

Memory Capacity for Sequences in a Recurrent Network with Biological Constraints

Christian Leibold^{1,3} and Richard Kempter^{1,2,3}

¹*Institute for Theoretical Biology, Humboldt-Universität zu Berlin, Germany*

²*Bernstein Center for Computational Neuroscience Berlin, Germany and*

³*Neuroscience Research Center, Charité, Medical Faculty of Berlin, Germany*

The CA3 region of the hippocampus is a recurrent neural network that is essential for the storage and replay of sequences of patterns that represent behavioral events. Here we present a theoretical framework to calculate a sparsely connected network's capacity to store such sequences. As in CA3, only a limited subset of neurons in the network is active at any one time, pattern retrieval is subject to error, and the resources for plasticity are limited. Our analysis combines an analytical mean-field approach, stochastic dynamics, and cellular simulations of a time-discrete McCulloch-Pitts network with binary synapses. To maximize the number of sequences that can be stored in the network, we concurrently optimize the number of active neurons, i.e. pattern size, and the firing threshold. We find that for one-step associations, i.e. minimal sequences, the optimal pattern size is inversely proportional to the mean connectivity c , whereas the optimal firing threshold is independent of the connectivity. If the number of synapses per neuron is fixed, the maximum number P of stored sequences in a sufficiently large, non-modular network is independent of its number N of cells. On the other hand, if the number of synapses scales as the network size to the power of $3/2$, the number of sequences P is proportional to N . In other words, sequential memory is scalable. Furthermore, we find that there is an optimal ratio r between silent and non-silent synapses of at which the storage capacity $\alpha = P/[c(1+r)N]$ assumes a maximum. For long sequences, the capacity of sequential memory is about one order of magnitude below the capacity for minimal sequences, but otherwise behaves similar to the case of minimal sequences. In a biologically inspired scenario the information content per synapse is far below theoretical optimality, suggesting that the brain trades off error tolerance against information content in encoding sequential memories.

I. INTRODUCTION

Recurrent neuronal networks are thought to serve as a physical basis for learning and memory. A fundamental strategy of memory organization in animals and man is the storage of sequences of behavioral events. One of the brain regions of special importance for sequence learning is the hippocampus (Brun et al., 2002; Fortin et al., 2002; Kesner et al., 2002). The recurrent network in the CA3 region of hippocampus, in particular, is critically involved in the rapid acquisition of single-trial or one-shot, episodic-like memory (Nakazawa et al., 2004), i.e., memory of the sequential ordering of events.

It is generally assumed that the hippocampus can operate in at least two states (Lörincz & Buzsàki, 2000). One state, called theta, is dedicated to fast, or one-shot, learning; the other state, referred to as sharp-wave ripple, is dedicated to the replay of stored sequences. Experiments by Wilson & McNaughton (1994), Nadasdy et al. (1999) and Lee & Wilson (2002) strongly corroborate the hypothesis that the hippocampus can replay sequences of previously experienced events. The sequences are assumed to be stored within the highly plastic synapses that recurrently connect the pyramidal cells of the CA3 region (Csicsvari et al., 2000).

In this paper we tackle the problem of how many sequences can be stored in a recurrent neuronal network such that their replay can be triggered by an activation of adequate cue patterns. This question is fundamental to neural computation, and many classical papers calculate the storage capacity of pattern memories. There, one can roughly distinguish between two major classes of network models: perceptron-like feed-forward networks in which associations occur within one time-step (Brunel et al., 1992; Gardner, 1987; Nadal & Toulouse, 1990; Willshaw et al., 1969) and recurrent networks that describe memories as attractors of a time-discrete dynamics (Amit et al., 1987; Golomb et al., 1990; Hopfield, 1982; Little, 1974). Also for networks that act as memory for sequences, capacities have been calculated in both the perceptron (Nadal, 1991) and the attractor case (Herz et al., 1991). An important result is that the capacity of sequence memory in Hopfield-type networks is about twice as large as that of a static attractor network (Düring et al., 1998).

The present paper describes sequence replay in a sparsely connected network by means of time-discrete dynamics, binary neurons and binary synapses. Our model for sequential replay of activity patterns is different from attractor-type models (Amit, 1988; Buhmann & Schulten, 1987; Sompolinsky & Kanter, 1986). In fact we completely dispense with fixed points of the network dynamics. Instead, we discuss transients that are far from equilibrium (August & Levy, 1999; Jensen & Lis-

man, 2005; Levy, 1996). In the case of a sequence consisting of a single transition between two patterns (a minimal sequence) the mathematical structure we choose is similar to the one of Nadal (1991) for an auto-associative Willshaw network (Willshaw et al., 1969). For longer sequences, our analysis resembles that of synfire networks (Diesmann et al., 1999), although we take expectation values as late as possible (Nowotny & Huerta, 2003).

Some of the previous approaches to memory capacity explicitly focus on questions of biological applicability. Golomb et al. (1990), for example, address the problem of low firing rates. Herrmann et al. (1995) explore the biological plausibility of synfire chains. Other approaches assess the dependence of storage capacity on restrictions to connectivity (Deshpande & Dasgupta, 1991; Gutfreund & Mezard, 1988; Maravall, 1999) and on the distribution of synaptic states (Brunel et al., 2004).

In this paper we propose a framework that allows for discussing how a combination of several biological constraints affects the performance of neuronal networks that are operational in the brain. An important constraint that supports dynamical stability at a low level of activity is a low mean connectivity. Another one is imposed by limited resources for synaptic plasticity, i.e., not every synapse that may combinatorially be possible can really be established. This constraint sets an upper bound to the maximum connectivity between two groups of neurons that are to be associated. Moreover, the number of synapses per neuron may be limited. Another important constraint for sequential memories is the length of replayed sequences, which interferes with dynamical properties of the network. Finally, the capacity of sequential memory is also influenced by the specific nature of a neuronal structure that reads out replayed patterns. This influence is often neglected by assuming a perfect detector for the network states.

In our approach we explicitly take into account that synapses are usually classified into activated and silent ones (Montgomery et al., 2001; Nusser et al., 1998). Activated synapses have a non-zero efficacy or weight and are essential for the network dynamics. Silent synapses, which do not contain postsynaptic AMPA receptors (Isaac et al., 1995), are assumed to not contribute to the network dynamics. Changing the state of synapses from the silent to the non-silent state and vice versa acts as a plasticity resource for the storage of sequences. Synaptic learning rules can set a fixed ratio between silent and non-silent synapses, which gives rise to an additional constraint.

We calculate the capacity of sequential memory in a constrained recurrent network by means of a probabilistic theory as well as a mean-field approach. Both theoretical models are compared to cellular simulations of networks of spiking units. We thereby describe the memory capacity for sequences in dependence upon five free parameters. The network size N , the mean connectivity c ,

and the ratio r between silent and non-silent synapses are three network parameters. In addition there are two replay parameters: the sequence length Q , and the threshold γ of pattern detection.

In this article, the number M of active neurons per pattern and the neuronal firing threshold θ are largely considered as dependent variables. It is shown how M and θ are to be optimized to allow for replaying a maximum number of sequences. Scaling laws are then derived by using the optimal values for M and θ , both being functions of the five free parameters.

II. MODEL OF A RECURRENT NETWORK FOR THE REPLAY OF SEQUENCES

In this section we specify notations to describe the dynamics and morphology of a recurrent network that allows for a replay of sequences of predefined activity patterns. A list of symbols that are used throughout this manuscript can be found in Appendix A.

A. Binary Synapses Connect Binary Neurons

Let us consider a network of N McCulloch-Pitts (McCulloch & Pitts, 1943) neurons that are described by binary variables x_k , $1 \leq k \leq N$. At each discrete point in time t neuron k can either be active, $x_k(t) = 1$, or inactive, $x_k(t) = 0$. The state of the network is then denoted by a binary N -vector $\mathbf{x}(t) = [x_1(t), \dots, x_N(t)]^T$. A neuron k that is active at time t provides input to a neuron k' at time $t + 1$ if there is an activated synaptic connection from k to k' . Neuron k' fires at time $t + 1$ if its synaptic input crosses some firing threshold $\theta > 0$.

In order to specify a neuron's input, we classify synapses into *activated* and *silent* ones. All activated connections contribute equally to the synaptic input. Silent synapses have no influence on the dynamics. Therefore, the synaptic input of neuron k' at time $t + 1$ equals the number of active neurons at time t that have an activated synapse to neuron k' . Silent synapses are assumed to act as a resource for plastic changes, although this paper does not directly incorporate plasticity rules.

The total number cN^2 of activated synapses in the network defines a mean connectivity $c > 0$, which later on will be interpreted as the probability of having an activated synapse connecting a particular neuron to another one. The connectivity through activated synapses in the network is described by the $N \times N$ binary matrix C where $C_{kk'} = 1$ if there is an activated synapse from neuron k to neuron k' , and $C_{kk'} = 0$ if there is a silent synapse or no synapse at all.

Similarly, the connectivity through silent synapses is denoted by c_s , and the total number of

silent synapses in the network is $c_s N^2$. Then, each neuron then has on average $(c + c_s)N$ morphological synapses, which in turn defines the *morphological* connectivity $c_m = c + c_s$. Experimental literature (Montgomery et al., 2001; Nusser et al., 1998) usually assesses the ratio c_s/c between the silent and non-silent connectivities. For the sake of convenience, we introduce the abbreviation $r = c_s/c$. We note that the four connectivity parameters c , c_m , c_s and r have two *independent* degrees of freedom.

B. Patterns and Sequences

A pattern or event is defined as a binary N -vector ξ where M elements of ξ are 1 and $N - M$ elements are 0. The network is in the state ξ at time t if $\mathbf{x}(t) = \xi$. An ordered series of events is called a sequence. A minimal sequence is defined as a series of two events, say a cue pattern ξ^A preceding a target pattern ξ^B . The minimal sequence $\xi^A \rightarrow \xi^B$ is said to be stored in the network if initialization with the cue $\mathbf{x}(t) \approx \xi^A$ at time t leads to the target $\mathbf{x}(t+1) \approx \xi^B$ one time step later. Typically, the network only approximately recalls or replays the events of a sequence; see section IV. Sequences of arbitrary length, denoted by $Q \geq 1$, are obtained by concatenating minimal sequences of length $Q = 1$.

In the next section we specify how to set up the connectivity such that a recurrent network can act as a memory for sequences.

III. EMBEDDING SEQUENCES AND STORAGE CAPACITY

For a minimal sequence $\xi^A \rightarrow \xi^B$ to be stored in the network, one requires an above-average connectivity through activated synapses from the cells that are active in the cue ξ^A to those that are supposed to be active during the recall of the target ξ^B . In what follows, we assume that all morphological synapses from neurons active in the cue pattern to cells active in the target pattern are switched on and none of them are silent. Such a network can be constructed similar to the one in Willshaw et al. (1969); see also Nadal & Toulouse (1990) and Buckingham & Willshaw (1993): let us therefore consider a randomly connected network, i.e., the probability of having a morphological synapse from one neuron to another one is c_m . Beginning with all synapses being in the silent state, one randomly defines pairs of patterns that are to be connected into minimal sequences. Then one converts those silent synapses into active ones that connect the M active neurons in a cue pattern to the M active neurons in the corresponding target pattern. Imprinting

of sequences stops, when the overall connectivity through activated synapses reaches the value c , i.e., the total number of activated synapses in the network attains a value of cN^2 .

Capacity of Sequential Memory

Let us now address the question of how many sequences can be concurrently stored using the above algorithm for a given mean connectivity c and morphological connectivity $c_m > c$. In so doing we define the capacity α of sequential memory as the maximum number P of minimal sequences that can be stored, normalized by the number $c_m N = (1 + r)cN$ of morphological synapses per neuron,

$$\alpha := \frac{P}{c_m N} . \quad (1)$$

The number P of minimal sequences that can be stored is assessed by extending the classical derivation of Willshaw et al. (1969): suppose that we have two groups of M cells that should be linked into a minimal sequence. For each morphological synapse in the network the probability that the presynaptic neuron is active in the cue pattern is M/N , and the probability that the postsynaptic neuron is active in the target pattern is also M/N . Then the probability that a synapse is *not* involved in this specific minimal sequence is $1 - M^2/N^2$. Given P stored minimal sequences, the probability that a synapse does not contribute to any of those sequences is $[1 - M^2/N^2]^P$, and therefore the probability of a synapse being in a non-silent state is $\mathcal{C} = 1 - [1 - M^2/N^2]^P$. For a mean connectivity c , on the other hand, the probability \mathcal{C} also equals the ratio between the number cN^2 of activated synapses and the total number $c_m N^2$ of synapses in the network, i.e., $\mathcal{C} = c/c_m$. Combining the two approaches we can derive P for any given pair of connectivities c and $c_m = c(1 + r)$ and find

$$\alpha = \frac{\log(1 - c/c_m)}{c_m N \log(1 - M^2/N^2)} . \quad (2)$$

Equation 2 is valid for all biologically reasonable choices of M , c and c_m and also can account for non-orthogonal patterns, as in Willshaw et al. (1969). A somewhat simpler expression for α can be obtained in the case $M/N \ll 1$. Independently of specific values of c and c_m we can expand $[1 - M^2/N^2]^P \approx 1 - P M^2/N^2$ to end up with

$$\alpha = \frac{c N}{c_m^2 M^2} \quad \text{for } 1 \ll M \ll N . \quad (3)$$

Equation 3 can also be interpreted through a different way of estimating the number P of minimal sequences that can be stored: P roughly equals the ratio between the total number cN^2 of activated

synapses in the network and the number $c_m M^2$ of activated synapses that link two patterns, i.e., $P = cN^2/(c_m M^2)$. This estimate, however, requires that different patterns are represented by different groups of neurons, i.e., that there is no overlap between the patterns, which is an excellent approximation for sparsely-coded patterns, $M/N \ll 1$.

Equations 2 and 3 for the capacity α of sequential memory, however, do not tell us whether embedded sequences can actually be replayed. In the next section we therefore introduce a method to quantify sequence replay.

IV. REPLAYING SEQUENCES

We consider a sequence as being stored in the network if and only if it can be replayed at a given quality. In order to be able to efficiently simulate replay in large networks, this section introduces a probabilistic ‘Markovian’ dynamics that approximates the deterministic cellular simulations well. Finally, we will define a measure to quantify the quality of sequence replay.

A. Capacity and Dynamical Stability

Let us design a network and patterns such that the number of sequences that can be concurrently stored is as large as possible. From equations 2 and 3 we see that the capacity α is maximized if the pattern size M is as small as possible. However, M cannot be arbitrarily small, which will be illustrated below and explained in detail in section V.

Examples of how sequence replay depends on network parameters are illustrated by simulations of a network of $N = 100\,000$ McCulloch-Pitts units at low connectivities $c = c_s = 0.05$. The choice $r = c_s/c = 1$ roughly resembles values experimentally obtained by Nusser et al. (1998), Montgomery et al. (2001), and Isope & Barbour (2002). Minimal sequences have been concatenated so that non-minimal sequences $[\xi^0 \rightarrow \xi^1 \rightarrow \dots \rightarrow \xi^Q]$ of length $Q = 20$ are obtained. In the simulations, the network is always initialized with the cue pattern ξ^0 at time step 0. The replay of non-minimal sequences at times $t > 0$ is then indicated through two order parameters: the number of correctly activated neurons (hits), $m_t := \mathbf{x}(t) \cdot \xi^t$, and the number of incorrectly activated neurons (false alarms), $n_t := \mathbf{x}(t) \cdot (\mathbf{1} - \xi^t)$, where $\mathbf{1} = [1, \dots, 1]^T$ and the symbol ‘ \cdot ’ denotes the standard dot product.

Figure 1 shows sequence replay in cellular simulations for two different pattern sizes ($M = 800$ and $M = 1600$). Sequence replay crucially depends on the value of the firing threshold θ . In

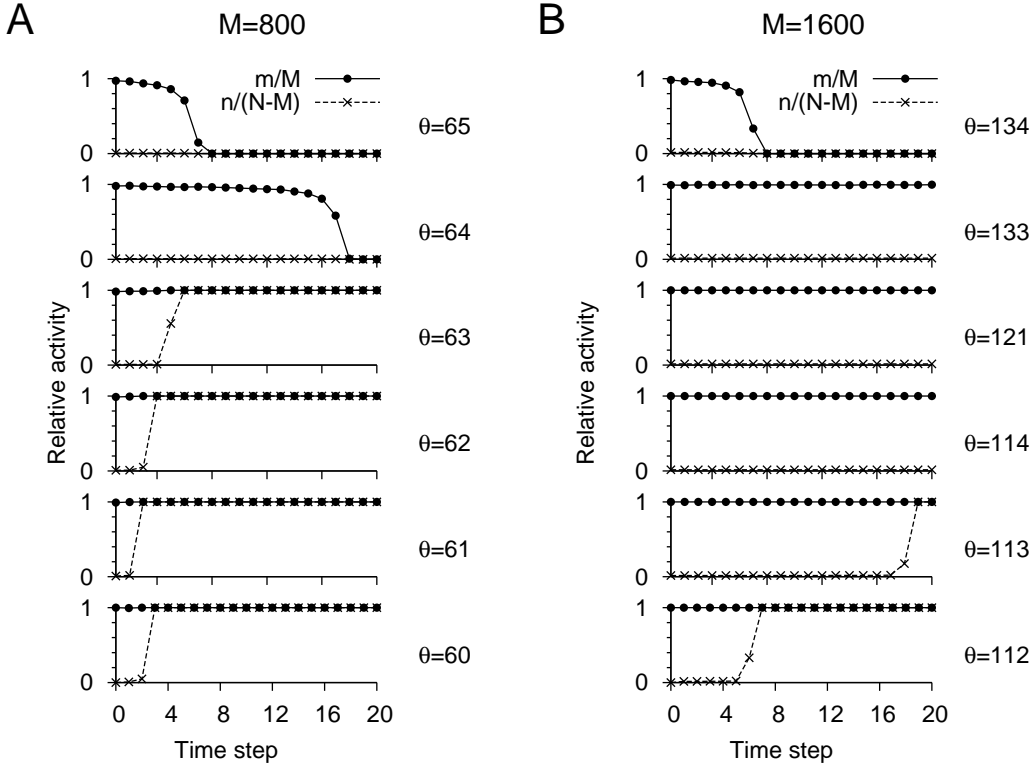


FIG. 1: Stability-capacity conflict. Sequence replay critically depends on both the firing threshold θ and the pattern size M . In all graphs we show the fraction m_t/M of hits (disks) representing sequence replay at time step t and the fraction $n_t/(N - M)$ (crosses) of false alarms in the rest of the network during the replay of a non-minimal sequence of length $Q = 20$. The network consists of $N = 10^5$ McCulloch-Pitts neurons with a mean connectivity of activated synapses of $c = 5\%$. The ratio of silent and activated synapses is $r = 1$. (A) For a pattern size $M = 800$, full replay is impossible. For high thresholds $\theta \geq 64$ the sequence dies out, whereas for low thresholds $\theta \leq 63$ the network activity explodes. (B) For a pattern size of $M = 1600$, sequence replay is possible for a broad range of thresholds θ between 114 and 133.

general, if the threshold is too high, the network becomes silent after a few iterations. On the other hand, if the threshold is too low, the whole network becomes activated within few time steps. Whether there exist values of θ at which a sequence can be fully replayed, however, also critically depends on the pattern size M . At a small pattern size of $M = 800$ there is no such firing threshold whereas for a pattern size $M = 1600$ there is a whole range of thresholds that allow replaying the full sequence. So, there is a conflict between the maximization of the capacity of the network, which requires M to be small, and the dynamical stability of replay, which becomes more robust for larger values of M ; cf. section VII

In section V we will derive a lower bound for the pattern size below which replay is impossible,

and we also determine the respective firing threshold. In connection with equation 2 these results enable us to calculate the maximum number of sequences that can be simultaneously stored in a recurrent network such that all stored sequences can be replayed. These calculations require a simultaneous optimization of pattern size M and threshold θ . A numerical treatment as shown in Figure 1, however, is infeasible for much larger networks. Therefore, we next introduce a numerically less costly approach.

B. Markovian Dynamics

Assessing the dynamics of large networks of neurons by means of computer simulations is mainly constrained by the amount of accessible memory. Simulations of a network of $N = 10^5$ cells with a connectivity of about $c = 5\%$, as the ones shown in Figure 1, require about 2 GB of computer memory. A doubling of neurons would therefore result in 8 GB and is thus already close to the limit of nowadays conventional computing facilities. Networks with more than 10^6 cells that need at least 200 GB are very inconvenient. It is therefore reasonable to follow a different approach for investigating scaling laws of sequential memory.

To be able to simulate pattern replay in large networks, we reduce the dynamical degrees of freedom of the network to the two order parameters defined in the previous section: the number m_t of correctly activated neurons (hits) and the number n_t of incorrectly activated neurons (false alarms) at time t ; see also Figure 2A. Furthermore, we take advantage of the fact that the network dynamics has only a one-step memory and, thus, reduce the full network dynamics to a discrete probabilistic dynamics governed by a transition matrix T (Gutfreund & Mezard, 1988; Nowotny & Huerta, 2003). The transition matrix is defined as the conditional probability $T(m_{t+1}, n_{t+1} | m_t, n_t)$ that a network state (m_{t+1}, n_{t+1}) follows the state (m_t, n_t) . We note that, due to this probabilistic interpretation the dynamics of (m_t, n_t) is stochastic, although single units behave deterministically. More precisely, we derive a dynamics for the probability distribution of (m_t, n_t) . How to interpret expectation values with respect to this distribution is specified next.

1. Reduced Connectivity Matrix

In the limit of a large pattern size M , the connectivities c and c_m can be interpreted as probabilities of having synaptic connections. In other words, the probability that in the embedded sequence $\xi^A \rightarrow \xi^B$ there is an activated synapse from a cell active in ξ^A to a cell active in ξ^B is c_m .




FIG. 2: Pattern size and connectivity matrix. (A) At some time t , the network is assumed to be associated with a specific event $\xi^t = \xi^A$ of size M . We therefore divide the network of N neurons into two parts. The first part consists of the M neurons that are supposed to be active if an event ξ^t is perfectly represented. The second part contains the $N - M$ neurons that are supposed to be inactive. The quantities m_t (hits) and n_t (false alarms) denote the number of active neurons in the two groups at time t . (B) The number m_{t+1} of hits and the number n_{t+1} of false alarms with respect to pattern $\xi^{t+1} = \xi^B$ at time $t + 1$ are determined by the state of the network at time t , $\mathbf{x}(t) = \xi^A$, and the connectivity matrix C . The average number of synaptic connections between the four groups of cells is described by the reduced connectivity matrix $\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix}$, which is defined in section IV B 1.

This probabilistic interpretation can be formalized by means of a reduction of the binary connectivity matrix C to four mean connectivities $\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix}$, which are average values over all P minimal sequences stored; see also Figure 2B. First, we define the mean connectivity c_{11} between neurons that are supposed to be active in cue patterns and those that are supposed to be active in their corresponding targets,

$$c_{11} = \frac{1}{P} \sum_{\{A,B\}} \frac{1}{N^2} \sum_{k,k'=1}^N \xi_k^A C_{kk'} \xi_{k'}^B. \quad (4)$$

Here the sum over $\{A, B\}$ is meant to be taken over the cue-target pairs of P different minimal sequences. By construction (see section III) c_{11} is at its maximum c_m . Second, the connectivity c_{10} describes activated synapses between cells that are active in cue patterns to cells that are supposed to be inactive in target patterns. Similarly, the mean connectivity c_{01} describes activated synapses from neurons that are supposed to be inactive in the cue to those that should be active in the target pattern. Finally, c_{00} denotes the mean connectivity between cells that are to be silent in both the cue and the target. The four mean connectivities are summarized in the reduced connectivity matrix $\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix}$; see also Figure 2B. The interpretation of the mean connectivities as probabilities

of having activated synaptic connections between two neurons can be considered as the assumption of *binomial statistics*. This assumption is a good approximation for Willshaw-type networks in the limit $N \gg M \gg 1$ (Buckingham & Willshaw, 1992).

Cues and targets of minimal sequences were assumed to be linked as tight as possible, which results in $c_{11} = c_m = c(1+r)$. The remaining three entries of the reduced connectivity matrix follow from normalization conditions: since every active neuron in a target pattern, e.g. ξ^B , receives, on average, cN activated synapses and those synapses originate from two different groups of neurons in a cue pattern, e.g. ξ^A , we have $cN = c_{11}M + c_{01}(N - M)$. Similarly, every inactive neuron in the target pattern receives, on average, $cN = c_{10}M + c_{00}(N - M)$ activated synapses. As a consequence of recurrence, every neuron of a cue pattern *projects*, on average, to cN postsynaptic neurons. From that we obtain two similar conditions with c_{10} and c_{01} interchanged and thus $c_{10} = c_{01}$.

All entries of the reduced connectivity matrix $\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix}$ can therefore be expressed in terms of the mean connectivity c , the ratio r of silent and non-silent connectivities, the pattern size M , and the network size N ,

$$\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix} = c \begin{pmatrix} 1+r & 1-r\frac{M}{N-M} \\ 1-r\frac{M}{N-M} & 1+r\frac{M^2}{(N-M)^2} \end{pmatrix}. \quad (5)$$

The assumption of binary statistics together with the reduced connectivity matrix enable us to specify the transition matrix T as it has been defined at the beginning of section IV B

Calculation of the capacity α for an arbitrary connectivity c_{11} , i.e., $c < c_{11} < c_m$, between cue and target patterns is somewhat more involved than in the case of section III where patterns were connected with the maximum morphological connectivity $c_{11} = c_m$. The scenario $c < c_{11} < c_m$ is outlined in Appendix B. For $1 \ll M \ll N$, however, equation 3 with c_m replaced by c_{11} turns out to be an excellent approximation.

2. Transition Matrix

Due to statistical independence of the activation of different postsynaptic neurons, the transition matrix can be separated,

$$T(m_{t+1}, n_{t+1} | m_t, n_t) = p(m_{t+1} | m_t, n_t) q(n_{t+1} | m_t, n_t), \quad (6)$$

where $p(m_{t+1} | m_t, n_t)$ is the probability that, at time $t + 1$, a number of m_{t+1} cells are correctly activated, and $q(n_{t+1} | m_t, n_t)$ is the probability of having n_{t+1} cells incorrectly active, given m_t and

n_t . Defining the binomial probability

$$b_{j,l}(x) = \binom{l}{j} x^j (1-x)^{l-j}, \quad (7)$$

with $0 \leq x \leq 1$ and $0 \leq j \leq l$ we obtain

$$p(m_{t+1}|m_t, n_t) = b_{m_{t+1}, M}(\rho_{m_t n_t}) \quad \text{and} \quad q(n_{t+1}|m_t, n_t) = b_{n_{t+1}, N-M}(\lambda_{m_t n_t}) \quad (8)$$

with $\rho_{m_t n_t}$ and $\lambda_{m_t n_t}$ denoting the probabilities of correct (ρ) and incorrect (λ) activation of a single cell, respectively. Both are specified by the reduced connectivity matrix $\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix}$ and the firing threshold θ ,

$$\rho_{m_t n_t} = \sum_{j,k; j+k \geq \theta} b_{j,M} \left(\frac{m_t}{M} c_{11} \right) b_{k, N-M} \left(\frac{n_t}{N-M} c_{01} \right), \quad (9)$$

$$\lambda_{m_t n_t} = \sum_{j,k; j+k \geq \theta} b_{j,M} \left(\frac{m_t}{M} c_{10} \right) b_{k, N-M} \left(\frac{n_t}{N-M} c_{00} \right). \quad (10)$$

Equations 9 & 10 can be understood as adding up the probabilities of all combinations of the number j of hits and the number k of false alarms that together cross the firing threshold θ .

The transition matrix T gives rise to probability distributions Λ_t for the number m_t of hits and the number n_t of false alarms. To be able to compare the Markovian dynamics with the network dynamics obtained from cellular simulations (see Figure 1) we calculate the expectation values $\langle m_t \rangle$ and $\langle n_t \rangle$ of hits and false alarms with respect to the probability distribution Λ_t for $t \geq 1$:

$$\langle m_t \rangle = \sum_{\mu=1}^M \sum_{\nu=1}^{N-M} \mu \Lambda_t(\mu, \nu | m_0, n_0) \quad (11)$$

$$\langle n_t \rangle = \sum_{\mu=1}^M \sum_{\nu=1}^{N-M} \nu \Lambda_t(\mu, \nu | m_0, n_0), \quad (12)$$

where

$$\Lambda_t(m_t, n_t | m_0, n_0) = \sum_{\{(m_1, n_1)\}} \cdots \sum_{\{(m_{t-1}, n_{t-1})\}} \prod_{\tau=1}^t T(m_\tau, n_\tau | m_{\tau-1}, n_{\tau-1}) \quad (13)$$

is the probability of having m_t hits and n_t false alarms at some time t , given that the network has been initialized with $m_0 = M$ hits and $n_0 = 0$ false alarms at time zero. Equation 13 can be derived from the recursive formula $\Lambda_t(\cdot|\cdot) = \sum_{\{(\cdot)\}} T(\cdot|\cdot) \Lambda_{t-1}(\cdot|\cdot)$, and the sums in equation 13 are meant to be over all pairs $(m_\tau, n_\tau) \in \{0, \dots, M\} \otimes \{0, \dots, N-M\}$, for $1 \leq \tau \leq t-1$.

An increase in numerical efficiency is gained from the fact that sums over binomial probabilities can be evaluated by means of the incomplete Beta function (Press et al., 1992). Moreover, numerical

treatment of the Markovian dynamics can take advantage of the separability of $T = pq$ (see equation 6). But still, for large numbers of N , computing and multiplying p and q in full is costly. We therefore reduced p and q to at most 5000 interpolation points where each of them is assigned to the same portion of probability. The reduced vectors are then used to calculate an iteration step $t \rightarrow t + 1$. Numerical results provided are thus estimates in the above sense and serve as approximations to the full Markov dynamics.

Figure 3 shows a numerical evaluation of the Markovian dynamics for the same parameter regime as used for the cellular simulations in Figure 1. We observe a qualitative agreement between the two approaches but also small differences regarding the upper and lower bounds for the set of firing thresholds allowing for infinite stable sequence replay. A further comparison is postponed to section V.

C. Quality of Replay and Detection Criterion

In the examples shown in Figures 1 and 3 it is obvious to tell the quality of sequence replay at a certain time step because we typically have to distinguish among only three different scenarios: (1) all neurons are silent, (2) all neurons are active, and (3) a pattern is properly represented. If, however, the network exhibits intermediate states one needs a quantitative measure of whether a particular sequence is actually replayed or not. For this purpose we specify the quality Γ at which single patterns ξ^t are represented by the actual network state \mathbf{x}_t . We consider Γ to be a function of the numbers m_t and n_t of hits and false alarms, respectively (see section IV B). The quality function

$$\Gamma(m_t, n_t) := m_t/M - n_t/(N - M). \quad (14)$$

is chosen such that a perfect representation of a pattern is indicated by $\Gamma = 1$. Random activation of the network, on the other hand, yields $|\Gamma| \ll 1$ in the generic scenario $1 \ll M \ll N$. The quality function weighs hits much stronger than false alarms, similar to the so-called “normalized winner-take-all recall” as introduced by Graham & Willshaw (1997) or Maravall (1999). Equation 14 is physiologically inspired by a downstream neuron receiving excitation from hits and inhibition from false alarms.

We say a pattern to be replayed correctly at time t if the detection criterion

$$\Gamma(m_t, n_t) \geq \gamma' \quad (15)$$

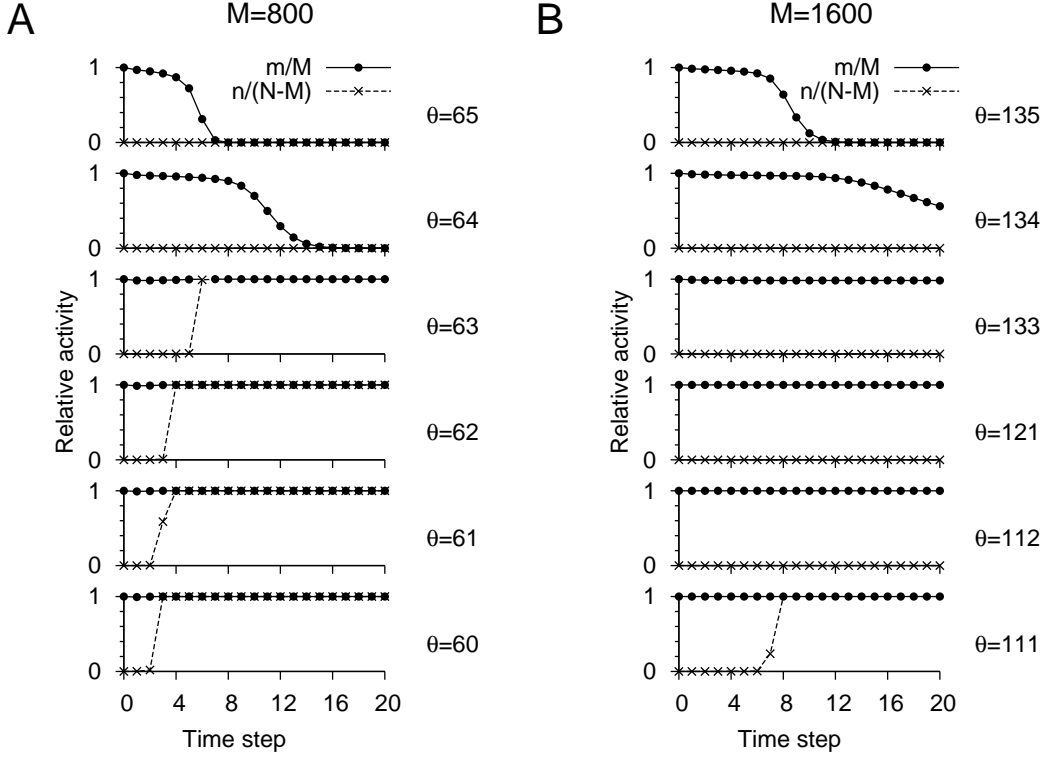


FIG. 3: Stability-capacity conflict for Markovian network dynamics. The expected fraction of hits $\langle m_t \rangle / M$ (disks) and false alarms $\langle n_t \rangle / (N - M)$ (crosses) are plotted as a function of time t after the network has been initialized with the cue pattern at $t = 0$. The parameters $N = 10^5$, $c = 5\%$, $r = 1$, and $Q = 20$ are the same as in Figure 1. (A) For a pattern size $M = 800$, full replay is impossible. For high firing thresholds $\theta \geq 64$ the sequence dies out, whereas for low thresholds $\theta \leq 63$ the network activity explodes, which is identical to Figure 1 although the time courses of hits and false alarms slightly differ. (B) For $M = 1600$, sequence replay is possible for thresholds $112 \leq \theta \leq 133$, whereas in Figure 1 we have obtained $114 \leq \theta \leq 133$.

is satisfied where γ' denotes the threshold of detection.

A sequence of Q patterns is said to be replayed if the final target pattern in the Q^{th} time step is correctly represented, i.e., $\Gamma(m_Q, n_Q) \geq \gamma'$. Here, we implicitly assume that all the patterns of a sequence are represented at least as proper as the last one.

Similarly to equation 15, we specify a detection criterion for sequence replay approximated by the Markovian dynamics,

$$\langle \Gamma(m_Q, n_Q) \rangle = \Gamma(\langle m_Q \rangle, \langle n_Q \rangle) \geq \gamma, \quad (16)$$

where the expectation values $\langle m_Q \rangle$ and $\langle n_Q \rangle$ are obtained from the Q -times iterated transition matrix T^Q for the initial condition $(m_0, n_0) = (M, 0)$.

The criteria 15 and 16 are obviously different. For $1 \ll M \ll N$, however, they are almost equivalent with $\gamma \approx \gamma'$ because the distribution of the quality measure Γ is typically unimodal and sharply peaked with variance below $1/(4M) + 1/[4(N - M)]$. Moreover, we will see in the next section that the specific value of the detection threshold does not affect scaling laws for sequential memory.

V. SCALING LAWS FOR MINIMAL SEQUENCES

The capacity α of sequential memory is proportional to M^{-2} ; see equation 3. In order to maximize α one therefore seeks a minimal pattern size M at which the replay of sequences serves a given detection threshold γ . In this section we assess this minimal pattern size for minimal sequences ($Q = 1$) and sparse patterns ($1 \ll M \ll N$). In particular, we explain why the minimal pattern size is independent of the network size N .

In the case $1 \ll M \ll N$, the reduced connectivity matrix in equation 5 can be approximated through $\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix} \approx c \begin{pmatrix} 1+r & 1 \\ 1 & 1 \end{pmatrix}$; neurons that are active in cue patterns are connected to neurons that should be active in target patterns with probability $c_m = c(1 + r)$. Otherwise, the connectivity is about c ; see Figure 2.

A. Hits and False Alarms in Pattern Recall

At some time t only those M neurons are supposed to be active that belong to the cue pattern ξ^A . We then require a particular minimal sequence $\xi^A \rightarrow \xi^B$ to be imprinted such that, at time $t + 1$, event ξ^B is recalled. We have assumed that the number j of inputs to each of the M ‘on’ neurons that should be active at time $t + 1$ is binomially distributed with probability distribution $b_{j,M}(c + c_s)$ (equation 7). In the same way the input distribution for the $N - M$ ‘off’ cells that should be inactive at time $t + 1$ is $b_{j,M}(c)$. As a result, a neuron that is supposed to be inactive receives, on average, input through cM activated synapses with a standard deviation $\sqrt{c(1 - c)M}$. To avoid unintended firing we require a firing threshold θ that is somewhat larger than cM . The larger the threshold the more noise-induced firing due to fluctuations in the number of synapses is suppressed. Let us take a threshold $\theta = cM + \kappa_+ \sqrt{c(1 - c)M}$ where κ_+ is a threshold parameter that determines the number of incorrectly activated neurons (Brunel et al., 2004), called false alarms. For $\kappa_+ = 1$, for example, we have $n_{t+1} \approx 0.16(N - M)$ false alarms; see Figure 4A. On the other hand, the threshold θ has to be small enough so that a neuron that is supposed to be

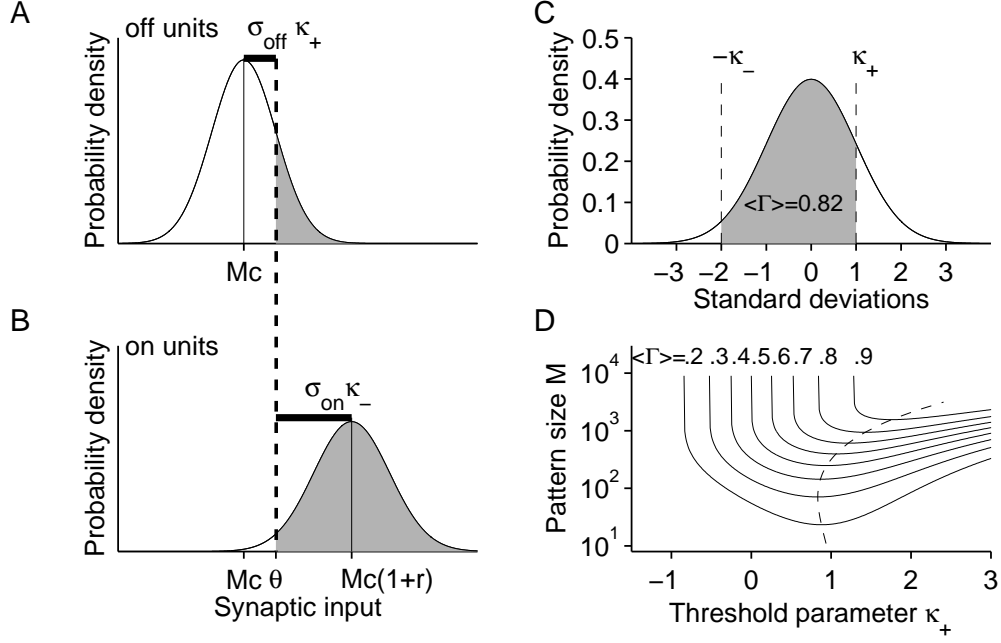


FIG. 4: Mean quality $\langle \Gamma \rangle$ of replay and threshold parameters κ_+ and κ_- . (A) Probability density of the number of synaptic inputs for ‘off’ units, which are supposed to be inactive during the recall of a target pattern. The vertical dashed line indicates the firing threshold θ . The gray area represents the probability $\langle n \rangle / (N - M)$ of having a false alarm. (B) Same as in A but for ‘on’ units, which are supposed to be active. The gray area represents the probability $\langle m \rangle / (N - M)$ of having a hit. (C) For $1 \ll M \ll N$, the binomial distributions in A and B can be approximated by normal distributions. The probability of hits minus that of false alarms equals the gray area under the normal distribution between $-\kappa_-$ and κ_+ . From equation 19 we see that this area can also be interpreted as the mean quality $\langle \Gamma \rangle$ of replay. (D) Pattern size M as a function of κ_+ for different replay qualities $\langle \Gamma \rangle$ at constant $r = 1$ and $c = 0.01$; see equations 17 and 20. The dashed line connects the minima of M .

active during event ξ^B is indeed activated. Each of these neurons receives, on average, $c_m M$ inputs with standard deviation $\sqrt{c_m(1 - c_m)M}$. A recall of ξ^B is therefore achieved by a threshold that is somewhat smaller than $c_m M$, i.e., $\theta = c_m M - \kappa_- \sqrt{c_m(1 - c_m)M}$ where κ_- is another threshold parameter that determines the number of correctly activated neurons, called hits. For $\kappa_- = 2$, for example, we have $m_{t+1} \approx 0.98 M$ hits; see Figure 4B.

The firing threshold θ is assumed to be the same for all neurons. Hence, combining above two conditions we find

$$cM + \kappa_+ \sqrt{cM(1 - c)} = c_m M - \kappa_- \sqrt{c_m M(1 - c_m)},$$

which then leads to expressions for the pattern size

$$M = \frac{1}{c} \left(\frac{\kappa_+ \sqrt{1-c} + \kappa_- \sqrt{[r+1][1-c(1+r)]}}{r} \right)^2 \quad (17)$$

and the firing threshold

$$\theta = cM + \kappa_+ \sqrt{c(1-c)M} . \quad (18)$$

The pattern size M in equation 17 is independent of the network size N and scales like c^{-1} for small values of c . Moreover, the firing threshold θ in equation 18 is independent of the network size N . For small mean connectivities c , the firing threshold θ is also independent of c . We emphasize that the validity of these scaling laws requires an almost perfect initialization of the cue pattern.

B. Optimal Pattern Size and Optimal Firing Threshold

We now argue that the firing-threshold parameters κ_+ and κ_- in equation 17 can be chosen such that M is minimal and, hence, the storage capacity is maximal. As indicated by the gray areas of the binomial distributions in Figures 4A and 4B, κ_+ and κ_- determine the mean numbers of false alarms $\langle n \rangle$ and hits $\langle m \rangle$, respectively. For $1 \ll M \ll N$, these binomial distributions are well approximated by Gaussians, and we have $\langle n \rangle / (N - M) = [1 - \text{erf}(\kappa_+ / \sqrt{2})] / 2$ and $\langle m \rangle / M = [\text{erf}(\kappa_- / \sqrt{2}) + 1] / 2$ where the error function $\text{erf}(x) := 2 / \sqrt{\pi} \int_0^x dt \exp(-t^2)$ is the cumulative distribution of a Gaussian. These approximations yield

$$\langle \Gamma(m, n) \rangle = \left[\text{erf}(\kappa_- / \sqrt{2}) + \text{erf}(\kappa_+ / \sqrt{2}) \right] / 2 , \quad (19)$$

which can be interpreted as the area under a normal distribution between $-\kappa_-$ and $+\kappa_+$; see Figure 4C.

From equation 19 we see that, for a given mean quality $\langle \Gamma \rangle$ of replay, the threshold parameters κ_+ and κ_- are not independent. More precisely, for some given detection criterion $\gamma = \langle \Gamma \rangle$ and $\kappa_+ > \sqrt{2} \text{erf}^{-1}(2\gamma - 1)$, equation 19 yields

$$\kappa_- = \sqrt{2} \text{erf}^{-1}[2\gamma - \text{erf}(\kappa_+ / \sqrt{2})] . \quad (20)$$

For fixed $\langle \Gamma \rangle = \gamma$ one therefore can choose κ_+ in equation 17 such that the pattern size M becomes minimal. At this minimal pattern size the capacity α in equation 3 reaches its maximum, and encoding of events is as sparse as possible. Let us therefore call this minimum value of M the

optimal pattern size M_{opt} for sequential memory. The dashed line in Figure 4D indicates that $M_{\text{opt}} := \min_{\kappa_+} M$ is located at values $\kappa_+ \gtrsim 1$. We also observe that the larger the detection threshold γ the larger is M_{opt} .

From equation 17 we find that for small connectivities $c \ll 1$, as they occur in biological networks, the minimum pattern size M_{opt} can be phrased as

$$M_{\text{opt}} = \frac{1}{c} [\mathcal{M}(r, \gamma) + \mathcal{O}(c)] \quad (21)$$

where $\mathcal{M}(r, \gamma)$ is a function of r and γ that has to be obtained by numerical minimization. Here, the order function $\mathcal{O}(c^k)$ is defined through $\lim_{c \rightarrow 0} c^{-k} \mathcal{O}(c^k) = \text{const.}$ for $k > 0$. At values $r = 1$ and $\gamma = 0.7$, for example, we have $\mathcal{M} = 6.1 \gg c$ corroborating the scaling law $M_{\text{opt}} \propto c^{-1}$.

For an optimal pattern size M_{opt} we can find the optimal firing threshold θ_{opt} from equation 18. In first approximation, θ_{opt} is independent of the connectivity c and the network size N , but depends on r and γ ,

$$\theta_{\text{opt}} = \mathcal{T}(r, \gamma) + \mathcal{O}(c) . \quad (22)$$

For example, $r = 1$ and $\gamma = 0.7$ account for $\theta_{\text{opt}} \approx 9.1 \gg c$.

The dependencies of M_{opt} and θ_{opt} on the connectivity c are indicated in Figure 5 through solid lines. Both M_{opt} and θ_{opt} increase with increasing detection threshold γ . These mean-field results match numerical simulations well: in cellular network simulations (open circles in Figure 5), M_{opt} and θ_{opt} were determined as the minimal M and the corresponding θ that account for replay at a fixed detection threshold $\gamma = 0.5$. The numerical evaluation of the Markovian network dynamics as defined in section IV B (full symbols in Figure 5) confirm the analytical results for a wider range of c and γ .

The lower bound M_{opt} for the pattern size in equation 21 enables us to determine an upper bound for the capacity α of sequential memory. Combining equations 3 and 21 we find

$$\alpha = cN \frac{1}{(1+r)^2 \mathcal{M}(r, \gamma)^2} + \mathcal{O}(c^2) . \quad (23)$$

We note that α is linear in the connectivity c and the network size N , decreases with increasing γ , and has a non-trivial dependency on the plasticity resources r that will be evaluated below. This scaling law for minimal sequences can now be used to study the storage of sequences in biologically feasible networks that face certain constraints.

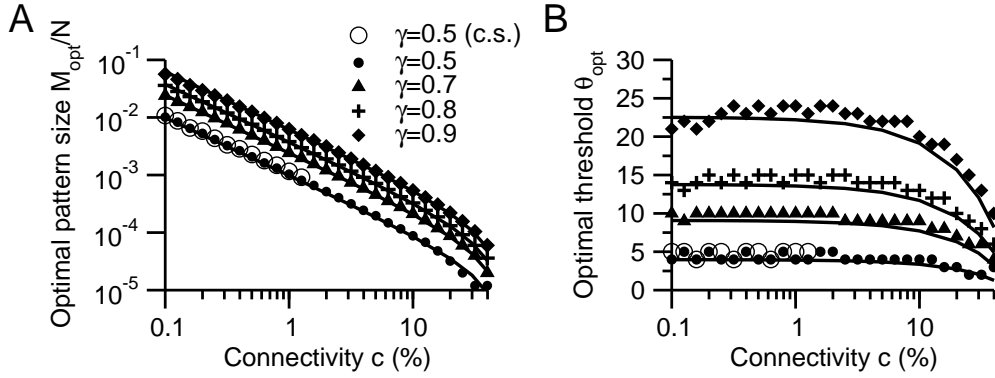


FIG. 5: Optimal pattern size M_{opt} and optimal firing threshold θ_{opt} . Lines depict results from the mean-field theory (equations 21 and 22). We also show numerical simulations (c.s.) of the network introduced in section IV A (open circles, $\gamma = 0.5$) and Markovian dynamics defined in section IV B (full symbols, $\gamma = 0.5, 0.7, 0.8, 0.9$). (A) For small connectivities c , the optimal pattern size M_{opt} scales like c^{-1} and increases with increasing γ . (B) The optimal threshold θ_{opt} is almost independent of the connectivity c for $c \lesssim 10\%$, and θ_{opt} increases with increasing γ . Further parameters: sequence length $Q = 1$, network size $N = 250\,000$, plasticity resources $r = 1$. For the Markovian dynamics, we used Brent’s method (Press et al., 1992) to numerically find a firing threshold θ as a root of the implicit equation $\langle m_Q \rangle / M - \langle n_Q \rangle / (N - M) = \gamma$, which is the detection criterion. By subsequently reducing M we end up with a minimal value M_{opt} for which the detection criterion $\langle \Gamma \rangle = \gamma$ can be fulfilled. The threshold root that is obtained at M_{opt} is called θ_{opt} .

VI. CONSTRAINED SEQUENCE CAPACITY

Biological and computational networks generally face certain constraints. Those constraints can lead to limiting values and interdependencies of the network parameters c , N and r . Some constraints and their implications on the optimization of the capacity α of sequential memory in equation 23 are discussed in this section.

A. Limited Number of Synapses per Neuron

A biological neuron may have a limited number cN of synapses. If cN is constant we find from equation 23 (for constant r and γ)

$$\alpha = \text{const.} \quad \text{and} \quad P = \text{const.}$$

Increasing the capacity α can therefore not be achieved by increasing N . Numerical results in Figure 6A (symbols) confirm this behavior for $c \ll 1$. For $r = 1$, the capacity α reaches its

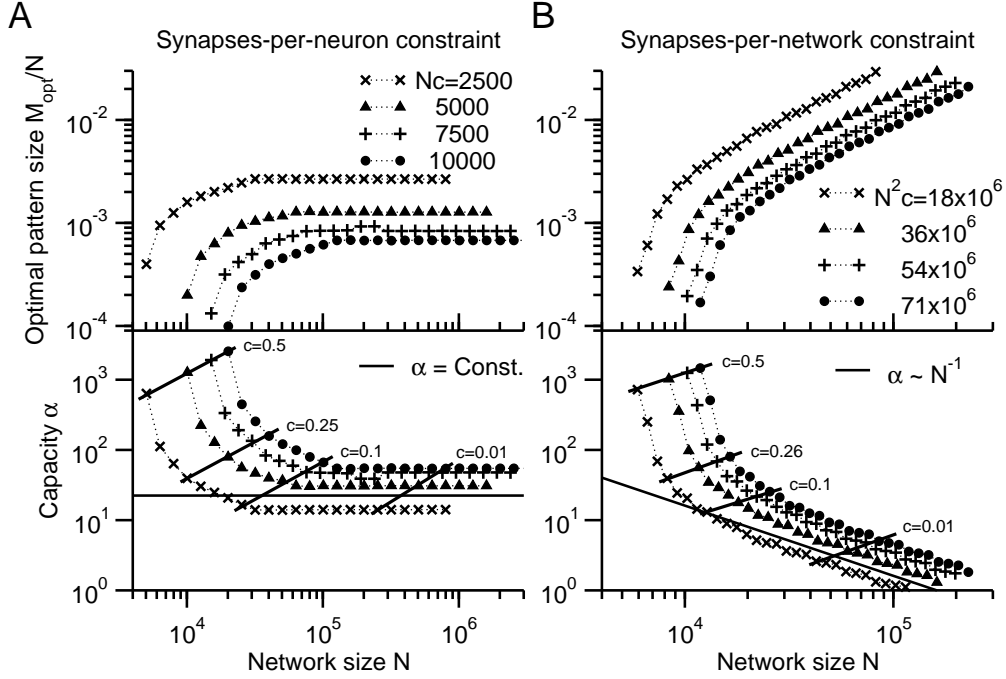


FIG. 6: Influence of constraints on the optimal pattern size M_{opt} and the capacity α of sequential memory. (A) Synapses-per-neuron constraint. For a fixed number cN of synapses per neuron we find $M_{\text{opt}} \propto N$ and $\alpha = \text{const.}$ as $N \rightarrow \infty$. The capacity α increases with increasing cN . Tilted solid lines connect symbols that refer to constant connectivities, e.g. $c = 0.01, 0.1, 0.25, 0.5$. (B) Synapses-per-network constraint. For a fixed number cN^2 of synapses in the network we find $M_{\text{opt}} \propto N^2$ and $\alpha \propto N^{-1}$ as $N \rightarrow \infty$. The capacity α increases with increasing cN^2 . For both constraints, $cN = \text{const.}$ and $cN^2 = \text{const.}$, there is an optimal network size at which the capacity α reaches its maximum. For $r = 1$ this maximum occurs at $c \approx 0.5$. A further increase in c is impossible since the morphological connectivity $c_m = c(1+r)$ cannot exceed 1. Other parameters are $Q = 1$ and $\gamma = 0.7$. Dotted lines link symbols and are *not* obtained by mean-field theory.

maximum at $c \approx 0.5$ where we have $c_m = c(1+r) = 1$, and the network can be considered an undiluted Willshaw one. For $c \rightarrow 0.5$, the scaling law $\alpha = \text{const.}$ (solid line) underestimates the storage capacity because the $\mathcal{O}(c^2)$ term in the mean-field equation 23 has been neglected.

In biologically relevant networks we typically have $c \ll 1$, and thus, for $cN = \text{const.}$, we face the scaling law $P = \alpha(1+r)cN = \text{const.}$ Therefore, the number cN of synapses a single neuron can support fully determines the network's computational power for replaying sequences in the sense that adding more neurons to the network does not increase α or P .

In the CA3-region of the rat hippocampus, for example, we have $c_m N \approx 12000$ recurrent synapses at each pyramidal cell; see Urban et al. (2001) for a review. The network size of CA3 is

$N \approx 2.4 \cdot 10^5$ (Rapp & Gallagher, 1996). From these numbers, $r = 1$ and $cN = \text{const.}$ we derive the connectivity $c \approx 0.025$. A comparison of these numbers with Figure 6A leads to estimates for the minimal pattern size being in the order of $M_{\text{opt}} \approx 200$ cells, a storage capacity of $\alpha \approx 15$ minimal sequences per synapse at a neuron and $P \approx 1.8 \cdot 10^5$ minimal sequences per CA3 network. The saturation of α and P at about $N = 10^5$ for $cN = 6000$ (see Figure 6A) may explain why the CA3 region has relatively few neurons ($N \lesssim 10^6$ in humans) despite its seminal importance for episodic memory.

B. Limited Number of Synapses in the Network

Numerical simulations of artificial networks are constrained by the available computer memory, which limits the number cN^2 of activated synapses in the network. For $cN^2 = \text{const.}$ we find from equation 23 (for constant r and γ)

$$\alpha \propto N^{-1} \quad \text{and} \quad P \propto N^{-2}.$$

Therefore an increase in both α and P can only be achieved by *reducing* the network size N at the expense of *increasing* the connectivity c . Numerical results in Figure 6B confirm this behavior for $c \ll 1$. The capacity α increases with increasing c and, for $r = 1$, assumes its maximum at the upper bound $c = 0.5$ when $c_m = 1$. For $c \rightarrow 0.5$, the scaling law $\alpha \propto N^{-1}$ (solid line) underestimates the storage capacity, similar as in Figure 6A.

We conclude that computer simulations of neural networks with constant cN^2 perform worse in storing sequences the more the connectivity resembles the biologically realistic scenario $c \ll 1$.

C. On the Ratio of Silent and Activated Synapses

In the previous two sections we have assumed a constant ratio r between the connectivity c_s through silent synapses and the connectivity c through non-silent synapses. The specific choice $r = 1$ was motivated by neurophysiological estimates from Nusser et al. (1998), Montgomery et al. (2001). We now focus on how the storage capacity α depends on this ratio r assuming that the total number $c_m N$ of morphological synapses per neuron is constant. Because of $c_m = c(1 + r)$, an increase in r increases c_s but reduces c . We note that this constraint is equivalent to a fixed number cN of activated synapses per neuron for constant r , a scenario which was evaluated in section VI A.

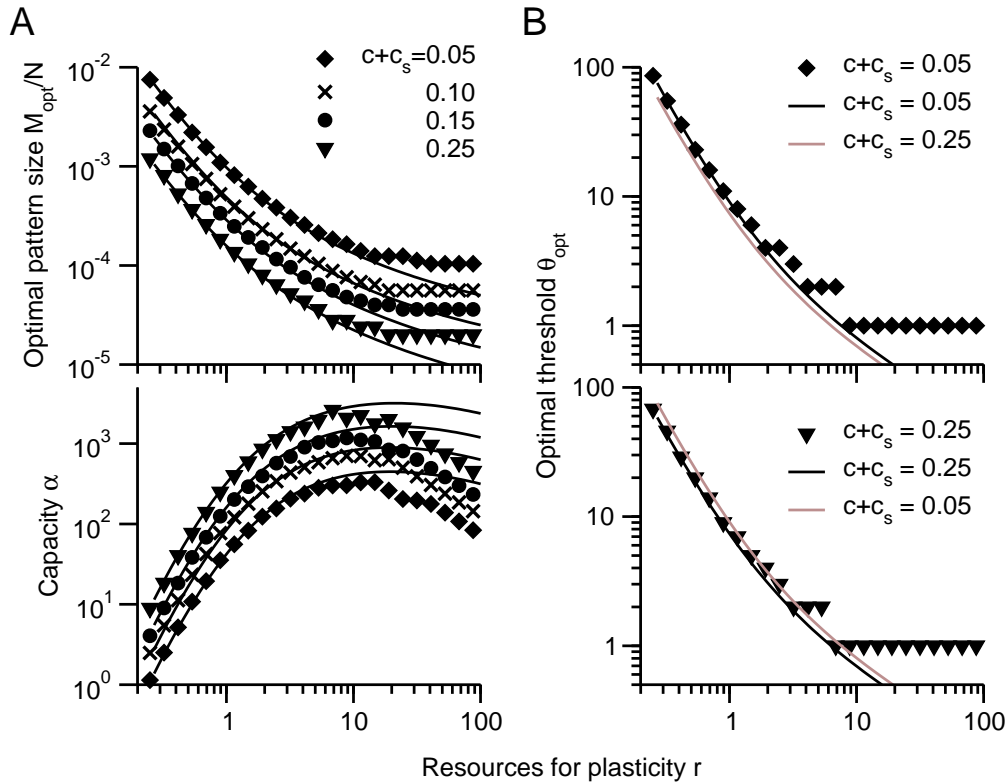


FIG. 7: Dependence of sequence replay on the resources r of synaptic plasticity for a constant total number $(c+c_s)N$ of synapses per neuron. Mean-field theory (solid lines) explains numerical results obtained from the Markovian dynamics (symbols) well as long as $r < 10$ and $\theta_{\text{opt}} \gtrsim 4$. Below $\theta_{\text{opt}} \lesssim 4$ the discreteness of θ_{opt} limits the validity of the mean-field theory. (A) The optimal pattern size M_{opt} (top) decreases with increasing r and saturates at values $(c+c_s)M_{\text{opt}} = 1$ (symbols). As a result, the capacity α (bottom) increases with r until M_{opt} has reached its lower bound, and α exhibits a maximum. A further increase in r reduces c but leaves M_{opt} constant and, thus, leads to a decrease of α ; see equations 2 and 3. (B) The optimal firing threshold θ_{opt} decreases with increasing r to its lower bound 1 (symbols). The weak dependence of θ_{opt} on $c_m = c+c_s$ is indicated by the gray lines. Further parameters for A and B: $N = 250\,000$, $\gamma = 0.7$, $Q = 1$.

For constant $(c+c_s)N$, numerical results in Figure 7A (symbols) show that the capacity α exhibits a maximum as a function of r . The maximum capacity occurs at a pattern size at which $(c+c_s)M_{\text{opt}} = c_{11}M_{\text{opt}} = 1$, i.e., the association from a cue pattern to an ‘on’ neuron of a target pattern is supported by a single spike, on average. For larger r , the optimal firing threshold θ_{opt} remains at its minimum value of one; see Figure 7B. An increase of r beyond its optimum reduces c but leaves M_{opt} constant and, thus, leads to a decrease of α ; see equations 2 and 3.

Values of r larger than 1 are thus beneficial for good memory performance, in our case the storage

capacities α . Similarly, Brunel et al. (2004) finds a high ratio r to be necessary for increasing the signal-to-noise ratio of read-out from a perceptron-like structure. These findings raise the question why values of r found in some experiments (Montgomery et al., 2001; Nusser et al., 1998) are in the range $r \lesssim 1$. We suppose that the specific value of r is due to the interplay between the recurrence in the hippocampus and the locality of synaptic learning rule; see Discussion. In contrast, Isope & Barbour (2002) report $r \approx 4$ at the cerebellar parallel fibers, which is a locally feed-forward system.

D. Scale-Invariance of Sequential Memory

Given the scaling laws $\alpha \propto cN$ of the storage capacity in equation 23, we can ask how the connectivity in a brain region should be set up in order to have scale-invariant sequential memory, which means

$$P \propto N .$$

From $P = \alpha c_m N \propto c^2 N^2$ (equation 23) we then find $c \propto N^{-1/2}$ or, equivalently, that the total number cN^2 of synapses in the network is proportional to $N^{3/2}$. Surprisingly, the latter result is in agreement with findings from Stevens (2001; 2005) in visual cortex and other brain areas. Thus, a $N^{3/2}$ -law for the number of synapses can generate a scalable architecture for associative memory.

To summarize this section, constraints have seminal influence on the scaling laws of the capacity of sequential memory, and different constraints lead to fundamentally different strategies for optimizing the performance of networks for replaying sequences.

VII. NON-MINIMAL SEQUENCES

In addition to constraints on intrinsic features of the network like a small connectivity or a limited number of synapses there are also constraints that may be imposed upon a sequence memory device from outside, for example a fixed detection threshold γ and a non-minimal sequence length Q .

A. Finite Sequences ($Q > 1$)

To determine the capacity α for non-minimal sequences $Q > 1$ in dependence upon the network size N , we apply the Markovian approximation as introduced in section IV B. As in the case $Q = 1$, replay of sequences is initialized with an ideal representation of the cue pattern, $(m_t, n_t) = (M, 0)$

for $t = 0$. Patterns that occur later in the sequence at $t \geq 1$, however, are not represented in an ideal way; typically, there is a finite number of false alarms n_t , and the number of hits m_t is generally below M ; see also Figure 3. The recall of patterns amidst a sequence therefore depends on noisy cues. As a consequence, for $Q > 1$ we expect that the dependence of the optimal pattern size M_{opt} and the optimal threshold θ_{opt} on the network size N is different as compared to the case $Q = 1$.

Assuming a constrained number of synapses per neuron ($cN = \text{const.}$), we nevertheless find that for $Q > 1$ the dependence of the optimal pattern size M_{opt} on N is almost linear for large N ; see Figure 8A. Accordingly, the capacity α is nearly independent of N ; see Figure 8B. Moreover, also the optimal firing threshold θ_{opt} is almost constant for large N ; see Figure 8C. These results for $Q > 1$ resemble the ones for $Q = 1$ shown in Figure 6, some of which are also indicated by disks in Figure 8. One reason for this correspondence is that patterns within a sequence are typically replayed at a high quality; see Figures 1 and 3.

Figure 8B also shows that α is a decreasing function of Q . We note that α still refers to the number of *minimal* sequences. Then, the maximum number of stored sequences of length Q is the Q -th fraction of $P = \alpha(1+r)cN$. Compared to $Q = 1$, the storage capacity α drops about an order of magnitude for $Q = 2, 4$ and soon, at $Q \gtrsim 8$, arrives at a baseline value for $Q \rightarrow \infty$ (solid line) that was obtained from a mean-field approximation to be explained below.

We thus conclude that non-minimal sequences impose no fundamental limit to the memory capacity for sequences. However, due to discrete time, our model cannot comprise temporal dispersion of synchronous firing, which may limit the replay of long sequences in biological networks (Diesmann et al., 1999).

B. Infinite Sequences ($Q \rightarrow \infty$)

A beneficial consequence of the weak dependence of the capacity α on the sequence length for $Q \gtrsim 8$ is that sequential memory for large Q can be more easily discussed in the framework $Q \rightarrow \infty$. Such a discussion requires to find the fixed-point distributions of the transition matrix T defined in equation 6. Assuming that the fixed-point distributions for hits m and false alarms n are unimodal and given the case $N \gg 1$ we can reduce the problem of finding fixed-point distributions of m and n to the much simpler problem of finding fixed points of the mean values $\langle m \rangle$ and $\langle n \rangle$. Let us

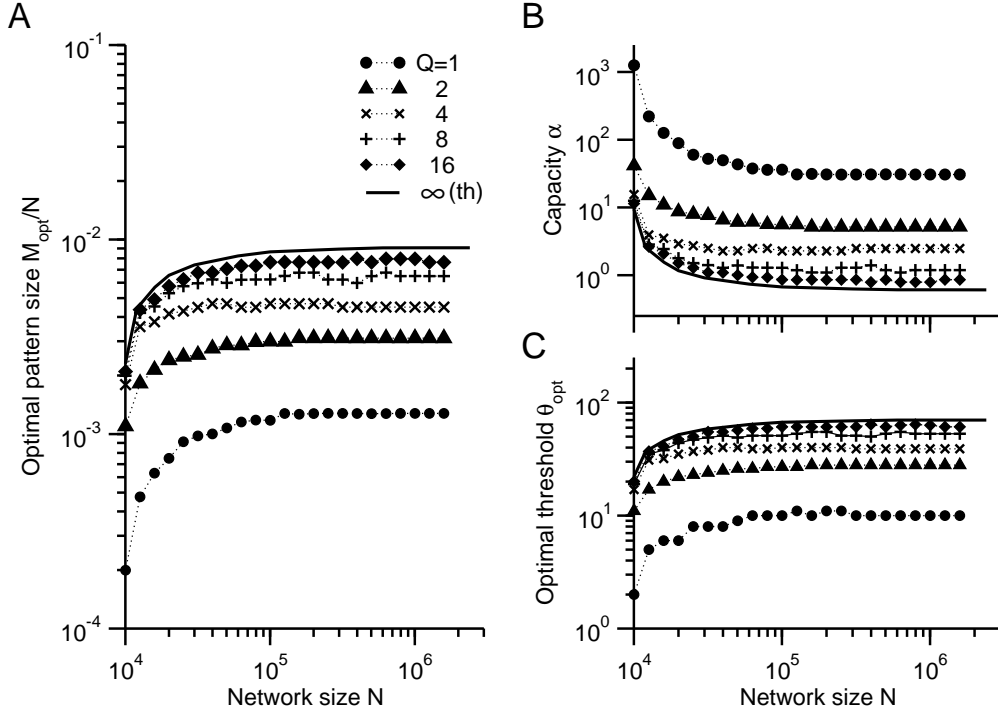


FIG. 8: Optimal event size M_{opt} , capacity α and optimal firing threshold θ_{opt} for non-minimal sequences in networks with a constrained number $c_m N = 10\,000$ of synapses per neuron. (A) The optimal pattern size M_{opt} increases half a magnitude between $Q = 1$ and $Q = 4$ (three bottom lines) and saturates for $Q \geq 8$ (three topmost lines). (B) The capacity α is almost constant for large N and decreases with increasing Q . (C) The optimal firing threshold θ_{opt} reflects the dependencies of the optimal pattern size M_{opt} . Numerical results (symbols) are obtained for $\gamma = 0.7$ and $r = 1$. Graphs for $Q \rightarrow \infty$ (lines) are obtained from the mean-field equation 24.

therefore introduce the iterated map

$$\begin{pmatrix} \langle m_{t+1} \rangle \\ \langle n_{t+1} \rangle \end{pmatrix} = T_{(\cdot)} \begin{pmatrix} \langle m_t \rangle \\ \langle n_t \rangle \end{pmatrix} \quad (24)$$

for the mean values of the order parameters. To specify the map $T_{(\cdot)}$ in accordance with the Markovian dynamics introduced in section IV B we define the mean synaptic inputs to on- and off units,

$$\mu_{\text{on}} = c_{11} \langle m \rangle + c_{01} \langle n \rangle \quad \text{and} \quad \mu_{\text{off}} = c_{10} \langle m \rangle + c_{00} \langle n \rangle ,$$

respectively, as well as the variances,

$$\begin{aligned} \sigma_{\text{on}}^2 &= c_{11} \langle m \rangle (1 - c_{11} \langle m \rangle / M) + c_{01} \langle n \rangle [1 - c_{01} \langle n \rangle / (N - M)] \\ \sigma_{\text{off}}^2 &= c_{10} \langle m \rangle (1 - c_{10} \langle m \rangle / M) + c_{00} \langle n \rangle [1 - c_{00} \langle n \rangle / (N - M)] \end{aligned} ,$$

which are determined by the reduced connectivity matrix $(\begin{smallmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{smallmatrix})$ from equation 5. A Gaussian approximation to binomial statistics then yields

$$T_{\langle \cdot \rangle} \left(\begin{matrix} \langle m \rangle \\ \langle n \rangle \end{matrix} \right) = \frac{1}{2} \left(\begin{matrix} M \left\{ 1 + \operatorname{erf} \left[(\mu_{\text{on}} - \theta) / \sqrt{2 \sigma_{\text{on}}^2} \right] \right\} \\ (N - M) \left\{ 1 + \operatorname{erf} \left[(\mu_{\text{off}} - \theta) / \sqrt{2 \sigma_{\text{off}}^2} \right] \right\} \end{matrix} \right). \quad (25)$$

Numerical iteration of equation 24 results in the fixed-points $(\langle m \rangle^*, \langle n \rangle^*)$ of the mean-field dynamics and their basins of attraction; see Figure 9A. The iterated map has two trivial fixed-points that are largely independent of the choice of firing threshold θ and pattern size M . These trivial fixed points represent complete activation of the network, on the one hand, and no activity at all, on the other hand. Shape and size of their basins of attraction (black and white areas in Figure 9A), however, are modulated by the specific values of M and θ . We also observe a third type of fixed point comprising a large number of hits and a small number of false alarms; numerics shows that we always find $\langle \Gamma \rangle \lesssim 1$ at this fixed point of infinite sequence replay. Its basin of attraction is plotted in gray and extends over a small interval of false alarm rates; note the logarithmic scale on the ordinates in Figure 9A.

In Figure 9A we see that the smaller the pattern size the narrower is the range of thresholds allowing for an infinite sequence replay. For a large enough pattern size, the range of possible thresholds is broad; see also Figures 1 and 3. The region in the (M, θ) space where infinite sequence replay can occur is summarized in Figure 9B. The wedge-shaped stability regions are not much affected by N but strongly depend on c .

The borders of such a stability region in Figure 9B can be described by upper and lower bounds for the thresholds, θ^{upper} and θ^{lower} , that can be approximated through linear functions of the pattern size M . The upper bound θ^{upper} is interpreted as an iso- $\langle \Gamma \rangle$ line that separates the region of a completely deactivated state with fixed point $\langle m \rangle^* = \langle n \rangle^* = 0$ from the region of stable sequence replay where $\langle m \rangle^* = M \langle \Gamma \rangle$ and $\langle n \rangle^* \ll M$. From the first line of equation 24 we then obtain

$$\theta^{\text{upper}} \approx c_{11} M \langle \Gamma \rangle - \operatorname{erf}^{-1}(2\langle \Gamma \rangle - 1) \mathcal{O}(\sqrt{\theta^{\text{upper}}}).$$

Thus, for large M , the bound θ^{upper} is an almost linear function of the pattern size with a slope $c_{11} \langle \Gamma \rangle \approx c_{11}$. Similarly, from the second line of equation 24 we obtain the boundary θ^{lower} between the region where $\langle m \rangle^* = M$ and $\langle n \rangle^* \ll N$ and the region of a completely activated state $\langle m \rangle^* / M = \langle n \rangle^* / (N - M) = 1$,

$$\theta^{\text{lower}} \approx c_{10} M + c_{10} N (1 - \langle \Gamma \rangle) + \operatorname{erf}^{-1}(2\langle \Gamma \rangle - 1) \mathcal{O}(\sqrt{\theta^{\text{lower}}})$$

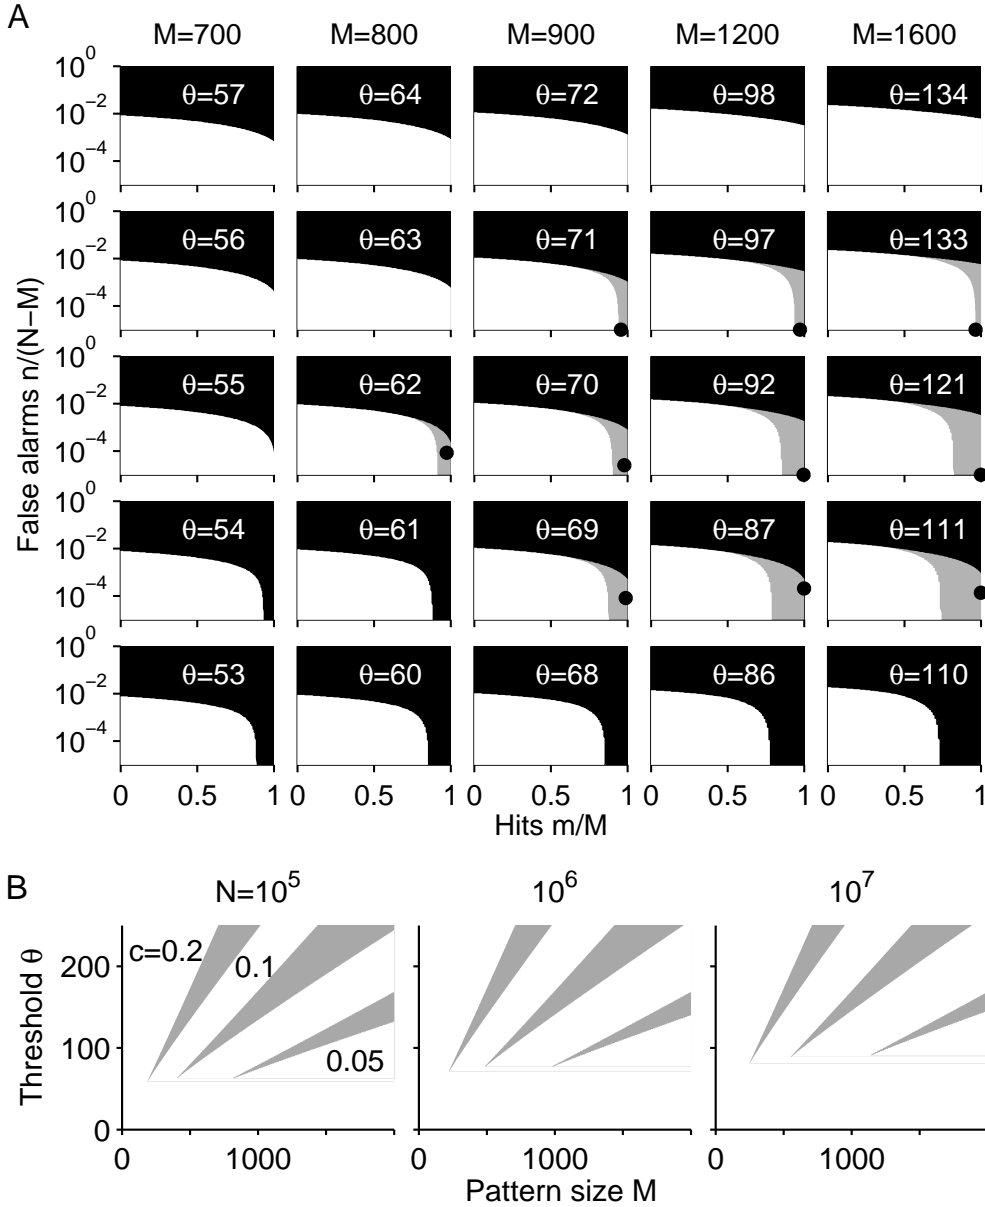


FIG. 9: Fixed points of infinite sequence replay. (A) Basins of attraction of the mean-field dynamics in equation 24 depend on pattern size M and firing threshold θ . The discrete dynamics of mean hit rates $\langle m \rangle / M$ and mean false alarm rates $\langle n \rangle / (N - M)$ exhibits two trivial fixed points: the first one is a completely deactivated state, $\langle m \rangle^* = \langle n \rangle^* = 0$, with basins of attraction represented by a white area. The second fixed point represents maximal activation, $\langle m \rangle^* / M = \langle n \rangle^* / (N - M) = 1$, with basins of attraction painted black. For few pairs of (M, θ) , we also observe non-trivial fixed points (black dots) corresponding to sequence replay. Their basins of attraction are depicted by gray areas. Parameters ($N = 10^5$, $c = 0.05$, $r = 1$) are the same as in Figures 1 and 3. (B) Regions of stable sequence replay in the (M, θ) space are plotted in gray; connectivities are $c = 0.05, 0.1, 0.2$ and network sizes are $N = 10^5, 10^6, 10^7$ for $r = 1$. The slopes of the upper and lower borders of these stability regions approximately equal the connectivities c_{11} and c_{10} , respectively.

The slope of θ^{lower} is about c_{10} , which for $r = 1$ is about half the slope of θ^{upper} . These predicted slopes agree with numerical results in Figure 9B. The size of the region of infinite sequence replay is therefore proportional to $c_{11} - c_{10} \propto r$. The larger the ratio r between silent and non-silent synapses, the larger are the stability regions and, hence, the more robust is sequence replay.

We emphasize that the above expressions for θ^{upper} and θ^{lower} are rough estimates that correspond to large pattern sizes M at which the distributions of synaptic inputs to ‘off’ and ‘on’ units do not overlap too much; see Figures 4A & B. Moreover, the optimal parameters M_{opt} and θ_{opt} at the tip of a stability region cannot be determined explicitly because we cannot assess the exact value of $\langle \Gamma \rangle$ analytically.

The mean field results in Figures 9A are largely consistent with the cellular simulations in Figure 1, but there are also important differences. Although cellular simulations have been obtained for finite sequences $Q = 20$ whereas mean-field results are valid for $Q \rightarrow \infty$, discrepancies at the edges of the stability regions also occur because random fluctuations in cellular simulations can kick the network into complete activation or deactivation. The edges of the wedge-shaped regions in Figure 9B therefore describe the behavior of cellular networks only approximately.

To summarize, the higher the capacity, the less robust is sequence replay against variations of the parameters M and θ . The wedge-shaped structures of the stability regions in Figure 9B indicate that the maximal sequence capacity and, hence, minimal M go along with a critical dependence of stability upon the firing threshold. In the limit of $M \rightarrow M_{\text{opt}}$ the network lives on the edge of dynamical (in)stability.

VIII. INFORMATION CONTENT FOR $N \rightarrow \infty$

The detection criterion we have proposed in section IV C permits a limited amount of errors. It is intuitively clear that these retrieval errors allow for an increase of the storage capacity α as compared to an error-less case. However, the more errors occur the more deteriorated is the representation of each of the patterns during replay. The common way of measuring the balance of these two opposing effects of retrieval errors is to calculate the information content I . The latter can be understood as the logarithm of the number of all possible ways of concurrently storing a number of P of associations or, more precisely (Nadal & Toulouse, 1990),

$$I = \lg_2 \left\{ \binom{N}{m+n} / \left[\binom{M}{m} \binom{N-M}{n} \right] \right\}^P. \quad (26)$$

Here, $\binom{N}{m+n} / \left[\binom{M}{m} \binom{N-M}{n} \right]$ is the number of patterns of size M that can be represented in a network of size N , given the hits m and false alarms n . We note that the number P of associations between patterns depends on the performance of the read-out device, and so does the information content.

The information content is often calculated as a function of the so-called coding ratio $f = M/N$, which is interpreted as a firing rate. In biological relevant networks the firing rate is low ($f \rightarrow 0$) while they are required to be operable in the limit $N \rightarrow \infty$. This asymptotic behavior of networks is extensively discussed in literature (e.g. Gardner 1987; Golomb et al. 1990; Willshaw et al. 1969). In what follows we will show that also in our framework we have $\lim_{N \rightarrow \infty} f = 0$. In this limit we will assess the information content I for $Q = 1$ and $Q \rightarrow \infty$.

From equation 26 we derive an approximation of I for $f \rightarrow 0$ given the retrieval criterion 15 we have applied throughout the present paper. For a specific value $0 < \gamma' < 1$ of the detection threshold and a fixed fraction $\eta := m/M > \gamma'$ of hits we can approximate I by evaluating equation 26 with $n = (\eta - \gamma')(N - M)$. Applying Stirling's formula for $\eta < 1$ and expanding the result up to $\mathcal{O}(f)$ we obtain

$$I/(c_m N^2) = \alpha f \left[(\eta - \gamma') |\lg_2(\eta - \gamma')| + \eta |\lg_2 \eta| + (1 - \eta) |\lg_2(1 - \eta)| \right] . \quad (27)$$

From equation 3 we know that the storage capacity scales like $\alpha \propto N/M^2$ and, thus, $I/N^2 \propto 1/M$. As a corollary, this shows that minimizing M not only maximizes α but also I .

In case $Q = 1$, a combined optimization of θ and M leads to M_{opt} being independent of network size N ; see section V. As a result we obtain $f \propto 1/N$. Accordingly, the information content per synapse $I/(c_m N^2) \propto 1/M$ is independent of network size N .

In order to also obtain the asymptotic behavior of I in the case of large sequence length we have assessed the optimal pattern size M_{opt} for $Q \rightarrow \infty$ as a function of network size N for a fixed connectivity c , i.e., without any constraint; see Figure 10. Numerics reveals a sub-logarithmic behavior, $M_{\text{opt}}(N) \propto (\ln N)^{0.82}$; the coding ratio $f \propto (\ln N)^{0.82}/N$ also falls below every bound as $N \rightarrow \infty$, i.e., coding becomes arbitrarily sparse. Together with equation 3 the unconstrained storage capacity diverges like

$$\alpha \propto N/(\ln N)^{1.64} .$$

From equation 27 we thus find for $N \rightarrow \infty$ the information content per synapse to decrease sub-logarithmically

$$I/N^2 \propto (\ln N)^{-0.82} .$$

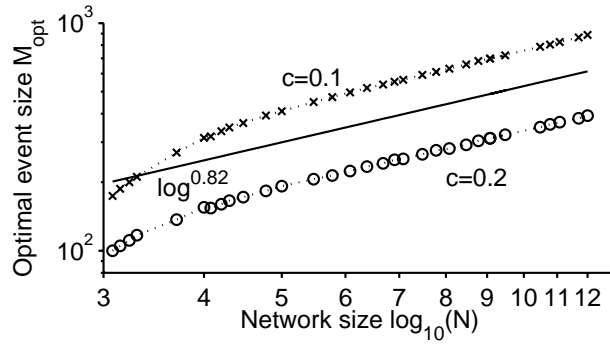


FIG. 10: The dependency of the optimal pattern size M_{opt} on the network size N is weak in the case of sequence length $Q \rightarrow \infty$; note the logarithmic scale on the abscissa. The connectivity c is fixed, i.e. no constraint is imposed (crosses: $c = 0.1$, circles: $c = 0.2$). The symbols represent optimal pattern sizes obtained from numerical solution of the fixed-point equation $(\langle m \rangle^*, \langle n \rangle^*) = T_{(\cdot, \cdot)}(\langle m \rangle^*, \langle n \rangle^*)$; see equation 24. The solid line illustrates the asymptotic behavior $M_{\text{opt}} \propto (\ln N)^{0.82}$ found from linear regression.

The information content per synapse converges to zero for $N \rightarrow \infty$, though very slowly. In fact, I/N^2 falls so slow that in the range of biologically reasonable network sizes $10^3 < N < 10^7$ the information content per synapse only varies by a factor of $(7/3)^{0.82} \approx 2$.

To summarize, combined optimization of M and θ provides an efficient algorithm to set up a sequential memory network for a broad range of network sizes. However, combining the results from sections VI and VII for biologically relevant parameter regimes, one obtains information contents $I < 3/2 \alpha f c_m N^2$ that are far below the theoretical maximum $c_m N^2$ of one bit per synapse.

IX. DISCUSSION

The present paper combines analytical and numerical methods to assess the capacity for storing sequences of activity patterns in a recurrent network of McCulloch-Pitts units. Results from mean-field theory are validated through simulations of cellular networks and a probabilistic dynamical description. Our approach is new in that we concurrently optimize the pattern size M and the firing threshold θ in order to maximize the storage capacity α . Within this framework we derive the capacity α in dependence upon five system parameters: network size N , mean connectivity c , synaptic plasticity resources r , sequence length Q , and detection threshold γ .

The storage capacity of a network crucially depends on the criterion for pattern detection. One typically requires that the quality of replay of patterns exceeds some detection threshold γ

(equations 15 and 16). Our retrieval criterion with $\gamma < 1$, which allows for errors in the replay of patterns, is fundamentally different from the error-free criterion $\gamma \rightarrow 1$ in the classical Willshaw et al. (1969) network where the storage capacity is subject to Gardner’s bound (Gardner, 1987). In the original Willshaw model as well as in our approach for minimal sequences ($Q = 1$) the network is initialized with a perfect representation of the cue pattern, $m_0 = M$ and $n_0 = 0$. The Willshaw model, however, requires a perfect retrieval of a target pattern in that the number of hits is maximal, $m_1 = M$, and that there is less than one false alarm on average, $\langle n_1 \rangle < 1$; furthermore, the firing threshold θ is set to the pattern size M . Then, binomial statistics yields the well-known logarithmic scaling laws for the optimal pattern size $M_{\text{opt}} \propto \log N$ and the capacity $\alpha_{\text{Willshaw}} \propto N/\log^2 N$ (Gardner, 1987; Willshaw et al., 1969); see also equation 3. In terms of the coding ratio $f = M_{\text{opt}}/N$, they find $\alpha_{\text{Willshaw}} \propto 1/(f|\ln f|)$ for $N \rightarrow \infty$. In contrast, in the present manuscript we optimize both the firing threshold θ and the pattern size M , and we use a read-out criterion that permits errors. Thus, the storage capacity diverges faster than α_{Willshaw} , viz. $\alpha \propto 1/f$; see equation 23.

An error-full representation of patterns is in agreement with the situation in the brain, for example in the hippocampal CA3 network. There, the recurrently connected pyramidal cells also have feed-forward connections to the pyramidal cells in CA1 via highly plastic synapses. It is generally assumed (Hasselmo, 1999) that these synapse are to be adjusted by CA3 activity and local learning rules, i.e., CA1 can learn replayed patterns. Read out in CA1 may therefore be successful even if the absolute number of false alarms in CA3 exceeds the number of hits. The detection criterion in equation 16 can be motivated by such downstream neurons that receive excitation from the correctly activated neurons and inhibition from the incorrectly activated ones, e.g. via a globally coupled network of interneurons.

For sequence length $Q = 1$, the concurrent optimization of M and θ leads to scaling laws for the replay of minimal sequences for biologically relevant connectivities $c \ll 1$: the optimal pattern size is inversely proportional to the mean connectivity, $M_{\text{opt}} \propto c^{-1}$, and the optimal firing threshold θ_{opt} is independent of c . Both θ_{opt} and M_{opt} are independent of the network size N . The above dependencies finally lead to the capacity of sequential memory that scales like $\alpha \propto cN$ (equation 23). Moreover, the number of associations that can be stored scales like $P \propto c^2 N^2$.

A main conclusion from the scaling laws $\alpha \propto cN$ and $P \propto c^2 N^2$ is that for a constrained number of synapses per cell (synapses-per-neuron constraint, $cN = \text{const.}$) the capacity α and the number P are constant, i.e., independent of the network size N ; see Figures 6 and 8. This means that

it is impossible to increase the computational power of the network by increasing N . One could argue, however, that taking two independent networks doubles P and therefore would account for a performance increase that is linear in N . The drawback of this strategy is that then each pattern can only be connected to only half of the other patterns, which are those located in the same network module.

A technically relevant constraint, e.g. in a computer simulation, is a constant total number of synapses in the network (synapses-per-network constraint, $cN^2 = \text{const.}$). From above scaling laws we conclude that α and P necessarily decrease with increasing network size; see Figure 6B.

One can also ask whether there *is* a scaling law for the connectivity that accounts for scale-invariant storage, i.e., $P \propto N$; see section VIC. In so doing we find scale invariance for $c \propto 1/\sqrt{N}$. As a result the total number of synapses then is proportional to $N^{3/2}$, which is in line with results by Stevens (2001; 2005).

For the synapses-per-neuron constraint, there is an optimal value for the ratio r between silent and non-silent synapses. For generic parameter regimes this optimal value is rather large ($r \approx 10$; see Figure 7). However, α exhibits a broad maximum as a function of r , and therefore the exact value of r is not critical for sequential memory. If one considers the network connectivity to be determined by local Hebbian or STDP-like learning rules ratios r that strongly deviate from 1 are implausible, since synaptic LTP at a specific pair of pre- and postsynaptic neuron can only be compensated for locally by LTD of another synapse at the very same pair of neurons (Bi & Poo, 1998; Gerstner et al., 1996; Kempster et al., 1999). One thus can argue that the functional benefit of a very high amount of plastic resources may no longer justify the expenses of non-local signaling in synaptic plasticity. In short, ratios $r \approx 1$ may be sufficient for an excellent performance of sequential memory.

This paper also shows that for long sequences, e.g. $Q > 8$, memory capacity becomes virtually independent of Q ; see Figure 8. For large Q , however, the optimal pattern size is necessarily such that the network is close to dynamical instability; see Figure 9. Yet from the point of view of maximizing storage capacity α , the strategy of avoiding dynamical instabilities by increasing pattern size M is problematic, since α is proportional to M^{-2} ; see equation 3. In order to approach the maximal storage capacity without the danger of complete activation or silencing of the network, one rather might introduce an activity-dependent stabilization mechanism that provides a negative feedback after a certain number of time steps. A biological realization that is at hand is a network of inhibitory interneurons (Battaglia & Treves, 1998; Bragin et al., 1995; Csicsvari et al., 2003;

Traub et al., 2000). This of course may come at the cost of limiting sequence length Q and/or reducing the detection threshold γ .

Our results for large sequence lengths Q are not immediately applicable to synfire chains (Abeles, 1991; Diesmann et al., 1999; Herrmann et al., 1995). The chief difficulty for translating our model into a more realistic network with continuous dynamics is to preserve the temporal separation between distinct patterns. The functional constraint of minimal sequence lengths is thus more likely a constraint on the temporal precision of network dynamics rather than on counting statistics. We speculate that, for biological networks, spike desynchronization restricts the applicability of our results to small values of Q .

The present framework is limited to orthogonal sequences, i.e., a particular pattern is not allowed to occur presynaptically in more than one minimal sequence. Non-orthogonal, or loop-like sequential memories can be taken into account by, e.g. generalizing the framework to neurons with more than one-step memory (Dehaene et al., 1987; Guyon et al., 1988) or adding “internal patterns” that represent repetitions (Amit, 1988) or context (Levy, 1996).

A possible neurophysiological application of our theory can be found in the hippocampus. During slow-wave sleep, low levels of the neuromodulator acetylcholine boost the impact of the excitatory feedback connections within CA3; see Hasselmo (1999) for a review. Slow-wave sleep goes along with a phenomenon called sharp-wave ripples, which is speculated to be a result of the replay of short sequences (Csicsvari et al., 2000; Draguhn et al., 2000). A sharp-wave ripple burst is a pulse-like incident of the local field potential in CA3 that is accompanied by 200 Hz oscillations. The latter are supposed to be generated by CA3 pyramidal cells (Behrens et al., 2005) and may reflect sequence replay (Lee & Wilson, 2002; Nadasdy et al., 1999; Wilson & McNaughton, 1994) occurring in time-slices of about 5 ms. The total duration of ripples of about 40 ms limits the number of putative events in a sequence to less than about 8. The temporal extent of a sharp wave may be controlled by inhibition (Maier et al., 2003), which would hint at dynamical stabilization of the network activity at a high level of storage capacity; see above.

In Figure 8 we have plotted the coding ratio $f = M_{\text{opt}}/N$ and storage capacity α as a function of network size for various sequence lengths. If we apply these results to the situation in the hippocampal CA3 region of rats and a sequence length of $Q = 8$, we find for a network size of $N = 240\,000$, a synapses-per-neuron constraint of $c_m N = 10\,000$ synapses per cell and plasticity resources $r = 1$, the optimal pattern size to be about 1500 cells. As a consequence, the storage capacity is about $\alpha = 1.2$ minimal sequence per synapse at a cell, which corresponds to about 1600

full sequences of length 8 stored in the network. Interestingly the firing threshold we obtain is 55, which is approximately the same as that assumed by Diesmann et al. (1999) for cortical synfire networks.

To summarize, the present paper provides a simple rule of how to choose pattern size and threshold in order to optimize storage capacity under biologically realistic constraints such as low connectivity and similar amounts of silent and non-silent synapses. From that, one can conclude that sequence completion in the recurrent network operates far below maximal information content. To put it more positively, information seems to be redundantly distributed over a large number of synapses, which seems consistent with the picture that memories are stored in a way that is robust against synaptic noise and some variability of morphological plasticity.

Acknowledgments

We thank D. Schmitz for valuable discussions on the hippocampal circuitry, A.V.M. Herz for ongoing support, and R. Gütig, R. Schaette, M. Stemmler, K. Thurley, and L. Wiskott for discussions and comments on the manuscript. This research was supported by the Deutsche Forschungsgemeinschaft (Emmy Noether Programm: Ke 788/1-3, SFB 618) and the Bundesministerium für Bildung und Forschung (Bernstein Center for Computational Neuroscience).

Appendix A: List of Symbols

Symbol	meaning (location of first use)
t	discrete time (section II A)
\mathbf{x}	binary network state vector (section II A)
θ	firing threshold (section II A)
c	mean connectivity of activated (section II A)
c_s	mean connectivity of silent synapses (section II A)
c_m	mean morphological connectivity (section II A)
$r = c_s/c$	ratio between silent and active connectivity (section II A)
$C = (C_{nn'})$	active connectivity matrix (section II A)
N	network size (section II A)
M	pattern size (section II B)
ξ	binary pattern vector (section II B)
Q	sequence length (section II B)
P	number of minimal sequences stored (section III)
$\alpha = P/(c_m N)$	capacity of sequential memory (section III)
m	number of hits (section IV A)
n	number of false alarms (section IV A)
$\begin{pmatrix} c_{11} & c_{10} \\ c_{01} & c_{00} \end{pmatrix}$	reduced connectivity matrix (equation 4)
T	transition matrix (equation 6)
b	binomial probability (equation 7)
p	conditional probability of hits (equation 8)
q	conditional probability of false alarms (equation 8)
ρ	conditional probability of one hit (equation 9)
λ	conditional probability of one false alarm (equation 10)
Γ	quality of replay (equation 14)
γ', γ	detection thresholds (equations 15 and 16)
$T_{(\cdot)}$	mean transition function (equation 25)
$f = M/N$	coding ratio (section VIII)
I	information content (section VIII)

Appendix B: Memory Capacity Revisited

Let us consider a naive network of size N that, initially, has no synapses at all. To imprint the first minimal sequence $\xi^A \rightarrow \xi^B$ in the network we need $M^2 c_{11}$ functional synapses in order to link two groups of M neurons at connectivity c_{11} ; see Figure 11. Let us first discuss the simpler case $c_{11} = c_m$. For the second sequence $\xi^C \rightarrow \xi^D$ fewer synapses are needed because we have to take into account that there are cells in pattern ξ^C that are already connected to cells in ξ^D because of some overlap with the first sequence. For random patterns, the probability that a neuron is active in a specific pattern is $f = M/N$, which is also called the coding ratio. As a result, the mean number of cells that are active in both of a given pair of patterns is Mf . Consequently, the Mf presynaptic cells that belong to both cue patterns ξ^A and ξ^C only have to be connected to the $M(1 - f)$ postsynaptic neurons of ξ^D that do *not* overlap with ξ^B . The number of new synapses needed is $Mf \cdot c_{11} \cdot M(1 - f)$. In order to complete the second minimal sequence, we are left with connecting the remaining $M(1 - f)$ presynaptic cells of ξ^C to *all* M postsynaptic cells of ξ^D . In summary, the second sequence consumes $Mf c_{11} M(1 - f) + M(1 - f) c_{11} M = M^2 c_{11} (1 - f^2)$ synapses. Similarly, the k^{th} minimal sequence consumes $M^2 c_{11} (1 - f^2)^{k-1}$ synapses that have not yet been accounted for. Summing up all contributions until we reach the limit $N^2 c$ of available

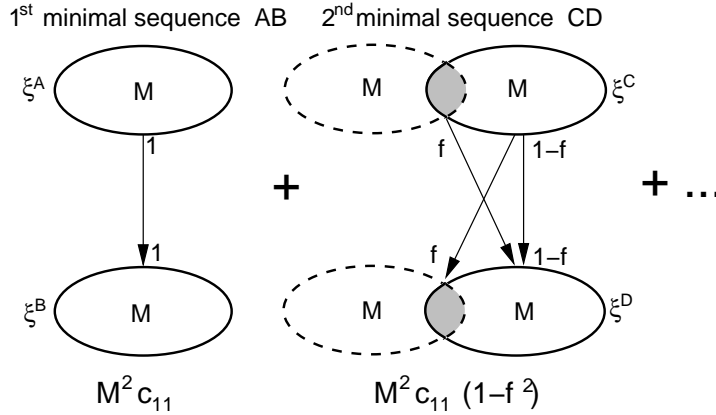


FIG. 11: Consumption of synapses by subsequently storing minimal sequences. The first minimal sequence $\xi^A \rightarrow \xi^B$ consumes $c_{11} M^2$ synapses. The patterns of a second minimal sequence $\xi^C \rightarrow \xi^D$ have some overlap f with ξ^A and ξ^B ; there are Mf cells (gray) both pre- and postsynaptically that contribute to both the first and the second minimal sequences. The number of synapses that are consumed by $\xi^C \rightarrow \xi^D$ is reduced by a factor of $(1 - f^2)$; see text.

non-silent synapses yields a condition on the maximal number P of minimal sequences,

$$N^2 c \stackrel{!}{=} M^2 c_{11} \sum_{k=1}^P (1 - f^2)^{k-1} = M^2 c_{11} f^{-2} [1 - (1 - f^2)^P] . \quad (28)$$

In case $c_{11} < c_m$ we need to take into account the probability c_{11}/c_m of having a morphological synapse from cue to target in the non-silent state. The transformation $f^2 \rightarrow f^2 c_{11}/c_m$ is sufficient to generalize the result in equation 28; solving the generalized version of equation 28 for P and normalizing the result by N we find the capacity to be

$$\alpha = \frac{\log(1 - c/c_m)}{c_m N \log(1 - f^2 c_{11}/c_m)} , \quad (29)$$

which can be approximated for $f \ll 1$ by

$$\alpha = \frac{c N}{c_m c_{11} M^2} .$$

Equation 29 is an extension to the results for the case $c_{11} = c_m = 1$, originally obtained by Willshaw et al. (1969), Nadal & Toulouse (1990) and Nadal (1991). In the main part of this paper, we discuss the scenario $c_{11} = c_m < 1$; see equations 2 and 3.

-
- Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex.*, New York: Cambridge University Press.
- Amit, D. J., Gutfreund, H., Sompolinsky, H. (1987). Information storage in a network with low levels of activity. *Phys. Rev. A*, *35*, 2293–2303.
- Amit, D. J. (1988). Neural networks counting chimes. *Proc. Natl. Acad. Sci. USA*, *85*, 2141–2145.
- August, D. A., & Levy, W. B. (1999). Temporal sequence compression by an integrate-and-fire model of hippocampal area CA3. *J. Comput. Neurosci.*, *6*, 71–90.
- Battaglia, F. P., & Treves, A. (1998). Stable and rapid recurrent processing in realistic autoassociative memories. *Neural Comput.*, *10*, 431–450.
- Behrens, C. J., van den Boom, L. P., de Hoz, L., Friedman, A., & Heinemann, U. (2005). Induction of sharp wave-ripple complexes in vitro and reorganization of hippocampal networks. *Manuscript in Revision*
- Bi, G.-q., & Poo M.-m. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, *18*, 10464–10472
- Bragin, A., Jando, G., Nadasdy, Z., Hetke, J., Wise, K., & Buzsaki, G. (1995). Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *J. Neurosci.*, *15*, 47–60.
- Brun, V. H., Otnass, M. K., Molden, S., Steffenach, H. A., Witter, M. P., Moser, M. B., Moser, E. I. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science*, *296*:2243–2246.

- Brunel, N., Nadal, J.-P., & Toulouse, G. (1992). Information capacity of a perceptron. *J. Phys. A: Math. Gen.*, *25*, 5017–5037.
- Brunel, N., Hakim, V., Isope, P., Nadal, J.-P., & Barbour, B. (2004). Optimal information storage and the distribution of synaptic weights: perceptron versus purkinje cell. *Neuron*, *43*, 745–757.
- Buckingham, J., Willshaw, D. (1992). Performance characteristics of the associative net. *Network*, *3*, 407–414.
- Buckingham, J., Willshaw, D. (1993) On setting unit thresholds in an incompletely connected associative net. *Network*, *4*, 441–459.
- Buhmann, J., & Schulten, K. (1987) Noise-driven temporal association in neural networks. *Europhys. Lett.*, *4*, 1205–1209.
- Csicsvari, J., Hirase, H., Mamiya, A., & Buzsaki, G. (2000). Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave associated population events. *Neuron*, *28*, 585–594.
- Csicsvari, J., Jamieson, B., Wise, K. D., & Buzsaki, G. (2003). Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron*, *37*, 311–322.
- Dehaene, S., Changeux, J.-P., Nadal, J.-P. (1987). Neural networks that learn temporal sequences by selection. *Proc. Natl. Acad. Sci. USA*, *84*, 2727–2731.
- Deshpande, V., & Dasgupta, C. (1991). An neural network for storing individual patterns in limit cycles. *J. Phys. A: Math. Gen.*, *24*, 5105–5119.
- Diesmann, M., Gewaltig, M. O., & Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, *402*, 529–533.
- Draguhn, A., Traub, R. D., Bibbig, A., & Schmitz, D. (2000). Ripple (approximately 200-Hz) oscillations in temporal structures. *J. Clin. Neurophysiol.*, *17*, 361–376.
- Düring, A., Coolen, A. C. C., Sherington, D. (1998). Phase diagram and storage capacity of sequence processing neural networks. *J. Phys. A: Math. Gen.*, *31*, 8607–8621.
- Fortin, N. J., Agster, K. L., & Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nat. Neurosci.*, *5*, 458–462.
- Gardner, E. (1987). Maximum storage capacity in neural networks. *Europhys. Lett.*, *4*, 481–485.
- Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature*, *383*, 76–81.
- Golomb, D., Rubin, N., & Sompolinsky, H. (1990). Willshaw model: Associative memory with sparse coding and low firing rates. *Phys. Rev. A*, *41*, 1843–1854.
- Graham, B., & Willshaw, D. (1997). Capacity and information efficiency of the associative net. *Network: Comput. Neural Syst.*, *8*, 35–54.
- Gutfreund, H., & Mezard, M. (1988). Processing of temporal sequences in neural networks. *Phys. Rev. Lett.*, *61*, 235–238.
- Guyon, I., Personnaz, L., Nadal, J.-P., & Dreyfus, G. (1988). Storage and retrieval of complex sequences in neural networks. *Phys. Rev. A*, *38*, 6365–6372.

- Hasselmo, M. E. (1999). Neuromodulation: acetylcholine and memory consolidation. *Trends Cogn. Sci.*, *3*, 351–359.
- Herrmann, M., Hertz, J. A., & Prügel-Bennett, A. (1995). Analysis of synfire chains. *Network: Comput. Neural Syst.*, *6*, 403–414.
- Herz, A. V. M., Li, Z., & van Hemmen, J. L. (1991). Statistical mechanics of temporal association in neural networks with transmission delays. *Phys. Rev. Lett.*, *66*, 1370–1373
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U. S. A.*, *79*, 2554–2558.
- Isaac, J. T., Nicoll, R. A., & Malenka, R. C. (1995). Evidence for silent synapses: implications for the expression of LTP. *Neuron*, *15*, 427–444.
- Isope, P., & Barbour, B. (2002). Properties of unitary granule cell → purkinje cell synapses in adult rat cerebellar slices. *J. Neurosci.*, *22*, 9668–9687.
- Jensen, O., & Lisman, J. E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends Neurosci.*, *28*, 67–72.
- Kempster, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E*, *59*, 4498–4514.
- Kesner, R. P., Gilbert, P. E., & Barua, L. A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behav. Neurosci.*, *116*, 286–290.
- Lee, A. K., & Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, *36*, 1183–1194.
- Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, *6*, 579–590.
- Little, W. A. (1974). Existence of persistent states in the brain. *Math. Biosci.*, *19*, 101–120.
- Lörincz, A., & Buzsàki, G. (2000). Two-phase computational model training long-term memories in the entorhinal-hippocampal region. *Ann. New York Acad. Sci.*, *991*, 83–111.
- Maier, N., Nimmrich, V., & Draguhn, A. (2003). Cellular and network mechanisms underlying spontaneous sharp wave-ripple complexes in mouse hippocampal slices. *J. Physiol.*, *550*, 873–887.
- Maravall, M. (1999). Sparsification from dilute connectivity in a neural network model of memory. *Network: Comput. Neural Syst.*, *10*, 15–39.
- McCulloch W. S., & Pitts, W. (1943). Logical calculus of ideas immanent in nervous activity. *Bull. of Math. Biophys.*, *5*, 115–133.
- Montgomery, J. M., Pavlidis, P., & Madison D. V. (2001). Pair recordings reveal all-silent synaptic connections and the postsynaptic expression of long-term potentiation. *Neuron*, *29*, 691–701.
- Nadal, J.-P., & Toulouse, G. (1990). Information storage in sparsely-coded memory nets. *Network*, *1*, 61–74
- Nadal, J.-P. (1991). Associative memory: on the (puzzling) sparse coding limit. *J. Phys. A: Math. Gen.*, *24*, 1093–1101,

- Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.*, *19*, 9497–9507.
- Nakazawa, K., McHugh, T. J., Wilson, M. A., & Tonegawa, S. (2004). NMDA receptors, place cells and hippocampal spatial memory. *Nature Rev. Neurosci.*, *5*, 361–372.
- Nowotny, T., & Huerta, R. (2003). Explaining synchrony in feed-forward networks: are McCulloch-Pitts neurons good enough? *Biol. Cybern.*, *89*, 237–241.
- Nusser, Z., Lujan, R., Laube, G., Roberts, J. D. B., Molnar, E., & Somogyi, P. (1998) Cell type and pathway dependence of synaptic AMPA receptor number and variability in the hippocampus. *Neuron*, *21*, 545–559.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1992). *Numerical Recipes in C: The Art of Scientific Computing*, New York: Cambridge University Press.
- Rapp, P. R., & Gallagher, M. (1996). Preserved neuron number in the hippocampus of aged rats with spatial learning deficits. *Proc. Natl. Acad. Sci. U.S.A.*, *93*, 9926–9930.
- Sompolinsky, H., & Kanter, I. (1986). Temporal associations in asymmetric neural networks *Phys. Rev. Lett.*, *57*, 2861–2864.
- Stevens, C. F. (2001). An evolutionary scaling law for the primate visual system and its basis in cortical function. *Nature* *411*, 193–195.
- Stevens, C. F. (2005). Private communication. The Salk Institute for Biological Studies, San Diego, CA.
- Traub, R. D., Bibbig, A., Fisahn, A., LeBeau, F. E. N., Whittington, M. A., & Buhl, E. H. (2000). A model of gamma-frequency network oscillations induced in the rat CA3 region by carbachol in vitro. *Europ. J. Neurosci.*, *12*, 4093–4106.
- Urban, N. N., Henze, D. A., & Barrionuevo, G. (2001). Revisiting the role of the hippocampal mossy fiber synapse. *Hippocampus*, *11*, 408–417.
- Willshaw, D. J., Bunetman, O. P., & Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature*, *222*, 960–962.
- Wilson, M. A., & McNaughton B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*, 676–679.