligned} L_{\rm PD}(t) &= -\sum_{i=1}^{N} h_i(t) S_i(t+1) \\ &= -\sum_{i,j=1}^{N} J_{ij} S_j(t) S_i(t+1) \ . \end{aligned} \tag{1.37}$$

Using the synaptic symmetry (1.24), the last line may also be written

$$L_{\rm PD}(t) = -\sum_{j=1}^{N} S_j(t) h_j(t+1) \ . \tag{1.38}$$

The difference $\Delta L_{\rm PD}(t) \equiv L_{\rm PD}(t+1) - L_{\rm PD}(t)$ is then

$$\begin{aligned} \Delta L_{\rm PD}(t) &= -\sum_{i=1}^{N} |h_i(t+1)| + \sum_{i=1}^{N} S_i(t) h_i(t+1) \\ &= -\sum_{i=1}^{N} [S_i(t+2) - S_i(t)] h_i(t+1) \end{aligned} \tag{1.39}$$

where (1.34) has been used to obtain the last equation.

Like $L_{\rm SD}$, the function $L_{\rm PD}$ is bounded. Evaluated along any solution of (1.34) and (1.35), $L_{\rm PD}$ is nonincreasing because the right-hand side of (1.39) is nonpositive; the product $S_i(t)h_i(t+1)$ is $\pm h_i(t+1)$ and thus smaller or at most equal to $|h_i(t+1)|$. Consequently, $\Delta L_{\rm PD}(t)$ has to approach zero as $t \to \infty$. $\Delta L_{\rm PD}(t)$ vanishes only if the system settles into a state with $S_i(t+2) = S_i(t)$ for all i, that is, a fixed-point solution (1.33) or a limit cycle of period two. In the latter case, some neurons switch between firing and quiescence at every time step, all other neurons remain in one activity state:

Assume that the synaptic couplings J_{ij} satisfy the symmetry condition (1.24). Then the dynamics of the Little model (1.34),(1.35) admit the Lyapunov function (1.36) and converge to fixed points (1.33) or period-two oscillations.

As will be shown in Section 1.3.5, the oscillating solutions can be excluded under additional assumptions for the synaptic couplings.

1.3.3 Continuous Time: Graded-Response Neurons

This section deals with the continuous-time dynamics of neural networks composed of analog neurons without signal delays. The network dynamics (1.14) reduce to a set of coupled ordinary differential equations,

$$C\frac{d}{dt}u_i = -R^{-1}u_i + \sum_{j=1}^{N} J_{ij}V_j + I_i^{\text{ext}} \ , \tag{1.40}$$

where

$$V_i = g_i(u_i) \ . \tag{1.41}$$

Since the dynamical variables u_i and V_i in (1.40) are taken at equal times, all temporal arguments have been omitted.

The input-output relation g_i will be called "sigmoid" if it is increasing, differentiable and grows in magnitude more slowly than linearly for large positive or negative arguments. The maximum slope of g_i will be referred to as the gain γ_i of neuron i. The nonlinearity is often modeled by a hyperbolic tangent, $g_i(u_i) = \frac{1}{2}[1 + \tanh(\gamma_i u_i)]$. In the high-gain limit $\gamma_i \to \infty$, one obtains a 0/1 representation of neural activity. It can be mapped onto Ising spins [14] through the identification $S_i = 2V_i - 1$.

Cohen and Grossberg [6] and Hopfield [7] studied the global behavior of networks with graded-response neurons, sigmoid response functions, and symmetric synapses. They used Lyapunov functions of the form

$$L_{\text{GR}} = -\frac{1}{2}\sum_{i,j=1}^{N} J_{ij}V_iV_j - \sum_{i=1}^{N} I_i^{\text{ext}}V_i + \sum_{i=1}^{N} R^{-1}G_i(V_i) \tag{1.42}$$

where the functions $G_i(V_i)$ are given by

$$G_i(V_i) = \int_0^{V_i} g_i^{-1}(x)dx \ . \tag{1.43}$$

The last expression is well defined because sigmoid nonlinearities are strictly monotone by definition. Since sigmoid functions grow less than linearly for large absolute arguments, the functions $G_i(V_i)$ increase faster than V_i^2 as $V_i \to \pm\infty$. The function L_{GR} is therefore bounded below.

Let us compute the time derivative of $L_{\rm GR}$ along a solution of the network dynamics. Using the synaptic symmetry (1.24), one obtains

$$\begin{aligned}\frac{d}{dt}L_{\rm GR}(t) &= -\sum_{i=1}^{N}[\sum_{j=1}^{N} J_{ij}V_j + I_i^{\rm ext} - R^{-1}u_i]\frac{dV_i}{dt} \\ &= -\sum_{i=1}^{N} C^{-1}\frac{du_i}{dt}\frac{dV_i}{dt} \\ &= -\sum_{i=1}^{N} C^{-1}\Big(\frac{du_i}{dt}\Big)^2\frac{dg_i}{du_i} \le 0\ . \end{aligned} \tag{1.44}$$

The formula proves that the function $L_{\rm GR}$ is nonincreasing along every trajectory. The time derivative vanishes only at equilibria, given by

$$V_i = g_i[R\sum_j J_{ij}V_j + RI_i^{\rm ext}], \tag{1.45}$$

or at network states, where $dg_i/du_i = 0$ for all i. If, however, the latter states do not satisfy (1.45), the system will continue to evolve according to (1.40),(1.41). The final result may be stated as follows:

Suppose that the synaptic efficacies in a network of graded-response neurons (1.40),(1.41) respect the symmetry condition (1.24) and that the input-output relations are sigmoid. Then the network dynamics admit the Lyapunov function (1.42) and relax to fixed-point solutions (1.45) only.

A comparison of the Lyapunov function $L_{\rm GR}$ with the Lyapunov function $L_{\rm SD}$ provides some hints about how to construct Lyapunov functions for systems with sigmoid input-output characteristics: the additional term $\sum_i R^{-1}G_i(V_i)$ dominates the quadratic term $-\frac{1}{2}\sum_{i,j} J_{ij}V_iV_j$ for large V_i if the g_i are sigmoid. Consequently, the function $L_{\rm GR}$ is bounded below even if the V_i are not.[15] Furthermore, the term $\sum_i R^{-1}G_i(V_i)$ is constructed in such a way that its partial derivative with respect to V_i supplies the term $R^{-1}u_i$ which makes it possible to insert the evolution equation (1.40) into (1.44). Similar ideas will be applied in Sections 1.3.4 and 1.3.5 to analyze discrete-time networks with sigmoid nonlinearities.

1.3.4 Iterated-Map Networks

Feedback networks with deterministic analog elements and synchronous discrete-time updating have been studied for a long time [32, 48, 49]. For vanishing signal delays and fixed inputs, the network dynamics (1.9),(1.10) become

[15] It should be noted that if a Lyapunov function is not globally bounded below, it might still be used for a local analysis.

$$V_i(t+1) = g_i[h_i(t)] \quad \text{for all } i \tag{1.46}$$

where

$$h_i(t) = \sum_{j=1}^{N} J_{ij} V_j(t) + I_i^{\text{ext}} \ . \tag{1.47}$$

Systems described by (1.46) and (1.47) have been called "iterated-map networks" [8]. Their fixed points coincide with those of graded-response networks (1.45) once one sets $R = 1$.

If the input-output functions g_i are threshold functions, $g_i(u_i) = \text{sgn}(u_i)$, one recovers the Little model (1.34),(1.35). This connection indicates that one may find a Lyapunov function for iterated-map networks by combining appropriate parts of the Lyapunov function for the Little model with that for networks of graded-response neurons.

Let us follow the approach of Marcus and Westervelt [8] and study the time evolution of the function

$$\begin{aligned} L_{\text{IM}}(t) = &- \sum_{i,j=1}^{N} J_{ij} V_i(t) V_j(t-1) - \sum_{i=1}^{N} I_i^{\text{ext}} [V_i(t) + V_i(t-1)] \\ &+ \sum_{i=1}^{N} [G_i(V_i(t)) + G_i(V_i(t-1))] \ , \end{aligned} \tag{1.48}$$

where $G_i(V_i)$ is defined as in (1.43).

Apart from a global time shift, the first term in (1.48) corresponds to L_{PD} as can be seen from equation (1.37); the other terms should be compared with the second and third term in (1.42). Notice that unlike L_{PD} (1.36), the function L_{IM} is written as an explicitly time-dependent function with temporal arguments t and $t-1$. In principle, one could use the evolution equations (1.46) and (1.47) and replace $V_i(t)$ by $g_i[\sum_{j=1}^{N} J_{ij} V_j(t-1) + I_i^{\text{ext}}]$ so as to obtain a description that involves a single time argument only. However, since we are mainly interested in the evaluation of L_{IM} along trajectories, the shorter definition (1.48) suffices.

Under the assumption of synaptic symmetry (1.24), the temporal difference $\Delta L_{\text{IM}}(t) \equiv L_{\text{IM}}(t+1) - L_{\text{IM}}(t)$ is

$$\Delta L_{\text{IM}}(t) = -\sum_i h_i(t) \Delta_2 V_i(t) + \sum_i [G_i(V_i(t+1)) - G_i(V_i(t-1))] \tag{1.49}$$

where

$$\Delta_2 V_i(t) \equiv V_i(t+1) - V_i(t-1) \tag{1.50}$$

is the change of V_i over 2 time steps.

The right-hand side of (1.49) is zero if $\Delta_2 V_i(t) = 0$ for all i. Let us analyze the case where $\Delta_2 V_i(t) \neq 0$ for at least some i. For sigmoid g_i,

g_i^{-1} is single valued and increasing. Consequently, G_i is strictly convex. Through a Taylor expansion of $G_i(V_i(t-1))$ around $V_i(t+1)$, one obtains

$$G_i(V_i(t+1)) - G_i(V_i(t-1)) < \Delta_2 V_i(t) G_i'(V_i(t+1)) \ . \tag{1.51}$$

For a graphical illustration of the inequality, see the left part of Figure 2.

Inserting the identity

$$G_i'(V_i(t+1)) = g_i^{-1}(V_i(t+1)) = h_i(t) \tag{1.52}$$

and (1.51) into (1.49), one arrives at the expression

$$\Delta L_{\rm IM}(t) \leq 0 \tag{1.53}$$

where the strict inequality holds if $\Delta_2 V_i(t) \neq 0$ for at least one neuron.

As demonstrated in the last section, the functions $G_i(V_i)$ increase faster than V_i^2 for large $|V_i|$. This result implies that $L_{\rm IM}$ is bounded below. As shown by (1.53), the function $L_{\rm IM}$ strictly decreases along any solution of (1.46),(1.47) unless $\Delta_2 V_i(t) = 0$ for all neurons. The derivation may be summarized in the following way:

Assume that the synaptic efficacies in an iterated-map network (1.46), (1.47) are symmetric (1.24) and that the nonlinearities are sigmoid. Then the network dynamics admit the Lyapunov function (1.48) and relax to fixed points solutions (1.45) or period-two oscillations.

In closing this section, let me briefly discuss antisymmetric synaptic couplings,

$$J_{ij} = -J_{ji} \ . \tag{1.54}$$

The derivation of Section 1.3.2 for the Little model (with no external input) shows that if (1.54) holds, one obtains

$$\Delta L_{\rm PD}(t) = -\sum_{i=1}^{N} [S_i(t+2) + S_i(t)] h_i(t+1) \ . \tag{1.55}$$

In this case the network approaches solutions that satisfy $S_i(t+2) = -S_i(t)$, that is, special limit cycles with period 4 [54].

It is left as an exercise to verify the same result for iterated maps without external input. Here, an additional condition is required, namely that the input-output characteristics have to be odd functions, $g_i(V_i) = -g_i(-V_i)$. The interested reader may also try to construct Lyapunov functions for more general systems. In particular, he or she could look at two problems. (i) What kind of time-varying external stimuli can be incorporated into the Lyapunov function of the Little model if one focuses on antisymmetric couplings? (ii) Are there Lyapunov functions for neural networks with McCulloch-Pitts neurons, antisymmetric couplings and sequential dynamics with fixed update order?

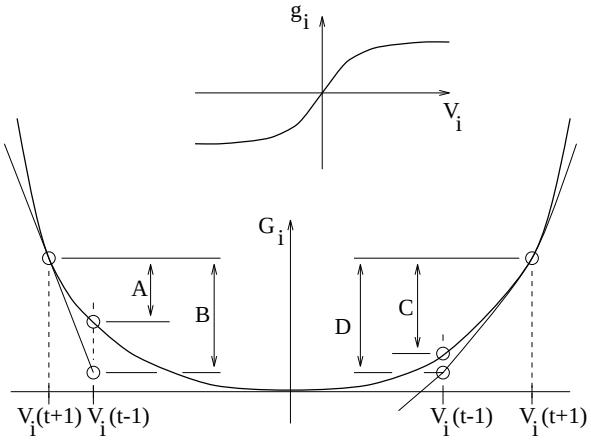

FIGURE 1.2. Illustration of the inequalities (1.51) and (1.58) for a sigmoid input-output function $g_i(V_i)$. The convex function $G_i(V_i)$ is defined in equation (1.43). The straight line on the left-hand side and the parabola on the right-hand side are tangent to $G_i(V_i)$. The inequality (1.51) is the statement $A < B$, the inequality (1.58) is the statement $C < D$.

1.3.5 Distributed Dynamics

In this section discrete-time updating schemes are considered that generalize beyond the Hopfield and Little models on both the single-neuron and network levels. Neurons are described by continuous variables with deterministic single-cell dynamics, that is they fall into class 1b and 2a in the scheme of Section 1.2.2. McCulloch-Pitts neurons with stochastic Glauber dynamics are discussed in Section 1.3.6. For the network dynamics, all choices of rules 3 and 4 are allowed that are fair sampling and do not lead to overlapping delays (rule 5a). The network dynamics are thus defined by a set of coupled nonlinear discrete-time equations

$$V_i(t+1) = \begin{cases} g_i(\sum_{j=1}^{N} J_{ij} V_j(t) + I_i^{\rm ext}) & \text{if } i \text{ is in } U(t), \\ V_i(t) & \text{otherwise.} \end{cases} \tag{1.56}$$

Here, $U(t)$ denotes the group of neurons updated at time t. The distributed dynamics (1.56) reduce to block-sequential algorithms studied by Goles-Chacc et al. [9] if one considers McCulloch-Pitts neurons and fixed update groups U_k, $k = 0, 1, \ldots, K-1$ with $U(t) = U_{t(\text{modulo } K)}$.

There are a number of reasons to study partially parallel network dynamics such as (1.56). First, one may achieve a better understanding of the essential ingredients needed to construct feedback networks that possess

fixed-point attractors only. Second, distributed dynamics map naturally on the architecture of parallel computers or computer networks. Third, the evolution equations (1.56) extend iterative methods that have been developed within the computer-science community to solve nonlinear systems of equations [55, 56, 57, 58] to systems with *noncontracting* functions and multiple solutions.

What can be said about the long-time behavior of neural networks with distributed dynamics? As in Sections 1.3.1 – 1.3.4, let us assume that the synaptic couplings are symmetric (1.24) and that the input-output characteristics are sigmoid. Consider again the Lyapunov function of networks with graded-response neurons (1.42). The function will now be called $L_{\rm DD}$ to distinguish its discrete-time evolution from the continuous-time evolution of Section 1.3.3.

The only neurons that may change their state at time t belong to the update group $U(t)$. Accordingly, $\Delta V_i(t) \equiv V_i(t+1) - V_i(t)$ vanishes for all other neurons. Using the symmetry (1.24) of the synaptic couplings, the change $\Delta L_{\rm DD}(t) = L_{\rm DD}(t+1) - L_{\rm DD}(t)$ is given by

$$\Delta L_{\rm DD}(t) = -\frac{1}{2}\sum_{i\in U(t)}\sum_{j\in U(t)} J_{ij}\Delta V_i(t)\Delta V_j(t) - \sum_{j=1}^{N}\sum_{i\in U(t)} J_{ij}V_j(t)\Delta V_i(t) - \sum_{i\in U(t)} I_i^{\rm ext}\Delta V_i(t) + \sum_{i\in U(t)}[G_i(V_i(t+1)) - G_i(V_i(t))] \ . \quad (1.57)$$

Since the functions $g_i(V_i)$ are assumed to be sigmoid, the auxiliary functions $G_i(V_i)$ are again strictly convex. Expanding $G_i(V_i(t))$ to second order around $V_i(t+1)$ and replacing the coefficient of the quadratic term with the smallest possible value, that is γ_i^{-1}, the following upper bound can be established (see also the right part of Figure 2):

$$G_i(V_i(t+1)) - G_i(V_i(t)) \le \Delta V_i(t) G_i'(V_i(t+1)) - \frac{1}{2}[\Delta V_i(t)]^2\gamma_i^{-1} \ . \quad (1.58)$$

Equality holds if and only if $V_i(t+1) = V_i(t)$. Inserting (1.52) and (1.58) into (1.57) gives

$$\Delta L_{\rm DD}(t) \le -\frac{1}{2}\sum_{i\in U(t)}\sum_{j\in U(t)} (J_{ij} + \delta_{ij}\gamma_i^{-1})\Delta V_i(t)\Delta V_j(t). \quad (1.59)$$

To facilitate the further discussion, let us define $W(t)$ as the number of neurons in the group $U(t)$ and symmetric matrices $\mathbf{U}(t)$ of dimension $W(t)\times W(t)$ as submatrices of the connection matrix $\mathbf{J}$, given by the synaptic strengths of those neurons that are updated at time t. For the Hopfield model (1.27),(1.28), where updating is one-at-a-time, $W(t) = 1$ for all t,

and $\mathbf{U}(t)$ reduces to the self-interaction term J_{ii} where i denotes the neuron being updated at time t. For the Little model (1.34),(1.35) or iterated-map analog networks (1.46),(1.47), the matrix is identical to $\mathbf{J}$ itself. As is obvious from these limiting cases, the structure of the set of matrices $\mathbf{U}(t)$ encodes the global dynamics.

The maximum neuron gain in the update group $U(t)$ will be denoted by $\gamma(t)$ and the minimum eigenvalue of the matrix $\mathbf{U}(t)$ by $\lambda_{\min}[\mathbf{U}(t)]$. Since for arbitrary symmetric matrices $\mathbf{A}$ and $\mathbf{B}$, $\lambda_{\min}[\mathbf{A}+\mathbf{B}] \geq \lambda_{\min}[\mathbf{A}] + \lambda_{\min}[\mathbf{B}]$, a sufficient condition for $\Delta L(t) \leq 0$ is given by

$$\lambda_{\min}[\mathbf{U}(t)] \geq -\gamma(t)^{-1} \ . \tag{1.60}$$

If the above condition holds for all t, $L_{\rm DD}(t)$ is strictly decreasing as long as $V_i(t+1) \neq V_i(t)$ for at least some i in the update group $U(t)$. As before, the function $L_{\rm DD}$ is bounded below. The network therefore relaxes asymptotically to a state where L does not vary in time if all directions in the space spanned by the neural activities are explored, that is, if the updating scheme is fair sampling. Since equality in (1.58) and (1.59) holds only if $V_i(t+1) = V_i(t)$, all solutions of (1.56) with time-independent $L_{\rm DD}$ are fixed-point solutions [10]. The result may be stated as follows:

Suppose the following three conditions hold: a) the updating rule is fair sampling, b) the neuron transfer functions are sigmoid, and c) the symmetric connection matrix satisfies (1.60) for all times. Then the distributed dynamics (1.56) admit the Lyapunov function (1.42) and converge to fixed points only.

For iterated-map networks, $\mathbf{U}(t)$ is constant in time and equals the set of all neurons. The criterion $\lambda_{\min}[\mathbf{J}] \geq -\gamma(t)^{-1}$ provides a sufficient condition to exclude two-cycles that exist in the general case as shown in Section 1.3.4: Lowering the neuron gain eliminates spurious oscillatory modes.

Neural networks with discrete elements correspond to the limit $\gamma_i \to \infty$ where (1.60) reduces to $\lambda_{\min}[\mathbf{U}(t)] \geq 0$. This implies in particular that there are no two cycles possible in the Little model if the whole connection matrix is nonnegative definite. The general remark from Section 1.3.1 about the convergence to solutions that are *not* minima of $L_{\rm DD}$ still holds in the discrete-neuron limit. This atypical behavior is, however, only possible because the g_i are piecewise constant functions in models with discrete neurons. For the generic case of continuous input-output characteristics, the network will always settle in a minimum as long as the initial conditions do not coincide with an unstable fixed-point of (1.56).

The convergence criterion (1.60) is less restrictive for smaller update groups than for larger ones because

$$\lambda_{\min}[\mathbf{U}_1] \geq \lambda_{\min}[\mathbf{U}_2] \quad \text{if} \quad U_1 \subset U_2 \ . \tag{1.61}$$

Note that (1.61) implies that the stability criterion for a fully parallel network where $\lambda_{\min}[\mathbf{J}] \geq -\gamma^{-1}$, is a sufficient condition for (1.60), and thus

sufficient to assure that the system (1.56) will converge to a fixed point for *any* fair sampling updating scheme.

Formula (1.61) has direct consequences for possible applications. Consider a high-dimensional optimization task such as the traveling salesman problem. It may be mapped onto a neural network architecture which then defines a fixed connection matrix $\mathbf{J}$ [59]. The computational time needed to find a good solution can easily be reduced on a parallel computer by increasing the size of update groups. However, the bounds given by (1.60) have to be met in order to assure convergence to fixed points, and will limit the maximal size of update groups. The goal of large updating groups will be achieved in an optimal way if one can form update groups of weakly or non-interacting neurons. All submatrices $\mathbf{U}(t)$ will have small off-diagonal elements in that case, and their eigenvalues will be close or identical to the diagonal elements, that is, the bounds (1.60) are largely independent of the size of the update groups. In principle, the search for optimal partitions of the above kind is itself a difficult optimization problem, but many applications exhibit an intrinsic structure (for example predominantly short-range interactions) which naturally leads to good choices for the updating groups.

1.3.6 Network Performance

The results obtained thus far demonstrate that the long-time behavior of neural networks with symmetric synaptic couplings is surprisingly robust with respect to alterations of model details at both the level of single neurons and the level of the overall network dynamics. All systems studied relax to fixed-point solutions under appropriate additional conditions on the synaptic efficacies and the input-output characteristics.

Various prescriptions for the storage of static patterns as fixed-point attractors have been discussed in the literature [22, 60, 61]. In what follows, I will focus on the Hebbian learning rule (1.23). A statistical mechanical analysis of performance measures, such as storage capacity and retrieval quality, can be carried out most readily for networks with McCulloch-Pitts neurons and block-sequential dynamics. It will also be assumed that the network can be partitioned into n fixed update blocks of equal size W such that there are no interactions within a group [10]. As emphasized before, such a situation can be arranged for many applications that map onto diluted or geometrically structured networks. In the limiting case $W = 1$, one recovers the Hopfield model.

To simplify the analysis, neurons are labeled by a double index S_{ia}. The first index $1 \leq i \leq W$ refers to the position within an update group, the second $1 \leq a \leq n$ labels the update group. The same notation applies to stored patterns ξ^{μ}_{ia} where the additional index μ, $1 \leq \mu \leq p$, labels the patterns. With these conventions, the Hebb rule (1.23) becomes

$$J_{ij}^{ab} = \begin{cases} \frac{1}{W(n-1)} \sum_{\mu=1}^{p} \xi_{ia}^{\mu} \xi_{jb}^{\mu} & \text{if } a \neq b \ , \\ 0 & \text{if } a = b \ . \end{cases} \tag{1.62}$$

The normalization factor N^{-1} in (1.23) has been changed to $[W(n-1)]^{-1}$ to guarantee the correct scaling behavior of L_{DD} in the thermodynamic limit $N \to \infty$.

Statistical mechanics may be used to analyze the emergent properties of feedback neural networks once it has been shown that under a stochastic update rule, the network relaxes to a Gibbsian equilibrium distribution generated by the Lyapunov function of the deterministic dynamics [22, 60, 62]. For Glauber dynamics (1.12) and a one-at-a-time or a parallel updating scheme such a relation exists as can be shown using the principle of detailed balance [28].

Although L_{DD} is identical to L_{SD} for two-state neurons, a block-sequential realization of Glauber dynamics need not approach a Gibbsian equilibrium distribution. However, in the special case of vanishing connection strength within all update groups (1.62), neurons "do not know" about the state of other neurons in the same group. Thus there is no formal difference between the block-sequential rule considered here and serial updating, where neurons change their state in consecutive order: every set of W successive updates of the latter dynamics is identical to one time step in the former case.

In what follows, I focus on the retrieval of unbiased random patterns where $\xi_{ia}^{\mu} = \pm 1$ with equal probability and study networks at a finite storage level $\alpha \equiv \frac{p}{N}$. The case of large cluster size, $W \to \infty$, with the number n of update groups kept finite will be analyzed; n has to be at least equal to two because according to (1.62), all neurons would be disconnected otherwise. Following the replica-symmetric theory of Amit, Gutfreund and Sompolinsky [63], a fixed number s of patterns is singled out, and it is assumed that the network is in a state highly correlated with these "condensed" memories. The remaining patterns are described collectively by a noise term. Notice that for coupling matrices of the form (1.62), both the overlaps m and spin-glass parameters q have to be defined as order parameters *on the level of the update groups*. For retrieval solutions, this requirement leads to the Ansatz

$$m_{\sigma a}^{\mu} \equiv W^{-1} \sum_{i=1}^{W} \xi_{ia}^{\mu} S_{ia}^{\sigma} = m \delta_{\mu,1} \tag{1.63}$$

and

$$q_{ab}^{\rho\sigma} \equiv W^{-1} \sum_{i=1}^{W} S_{ia}^{\rho} S_{ib}^{\sigma} = \delta_{ab} [\delta_{\rho\sigma}(1-q) + q] \tag{1.64}$$

for a k-fold replicated network, $1 \leq \rho, \sigma \leq k$. The resulting fixed-point equations are

$$m = \langle\langle \tanh[T^{-1}\{m + \sqrt{\alpha r} z\}] \rangle\rangle \tag{1.65}$$

and

$$q = \langle\langle \tanh^2[T^{-1}\{m + \sqrt{\alpha r} z\}] \rangle\rangle \tag{1.66}$$

where

$$r \equiv \frac{q}{[1 - T^{-1}(1-q)]^2} - \frac{q(n-1)}{[n - 1 + T^{-1}(1-q)]^2} \ . \tag{1.67}$$

Double angular brackets represent an average with respect to both the condensed patterns and the normalized Gaussian random variable z[10].

Equations (1.65) – (1.67) closely resemble their counterparts for the Hopfield model [63] and become identical to them in the limit of large n. On a formal level, the same holds for $n = 1$ but as explained before, this case does not correspond to a physical situation. For a general number of update groups there exists a first-order phase transition at $T = 0$ between the retrieval state and a spin-glass phase as α is varied. The critical storage level is denoted by α_c, the corresponding overlap by m_c.

The relative information content I_R, measured *per synapse* and *relative* to that of the Hopfield model,

$$I_R(n) \equiv \frac{I_n(\text{block-sequential})}{I(\text{random-sequential})} = \frac{n \cdot \alpha_c(n)}{(n-1) \cdot \alpha_c(\text{Hopfield})} \ , \tag{1.68}$$

is a third performance measure. A comparison between various network architectures in terms of all three measures is given in Table 1.

The performance of block-sequential updating schemes is quantitatively similar to that of the Hopfield model where $\alpha_c = 0.138$ and $m_c = 0.97$ [63]: the capability to retrieve stored random patterns is slightly lower when measured in terms of patterns per neuron, as indicated in the second column of Table 1, and slightly higher when measured in terms of patterns per synapse, as shown in the last column. Notice, in particular, that the information content increases with decreasing network connectivity, that is, for small n.

The results demonstrate that feedback networks can be used to store large amounts of information: the number of patterns (each of size N) that can be memorized grows linearly with N so that the information stored per synapse remains at a constant value of roughly 0.1 bits per synapse.[16] Stored patterns can be retrieved from noisy or incomplete data as long as the storage level remains below the critical level α_c. Compared to sequential or fully synchronous update schemes partially parallel schemes offer a

[16] This number is increased significantly by more elaborate learning rules [64].

potentially large advantage in terms of computational costs when implemented on a parallel computer allowing for a speed-up that may be as large as the number of processors without sacrificing network stability.

n	α_c	m_c	I_R
2	0.100	0.93	1.45
3	0.110	0.95	1.20
4	0.116	0.96	1.12
5	0.120	0.96	1.09

(1.69)

TABLE 1.1 Numerical solution of the saddle-point equations at $T = 0$. Displayed are the storage capacity α_c, the retrieval overlap m_c, and the relative information content I_R as functions of the number n of update groups.

1.3.7 Intermezzo:Delayed Graded-Response Neurons

The dynamical description of section 1.3.3 neglects any time lags due to finite propagation velocities of neural signals. As a first step towards the general formulation (1.14), one may study models where the communication time between neurons is modeled by one fixed delay τ,

$$C\frac{d}{dt}u_i(t) = -R^{-1}u_i(t) + \sum_{j=1}^{N} J_{ij}V_j(t-\tau) + I_i^{\rm ext}(t) \qquad (1.70)$$

with

$$V_i = g_i(u_i)\ . \qquad (1.71)$$

A mathematical analysis of this model is quite complicated. Because of the discrete delay, the initial condition for each neuron has to be specified as a function over a time-interval of length τ. Consequently, equations (1.70),(1.71) describe an infinite-dimensional dynamical system even in the scalar case ($N = 1$) which will be discussed in detail in Section 1.4.2.

Obviously, fixed-point solutions of (1.70),(1.71) do not depend on the time lag and are thus identical with those of the original model without delays, described by equations (1.40),(1.41). However, equilibria that are stable without delays may become unstable for large enough time lags as can been verified through a local stability analysis [65].

Global results about (1.70),(1.71) have been obtained under conditions that exclude nontrivial fixed-point solutions. A proof based on a Lyapunov functional shows that in this case there are no limit cycles either [66].

The lack of stronger global analytical results illustrates the limits of Lyapunov's direct method. It is often very hard or impossible to find a Lyapunov function for a *given* dynamical system under conditions that admit

interesting applications — multiple fixed points in the present example. On the other hand, there are many cases where one can find Lyapunov functions as soon as one enlarges the class of systems studied. In the present case, one could replace the single discrete lag in (1.70) by a distributed delay as in (1.14). At a first glance, this seems to complicate the analysis even further. However, there exist nontrivial delay distributions for which the dynamics generated by (1.14) admit global Lyapunov functionals [67].

The remark applies also to systems with synaptic couplings $J_{ij}(\tau)$ that are of the form $J_{ij}\varepsilon(\tau)$ where $\varepsilon(\tau)$ satisfies a linear ordinary differential equation in τ. For instance, if $\tau_{\max} = \infty$ and $\varepsilon(\tau) = \exp(-\tau)$, one may rewrite the dynamical equations as a set of $2N$ ordinary differential equations. The example demonstrates that unlike networks with discrete time lags, networks with distributed delays need not represent infinite dimensional dynamical systems. Models with delay distributions that are "reducible" in this sense have been studied extensively in the applied mathematics literature [68]. For a neurobiologically motivated system of two limit-cycle oscillators with reducible signal delay, a Lyapunov function is given in reference [69].

1.4 Periodic Limit Cycles and Beyond

Natural stimuli provide information in both space and time. Recurrent neural networks with delayed feedback can be programmed to recognize and generate such pattern sequences or "temporal associations" [70, 71, 72, 73, 74, 75, 76].[17] Recurrent networks with a broad distribution of signal delays and a Hebbian learning rule such as (1.22) are well suited to *learn* pattern sequences as well. [47, 77, 78, 79, 80, 81]. These systems are characterized by a high degree of compatibility between the network architecture, the task of learning spatio-temporal associations, and the learning algorithm. As in networks with fixed-point attractors, an initial state or "stimulus" lying in the basin of attraction of a stored "memory" will spontaneously evolve towards this attractor. In the present context, however, memories are spatio-temporal patterns of neural activity.

This section demonstrates that one can understand the computation of certain networks with signal delays as a down-hill march on an abstract spatio-temporal energy landscape. The result allows the application of techniques developed in the last sections.

---

[17] A detailed discussion can be found in reference [33].

1.4.1 Discrete-Time Dynamics

Let us focus on a synchronous discrete-time dynamics with deterministic McCulloch-Pitts neurons. For vanishing external inputs, the network dynamics (1.7),(1.8) become

$$S_i(t+1) = \mathrm{sgn}[h_i(t)] \quad \text{for all } i \tag{1.72}$$

with

$$h_i(t) = \sum_{j=1}^{N} \sum_{\tau=0}^{\tau_{\max}} J_{ij}(\tau) S_j(t-\tau) \ . \tag{1.73}$$

In the following, it is assumed that the synaptic couplings $J_{ij}(\tau)$ satisfy the extended symmetry $J_{ij}(\tau) = J_{ij}(D-(2+\tau))$. As was shown in Section 1.2.4, this symmetry arises if the network is taught cyclic pattern sequences of equal duration D.

The construction of a Lyapunov function for the retrieval dynamics (1.72), (1.73) is facilitated by the following consideration: If the network has learned cyclic associations with common length D, every correct retrieval solution corresponds to a D-periodic limit cycle. D-periodic oscillatory solutions of a discrete-time network, however, can always be interpreted as static states in a *fictitious* system of size $D \times N$ [50, 51].

Let us consider such a "D-plicated" network with D columns and N rows. The neural activities are denoted by S_{ia} where $1 \leq i \leq N$ and $0 \leq a < D$. To reproduce the synchronous dynamics of the original system, neurons S_{ia} with $a = t$ (modulo D) are updated at time t.

The time evolution of the new network is block sequential: synchronous within single columns and sequential with respect to these columns. In terms of the original variables S_i, the new activities S_{ia} are therefore given by $S_{ia}(t) \equiv S_i(a+n_t)$ for $a \leq t$ (modulo D) and $S_{ia}(t) \equiv S_i(a+n_t-D)$ for $a > t$ (modulo D), where n_t is defined through $t \equiv n_t + t$ (modulo D). The update rule reads

$$S_{ia}(t+1) = \begin{cases} \mathrm{sgn}[\sum_{j=1}^{N} \sum_{b=0}^{D-1} J_{ij}^{ab} S_{jb}(t)] & \text{if } a = t \text{ (modulo } D), \\ S_{ia}(t) & \text{otherwise.} \end{cases} \tag{1.74}$$

The synaptic couplings J_{ij}^{ab} are defined as

$$J_{ij}^{ab} = J_{ij}\big((b-a-1) \text{ (modulo } D)\big) \ . \tag{1.75}$$

Notice that the time evolution (1.74) of the equivalent fictitious system is the same as a block-sequential updating of a network with $D \times N$ McCulloch-Pitts neurons and block size N, as is illustrated in Figure 3. Section 1.3.5 shows how to guarantee that such a system relaxes to fixed points only: through synaptic symmetry together with the condition $\lambda_{\min}[\mathbf{U}(t)] \geq 0$.

Synaptic symmetry in the fictitious system, $J_{ij}^{ab} = J_{ji}^{ba}$, is equivalent to the extended symmetry (1.26) for the original couplings $J_{ij}(\tau)$. The second condition, $\lambda_{\min}[\mathbf{U}(t)] \geq 0$, is equivalent to $\lambda_{\min}[\mathbf{J}(D-1)] \geq 0$. This condition can be satisfied by setting $\tau_{\max} = D-2$.

It is left as an exercise for the interested reader to show that the Lyapunov function $L_{\rm DD}$, formulated for the equivalent fictitious system, may be rewritten in terms of the original time-delay network as

$$L_{\rm TD}(t) = -\frac{1}{2}\sum_{i,j=1}^{N}\sum_{a,\tau=0}^{D-1} J_{ij}(\tau) S_i(t-a) S_j\big(t-(a+\tau+1)\ (\text{modulo } D)\big)\ . \tag{1.76}$$

One may once again calculate the difference $\Delta L_{\rm TD}(t) \equiv L_{\rm TD}(t+1) - L_{\rm TD}(t)$ and arrives, as expected, at

$$\Delta L_{\rm TD}(t) = -\sum_{i=1}^{N}\big[S_i(t+1) - S_i(t+1-D)\big] h_i(t) \leq 0\ . \tag{1.77}$$

The derivation may be summarized as follows:

Suppose that the synaptic efficacies of the time-delay network (1.72), (1.73) satisfy the extended symmetry condition (1.26). Then the retrieval dynamics are governed by the Lyapunov function (1.76). The network relaxes to a fixed-point solution or a limit cycle with $S_i(t) = S_i(t-D)$, that is, an oscillatory solution with the same period as that of the taught cycles or a period which is equal to an integer fraction of D.

Due to the equivalence of (1.72),(1.73) with a block-sequential update rule for the fictitious system, one may apply the quantitative analysis of Section 1.3.6 to time-delay networks that store temporal associations. There is, however, a slight technical difficulty that has to be handled properly. Storing one D-periodic pattern sequence in the original model corresponds to memorizing D static patterns of size $D \times N$ in the equivalent system, each shifted by one column (modulo D) with respect to the next pattern. This complication arises because every sequence may be occuring with its first pattern recalled at some time t, or at time $t+1$, or at time $t+2$, and so on. In the equivalent D-plicated system, each of these time-shifted cyclic temporal associations corresponds to a new pattern.

For generic temporal associations, the analysis becomes rather complicated due to nontrivial correlations between shifted copies of the same pattern. If, however, each pattern of a sequence lasts for one time step only, all relevant correlations are the same as if one had stored D unrelated patterns. This implies that the results of Section 1.3.6 cover also the storage of pattern sequences where each pattern lasts for one unit of time.

As an example, take $D = 2$. With the maximal delay $\tau_{\max}$ set to $D-2$, $\tau_{\max}$ is zero, and one has recovered the Little model. According to Table 1, $0.100N$ two-cycles of the form $\psi_{i1}^{\mu} \rightleftharpoons \psi_{i2}^{\mu}$ may be recalled as compared to $0.138N$ static patterns [82]: a 1.45-fold increase of the information content

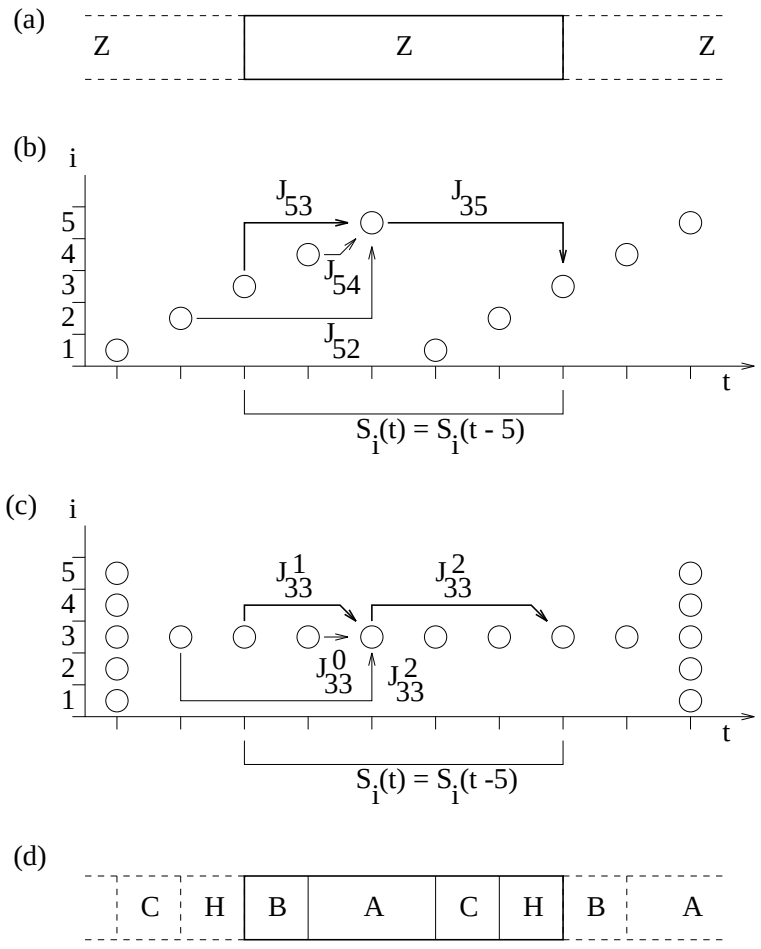

FIGURE 1.3. Schematic drawing of the dynamics of a time-delay network (c and d) and its equivalent fictitious system with block-sequential time evolution (a and b). Horizontal axes represent time, vertical axes in (b) and (c) denote the index of neurons. (a) The pattern "Z" is retrieved in the fictitious network with five update groups that are represented in (b) by five neurons. (c) Time evolution of one neuron in a network with signal delays and discrete-time dynamics. The system recalls the cyclic pattern sequence "BAACH" as shown in (d).

per synapse. At the same time, the retrieval overlap drops slightly from 0.97 to 0.93.

The performance of networks with distributed delays and $D = 4$ is displayed in Table 2.

τ	=	0	1	2	3	α_c	m_c	I_R
$\varepsilon(\tau)$	=	1/3	1/3	1/3	0	0.116	0.96	1.12
$\varepsilon(\tau)$	=	1/2	0	1/2	0	0.100	0.93	1.45
$\varepsilon(\tau)$	=	0	1	0	0	0.050	0.93	1.45

(1.78)

TABLE 1.2. Influence of the weight distribution on the collective network properties. The storage capacity α_c, the critical overlap m_c, and the relative information content I_R are displayed for some choices of $\varepsilon(\tau)$ for $D = 4$.

As shown in Table 2, the uniform distribution leads to the largest α_c but smallest I_R. The other two networks have the same value of I_R as the (unique) $D = 2$ system due to the particular structure of their eigenvalue spectrum. Furthermore, one obtains $I_R = 1.45$ independently of D for all networks with a minimal connectivity where only one synapse links two neurons.[18] Simulation data show slightly higher values of α_c, possibly indicating effects of replica symmetry breaking as in the Hopfield model [63].

In passing, let me remark that each cycle consists of D patterns so that the storage capacity for *single* patterns is $\bar{\alpha}_c = D\alpha_c$. During the recognition process, however, each pattern will trigger the cycle it belongs to and cannot be retrieved as a static memory.

If static patterns instead of temporal associations are learned, the synaptic strengths do not depend on the delay, see also equation (1.23). The synaptic couplings still satisfy the extended symmetry, and with $\tau_{\max} = D - 2$ one recovers the Lyapunov function for networks with McCulloch-Pitts neurons and "multiple-time-step parallel dynamics" [83],

$$L_{\mathrm{MTS}}(t) = -\frac{1}{2} \sum_{i,j=1}^{N} J_{ij} \sum_{a=0}^{D-2} S_i(t-a) \sum_{b=0}^{D-2} S_j(t-b) \; . \tag{1.79}$$

The evolution equations (1.72),(1.73) may be generalized to analog systems with periodic external inputs. Using the "cooking recipes" of Sections 1.3.1 – 1.3.4, it is possible to construct a Lyapunov function for that case as well [84].

The learning rule (1.26) may also be utilized to store cycles of correlated real-valued pattern sequences. Numerical studies have been performed for

[18] This case is possible if D is an even number.

low-dimensional trajectories (small N) with high numbers of data points (large D). For many examples, good retrieval could be obtained without any need for highly time-consuming supervised learning schemes. However, algorithms of the latter kind facilitate the learning of more sophisticated real-world tasks. Here Lyapunov functions are of great help since they permit the application of mean-field techniques [85] to a wide class of supervised learning strategies such as spatio-temporal extensions of the "Boltzmann Machine" concept [86] and contrastive-learning schemes [87].

In closing this section, let me mention that an analysis of the storage capacity along Gardner's approach [88] has been given in reference [89]. Analytical results on highly diluted systems with time lags have also been obtained [90].

1.4.2 Continuous-Time Dynamics

The global dynamics of certain networks with graded-response neurons and delayed interactions may be studied in a manner similar to that of Section 1.4.1 [67]. In the following, I will focus on the simplest case, a single neuron (or a homogeneous assembly of neurons) coupled to itself through one inhibitory feedback loop with delay τ. Equation (1.14) reduces to

$$C\frac{d}{dt}u(t) = -R^{-1}u(t) - g[u(t-\tau)] \ , \tag{1.80}$$

where g satisfies the condition

$$ug(u) > 0 \quad \text{for } u \neq 0 \quad \text{and} \quad g(0) = 0 \ . \tag{1.81}$$

Solutions of this seemingly simple scalar equation include a fixed point $u(t) = 0$ and, depending on the graph of g, periodic limit cycles and chaotic trajectories [91]. Such a diversity of temporal phenomena is possible since due to the discrete delay, equation (1.80) describes an infinite dimensional dynamical system as was already mentioned in Section 1.3.7.

Various aspects of the scalar delay differential equation (1.80) have been discussed in the mathematics literature. Most articles have concentrated on periodic solutions, in particular on those that are "slowly oscillating," that is, periodic solutions with zeros spaced at distances larger than the time lag τ. Results about their existence, uniqueness and local stability have been obtained by Kaplan and Yorke [92], Nussbaum [93], and Chow and Walther [94], respectively.

The global analysis of (1.80) is simplified significantly if one neglects the transmembrane current $R^{-1}u(t)$ and if g is an odd sigmoid function. Without loss of generality, one may set $C = \tau = 1$ and study the evolution equation

$$\frac{d}{dt}u(t) = -g[u(t-1)] \ . \tag{1.82}$$

Consider the auxiliary function $L_{\rm DDE}(t)$,

$$\begin{aligned} L_{\rm DDE}(t) = &- \frac{1}{2}\int_0^1 \int_{t+\tau-1}^{t+1} \dot{u}(s)\dot{u}(s-\tau)\; ds\; d\tau \\ &+ \frac{1}{2}\int_1^2 \int_{t+\tau-1}^{t+1} \dot{u}(s)\dot{u}(s-\tau)\; ds\; d\tau \\ &+ \int_{t-1}^{t+1} G(\dot{u}(s))\; ds + \frac{1}{4}[u(t+1)+u(t-1)]^2 \end{aligned} \tag{1.83}$$

where $G(x)$ is defined as in (1.43).[19] For bounded nonlinearities g, all solutions of (1.82) are bounded. They are differentiable for $t > 1$. Consequently, $L_{\rm DDE}(t)$ is bounded below for $t > 2$. It follows that for $t > 1$, the time derivative of $L_{\rm DDE}(t)$ along a solution of (1.82) is well defined and given by

$$\begin{aligned} \frac{d}{dt}L_{\rm DDE}(t) = &[\dot{u}(t+1)+\dot{u}(t-1)][u(t)-\frac{1}{2}u(t+1)-\frac{1}{2}u(t-1)] \\ &+G(\dot{u}(t+1))-G(\dot{u}(t-1)) \\ &+\frac{1}{2}[u(t+1)+u(t-1)][\dot{u}(t+1)+\dot{u}(t-1)] \\ = &\, u(t)[\dot{u}(t+1)+\dot{u}(t-1)]+G(\dot{u}(t+1))-G(\dot{u}(t-1))\;. \end{aligned} \tag{1.84}$$

Because the input-output characteristic is assumed to be an odd sigmoid function, g^{-1} is odd, single valued and monotone increasing. Consequently, the function G is even and strictly convex. In particular, the equality $G(\dot{u}(t-1)) = G(-\dot{u}(t-1))$ holds. Performing a Taylor expansion as in (1.51), one therefore obtains

$$\begin{aligned} G(\dot{u}(t+1))-G(\dot{u}(t-1)) &\le [\dot{u}(t+1)+\dot{u}(t-1)]g^{-1}(\dot{u}(t+1)) \\ &\le -[\dot{u}(t+1)+\dot{u}(t-1)]u(t)\;. \end{aligned} \tag{1.85}$$

Equality in (1.85) holds if and only if $\dot{u}(t+1) = -\dot{u}(t-1)$. Taking the evolution equation (1.82) and the strict monotonicity of g into account, the last equation may also be written $u(t) = -u(t-2)$.

Inserting (1.85) into (1.84), one finally arrives at

$$\frac{d}{dt}L_{\rm DDE}(t) \le 0 \quad \text{for } t \ge 2 \tag{1.86}$$

---

[19] $L_{\rm DDE}$ has been introduced as an explicitly time-dependent function for simplicity and has been written in terms of both u and $\dot{u}$ for the same reason. The initial function may, however, not be differentiable. This (purely technical) difficulty can be avoided if $\dot{u}(s)$ is replaced by $-g(u(s-1))$. $L_{\rm DDE}$ may then be properly defined as a functional in the space of continuous functions from the interval $[-2,0]$ to the real numbers [95].

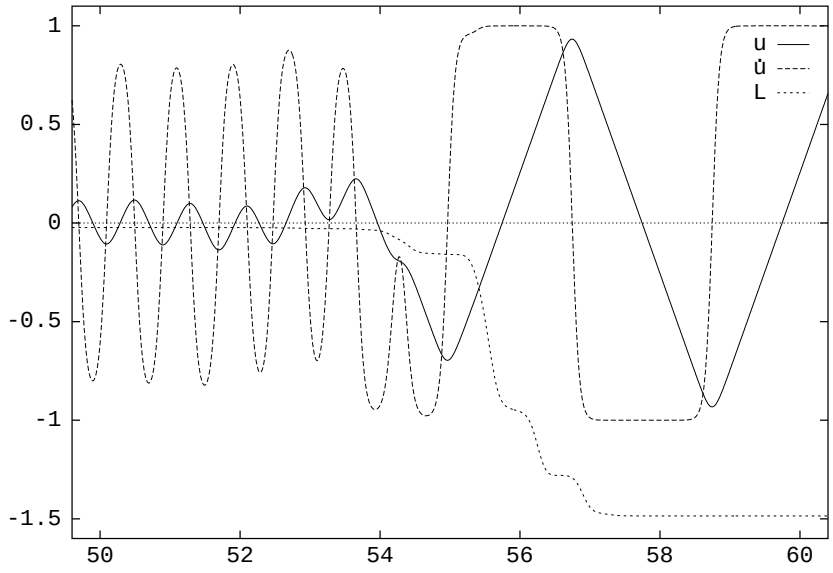

FIGURE 1.4. Time evolution of a single neuron with delayed feedback according to the evolution equation (1.82). The input-output characteristic is $g(u) = \tanh(5u)$. The state variable u is plotted as a solid line, its derivative $\dot{u}$ as a dashed line, and the Lyapunov function $L_{\rm DDE}$ as a dotted line. Notice that $L_{\rm DDE}(t)$ approaches a constant value as required for a Lyapunov function whereas u relaxes towards a periodic oscillatory solution with period 4.

where equaility holds if and only if $u(t) = -u(t-2)$.[20] An illustration is given in Figure 4. According to (1.86), $L_{\rm DDE}(t)$ is nonincreasing along every solution for $t > 2$. The overall result may be summarized in the following way:

Suppose that the function g is odd, bounded and sigmoid. Then the evolution equation (1.82) admits the Lyapunov function (1.83). Solutions of (1.82) converge either to the trivial fixed point $u = 0$ or to a periodic limit cycle that satisfies

$$u(t) = -u(t-2) \; . \tag{1.87}$$

Notice that the period P of the limit cycles does *not* depend on the graph of g; according to (1.87), it is always given by $P = \frac{4}{4k+1}$ where k is a nonnegative integer.[21] On the other hand, it is well known that for the general equation (1.80), the period of a periodic solution is influenced by the ratio of RC to τ and the shape of g [96]. This fact implies that the above methods can probably not be extended to study delay differential equations of the type (1.80). There is, however, another way to analyze this equation [97]. To facilitate the discussion, let t_i, $i \in I\!N$ with $t_i < t_{i+1}$ denote the times of consecutive zero crossings $u(t) = 0$ of a solution of (1.80). One may then prove the following proposition:

Assume that the function g is bounded and satisfies the condition (1.81). For every solution $u(t)$ of (1.80), the number $n(i)$ of zero crossings in the interval $[t_i - \tau, t_i)$ is a nonincreasing function of i.

This result means that a solution of (1.80) oscillates more and more slowly around zero as time proceeds. For long times it approaches a solution with constant $n = n(i)$, possibly $n = 0$. In particular, if the system is initialized with a solution that has n zero crossings in the interval $[-\tau, 0)$ it can never reach an oscillation with more than n zero crossings in any one of the intervals $[t_i - \tau, t_i)$.

Let me briefly sketch the proof. The reader is also referred to Figure 5. If g is bounded and satisfies the condition (1.81), solutions of (1.80) exist for all positive t and are continuous [98]. Assume without loss of generality that at time t_j, $\frac{d}{dt}u(t_j) > 0$. According to (1.80) and (1.81), this means that $u(t_j - \tau) < 0$ because $u(t_j) = 0$ by definition. The same argument may be used at time t_{j+1} where it implies that $u(t_{j+1} - \tau) > 0$ because

---

[20] The curious reader is invited to compare this result and its derivation with that for the Little model with antisymmetric couplings (1.54).

[21] Further results derived with the help of $L_{\rm DDE}$ can be found in reference [95]. One proof is well suited to highlight the potential of Lyapunov functions — once they are found: It can be shown that for large enough $g'(0)$, the global minimum of $L_{\rm DDE}$ is always achieved on a slowly oscillating solution (otherwise on the trivial fixed point $u(t) = 0$). This immediately implies that those solutions have to be asymptotically stable (except for global phase shifts), a conclusion that previously required elaborate analytical techniques.

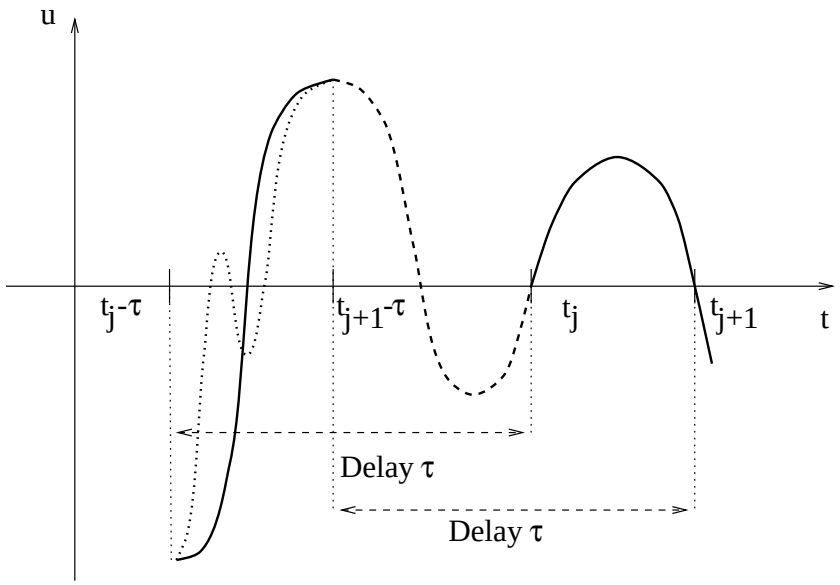

FIGURE 1.5. Time evolution of a single graded-response neuron with delayed self inhibition modeled by the delay differential equation (1.80). There are zero crossings of the solution $u(t)$ at time t_j, t_{j+1}, and at various earlier (and later) times. In the interval $[t_j - \tau, t_{j+1} - \tau)$ two possible solutions are drawn. They have one and three zero crossings, respectively.

$u(t_{j+1}) = 0$ and $\frac{d}{dt}u(t_{j+1}) < 0$. Together with the continuity of $u(t)$, this implies that there is an odd number $k(j) \geq 1$ of zero crossings in the interval $[t_j - \tau, t_{j+1} - \tau)$.

Denote the number of zero crossings in the interval $[t_{j+1} - \tau, t_j)$ by $l(j)$.[22] It follows that $n(j) = l + k(j)$ and $n(j+1) = l(j) + 1$. Since $k(j) \geq 1$, both relations may be combined to the statement $n(j) \geq n(j+1)$ which proves the proposition.

The number of zero crossings in any interval is nonnegative — the function $n(i)$ is bounded below. Since it is nonincreasing along every solution of (1.80), it is an integer-valued Lyapunov function. Accordingly, solutions of (1.80) relax to solutions with constant $n(i)$. Notice that those solutions may be periodic but could — at least in principle – also be aperiodic. This is a surprising result. It highlights the generality of Lyapunov's second method in a rather illuminating way.

[22] It is understood that $l(j) = 0$ if $t_{j+1} - \tau \geq t_j$.

1.5 Synchronization of Action Potentials

While it may frequently be the case that mean firing rates are an adequate description of neural information, there are many instances where the detailed timing and organization of action potentials matters. An important example is given by the stimulus-dependent synchronization of action potentials [15, 16, 17].

Due to the inherent limitations of descriptions based on discrete-time dynamics or mean-firing rates, realistic synchronization processes are not captured by the networks discussed in Sections 1.3 and 1.4. They may, however, be studied using networks with integrate-and-fire neurons, whose time evolution has been introduced in section 1.2.3.

Networks of that type often show globally synchronized neurons when all-to-all couplings are used.[23] Note that throughout this section, terms such as "synchronized neurons" always refer to the time of spike generation. According to this definition, a periodic network state (also called a "phase-locked solution") may or may not be "globally synchronized." A global analysis for networks described by (1.13), (1.18), and (1.20) has been given in reference [102]. With excitatory all-to-all couplings of equal strength, nonzero leakage currents, uniform external inputs, and a reset to zero after spike generation ($\gamma = 0$), the size of synchronized clusters is a nondecreasing function of time — a (discrete-valued) Lyapunov function! The proof then shows that such systems indeed approach a globally synchronized solution where all neurons fire in unison.

Networks with more general nonuniform interaction admit richer dynamical behaviour [25, 103, 104]. Equipped with excitatory finite-range couplings, one class of networks relaxes to phase-locked clusters of (locally) synchronized neurons [105, 106]. The shapes and relative phases of the clusters encode information about the initial stimulus. This result is in accordance with the hypothesis that synchronized cortical neurons are used to bind stimulus features together [107].

1.5.1 Phase Locking

Global results for locally coupled networks with integrate-and-fire neurons have been obtained in the limiting case $R \to \infty$ of perfectly integrating cells and uniform positive input currents $I_i^{\text{ext}} = I > 0$. In this situation, external information is encoded in the initial conditions $u_i(t = 0)$, not in the input currents. This choice is reminiscent of the experimental paradigm of "stimulus induced oscillations" [15]. Due to the constant positive input current I, each model cell fires regularly if there is no further synaptic

[23] Doubts about the structural stability of simple integrate-and-fire models have been raised because some model variants do not exhibit system-wide synchronization with all-to-all couplings [99, 100, 101].

input from other cells. Thus I^{-1} represents the spontaneous firing rate of an isolated neuron. By rescaling time, the capacitance C and input I in (1.13) can be taken as unity. The overall dynamics may then be summarized by the following update rules:

(i) Initialize the $u_i(t=0)$ in $[0,1]$ according to the external stimulus.

(ii) If $u_i \geq 1$ and if neuron i is next in the update scheme then

$$u_i \to u_i' = \gamma(u_i - 1) \tag{1.88}$$

and

$$u_j \to u_j' = u_j + J_{ji} \ . \tag{1.89}$$

(iii) Repeat step (ii) until $u_i < 1$ for all i.

(iv) If the condition of step (ii) does not apply then

$$\frac{d}{dt} u_i = 1 \quad \text{for all } i \ . \tag{1.90}$$

Under the condition that all neurons have the same total incoming synaptic strength,

$$\sum_j J_{ij} = A \ , \tag{1.91}$$

and the same total outgoing synaptic synaptic strength,

$$\sum_i J_{ij} = A \ , \tag{1.92}$$

one may proof that the simple function $L_{\rm IAF}$,

$$L_{\rm IAF} = -\sum_i u_i \ , \tag{1.93}$$

that is, the (negative) total membrane potential, plays the role of a Lyapunov function for the system defined by (i)–(iv) as shown in reference [106]:

Assume that $\gamma = 1$ and that the synapses satisfy $J_{ij} \geq 0$ and the constraints (1.91) and (1.92) with $A < 1$. Then the dynamics generated by (1.88) – (1.90) admit the Lyapunov function (1.93) and converge to cyclic attractors with period $P_{\rm IAF} = 1 - A$. On the limit cycle, each neuron fires exactly once in a period.

Notice that synaptic symmetry has *not* been required! This distinguishes the present model from the networks discussed in the previous sections.

Depending on the initial conditions, the limit cycles can contain events in which one neuron fires alone, and others in which many neurons fire in synchrony. In networks with excitatory short-range connections only, regions with small variability of the initial conditions are smoothed out and represented by locally synchronized clusters of neurons whose firing times encode the stimulus quality. Regions with high variability, on the other hand, give rise to spatially uncorrelated firing patterns. Through appropriate choice of coupling strengths, more complex computations can be preformed as demonstrated by numerical simulations [106].

In order to proof the proposition, let us first show that no neuron fires more than once in any interval of length $P_{\rm IAF}$.

Lemma: Let $n_i(t,t')$ denote the number of times neuron i fires in $[t,t')$. If the conditions of the proposition hold then $n_i(t,t+P_{\rm IAF}) \leq 1$.

Starting at time t, if some neuron fires twice before $t+P_{\rm IAF}$, then some neuron k must first fire twice, and at time $t' < t+P_{\rm IAF}$. For that to happen, the total change in u_k from t to t' due to the synaptic currents and the external input must be greater than 1. Thus one requires that for neuron k,

$$(t'-t)\frac{(1-A)}{P_{\rm IAF}} + \sum_j J_{kj} n_j(t,t') > 1 \ . \tag{1.94}$$

However, by hypothesis $(t'-t) < P_{\rm IAF}$, and since k is the first neuron to fire twice, the number $n_j(t,t')$ of firings of each of the other neurons up to t' is less than or equal to 1. For J_{ij} nonnegative the left-hand side of equation (1.94) is less then $(1-A)+A = 1$. The contradiction shows that k cannot have fired twice.

Returning to the proof of the proposition, let us consider the change of $L_{\rm IAF}$ in a time interval of length $P_{\rm IAF}$, $\Delta L_{\rm IAF}(t) \equiv L_{\rm IAF}(t+P_{\rm IAF}) - L_{\rm IAF}(t)$. It is

$$\Delta L_{\rm IAF}(t) = -(1-A)N - \sum_{i,j} J_{ij} n_j(t,t+P_{\rm IAF}) + \sum_i n_i(t,t+P_{\rm IAF}) \ . \tag{1.95}$$

The first term comes from the constant input current, the second term from the effect of the firing of other neurons, and the third term comes from i itself firing. Using the condition (1.91), one finds

$$\Delta L_{\rm IAF}(t) = -(1-A)[N - \sum_i n_i(t,t+P_{\rm IAF})] \tag{1.96}$$

Due to the lemma, $n_i(t,t+P_{\rm IAF}) \leq 1$ for all t. The change of $L_{\rm IAF}$ in each time interval $P_{\rm IAF}$ is thus nonpositive. Since $L_{\rm IAF}$ is bounded, the system performs a downhill march on the energy landscape generated by the Lyapunov function $L_{\rm IAF}$ — if the function is measured after time steps of length $P_{\rm IAF}$. The difference $\Delta L_{\rm IAF}(t)$ vanishes if and only if $n_i(t,t+P_{\rm IAF}) =$

1 for all i, that is, on periodic limit cycles where every neurons fires exactly once in a time interval of length $P_{\rm IAF}$ [106].[24]

To avoid the unfamiliar evaluation of the Lyapunov function $L_{\rm IAF}$ at the discrete times $t + kP_{\rm IAF}$, $k \in I\!N$, one may alternatively use the functional

$$\tilde{L}_{\rm IAF} = \int_0^{P_{\rm IAF}} L_{\rm IAF}(s)ds \ . \tag{1.97}$$

Along solutions, $\tilde{L}_{\rm IAF}$ is differentiable with $\frac{d}{dt}L_{\rm IAF}(t) = \Delta L_{\rm IAF}(t)$ for all $t \geq 0$, so that the previous conclusions are reached again. For an illustration, see Figure 6.

1.5.2 Rapid Convergence

The results of the previous section prove that specific networks of integrate-and-fire neurons approach phase-locked solutions. Numerical simulations of these and more general networks [102, 106, 108, 109, 110, 111] indicate that the convergence process takes place in a very short time — see also Figure 6.[25] This observation can be substantiated under certain conditions [105, 106]:

Assume that the synapses satisfy $J_{ij} \geq 0$ and the condition (1.91) with $A < 1$. Then all solutions of (1.88) – (1.90) converge to cyclic attractors with period $P_{\rm IAF} = 1 - A$. The limit cycles are reached as soon as every neuron has fired once. On the limit cycle, each neuron fires exactly once in a period.

Notice that although the conditions on γ and on the sum of outgoing synaptic strength have been dropped, the conclusions are now stronger than in the previous proposition. However, the proof given is *not* based on a Lyapunov function so that the concept of a down-hill march on an energy landscape generated by the Lyapunov function is not available anymore. The lack of a Lyapunov function might also be a drawback when stochastic extensions are considered in the future.[26]

Let $t_{\rm max}$ denote the first time where every neuron has fired at least once. Some cells may have fired repeatedly before $t_{\rm max}$, depending on the

[24] A related proof has been given in reference [35].

[25] In general, clusters of locally synchronized neurons will slowly reorganize. The models analyzed in this article are an exception in that they do not show such slow relaxation phenomena.

[26] The sentence reflects the author's hope that it might be possible to construct simple stochastic dynamics of integrate-and-fire neurons such that the Lyapunov function of the noiseless dynamics determines a Gibbs distribution for the stochastic extension. Equilibrium statistical mechanics could then be applied to analyze the collective phenomena in networks of integrate-and-fire neurons in the same spirit as this has been done for neural network models discussed in Sections 1.3 and 1.4. Regrettably, such evolution equations have not been found yet.

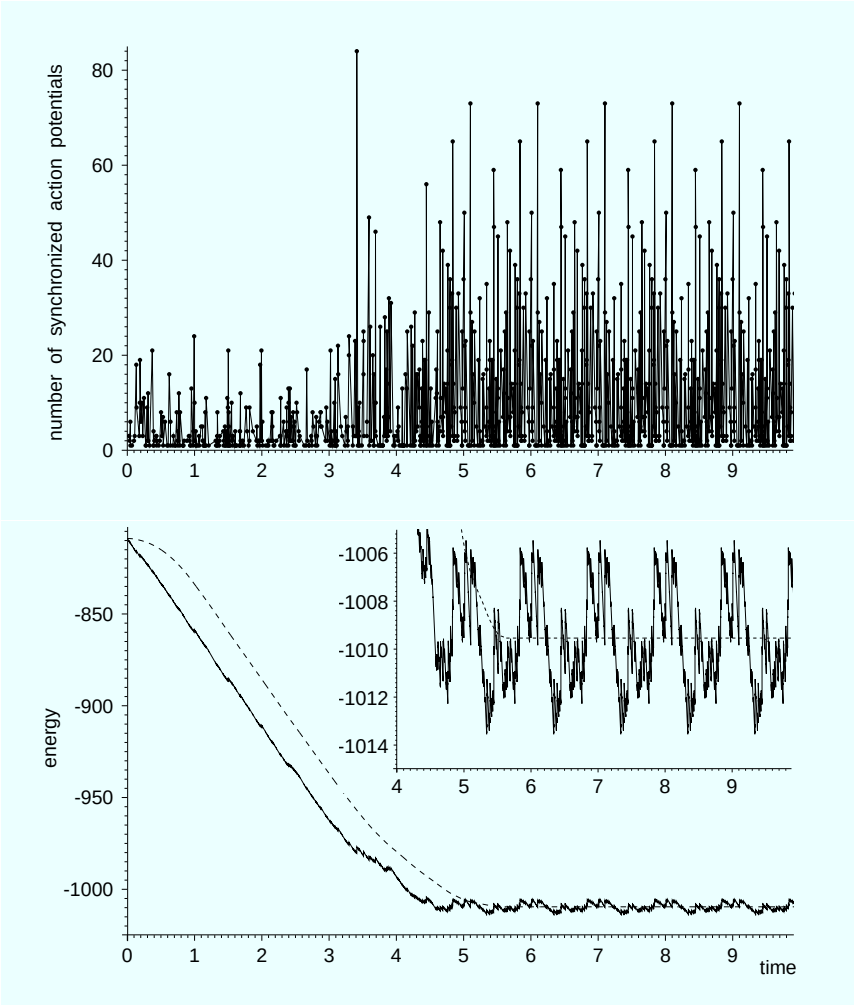

FIGURE 1.6. Rapid local synchronization of action potentials. Shown are results from numerical simulations of a planar network with 40×40 integrate-and-fire neurons ($R^{-1} = 0$, $\gamma = 1$), periodic boundary conditions and nearest-neighbor interactions of strength $J_{\rm nn} = 0.24$. Each dot in the upper trace represents the number of simultaneous action potentials as a function of time. The lower trace depicts the time evolution of the Lyapunov function $L_{\rm IAF}$ (solid line) and the Lyapunov functional $\tilde{L}_{\rm IAF}$ (dashed line). The inset verifies that, as predicted, the latter approaches a constant value.

parameter values and initial conditions. Let t_i denote the last time when neuron i fires before $t_{\max}$, $t_{\min}$ the minimum of all these times t_i, and k a cell that fires at $t_{\min}$ for the last time.

By definition, every cell discharges at least once in the interval $[t_{\min},t_{\max}]$. This implies in particular that every neuron j from which cell k receives synaptic input emits one or more action potentials in that interval. Each spike adds J_{kj} to u_k. The total change of u_k in $[t_{\min}, t_{\max}]$ is thus equal or greater than $A + t_{\max} - t_{\min}$. This number has to be smaller than 1 because otherwise, neuron k would fire a second time in the interval $[t_{\min}, t_{\max}]$ in contradiction to the assumption. It follows that $t_{\max} - t_{\min} < P_{\rm IAF}$.

Going back to Section 1.5.1, one notices that the condition on the sum of *outgoing* synaptic strengths (1.92), although essential for the proof of the main proposition, is not required for the proof of the lemma: The lemma is also valid under the weaker conditions of the present section. Evaluated at time $t = t_{\max} - P_{\rm IAF}$ and combined with the previous results, the lemma implies that every cell fires exactly once in $[t_{\min}, t_{\max}]$ and no cell fires in $(t_{\max} - P_{\rm IAF}, t_{\min})$. Since $t_{\max} \leq 1$, the last result proves that in finite time $t_{\max} - P_{\rm IAF}$, all limit cycles are approached in the sense that $u_i(t) = u_i(t + P_{\rm IAF})$ for $t \geq t_{\max} - P_{\rm IAF}$. The argument also shows that the attractors are reached as soon as every neuron has fired once.

The proof does not depend on the details of the reset mechanism. This means that it covers not only the present model with arbitrary $0 \leq \gamma \leq 1$ but also all schemes where a neuron i firing at time t is relaxed to some value between 0 and $u_i(t^-) - 1$. Perhaps surprisingly, this allows stochastic updatings during the transient phase.

In all model variants except from the limiting case $\gamma = 1$ limit cycles with period $P_{\rm IAF}$ and one spike per cycle cannot occur if a neuron is driven above threshold. In events with multiple neurons firing "at the same time," the potentials have to be fine-tuned such that if neuron i is triggered by neuron j, $u_i(t^-) = 1 - J_{ij}$. This implies that although every *firing sequence* of the model with $\gamma = 1$ can be realized in these models, the volume of all attractors is greatly reduced when measured in the space of the dynamical variables u_i.

1.6 Conclusions

The examples presented in this article demonstrate that Lyapunov's direct method has widespread applications within the theory of recurrent neural networks. With respect to the list of levels of analysis sketched in the Introduction, it has been shown that Lyapunov's method is most helpful on the second level which deals with questions about the type of attractors possible in a neural network.

Combined with powerful techniques from statistical mechanics, Lyapunov's approach allows not only for a qualitative understanding of the global

dynamics but also for quantitative results about the collective network behavior. As shown in Sections 1.3, 1.4, and 1.5, Lyapunov's method applies to the retrieval of static patterns in networks with instantaneous interactions, to the recall of spatio-temporal associations in networks with signal delays, and to synchronization processes in networks of integrate-and-fire neurons.

There remain numerous interesting questions about the global dynamics of feedback neural networks. These include questions concerning the convergence of network models with discrete-time dynamics, symmetric couplings and *overlapping* delays (see Figure 1d). Numerical simulations suggest that such systems relax to fixed-point solutions [112] but the analytic results from the computer-science literature [55, 56, 57, 58] only cover the case where a single pattern is stored in the network.

With regard to networks with transmission delays, it would be interesting to know more about the global dynamics generated by equations (1.70),(1.71) under conditions that admit multiple fixed-point attractors. With a similar interest in mind, one could try to perform a statistical mechanical analysis of the system (1.72),(1.73) with delay-independent symmetric couplings (1.23) to study the influence of signal delays on the collective properties of networks that store static patterns.

In the proofs concerning integrate-and-fire neurons, synaptic strengths were assumed to be excitatory. There is, however, strong numerical evidence that inhibition does not change the overall results [106]. If the synaptic couplings continue to satisfy the condition (1.91) with $A < 1$ and if the network parameters are chosen such that there are no run-away solutions and no solutions with neurons that are permanently below threshold, then all simulations of the dynamics generated by (1.88) – (1.90) approach periodic limit cycles of period $P_{\rm IAF} = 1 - A$. For leaky integrate-and-fire models (finite R), the same is true but the period is given by the period $P_{\rm LIAF}$ of the globally synchronized solution in such a system,

$$P_{\rm LIAF} = RC[\ln(RI - A) - \ln(RI - 1)] \; . \tag{1.98}$$

This observation gives hope that further understanding of integrate-and-fire models is possible although the mathematical situation is more complicated than in the cases discussed in Section 1.5. A convergence proof based on Lyapunov functions such as (1.93) is possible because *every* periodic solution of the model has the *same* period. This is not the case for models for finite R as shown by the following counterexample. Consider a spatio-temporal "checkerboard" pattern, where the "black" sites fire at even multiples of $\Delta/2$, the "white" sites at odd multiples of $\Delta/2$. A self-consistent calculation of the firing pattern leads to an implicit equation for Δ,

$$A{\rm e}^{-\frac{\Delta}{2RC}} + RI[1 - {\rm e}^{-\frac{\Delta}{RC}}] = 1 \; . \tag{1.99}$$

Except from the limiting case $R \to \infty$, Δ differs from the period of the globally synchronized solution. A linear stability analysis verifies that the checkerboard pattern is unstable but its mere existence indicates that it will be difficult to find Lyapunov functions for leaky integrate-and-fire models.

More generally, one may ask which conditions in the proofs of Sections 1.3, 1.4, and 1.5 can be violated without changing the desired emergent network behavior. These questions deal with the structural stability of neural networks, the fifth level of analysis, and have to be answered if one wants to evaluate the biological relevance of specific networks. In order to keep the article within reasonable bounds, this topic has not been discussed here. A particularly important issue, the convergence of "conventional" recurrent neural networks (of the type studied in Section 1.3) without synaptic symmetry, has been studied extensively in the literature [113, 114]. In passing let me note that one may always generate specific asymmetric networks through appropriate transformations of both the coupling matrix and dynamical variables of systems with symmetric interactions.

There are a number of other topics related to the main theme of this article that could not be included. Let me briefly list two of these issues.

First, one may design dynamical systems such that they perform a downhill march on an energy landscape that encodes some optimization task [59]. Various biologically motivated examples can be found in the computer-vision literature [115, 116].

Second, one may construct feedback networks that possess desired attractors but no spurious stable states [117, 118]. The construction of such artificial associative memories is greatly facilitated if one deliberately lifts modeling restrictions that would otherwise be naturally imposed by biological constraints.

Let me close with a general comment: "associative computation" means that many different inputs are mapped onto few output states. The time evolution of a dynamical system that performs such a computation is characterized by a contraction in its state space, that is, it is dissipative.[27] This observation suggests that many dynamical systems that have been used as models for associative computation may admit Lyapunov functions. As emphasized in Section 1.3.7, minor modifications of the models may be needed to satisfy technical requirements.

In view of the many Lyapunov functions already found, I would like to conclude with a remark from the monograph of Rouche, Habets and Laloy [3]: "Lyapunov's second method has the undeserved reputation of being mainly of theoretical interest, because auxiliary functions are so difficult to construct. We feel this is the opinion of those people who have not really tried ..."

[27] The threshold operation of a two-state neuron might be interpreted as a special realization of this contraction process.

Acknowledgments: Most of the author's own results presented in this article were obtained in collaborations with Bernhard Sulzer, Charlie Marcus, John Hopfield, Leo van Hemmen, Reimer Kühn, and Zhaoping Li. Discussions with David MacKay and Hans-Otto Walther were of great help. The work on integrate-and-fire neurons was stimulated by a series of helpful conversations with John Rundle and benefitted from valuable comments by Tom Heskes and powerful computing resources provided by Klaus Schulten. Atlee Jackson, Ken Wallace, Reimer Kühn, Ron Benson, and Wulfram Gerstner contributed with critical remarks on earlier drafts of the manuscript. Burkhard Rost, John Hopfield, Li-Waj Tang, and Tanja Diehl were sources of inspiration when it came to understanding complex biological systems of various kinds.

It is a great pleasure to acknowledge the continuous support from the above people. I would also like to thank these colleagues and friends for sharing happy and exciting times and for helping through the difficult ones. This work has been made possible through grants from the Deutsche Forschungsgemeinschaft, the Beckman Institute, and the Commission of the European Communities under the Human Capital and Mobility Programme.

1.7 References

[1] M. Abeles, *Corticonics: Neural Circuits of the Cerebral Cortex* (Cambridge University Press, Cambridge, 1991)

[2] J. LaSalle and S. Lefschetz, *Stability by Ljapunov's Direct Method* (Academic Press, New York, 1961)

[3] N. Rouche, P. Habets and M. Laloy, *Stability theory by Liapunov's direct method* (Springer Verlag, New York, 1977)

[4] J.J. Hopfield, Proc. Natl. Acad. Sci. USA **79**, 2554-2558 (1982)

[5] W.A. Little, Math. Biosci. **19**, 101-120 (1974)

[6] M.A. Cohen and S. Grossberg, IEEE Transactions SMC **13**, 815-826 (1983)

[7] J.J. Hopfield, Proc. Natl. Scad. Sci. USA **81**, 3088-3092 (1984)

[8] C.M. Marcus and R.M. Westervelt, Phys. Rev. A **40**, 501-504 (1989)

[9] E. Goles-Chacc, F. Fogelman-Soulie, and D. Pellegrin, Disc. Appl. Math. **12**, 261-277 (1985)

[10] A.V.M. Herz and C.M. Marcus, Phys. Rev. E **47** 2155-2161 (1993)

[11] E.R. Kandel and J.H. Schwartz, *Principles of Neural Science* (Elsevier, New York, 1985)

[12] A.L. Hodgkin and A.F. Huxley, J. Physiol. (London), **117**, 500-544 (1952)

[13] W.C. McCulloch and W. Pitts, Bull. Math. Biophys. **5**, 115-133 (1943)

[14] E. Ising, Z. Phys. **31**, 253 (1925)

[15] R. Eckhorn, Biol. Cybern. **60**, 121-130 (1988)

[16] C.M. Gray and W. Singer, Proc. Natl. Acad. Sci. USA **86**, 1698-1702 (1989)

[17] G. Laurent and H. Davidowitz, Science **265**, 1872-1875 (1994)

[18] W. Gerstner and J.L van Hemmen, this volume, Chapter ??

[19] M.A. Wilson and J.M. Bower, in: *Methods in Neuronal Modeling: From Synapses to Networks* edited by C. Koch and I. Segev (MIT Press, Cambridge, MA, 1989) pp. 291-334

[20] R.D. Traub and R. Miles, *Neural Networks of the Hippocampus* (Cambridge University Press, Cambridge, 1991)

[21] V. Braitenberg, in: *Brain Theory* edited by G. Palm und A. Aertsen (Springer Verlag, Berlin, Heidelberg, New York, 1986) pp. 81-96

[22] D.J. Amit, *Modeling Brain Function: The World of Attractor Neural Networks* (Cambridge University Press, Cambridge 1989)

[23] E. de Schutter and J.M. Bower, Proc. Natl. Acad. Sci. USA **91**, 4736-4740 (1994)

[24] P.C. Bressloff, Phys. Rev. A **44**, 4005-4016 (1991)

[25] W. Gerstner and J.L. van Hemmen, Biol.Cybern. **67**, 195-205 (1992); Network **3**, 139-164 (1992)

[26] B. Katz, *Nerve, Muscle, and Synapse* (McGraw-Hill, New York 1966)

[27] Y. Burnod and H. Korn, Proc. Natl. Acad. Sci. USA **86**, 352-256 (1989)

[28] P. Peretto, Biol. Cybern. **50**, 51-62 (1984)

[29] R.J. Glauber, J. Math. Phys. **4**, 294-307 (1963)

[30] H. Horner, in *Computational Systems — Natural and Artificial*, edited by H. Haken (Springer-Verlag, Berlin Heidelberg, 1988) pp. 118-132

[31] K.M. Chandy, Sci. Comp. Pr. **14**, 117-132 (1990)

[32] E. Caianiello, J.Theor.Biol. **1**, 204-235 (1961)

[33] R. Kühn and J.L. van Hemmen, in *Physics of Neural Networks*, edited by E. Domany, J.L. van Hemmen, and K. Schulten (Springer, Berlin Heidelberg, 1991) pp. 213-280

[34] A. Gabrielov, Physica A **195**, 253-274 (1993)

[35] A. Gabrielov, W.I. Newman, and L. Knopoff, Phys. Rev. E **50**, 188-196 (1994)

[36] P. Bak, C. Tang, and K. Wiesenfeld, Phys. Rev. Lett. **59**, 381-384 (1987)

[37] S. Dunkelmann and G. Radons, in *Proceedings of the International Conference on Artificial Neural Networks*, edited by M. Marimnaro and P.G. Morasso (Springer, London, 1994), pp. 867-871

[38] R.W. Kentridge, in *Computation and Neural Systems*, edited by F.H. Eeckman and J.M. Bower (Kluwer, 1994), pp. 531-535

[39] M. Usher, M. Stemmler, C. Koch, and Z. Olami, Neural Computation **6**, 795-836 (1994)

[40] J.J. Hopfield, Physics Today **46**, 40-46 (1994)

[41] J.B. Rundle, A.V.M. Herz, and J.J. Hopfield, preprint (1994)

[42] D.O. Hebb, *The Organization of Behavior* (Wiley, New York, 1949)

[43] J.H. Byrne and W.O. Berry (editors), *Neural Models of Plasticity* (Academic Press, San Diego, 1989)

[44] F. Edwards, Nature **350**, 271-272 (1991)

[45] T.H. Brown, A.H. Ganong, E.W. Kairiss, C.L. Keenan, and S.R. Kelso, in: *Neural Models of Plasticity*, edited by J.H. Byrne and W.O. Berry (Academic Press, San Diego, 1989) pp. 266-306

[46] D.W. Dong and J.J. Hopfield, Network **3**, 267-283 (1992)

[47] A.V.M. Herz, B. Sulzer, R. Kühn, and J.L.van Hemmen, Europhys. Lett. **7**, 663-669 (1988); Biol. Cybern. **60**, 457-467 (1989)

[48] S. Grossberg, J. Math. Anal. Appl. **21**, 643-694 (1968)

[49] S.I. Amari, IEEE Trans. Comp. C **21**, 1197-1206 (1972)

[50] Z. Li and A.V.M. Herz, in *Proceedings of the XI. Sitges Conference, "Neural Networks"*, edited by L. Garrido, Lecture Notes in Physics Vol. 368 (Springer Verlag, Berlin Heidelberg, 1990), pp. 287-302

[51] A.V.M. Herz, Z. Li and J.L. van Hemmen, Phys. Rev. Lett. **66**, 1370-1373 (1991)

[52] D.J. Amit, N. Brunel, and M.V. Tsodyks, J. Neurophysiol. **14**, 6445 (1994)

[53] K. Sakai and Y. Miyashita, Nature **354**, 152-155 (1991)

[54] H. Gutfreund, J.D. Reger and A.P. Young, J. Phys. A **21**, 2775-2797 (1988)

[55] D. Chazan and W. Miranker, Lin. Alg. Appl. **2** 199 (1969)

[56] G.M. Baudet, J. Assoc. Comp. Mach. **25** 226 (1978)

[57] D. Mitra, SIAM J. Sci. Stat. Comput. **8**, 43-58 (1987)

[58] P. Tseng, D.P. Bertsekas, and J.N. Tsitsiklis, Siam J. Control **28**, 678-710 (1989)

[59] J.J. Hopfield and D.W. Tank, Biol. Cybern. **52**, 141-152 (1985)

[60] J. Hertz, A. Krogh, and R.G. Palmer, *Introduction to the Theory of Neural Computation* (Addison-Wesley, Redwood City, 1991)

[61] E. Domany, J.L. van Hemmen, and K. Schulten (editors), *Physics of Neural Networks* (Springer, Berlin Heidelberg, 1991)

[62] J.L. van Hemmen and R.Kühn, in *Physics of Neural Networks*, edited by E. Domany, J.L. van Hemmen, and K. Schulten (Springer, Berlin Heidelberg, 1991), pp. 1-106

[63] D.J. Amit, H. Gutfreund and H. Sompolinsky, Phys. Rev. Lett. **55**, 1530-1533 (1985)

[64] B.M. Forrest and D.J. Wallace, in *Physics of Neural Networks*, edited by E. Domany, J.L. van Hemmen, and K. Schulten (Springer, Berlin Heidelberg, 1991), pp. 121-148

[65] C.M. Marcus and R.M. Westervelt, Phys. Rev. **A**, 347-359 (1989)

[66] T.A. Burton, Neural Networks **6**, 677-680 (1993)

[67] A.V.M. Herz, in *Proceedings, SFB Workshop, Riezlern 1991*, edited by U. Krüger (SFB 185, Frankfurt, 1992), pp. 151-164

[68] N. MacDonald, *Biological Delay Systems: Linear Stability Theory* (Cambridge University Press, Cambridge 1989)

[69] E. Niebur, H.G. Schuster, and D.M. Kammen, Phys. Rev. Lett. **67**, 2753-2756 (1991)

[70] D. Kleinfeld, Proc. Natl. Acad. Sci. USA **83**, 9469-9473 (1986)

[71] H. Sompolinsky and I. Kanter, Phys. Rev. Lett. **57**, 2861-2864 (1986)

[72] H. Gutfreund and M. Mezard, Phys. Rev. Lett. **61**, 235-238 (1988)

[73] D.J. Amit, Proc. Natl. Acad. Sci. USA **85**, 2141-2145(1988)

[74] U. Riedel, R. Kühn, and J.L. van Hemmen, Phys. Rev. **A 38**, 1105-1108(1988)

[75] D. Kleinfeld and H. Sompolinsky, Biophys. J. **54**, 1039-1051 (1988)

[76] A.C.C. Coolen and C.C.A.M. Gielen, Europhys. Lett. **7**, 281-285 (1988)

[77] K. Bauer and U. Krey, Z. Phys. B **79**, 461-475 (1990)

[78] M. Kerszberg and A. Zippelius, Phys. Scr. T **33**, 54-64 (1990)

[79] M. Bartholomeus and A.C.C. Coolen, Biol. Cybern. **67**, 285-290 (1992)

[80] T.M. Heskes and S. Gielen, Neural Networks **5**, 145-152 (1992)

[81] B. de Vries and J.C. Principe, Neural Networks **5**, 565-576 (1992)

[82] J.F. Fontanari and R. Köberle, Phys. Rev. A **36**, 2475 (1987)

[83] C.M. Marcus and R.M. Westervelt, Phys. Rev. A. **42**, 2410-2417 (1990)

[84] A.V.M. Herz, Phys. Rev. A **44**, 1415-1418 (1991)

[85] C. Peterson and J.R. Anderson, Complex Systems **1**, 995 (1987)

[86] D.H. Ackley, G.E. Hinton, and T.J. Sejnowski, Cognitive Sci. **9**, 147-169 (1985)

[87] P. Baldi and F. Pineda, Neural Comp. **3**, 526-545 (1991)

[88] E. Gardner, Europhys. Lett. **4** 481-485 (1987)

[89] K. Bauer and U. Krey, Z. Phys. B **84** 131-141 (1991)

[90] G. Mato and N. Parga, Z. Phys. B **84**, 483-486 (1991)

[91] M.C. Mackey and L. Glass, Science **197**, 287-289 (1977)

[92] J.L. Kaplan and J.A. Yorke, J. Math. Anal. Appl. **48**, 317-324 (1974)

[93] R.D. Nussbaum, J. Differential Equations **34**, 25-54 (1979)

[94] S.-N. Chow and H.-O. Walther, Trans. Amer. Math. Soc. **307**, 127-142 (1988)

[95] A.V.M. Herz, to appear in J. Differential Equations

[96] G.S. Jones, J. Math. Anal. Appl. **4**, 440-469 (1962)

[97] J. Mallet-Paret, J. Differential Equations **72**, 270-315 (1988)

[98] J.K. Hale, *Theory of Functional Differential Equations* (Springer Verlag, New York, Heidelberg, Berlin 1977)

[99] L.F. Abbott and C. van Vreeswijk, Phys. Rev E **48**, 1483-1490 (1993)

[100] M. Tsodyks, . Mitkov, and H. Sompolinsky, Phys. Rev. Lett. **71**, 1280-1283 (1993)

[101] C. van Vreeswijk and L.F. Abbott, SIAM J. Appl. Math. **53**, 253-254 (1993)

[102] R.E. Mirollo and S.H. Strgatz, SIAM J. Appl. Math. **50**, 1645-1662 (1990)

[103] L.F. Abbott, J. Phys A **23**, 3835-3859 (1990)

[104] A. Treves, Network **4**, 259-284 (1993)

[105] A.V.M. Herz and J.J. Hopfield, preprint (1994)

[106] J.J. Hopfield and A.V.M. Herz, preprint (1994)

[107] C. von der Malsburg, Internal Report 81-2, MPI for Biophysical Chemistry, Göttingen (1981)

[108] M. Usher, H. Schuster, and E. Niebur, Phys. Rev. Lett. **71**, 1280-1283 (1993)

[109] P. Bush and T. Sejnowski, preprint (1994

[110] D. Hansel, G. Mato, and C. Meunier, preprint (1994)

[111] M. Tsodyks and T. Sejnowski, preprint (1994)

[112] G. Sawitzki, *The NetWork Project*, StatLab, Universität Heidelberg, 1989. Republished on Apple Developer CD Series disk *IV* (1990); G. Sawitzki, R. Kühn, and J.L. van Hemmen (private communication)

[113] M.W. Hirsch, Neural Networks **2**, 331-349 (1989)

[114] E.K. Blum and X. Wang, Neural Networks **5**, 577-587 (1992)

[115] T. Poggio, V. Torre, and C. Koch, Nature **317**, 314-319 (1985)

[116] A.L. Yuille, Biol. Cybern. **61**, 115-123 (1989)

[117] C.M. Bachmann, L.N Cooper, A. Dembo, and O. Zeitouni, Proc. Natl. Acad. Sci. USA **84**, 7529-7531 (1987)

[118] M.A. Cohen, Neural Networks **5**, 83-103 (1992)