related to the network's ability to maintain a persistent state. We discuss generalizations to other models of neural firing, and consequences for coding and computation in continuous attractor networks.

doi:

I-26. Analyzing the trade-off between plasticity and stability in recurrent neuronal networks

Friedemann Zenke FRIEDEMANN.ZENKE@EPFL.CH
Guillaume Hennequin GUILLAUME.HENNEQUIN@EPFL.CH
Henning Sprekeler HENNING.SPREKELER@EPFL.CH
Tim Vogels TIM.VOGELS@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH

École Polytechnique Fédérale

Over the last 30 years, Bienenstock Cooper Munro (BCM) type learning rules have shaped our understanding of synaptic plasticity. While they excel at explaining the emergence of receptive fields and stimulus selectivity in networks with feed-forward architecture, their impact on recurrent scenarios is less distinctive. Here, we analyze general BCM-type synaptic plasticity rules with a homeostatic sliding threshold in the framework of recurrent networks of rate-based and spiking neurons. We begin by considering the effects of learning rate and homeostatic timescales on network stability in a non-linear firing rate model. We show how a sensible choice of timescales leads to stable weight dynamics, but other seemingly sensible parameter choices will inevitably lead to catastrophic run-away potentiation. We discuss under which conditions a stable fixed-point in a regime of Hebbian learning exists. We then study the network's response to perturbations and quantify the critical point whereupon network stability is compromised. By viewing perturbation as a consequence of pattern storage in synaptic connections, we quantify the number of such patterns that can be learned safely in a given time. Our model could provide simple explanations as to why memory intake capacity is limited and why learning becomes increasingly inefficient during intensive learning periods. We confirm these findings in numerical simulations of spiking neural networks and show that our analytical results apply to synapses with triplet based STDP rules.

doi:

I-27. The role of horizontal long-range connections in shaping the dynamics of multi-electrode array data

Alexander Lerchner[1] LERCHNER@GATSBY.UCL.AC.UK
Krishna V. Shenoy[2] SHENOY@STANFORD.EDU
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]University College London
[2]Stanford University

Recordings from electrode arrays in primate motor- and premotor cortex indicate that the neural state does not change abruptly but traces out smooth trajectories in firing rate space ("neural space"). Specifically, single-trial analysis shows a positive correlation between the distance that must be traversed in neural space and the reaction time of the subject. Moreover, there is evidence for inertial dynamics, as if the direction and magnitude of motion in neural space cannot change arbitrarily quickly. Given current models and theories of spiking networks, it is not clear what network properties might underlie such dynamics. We model the piece of cortex that is covered by the electrode array as a collection of unstructured local balanced spiking networks—one for each electrode position. The local networks are then interconnected via "horizontal long-range" connections that diminish in strength with increasing distance. We find that the dynamics in neural state space depend strongly on the existence and

type of horizontal long-range connections. When such connections are absent, or when they consist of equally strong excitatory and inhibitory connections, the network displays dynamics that are essentially abrupt with no distance-time correlation. For the case of no lateral connections, this behavior is in agreement with balanced network theory for unstructured cortical networks. When the long-range connections are excitatory only, however, the model exhibits “slow” and smooth dynamics that are in good agreement with experimental observations. Furthermore, the same network configuration exhibits a dependence on initial firing rate velocity as observed in experiments. Anatomical data shows that horizontal long-range connections in cortex are excitatory only, while the influence of inhibitory neurons is primarily local. Our results suggest that such an arrangement may necessarily entail a smooth evolution of network states in multi-electrode array data.

doi:

I-28. A firing-rate model that better approximates the population dynamics of spiking networks

Evan S Schaffer ESS2129@COLUMBIA.EDU
Larry F Abbott LFA2103@COLUMBIA.EDU

Columbia University

Firing-rate models provide an attractive approach for studying large neural networks because they can be simulated rapidly and are amenable to mathematical analysis. Traditional firing-rate models have the shortcoming of using exponential decay with a single time constant (typically a membrane or synaptic time constant) to describe all changes in rate, which is only appropriate in the high-noise, asynchronous regime. This is a serious limitation because transient synchronization of subgroups of neurons is commonly seen in a variety of neural systems, and it may be an important mechanism for generating rapid responses. To address this issue without losing the advantages associated with modeling firing-rates, we have developed a model based on a Fokker-Planck description of the membrane-potential distribution dynamics of a population of spiking neurons. Whereas most methods for approximating solutions to such a Fokker-Planck equation lead to considerably more complex equations than are practical for large networks, we show that the firing-rate for a population of Quadratic Integrate-and-Fire (QIF) model neurons can be approximated in a surprisingly simple form. Using this approach, we compare a large randomly connected network of excitatory (E) and inhibitory (I) neurons to a network of two rate units. We find that the firing-rate network provides a good approximation to the time-varying activity of the spiking network across a wide range of parameters, including changes in synaptic strength, as can occur during learning. Most surprisingly, we also find that the simple rate network can approximate the phase diagram of a large E-I spiking network, predicting the bifurcation between synchronous and asynchronous states. Finally, although our model was derived from the QIF model, we show that a rate model of this form is a good description of firing-rate dynamics in other spiking models.

doi:

I-29. Learning sparse representations through learned inhibition

Paul Hollensen HOLLENSE@CS.DAL.CA
Thomas Trappenberg TT@CS.DAL.CA

Dalhousie University

As already pointed out by Barlow (Barlow, 1961), learning sparse representation is a key ingredient in biological systems. Sparse representations offer a number of benefits to biological systems. Primarily, they lead to the exploitation of statistical regularities in the incoming signal by forcing the signal to be represented with a relatively small number of active neurons or populations. This eases subsequent interpretations by suppressing noise