Memory Formation and Recall in Recurrent Spiking Neural Networks

Friedemann Zenke
I&C, EPFL

***Abstract*—Performance of a memory system which we consider as useful in an everyday scenario might best be characterized by fast learning and slow forgetting. How is this idea compatible with synaptic plasticity in spiking neural networks? The reviewed literature focuses on how plasticity can perform efficiently under the above performance criteria. Yet it does not address homeostatic self control mechanisms that maintain the network in a state of constant vigilance for either acquiring or recalling memories. Here a line of research is proposed that focuses on biologically motivated plasticity mechanisms in the context of large recurrent networks.**

I. Introduction

My project is dedicated to identifying the necessary mechanisms responsible for learning, recalling and stabilizing memories in recurrent spiking neural networks. This research necessarily brings together the different aspects of modeling networks of spiking neurons, the plasticity of their connections and finally the maintenance mechanisms in the background that make sure everything works smoothly.

In the past a lot of attention was given to attractor networks and they still are a field of active study. The common ground of all these models is that memories are stored in the connections. Memory recall proceeds via the network dynamics, which means the system state tends towards the stored patterns of activity. One such study is presented in Section II-A. The work belongs to the class of non-spiking models, which is generally theoretically more tractable. Learning in this type of framework happens once, because for analytical ease the synaptic weights remain fixed. This strong separation of learning and recall contradicts our everyday experience. We can form new memories or forget old ones as we go. Furthermore real world stimuli are subject to change over time and so do memories. The articles reviewed in Sections II-B and II-C concentrate on this idea. How can synapses be modified effectively to guarantee fast acquisition and slow forgetting?

Finally the focus of the three reviews is biased towards topics most relevant for the proposal outlined in Section III.

Proposal submitted to committee: July 30th, 2010; Candidacy exam date: June 29th, 2010; Candidacy exam committee: Exam president: R. Urbanke, thesis director: W. Gerstner, co-examiner: M. Hasler.

This research plan has been approved:

Date: _____

Doctoral candidate: _____
(name and signature)

Thesis director: _____
(name and signature)

Thesis co-director: _____
(if applicable) (name and signature)

Doct. prog. director: _____
(R. Urbanke) (signature)

EDIC-ru/05.05.2009

II. Journal Articles

A. Willshaw model: Associative memory with sparse coding and low firing rates

The article written by Golomb et al. [12] constitutes a theoretical study of the Willshaw model [24]. Unlike earlier publications that focused on optimal storage capacity this analysis concentrates on the robustness of memory retrieval of low activity patterns. Biological plausibility is approached by including a lower activity bound and the inclusion of inhibitory feedback.

1) Mean field theory: The authors consider a model of an associative memory with N binary units. P sparsely coded patterns V^μ are stored in the binary weight matrix by using Willshaw's learning rule. That means units receive input from all other units that were mutually active in any of the patterns stored. This is formalized as follows

$$J_{ij} = \frac{1}{Nf}\Theta\left(\sum_{\mu=1}^{P} V_i^\mu V_j^\mu\right)$$

with the coding level $f \equiv \frac{1}{N}\sum_i V_i^\mu$ for all μ. Θ is the Heaviside step function. J_{ij} can be interpreted as the strength of the synaptic coupling of the two binary units V_i and V_j. The input to unit V_i is then defined as

$$h_i = \sum_j J_{ij} V_j - \frac{K}{Nf}\sum_j V_j + \theta$$

where K characterizes the strength of the inhibitory feedback from all N units. θ is a negative threshold term that can be

interpreted as external input. The system evolves according to stochastic asynchronous dynamics. The noise level is characterized by the finite temperature T.

Golomb et al. chose a statistical mechanical approach. By exploiting the symmetry of the synaptic matrix it is possible to introduce the energy function H defined as

$$H = \underbrace{-\frac{1}{2}\sum_{ij} J_{ij}V_iV_j}_{\text{unit interactions}} + \overbrace{\frac{K}{2Nf}\left(\sum_i V_i\right)^2}^{\text{global inhibition}} \underbrace{-\theta\sum_i V_i}_{\text{external input}} \quad (1)$$

The authors apply the methods of mean field theory under the constraint that the fraction of non-zero connections C is finite. To do so they replace J_{ij} by an effective interaction strength J^0_{ij}. This can be done by averaging over all patterns. An important assumption that is made here is that the stored patterns are random and independent of each other. By additionally assuming a large initial proximity of the system state to one specific pattern (without loss of generality this is chosen to be $\mu = 1$) one can approximate

$$J^0_{ij} \approx \frac{1}{Nf}\left(CV^1_iV^1_j + 1 - C\right) \quad (2)$$

Depending on i and j this expression can only take two values. If both units are active in pattern V^1 they are surely connected. This is guaranteed by the definition of the learning rule. Otherwise if one or both units are not part of pattern V^1 they might still be connected if the two units are active together in any of the other $P-1$ patterns. The probability that any two arbitrary units V_i and V_j are active together in any random pattern is f^2. Suppose we are interested in the probability of $J_{ij} = 0$. For $P-1$ patterns this probability is equal to $(1-f^2)^{P-1}$. For large P this is approximately equal to $C = \exp(-Pf^2)$. Therefore V_i and V_j are connected with probability $C-1$. Consequently Expression 2 simplifies the interaction strength to the question of whether two units are both coding for pattern V^1 or not.

The central idea of mean-field calculations is to use simplifications like Equation 2 to derive expressions that depend on macroscopic quantities only. Although in principle an approximation these theories usually become exact in the $N \to \infty$ limit where mean values become equal to their expectation value.

Due to stochastic nature of the system and because there exists an energy function H Golomb et al. can apply the same methods as for thermodynamic systems. The dynamical evolution of these systems tends to minimize the free energy. The idea of central importance behind this is the fact that the probability of finding an equilibrated thermodynamical system in a given state is maximal for configurations that minimize the free energy. In other words the partition function only has significant contributions around the local minima of the free energy. Therefore it is very desirable to find macroscopically meaningful parameters minimizing the free energy. For an associative memory *meaningful* parameters generally characterize the overlap or other measures of proximity of the dynamics to specific patterns of activity. By finding the local minima of the free energy Golomb et al. obtain the following mean field equations

$$V^+ = \frac{1}{1+\exp^{-\beta h^+}} \quad (3)$$

$$fV^- = \frac{1}{1+\exp^{-\beta h^-}} \quad (4)$$

where β measures the degree of stochasticity. It enters as an inverse temperature $\beta = 1/T$. The quantities $V^\pm$ have an intuitive interpretation. V^+ is the thermal average of the overlap of the *on* units (i.e. a unit that is supposed to be on in the active pattern) with pattern V^1. V^- on the other hand is the thermal mean overlap of the *off* units with the system state. The variable h^+ is the local field experienced by an *on* unit and h^- is the input of the *off* unit respectively. The local fields for the *on* units are given by

$$h^+ = \overbrace{(1-K)V^+}^{\text{on units}} + \overbrace{(1-C-K)V^-}^{\text{off units}} + \overbrace{\theta}^{\text{external}} \quad (5)$$

and for the *off* units respectively

$$\begin{aligned} h^- &= \overbrace{(1-C-K)V^+}^{\text{on units}} + \overbrace{(1-C-K)V^-}^{\text{off units}} + \overbrace{\theta}^{\text{external}} \\ &= h^+ - CV^+ \end{aligned} \quad (6)$$

Here V^+ is equal to the excitatory feedback from other *on* units. The same units additionally cause an inhibitory feedback of magnitude $-KV^+$. Both are represented by the first term on the right hand side RHS of Equation 5. The second term represents the contribution of *off* units that project onto an arbitrary *on* unit. Their contribution represents random crosstalk from other patterns. By using the same arguments as above one can see that the fraction of spurious connections from *off* units is approximately $1-C$. Likewise the *off* units also contribute inhibitory input to the *on* units with magnitude $-K$. The last term on the RHS of Equation 5 represents the global positive input which can be interpreted as external afferents. The local field at *off* units looks very similar to the one at the *on* units (Equation 6). The only difference is that here input from *on* units is random. Its probability is the same as the probability of having a spurious connection $1-C$. Therefore $(1-K)V^+$ needs to be replaced by $(1-C-K)V^+$. This directly yields Equation 6.

2) Low activity solutions: Golomb and colleagues focus their analysis on solutions of Equations 5 and 6 with low global activity $f \to 0$ where f vanishes polynomially as $f \propto N^{-x}$ with $0 < x < 1$. The most important result of this analysis is the prediction of different phases of the network depending on the variable C, the amount of external input θ, the strength of inhibition K and the temperature T. The different phases can be categorized in two groups. On the one hand there are so called *retrieval states* which are characterized by $V^+ \gg fV^-$. That means units coding for the memory pattern to be retrieved are noticeably more active than all the other units. On the other hand there exists a symmetric state with $V^+ \approx fV^-$, where all units of the network display the same averaged activity. Both phases exhibit an overall

network activity $V = \sum_i V_i$ that is relatively low. This is a prominent property also found in many biological systems. While the retrieval state is highly correlated with one of the stored memories and allows for memory recall. The latter state is not correlated to any of the patterns and can be interpreted as background activity.

Golomb et al. systematically analyze the solutions for different parameter ranges. When only looking at the retrieval states one fundamental distinction can be made from the sign of the input to the *on* units h^+. For stable retrieval solutions it is necessary that $h^+ > 0$. In this case the *on* units receive positive feedback and are active with high probability. For the the simple case $T = 0$ it is easily verified that $V^+ = 1$ and $V^+ = 0$ solve Equations 3 and 4 if additionally

$$C > h_0 > 0$$

is fulfilled, where the short notation $h_0 = \theta + 1 - K$ was used. This inequality puts external input, inhibitory feedback and the number of zero-bonds in relation. Such a retrieval state ($T = 0$) represents a state of frozen activity, where the pattern is recalled completely and without erroneous bits. If C is kept constant and at the same time h_0 is increased (this can be achieved by either increasing external input or decreasing feedback inhibition) the retrieval state changes its character. For $h_0 > C$ the *on* units are still active with $V^+ = 1$ while the *off* units receiving connections from *on* units begin to turn on occasionally with probability

$$V^- = \frac{h_0 - C}{K + C - 1}, \quad C < h_0$$

Due to the stochastic asynchronous network dynamics the retrieval state actually exhibits randomly fluctuating *off* units even at $T = 0$. This is fundamentally different from Hopfield networks. It is important to note that this effect truly only arises from to the asynchronous update which inherently contains a source of stochasticity. For a synchronous update the interplay of excitation of the *off* cells and global inhibition would most likely result in an oscillatory network behavior.

Looking at the retrieval states for $T > 0$ Golomb et al. show that for small temperatures i.e. $C > \bar{T} + h_0$ with $\bar{T} = T(-\ln f)$ the change from the behavior described above is negligible. Graphically speaking for small values of T and β the sigmoidal Functions 3 and 4 still look very much like sigmoidal functions. However for higher temperatures $C < \bar{T} + h_0$ the system enters what Golomb and colleagues call the *weak memory regime* where V^+ is still close to 1 but the background activity fV^- becomes substantial. This behavior can be interpreted as a form of escape noise in the spiking picture. Although the mean input h^- is kept negative close to threshold (which is at zero in this model) the decreased slope of the sigmoid still assures a finite firing probability. The threshold for this *weak memory regime* is indeed smooth and Golomb et al. chose it such that it characterizes the point where $V^- \approx 1$.

When temperature is increased even further the weak memory regime destabilizes and the symmetric state eventually becomes the only stable state of the network. This transition depends on the loading level and occurs at the point where $\bar{T} = C$. In this phase memory recall is not possible any more. A similar phase transition is also found in the Hopfield model [2].

Furthermore Golomb et al. show that for finite values of f the symmetric background state actually becomes instable at small temperatures. Then the network always evolves towards a retrieval state. It is also worth mentioning that there are retrieval phases at low temperatures with $h_0 < 0$. Since the net input to all units is negative unit activation can only be caused by thermal fluctuations. These retrieval states are characterized by a low local activity $V^+ < 1$ of *on* units and a quiescent *off* population ($V^- = 0$).

3) Finite size corrections: Since the mean-field arguments rely mainly on the scaling behavior of N, f and P the theory makes predictions for arbitrary f and N. However the mean-field approach is only exact in the limit $N \to \infty$ therefore small N are actually problematic. The authors analyze the finite size effects by estimating the deviations from the idealized J_{ij}^0 The fluctuations of the connection strength J_{ij} result in fluctuations of the input strength $h_i^{\pm}$. For large N these fluctuations can be well approximated as zero-mean Gaussian variables. Fluctuations of the input result in inhomogeneous unit activity. In the Hopfield picture this means that the so called cross talk term becomes significant [13]. If it becomes large enough it eventually can flip a bit in the retrieval pattern and impair the network's performance as an associative memory. The authors propose first order finite size corrections to their mean field solutions to account for fluctuations. Furthermore they show with simulations that the corrections qualitatively improve their results.

In the mean-field calculations $P \to \infty$ was implicit. For a finite network and finite f also P is necessarily finite. Using the formalism developed for finite size corrections the authors derive the maximum storage capacity for various different limiting cases of f and h_0.

In the present model input fluctuations have another striking effect. Indeed retrieval phases that are highly correlated to the stored memories still exist with input fluctuations. Nevertheless in strong contrast to the $N \to \infty$ case here the dynamics are frozen. This behavior is found at low temperatures $\bar{T} \ll \Delta^-$ where Δ^- denotes the standard deviation of the input fluctuations of the *off* units. At the same time the thermal averages deviate only slightly from the values $V^{\pm}$ obtained from the mean field equations.

4) Indirect inhibition: Up to now inhibition entered the energy function as a simple input term $-K$ proportional to the squared global network activity $V = \sum_i V_i$. In that sense inhibitory feedback was direct. Golomb et al. show that inhibition can be realized indirectly by adding a population of inhibitory neurons to their network. Thus the model comes closer to biology where inhibition is indeed realized by dedicated neurons. The inhibitory population in the Golomb model receives uniform input from the excitatory units and has uniform projections to itself and to the excitatory units.

The local fields now change to

$$h_i^{\mathsf{ex}}(t) = \sum_{j=1}^{N} J_{ij} V(t)_j - f^{-1} K U(t) + \theta \quad (7)$$

$$h_i^{\mathsf{in}}(t) = [V(t) - U(t)]/f \quad (8)$$

where $U = 1/N \sum_i U_i$. Equation 7 describes the local field on excitatory units and Equation 8 the uniform field for inhibitory units respectively. Once again it is assumed that the network state is correlated to a particular memory ($\mu = 1$) initially. The mean field equations become

$$\begin{aligned} \frac{\partial}{\partial t} V^+ &= -V^+ + g(\beta h^{\mathsf{ex}+}) \\ \frac{\partial}{\partial t} V^- &= -V^- + g(\beta h^{\mathsf{ex}-}) \\ \frac{\partial}{\partial t} U &= -U + g(\beta h^{\mathsf{in}}) \end{aligned}$$

with $g = (1 + \exp(-x))^{-1}$. The fact that temporal dynamics explicitly enter the mean-field equations allows to have temporally periodic solutions. However Golomb et al. limit their analysis to the stable solutions where the mean-field equations again take the form of Equation 3. Golomb and colleagues show that the problem can be reduced to the same set of mean field equation like in the case discussed above. Based on this they conclude that a system with indirect inhibitory feedback via an inhibitory population has the same stationary solutions as a system with the direct inhibitory coupling as discusses above. They also point out that this argument is limited to the large N limit only. For finite N fluctuations in excitation lead mostly to excitatory runaway behavior that destabilizes phases with low activity ($h_0 < 0$). In these cases the indirect inhibition simply does not react fast enough.

The authors validate their findings from the mean-field theory, the finite size corrections, capacity considerations and finally the indirect inhibition with numerical simulations. They show that their analytical results qualitatively agree well with simulations.

5) Summary and Conclusion: In summary the paper illustrates in a simplistic model robust retrieval of patterns previously stored in the connection matrix. The model relies on a Hebbian learning rule, binary excitatory units, low activity patterns and global inhibitory feedback. The authors illustrate under which conditions memory retrieval is possible. Here the accessible network phases depend in a simple (but non-trivial) way on network load, represented by C, the temperature T, strength of inhibition K and external afferents θ. The paper underlines the importance of the level of stochasticity to maintain retrieval states at low local firing rates.

Although the model discussed in this publication is not a spiking neural network, a lot of the knowledge can be extrapolated to the spiking case. Indeed it has been shown by [1], [19] that spiking neural network models exist that resemble to some extent the results of Golomb et al. The article can therefore be seen as a proof of principle study for related work done on spiking neural networks. However, in spiking models generally a lot of fine tuning is required. While the level of stochasticity and the external input remain free parameters in Golomb's model, the dynamical state of recurrent neural networks is generally very sensitive to changes of them. Here additional care has to be taken by the wide variety of dynamical regimes expressed by these networks. For spiking neural networks the accessible dynamical states are ranging from regular types of behavior to asynchronous irregular (AI) states [4], [22]. Biological networks generally exhibit low global and local firing rates. These effects are not completely described by the model of Golomb et al. as a non-spiking model. In the Golomb model local low activity retrieval states exist at the cost of background activity in the remaining *off* units. If the *off* neurons are to display some level of background activity retrieval states are accompanied with high local activity which is close to saturation.

Another important point is that asynchrony in the Golomb model is generated inherently by the dynamic update rule, while plasticity is not included at all (except at the beginning when the weight matrix is set). Since activity in many parts of the brain seems to be very similar to the AI state a biologically plausible model should be able to reproduce AI behavior. Although the AI state is the basis for speculation its universality seems to suggest that it plays an important role. If one accepts this assumption it becomes a fundamental criterion when considering plasticity in recurrent networks. Plasticity should ultimately serve the conservation of this state. It is at the heart of homeostatic plasticity that functional alterations and memory formation can proceed without destroying what is considered as *healthy* network activity.

B. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates

In this article [8] Fusi addresses the questions of learning in an associative memory that does not necessarily distinguish between learning and recall. This idea of on-line learning is biologically motivated by cortical reverberations that have been found experimentally in the cortex. Such reverberations are likely to result from a feedback mechanism between multiple cells. They only occur for stimuli that are already known. Most real learning systems exhibit the palimpsest property. That means old memories are forgotten for the benefit of new ones. This article deals with the question of how quickly previously stored memories are forgotten and how the memory span can be extended.

The three underlying assumptions that serve as a basis of the theoretical considerations are:

1) Locality in time and space - i.e. a synapse can only access and process information that is currently available about the two neurons it connects.
2) All internal variables of a synapse are bounded quantities.
3) Modifications of synapses are finite and long-term modifications cannot be infinitesimal.

A synaptic system which fulfills these requirements automatically exhibits the palimpsest property. This results in a finite memory span during which memories can be recalled. Memories older than that are lost.

To approach the problem analytically Fusi makes use of assumptions 2 and 3 to motivate discrete synaptic systems. The synapse dynamics in a network subject to a Hebbian learning rule and exposed to a stream of random stimuli is mapped onto a random walk. The transition matrix is defined as

$$\mathbf{M} = \sum_{\xi_{\mathsf{pre}},\xi_{\mathsf{post}}} p(\xi_{\mathsf{pre}},\xi_{\mathsf{post}})\mathbf{Q}(\xi_{\mathsf{pre}},\xi_{\mathsf{post}}) \tag{9}$$

where $p(\xi_{\mathsf{pre}},\xi_{\mathsf{post}})$ is the probability of finding the pre- and postsynaptic activities ξ_{pre} or ξ_{post} respectively for a specific stimulus. $\mathbf{Q}$ represents the learning rule and therefore the transitions probabilities from one synaptic state to the other, given the pre- and postsynaptic activity. The study now focuses on an ongoing stream of stimuli that ought to be stored. For the analysis one of these stimuli is selected and tracked as it fades over time. Fusi supposes that that specific stimulus imposes the activities ξ^1_{pre} or ξ^1_{post} which results in a modified probability density of $\rho^1_K(\xi^1_{\mathsf{pre}},\xi^1_{\mathsf{post}})$. The conditional probability function ρ of finding an arbitrary synapse in state J after the presentation of p stimuli is then given by

$$\rho^p_J(\xi^1_{\mathsf{pre}},\xi^1_{\mathsf{post}}) = \sum_{K=1}^{n_s} \rho^1_K(\xi^1_{\mathsf{pre}},\xi^1_{\mathsf{post}})(\mathbf{M}^{p-1})_{KJ}$$

where the sum runs over all synaptic states. The possible transitions caused by a stream of $p-1$ random stimuli are given by $\mathbf{M}^{p-1}$. Under the assumption that all synaptic states are eventually accessible from any given state (ergodicity) there exists an asymptotic distribution of weights. The asymptotic distribution ρ^∞_J is independent of the initial distribution. It satisfies

$$\rho^p_J(\xi^1_{\mathsf{pre}},\xi^1_{\mathsf{post}}) \to \rho^\infty_J$$

Informally this means that the weight distribution $W(p)$ after p steps should be sufficiently different from the asymptotic distribution. This is to ensure that the system is still able to contain specific information about the first stimulus. This idea introduces the notion of a memory span and is formalized as follows

$$\Delta W = \left| \sum_{ij} [W_{ij}(p) - W_{ij}(\infty)] \right| > \delta \tag{10}$$

where δ is a parameter characterizing the amount of change left of the memory being tracked. Based on the assumption that the number of synapses is significantly larger than the number of synaptic states n_s one can estimate ΔW as

$$\Delta W \approx N^2 \left| \sum_{J} \left[\rho^p_J(\xi^1_{\mathsf{pre}},\xi^1_{\mathsf{post}}) - \rho^\infty_J\right] \right| > \delta \tag{11}$$

where it was additionally assumed that the number of synapses scales with $\sim N^2$. Here replacing the sum over i and j by a sum over the synaptic states can be seen as a form of self-averaging. This is only justified because $N^2 \gg n_s$.

In the following Fusi uses the ergodic property. According to the ergodic theorem the transition probabilities $\mathbf{M}^p$ converge to their limit at least as

$$|(\mathbf{M}^p)_{KJ} - \mathbf{M}^\infty)_{KJ}| < (1 - \Lambda\ _{min})^{(p/n_0)-1}$$

where $\Lambda_{\mathsf{min}} = \min_{KJ}(\mathbf{M}^p)_{KJ}$. By using Equation 9 and substituting this as an upper bound into Equation 11 one finally obtains

$$\Delta W < N^2 n_s W_{\mathsf{max}} (1 - \Lambda_{\mathsf{min}})^{(p/n_0)-1} \tag{12}$$

where W_{max} denotes the maximum efficacy. This expression represents an upper bound for p

$$p < \frac{-n_o \log(n_s N^2/\delta)}{\log(1 - \Lambda_{\mathsf{min}})} \tag{13}$$

The scaling of p therefore depends only on the number of neurons N, the number of synaptic states n_s, the mean number of presentations to visit all synaptic states n_0 and Λ_{min}. It is an important result of the article that Equation 13 scales only logarithmically and thus very poorly in N^2 and n_s. As a consequence Fusi proposes to chose Λ_{min} appropriately small to achieve large p. This idea is at the heart of stochastic learning. Since Λ_{min} scales with the transition probabilities between the synaptic states it can be decreased by reducing these very transition probabilities. Doing so assures that only a finite subset of all synapses is modified upon stimulus presentation.

How can such a stochastic update be achieved? Motivated by the results described above the publication turns to the more practical part. Accepting the fact that high capacity can only be achieved efficiently for small synaptic transition probabilities Fusi suggests that the irregular activity as found in biological networks itself could serve as a noise source to produce stochastic on-line selection of synapses to be updated. To illustrate this point a bistable synaptic system is introduced that makes use of the stochasticity generated by the network. The synapse relies on an internal state variable and a threshold, while the dynamics are biased by a Hebbian learning term. The main results concerning this analysis are that transition probabilities are strongly firing rate dependent and very long switching times can occur for low levels of activity. This feature prevents forgetting in presence of spontaneous network activity without having to introduce long synaptic time constants.

In summary this article [8] develops a framework that gives insight into the limiting behavior of storage capacity for synaptically stored memories. The most important fact is that in an on-line learning scenario memories are overwritten constantly by new memories. To protect memories from ongoing plasticity the concept of stochastic learning is proposed that reduces the overall rate of plasticity events. As a direct consequence memories are overwritten less rapidly and memory lifetimes are increased. However if successful storage of a memory needs a certain amount of synaptic change, this number is also going to be obtained slower as a direct consequence of stochastic plasticity. This means that the memory lifetime is increased at the expense of initial signal to noise ratio (SNR). Fusi leaves this question open in this publication but solves it in [9]. This article is reviewed in the following section.

Another important point is that the underlying ideas are based on three axiomatic assumptions that do not depend in detail on the underlying architecture. Therefore the results remain qualitative. Here the ignorance of the absolute scale of

storage capacity is put in the parameter δ (see Equation 10). To gain quantitative insight into the capacity of such a system one necessarily has to appoint a specific network architecture and a specific notion of memory recall. Although the main theoretical results of the article do not depend on architectural details the proposition of a stochastic synaptic update based on spiking background activity does. Nevertheless the question of the self control mechanisms that lead to, and ultimately maintain this state of activity in the network, is not addressed. A logical continuation of this work should therefore include a defined architecture for a spiking neural network that includes homeostatic plasticity and makes use of stochastic synaptic updates. A specific synaptic plasticity model for spiking neural networks that already includes two of these features is for example the TagTriC model [6] which serves as one important cornerstone for my work.

C. Cascade Models of Synaptically Stored Memories

In the article reviewed here [9] Fusi and colleagues address the problem of memory lifetimes from a different perspective. As was already summarized in Section II-B, preservation of memories requires as little plasticity as possible. Controversially at the same time the effective storage of new stimuli needs a high degree of plasticity. Fusi and colleagues propose one possible solution to this. By dint of introducing a wide range of timescales manifested in what they call meta-plastic states they show that a synaptic model can indeed combine fast memory storage with long retention times.

For the qualitative analysis of memory lifetimes three important probabilities are introduced. r denotes the probability per unit time that a plasticity event occurs. $f_\pm$ denotes the probability for weakening (-) or strengthening (+) respectively of a synapse. Finally q stands for the probability that a candidate event for synaptic strengthening or depression given by the the rates $rf_\pm$ is actually applied.

The analysis follows the same idea like [8]. From a continuous stream of incoming stimuli one memory is chosen and its SNR is tracked over time to see how it fades. The "signal" in this framework is proportional to $\sim qN_{\mathrm{syn}}$ where N_{syn} is the number of synapses that have been affected by a given memory storage event. If the rate of synaptic change r is constant and the system is continuously exposed to new stimuli then similarly to nuclear decay the probability at time t that a particular synapse has not been modified again is given by the function $\exp(-qrt)$. Synaptic noise due to ongoing plasticity is assumed to be in the order of $\sim \sqrt{N_{\mathrm{syn}}}$. The SNR then behaves like $\sim qN_{\mathrm{syn}}\exp(-qrt)$ and is equal to one at time

$$t_{\max} \sim \frac{\ln(q\sqrt{N_{\mathrm{syn}}})}{qr}$$

This qualitative expression immediately shows that forgetting proceeds exponentially in time and scales logarithmically in N_{syn}. In [8] this problem was approached by allowing q to be very small. This indeed solves the problem of short memory lifetimes but at the expense of a high initial SNR. In fact it can be shown that the SNR is in the order of 1 for optimal q. By this Fusi and colleagues illustrate that a strong initial memory signal is incompatible with long memory lifetimes. The effect is due to the unavoidable nature of exponential decay[1] of the memory trace. The unwanted behavior can be loosened by requiring power-law forgetting instead. This means that the memory signal decays as $t_{\max} \sim t^{-k}$ for some positive number k.

The novel idea outlined in the publication is that such power-law behavior can be mimicked by introducing memory traces that decay exponentially using a wide range of different decay time constants. This idea forms the basis for meta-plasticity. While short lived synaptic states contribute to a high initial SNR, the long lived synaptic states assure a long memory lifetime. In the *Cascade Model* synapses with the same strengths can be in different states of meta-plasticity. These states are called meta-plastic states. That means that a synapse can undergo meta transitions that only change the modification probability of the synapse q but not the synaptic efficacy. In the example given in the article q can take values $q_i = x^{i-1}$ with $i = 1..n-1$, where states with large i can be thought of as better protected from future change, or figuratively further down in the cascade. Candidate events for changing synaptic efficacy are executed with probability q_i. If executed they result in a change of the efficacy and a reset of the meta-plastic state to $i = 1$. Therefore the synapse is more susceptible to future plasticity. On the other hand candidate events that would leave the synaptic efficacy unchanged can result in a change of the meta-plastic state and therefore in an increase of the variable i which manifests in a more consolidated synaptic state.

The cascade model can be seen as an extension of the ideas brought forward in [8]. Transitions from one synaptic state to the other are stochastic but the susceptibility to change is directly dependent of the pre-history of the synapse itself. It is reflected in the meta-plastic state. In this manner the cascade model establishes a notion of importance for a specific synapse being in a designated efficacy or plasticity state.

It is demonstrated with mean field calculations and numerical simulations that indeed this meta-plastic modification rule leads to power-law forgetting over extended periods of time. It is shown that even for a small number of meta-states the model outperforms comparable switch like synapses without meta-plasticity.

Although the stochastic learning seems to give a solution for long memory lifetimes it does not support a high degree of initial plasticity if the transition probabilities for synaptic state transitions are fixed. The cascade model solves this by introducing, on top of stochastic learning, a synaptic model that makes use of meta-plastic states. Similarly to [8] it is a general concept that is introduced in this publication rather than a specific implementation. However, since the underlying arguments are formulated in a probabilistic way the transfer to spike timing dependent plasticity (STDP) or stochastic STDP in particular is straight forward.

[1] Here the word decay simply illustrates the fact that the synaptic state might be altered due to ongoing plasticity.

III. Research Proposal

In summary [12] shows how a very simple learning rule can lead to successful memory retrieval in a non spiking network, while [8], [9] give important hints of how a synaptic update rule should look like to perform well in on-line learning tasks. The main ideas can readily be generalized to the spiking network case. However, what the articles do not address is the question of what mechanisms do have to be present in the plasticity rules to preserve network integrity, AI background activity and ultimately achieve longterm stability of the memory system. A logical continuation therefore is to study memory formation in large spiking neural network models.

Plasticity in spiking neural networks often comes in the flavor of spike timing dependent plasticity (STDP). It is a form of Hebbian plasticity that takes into account causality of pre- and postsynaptic spike times. STDP is believed to form along with other forms of plasticity the basis of learning in neural circuits in the brain [20]. Recently homeostatic properties of STDP have received increasing interest [21], [14], [17], [23]. However up to now it remains an open question how plasticity forms functional neural circuits and achieves the necessary degree of fine tuning to preserve network integrity.

There is one class of STDP learning rules that takes into account interactions of spike triplets [18]. These rules seem to excel in describing experimental data and have the interesting property that they regain temporal symmetry in the high rate limit [5]. This property allows for learning of rate coded patterns in an on-line fashion as proposed by Fusi [8]. Evidence for such behavior has been reported experimentally in different associative areas in the brain (see [10], [16], [25] for some examples).

Concerning STDP learning rules this study can rely on previous studies in the field (in particular [6], [5]). The challenge is now to port the most important concepts identified earlier to spiking network models that exhibit plausible behavior (e.g. AI background activity). In addition these models should be to a certain extent self stabilizing and perform the given task well enough that they can be fitted to experimental data. This line of research is both novel and tries to make a logical and important step forward in understanding fundamental principles of neural computation.

In preliminary studies I performed numerical simulations on recurrent networks subject to STDP. The results show that a behavior similar to the one reported experimentally can be reproduced. However the results suggest also that there might be shortcomings in homeostatic control. What is seen consistently during simulation experiments is that memories can be formed reliably in the framework of a simple rate code. Nevertheless during activation of a memory new cells are recruited that where initially not coding for the memory. There is good reason to believe that this can be understood theoretically in the framework of [3]. Since cortical reverberations seem to require a critical minimum number of cells to give sufficient excitatory feedback this mechanism could in principle be very useful. However, unfortunately the process does not stop and eventually develops runaway behavior. It is therefore essential to identify a mechanism that is capable of regulating the number of units coding for a pattern. Such a mechanism should act on the same timescale as plasticity itself. Therefore one promising candidate mechanism relies on inhibitory plasticity. This could be either indirect, via plastic excitatory synapses of inhibitory inter-neurons, or direct via plastic inhibitory synapses. Although inhibitory plasticity has been found in biological systems [11] it has not been studied in depth. This seems to be a promising direction for further research.

In a second line of research I am working with a particular synaptic model developed by Clopath et al. [6]. The model includes synaptic tagging as has been reported experimentally [7] and uses a mechanism similar to metaplasticity as described in Section II-C. In its original form the synaptic changes do not depend on the postsynaptic spikes directly but on the membrane potential [5]. I performed numerical simulations to study the behavior of the model in networks which can best be described as to be in the high conductance state [15]. In this context the membrane potential is subject to strong fluctuations close to firing threshold. This in turn affects the way synaptic weights change in different dynamical states of the cell. The effect is qualitatively different from other STDP models and the implications are not well understood. Further research needs to be conducted in this area.

In specific I am proposing to study (1) homeostatic STDP in spiking networks. A first partial aim is to make a network learn repeated rate coded stimuli from a continuous stream of stimulation and demonstrate stability over many recall cycles. A second aim is to demonstrate the necessity of a homeostatic mechanism beyond mechanisms already included in [6]. Here I am proposing in particular to systematically investigate and explore the homeostatic effects of direct or indirect inhibitory plasticity.

One long term aim (2) is to study encoding of spatio-temporal patterns in recurrent networks with plasticity and realistic synaptic delays. Side projects along these lines of research include to investigate the evolution and the asymptotic distribution of the synaptic weights under triplet STDP for AI background activity.

These studies are to be accompanied directly (3) by systematic comparative studies of spike- and voltage-based STDP rules as well as models accounting for synaptic tagging and metaplasticity.

First results on (1) and (3) are expected by the end of the first year. The detailed study of (1) is likely to continue during the second year. The studies (2) require advances in software parallelization which will start by the second year. First results are expected towards the end of the second year.

References

[1] Amit, D. J., and Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex 7*, 3 (May 1997), 237–252.

[2] Amit, D. J., Gutfreund, H., and Sompolinsky, H. Statistical mechanics of neural networks near saturation. *Annals of Physics 173*, 1 (Jan. 1987), 30–67.

[3] BIENENSTOCK, E., COOPER, L., AND MUNRO, P. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience 2*, 1 (Jan. 1982), 32.
[4] BRUNEL, N. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience 8*, 3 (May 2000), 183–208.
[5] CLOPATH, C., BÜSING, L., VASILAKI, E., AND GERSTNER, W. Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nature Neuroscience 13*, 3 (Mar. 2010), 344–52.
[6] CLOPATH, C., ZIEGLER, L., VASILAKI, E., BÜSING, L., AND GERSTNER, W. Tag-trigger-consolidation: a model of early and late long-term-potentiation and depression. *PLoS Computational Biology 4*, 12 (Dec. 2008), e1000248.
[7] FREY, U., AND MORRIS, R. G. M. Synaptic tagging and long-term potentiation. *Nature 385*, 6616 (Mar. 1997), 533–6.
[8] FUSI, S. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biological Cybernetics 87*, 5 (Dec. 2002), 459–470.
[9] FUSI, S., DREW, P. J., AND ABBOTT, L. F. Cascade models of synaptically stored memories. *Neuron 45*, 4 (Feb. 2005), 599–611.
[10] FUSTER, J. M., AND JERVEY, J. P. Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. *Science (New York, N.Y.) 212*, 4497 (May 1981), 952–5.
[11] GAIARSA, J.-L., AND BEN-ARI, Y. *The Dynamic Synapse: Molecular Methods in Ionotropic Receptor Biology*. CRC/Taylor & Francis, 2006, ch. 2.
[12] GOLOMB, D., RUBIN, N., AND SOMPOLINSKY, H. Willshaw model: Associative memory with sparse coding and low firing rates. *Physical Review A 41*, 4 (Feb. 1990), 1843–1854.
[13] HERTZ, J., KROGH, A., AND PALMER, R. Introduction to the theory of neural computation.
[14] KEMPTER, R., GERSTNER, W., AND VAN HEMMEN, J. L. Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Computation 13*, 12 (Dec. 2001), 2709–41.
[15] KUMAR, A., SCHRADER, S., AERTSEN, A., AND ROTTER, S. The high-conductance state of cortical networks. *Neural Computation 20*, 1 (Jan. 2008), 1–43.
[16] MIYASHITA, Y. Neural correlate of visual associative long-term memory in the primate temporal cortex. *Nature 335*, 27 (Oct. 1988), 817–820.
[17] MORRISON, A., AERTSEN, A., AND DIESMANN, M. Spike-timing-dependent plasticity in balanced random networks. *Neural Computation 19*, 6 (June 2007), 1437–67.
[18] PFISTER, J.-P., AND GERSTNER, W. Triplets of spikes in a model of spike timing-dependent plasticity. *The Journal of neuroscience : the official journal of the Society for Neuroscience 26*, 38 (Sept. 2006), 9673–82.
[19] RENART, A., MORENO-BOTE, R., WANG, X.-J., AND PARGA, N. Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural Computation 19*, 1 (Jan. 2007), 1–46.
[20] SJÖSTRÖM, J., AND GERSTNER, W. Spike-timing dependent plasticity. *Scholarpedia 5*, 2 (2010), 1362.
[21] SONG, S., MILLER, K. D., AND ABBOTT, L. F. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience 3*, 9 (Sept. 2000), 919–26.
[22] VOGELS, T. P., RAJAN, K., AND ABBOTT, L. F. Neural network dynamics. *Annual Review of Neuroscience 28*, c (Jan. 2005), 357–76.
[23] WATT, A. J. Homeostatic plasticity and STDP: keeping a neuron's cool in a fluctuating world. *Frontiers in Synaptic Neuroscience 2*, June (June 2010), 1–16.
[24] WILLSHAW, D., BUNEMAN, O., AND LONGUET-HIGGINS, H. Non-holographic associative memory. *Nature 222*, 7 (June 1969), 960–962.
[25] YAKOVLEV, V., FUSI, S., BERMAN, E., AND ZOHARY, E. Inter-trial neuronal activity in inferior temporal cortex: a putative vehicle to generate long-term visual associations. *Nature Neuroscience 1*, 4 (Aug. 1998), 310–7.