

Suggested Project Manuscripts - Phys178/278 Winter 2021

1. Learning in layered networks with biologically plausible learning rules

Krotov and Hopfield 2019

PNAS

Unsupervised learning by competing hidden units

<https://www.pnas.org/content/116/16/7723.short>

It is essential for learning good feature detectors in early layers of artificial neural networks, so that these detectors are useful for the task performed by the higher layers of that neural network. At the same time, the traditional form of backpropagation is biologically implausible. In the present paper we propose an unusual learning rule, which has a degree of biological plausibility and which is motivated by Hebb's idea that change of the synapse strength should be local—i.e., should depend only on the activities of the pre- and postsynaptic neurons. We design a learning algorithm that utilizes global inhibition in the hidden layer and is capable of learning early feature detectors in a completely unsupervised way. These learned lower-layer feature detectors can be used to train higher-layer weights in a usual supervised way so that the performance of the full network is comparable to the performance of standard feedforward networks trained end-to-end with a backpropagation algorithm on simple tasks.

Reproducing all the math is challenging but suitable for a physics graduate students. All students can download the simulation code and demonstrate the utility of this approach to learning in a feedforward model

AND/OR

Song Mei, Andrea Montanari and Phan-Minh Nguyen 2018

A mean field view of the landscape of two-layer neural networks

PNAS

<https://www.pnas.org/content/115/33/E7665.short>

Multilayer neural networks are among the most powerful models in machine learning, yet the fundamental reasons for this success defy mathematical understanding. Learning a neural network requires optimizing a nonconvex high-dimensional objective (risk function), a problem that is usually attacked using stochastic gradient descent (SGD). Does SGD converge to a global optimum of the risk or only to a local optimum? In the former case, does this happen because local minima are absent or because SGD somehow avoids them? In the latter, why do local minima reached by SGD have good generalization properties? In this paper, we consider a simple case, namely two-layer neural networks, and prove that—in a suitable scaling limit—SGD dynamics is captured by a certain nonlinear partial differential equation (PDE) that we call distributional dynamics (DD). We then consider several specific examples and show how DD can be used to prove convergence of SGD to networks with nearly ideal generalization error. This description allows for “averaging out” some of the complexities of the landscape of neural networks and can be used to prove a general convergence result for noisy SGD.

Reproducing all the math is challenging but suitable for a physics graduate students.

2. Synaptic connectivity in relation to memory storage

Brunel 2016

Is cortical connectivity optimized for storing information?

Nature Neuroscience

<https://www.nature.com/articles/nn.4286>

Cortical networks are thought to be shaped by experience-dependent synaptic plasticity. Theoretical studies have shown that synaptic plasticity allows a network to store a memory of patterns of activity such that they become attractors of the dynamics of the network. Here we study the properties of the excitatory synaptic connectivity in a network that maximizes the number of stored patterns of activity in a robust fashion. We show that the resulting synaptic connectivity matrix has the following properties: it is sparse, with a large fraction of zero synaptic weights ('potential' synapses); bidirectionally coupled pairs of neurons are over-represented in comparison to a random network; and bidirectionally connected pairs have stronger synapses on average than unidirectionally connected pairs. All these features reproduce quantitatively available data on connectivity in cortex. This suggests synaptic connectivity in cortex is optimized to store a large number of attractor states in a robust fashion.

Brunel shows that perceptron theory to say something about cortical connectivity. Reproducing all the math is hard but suitable for a physics graduate students. All students will want to draw a connection between the claims and experimental findings; a place to start is with "*Highly nonrandom features of synaptic connectivity in local cortical circuits*" (S Song, PJ Sjöström, M Reigl, S Nelson, DB Chklovskii. PLoS Biol 3 (3), e68/

AND/OR

Battista and Monasson 2020

Storage of correlated patterns in recurrent networks

Physical Review Letters

<https://doi.org/10.1103/PhysRevLett.124.048302>

Recurrent neural networks (RNN) are powerful tools to explain how attractors may emerge from noisy, high-dimensional dynamics. We study here how to learn the $\sim N^2$ pairwise interactions in a RNN with N neurons to embed L manifolds of dimension $D \ll N$. We show that the capacity, i.e., the maximal ratio L/N , decreases as $|\log \epsilon| - D$, where ϵ is the error on the position encoded by the neural activity along each manifold. Hence, RNN are flexible memory devices capable of storing a large number of manifolds at high spatial resolution. Our results rely on a combination of analytical tools from statistical mechanics and random matrix theory, extending Gardner's classical theory of learning to the case of patterns with strong spatial correlations.

A true "theory" paper. Reproducing all the math is hard but suitable for a physics graduate students.

3. Probing networks through perturbations to single neuronal outputs

Sadeh and Clopath 2020

Relationship between connectivity and effects of perturbation

PNAS

<https://doi.org/10.1073/pnas.2004568117>

To unravel the functional properties of the brain, we need to untangle how neurons interact with each other and coordinate in large-scale recurrent networks. One way to address this question is to measure the functional influence of individual neurons on each other by perturbing them in vivo. Application of such single-neuron perturbations in mouse visual cortex has recently revealed feature-specific suppression between excitatory neurons, despite the presence of highly specific excitatory connectivity, which was deemed to underlie feature-specific amplification. Here, we studied which connectivity profiles are consistent with these seemingly contradictory observations, by modeling the effect of single-neuron perturbations in

large-scale neuronal networks. Our numerical simulations and mathematical analysis revealed that, contrary to the prima facie assumption, neither inhibition dominance nor broad inhibition alone were sufficient to explain the experimental findings; instead, strong and functionally specific excitatory–inhibitory connectivity was necessary, consistent with recent findings in the primary visual cortex of rodents. Such networks had a higher capacity to encode and decode natural images, and this was accompanied by the emergence of response gain nonlinearities at the population level. Our study provides a general computational framework to investigate how single-neuron perturbations are linked to cortical connectivity and sensory coding and paves the road to map the perturbome of neuronal networks in future studies.

All students will want to draw a connection between the claims and experimental findings; a place to start is with "*Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex*" (M London, A Roth, L Beeren, M Häusser, PE Latham Nature 466 (7302), 123-127)

4. Learning

Lu and Bassett 2020

Invertible generalized synchronization: A putative mechanism for implicit learning in neural systems

Chaos 30, 063133; doi: 10.1063/5.0004344

<https://doi.org/10.1063/5.0004344>

Regardless of the marked differences between biological and artificial neural systems, one fundamental similarity is that they are essentially dynamical systems that can learn to imitate other dynamical systems whose governing equations are unknown. The brain is able to learn the dynamic nature of the physical world via experience; analogously, artificial neural systems such as reservoir computing networks (RCNs) can learn the long-term behavior of complex dynamical systems from data. Recent work has shown that the mechanism of such learning in RCNs is invertible generalized synchronization (IGS). Yet, whether IGS is also the mechanism of learning in biological systems remains unclear. To shed light on this question, we draw inspiration from features of the human brain to propose a general and biologically feasible learning framework that utilizes IGS. To evaluate the framework's relevance, we construct several distinct neural network models as instantiations of the proposed framework. Regardless of their particularities, these neural network models can consistently learn to imitate other dynamical processes with a biologically feasible adaptation rule that modulates the strength of synapses. Further, we observe and theoretically explain the spontaneous emergence of four distinct phenomena reminiscent of cognitive functions: (i) learning multiple dynamics; (ii) switching among the imitations of multiple dynamical systems, either spontaneously or driven by external cues; (iii) filling-in missing variables from incomplete observations; and (iv) deciphering superimposed input from different dynamical systems. Collectively, our findings support the notion that biological neural networks can learn the dynamic nature of their environment through the mechanism of IGS.

Questions:

1. From Section V of the paper: " A recent work has shown that a reservoir computer can be successfully trained to solve the missing-variable-inference problem for chaotic systems. Here, with our biological plausible learning framework, we consider whether the experience of learning the full dynamical system can, thereafter, immediately allow the neural network to fill in missing variables from a sparse measurement of the learned dynamical system. " Train the RNN on all variables of an input ODE system $s(t)$ of your own choosing. Appendix B provides some examples. Provide the trained RNN with a stimulus that is a possibly incomplete-dimensional version of $s(t)$ (so if $s(t)$ is 4-dimensional, try 1, 2, or 3 dimensions).

How well does the RNN output the correct trajectory for all dimensions of $s(t)$ on average over multiple trials (use a metric like net mean squared error to quantify how well the RNN predicts)? If the stimulus is provided over only a limited window of time, does the RNN continue to output the correct trajectory after the stimulus has ended?

2. From Section VI of the paper: " in real environments, complex adaptive systems—including the human brain—often process mixed sensory input given by a superposition of multiple input systems.³⁷ Recent studies have shown that such separation problems can be successfully solved by various systems, including a reservoir computer and even a tank of water. It is therefore natural to ask whether the encoding mechanism (invertible generalized synchronization) also allows our more biologically plausible learning system to execute such a function. " Try to reproduce this finding, with your own training data.
3. Try different connection topologies (apart from Erdos-Renyi, the randomly connected network) to see how well the network learns from data and predicts forward in time. See Appendix C.

AND/OR

Vladimir Itskov, Carina Curto, Eva Pastalkova and György Buzsáki 2011

Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus

Journal of Neuroscience

<https://www.jneurosci.org/content/31/8/2828.short>

Hippocampal neurons can display reliable and long-lasting sequences of transient firing patterns, even in the absence of changing external stimuli. We suggest that time-keeping is an important function of these sequences, and propose a network mechanism for their generation. We show that sequences of neuronal assemblies recorded from rat hippocampal CA1 pyramidal cells can reliably predict elapsed time (15–20 s) during wheel running with a precision of 0.5 s. In addition, we demonstrate the generation of multiple reliable, long-lasting sequences in a recurrent network model. These sequences are generated in the presence of noisy, unstructured inputs to the network, mimicking stationary sensory input. Identical initial conditions generate similar sequences, whereas different initial conditions give rise to distinct sequences. The key ingredients responsible for sequence generation in the model are threshold-adaptation and a Mexican-hat-like pattern of connectivity among pyramidal cells. This pattern may arise from recurrent systems such as the hippocampal CA3 region or the entorhinal cortex. We hypothesize that mechanisms that evolved for spatial navigation also support tracking of elapsed time in behaviorally relevant contexts.

1. Give a brief summary of what the authors observed experimentally - temporal encoding by cell assembly sequences.
2. Construct a network of rate neurons that have Mexican-hat connectivity and threshold adaptation as described in the Methods, and show that different initial conditions produce different activity trajectories (Fig. 2).
3. Simulate the same initial conditions multiple times to show that trajectories are reliable (quantify with their measure) across trials despite being driven by noisy, temporally and spatially unstructured inputs (Fig. 3). Explain why this result explains the experimentally observed phenomenon.

5. Spiking statistics of cortical neurons

Roxin, Brunel, Hansel, Mongillo and van Vreeswijk 2011

On the distribution of firing rates in networks of cortical neurons

Journal of Neuroscience 31:16217-26.

<https://doi.org/10.1523/JNEUROSCI.1677-11.2011>

The distribution of *in vivo* average firing rates within local cortical networks has been reported to be highly skewed and long tailed. The distribution of average single-cell inputs, conversely, is expected to be Gaussian by the central limit theorem. This raises the issue of how a skewed distribution of firing rates might result from a symmetric distribution of inputs. This article argues that skewed rate distributions are a signature of the nonlinearity of the *in vivo* $f - I$ curve. During *in vivo* conditions, ongoing synaptic activity produces significant fluctuations in the membrane potential of neurons, resulting in an expansive nonlinearity of the $f - I$ curve for low and moderate inputs.

Investigate the effects of single-cell and network parameters on the shape of the $f - I$ curve and on the distribution of firing rates in randomly connected networks.

1. Analytical approach – deriving the firing rate distribution
Change sigma, Fig. 3A to 3C Explore different mean firing rate, Fig. 4
2. Numerical approach – compute the firing rate distributions by simulating a randomly connected network of integrate-and-fire neurons (Fig. 5)

6. Generation of stable heading representations in diverse visual scenes

Kim, Hermundstad, Romani, Abbott and Jayaraman 2019

Generation of stable heading representations in diverse visual scenes

Nature. 576:126-31 (2019)

<https://doi.org/10.1038/s41586-019-1767-1>

Many animals rely on an internal heading representation when navigating in varied environments. How this representation is linked to the sensory cues that define different surroundings is unclear. In the fly brain, heading is represented by ‘compass’ neurons that innervate a ring-shaped structure known as the ellipsoid body. Each compass neuron receives inputs from ‘ring’ neurons that are selective for particular visual features; this combination provides an ideal substrate for the extraction of directional information from a visual scene. Here we combine two-photon calcium imaging and optogenetics in tethered flying flies with circuit modelling, and show how the correlated activity of compass and visual neurons drives plasticity, which flexibly transforms two-dimensional visual cues into a stable heading representation. We also describe how this plasticity enables the fly to convert a partial heading representation, established from orienting within part of a novel setting, into a complete heading representation. Our results provide mechanistic insight into the memory-related computations that are essential for flexible navigation in varied surroundings.

- Investigate the experimental observations regarding the ‘compass’ neurons of fly brain.
- Using the ring model provided in this article, run the simulation and study
 - how the correlated activity of compass and visual neurons drives plasticity, which flexibly transforms two-dimensional visual cues into a stable heading representation;
 - how network plasticity enables the flexible generation of a stable compass-neuron heading representation in different visual scenes.
- Try some of the untested ideas proposed in the Section “Other aspects of plasticity rules to explore in the future” in the SI texts.

7. Coupled oscillators

Cohen, Holmes and Rand 1982

The nature of the coupling between segmental oscillators of the lamprey spinal generator for locomotion: A mathematical model

Journal Mathematical Biology 13:345–369

<https://doi.org/10.1007/BF00276069>

We present a theoretical model which is used to explain the intersegmental coordination of the neural networks responsible for generating locomotion in the isolated spinal cord of lamprey. A simplified mathematical model of a limit cycle oscillator is presented which consists of only a single dependent variable, the phase $\theta(t)$. By coupling N such oscillators together we are able to generate stable phase locked motions which correspond to traveling waves in the spinal cord, thus simulating “fictive swimming”. We are also able to generate irregular “drifting” motions which are compared to the experimental data obtained from cords with selective surgical lesions.

This study introduced a theoretical model which is used to explain the intersegmental coordination of the neural networks responsible for generating locomotion, i.e., swimming, in the isolated spinal cord of *lamprey*. The electromyographic activity of the myotomal muscle of a fish exhibits a stereotyped temporal pattern. In spinal cord the ventral root (VR) output pattern underlying the muscle activity is believed to have three important features: (1) The activity of the two ventral roots of a single segment strictly alternates in time; (2) the duration of the activity of a VR is a constant proportion of the period of the cycle; (3) there is a delay between the bursts of any two ipsilateral ventral roots and that delay is proportional to the period (Fig 1.2).

Here, the model assumes that each segment of the cord consists of a pair of neural networks which can generate oscillatory activity. Pairs of oscillators (in ventral root) are assumed to be coupled together to form the central pattern generators (CPG) which then generates the complete stable pattern.

Each segmental oscillator of spinal cord is approximated as a limit cycle oscillator which consists of only a single dependent variable, the phase $\theta(t)$. By considering set of N coupled limit cycle oscillators, demonstrate that,

- how stable phase locked motions which correspond to traveling waves in the spinal cord can be generated, thus simulating “fictive swimming”;
- bidirectional coupling between the oscillators can generate a stable traveling wave.

There is a delay between the bursts of any two ipsilateral ventral roots and that delay is proportional to the period. The third feature implies that the delay occupies a constant phase of the cycle, i.e. there is a constant phase coupling between the two segments.

AND/OR

Jeong, Ko and Moon 2002

Time-delayed spatial patterns in a two-dimensional array of coupled oscillators

Physical Review Letters 89 (15)

<https://doi.org/10.1103/PhysRevLett.89.154104>

We investigated the effect of time delays on phase configurations in a set of two-dimensional coupled phase oscillators. Each oscillator is allowed to interact with its neighbors located within a finite radius, which serves as a control parameter in this study. It is found that distance-dependent time delays induce various patterns including traveling rolls, squarelike and rhombuslike patterns, spirals, and targets. We analyzed the stability boundaries of the emerging

patterns and briefly pointed out the possible empirical implications of such time-delayed patterns.

Spatiotemporal patterns arise in numerous physical, chemical, and biological systems. The brain, one of the most complex systems, is now also known to generate spatiotemporal patterns such as plane waves and spirals. This study investigated the effects of time-delayed interactions in an ensemble of two-dimensional coupled phase oscillators. Each oscillator is allowed to interact with its neighbors located within a finite radius r_0 , and the coupling signal propagates at speed v .

- Demonstrate that distance-dependent time delays induce various patterns including traveling rolls, square like and rhombus like patterns, spirals, and targets.
- Explore the parametrical space of coupling length r_0 , coupling coefficients K , and signal propagation speed v . and analyze the stability boundaries between the synchronized planar solutions (when all oscillators are in the same phase) and the emerging patterns.

AND/OR

Ermentrout and Kopell 1984

Frequency plateaus in a chain of weakly coupled oscillators, I

SIAM J. Math. Anal., 15(2), 215–237

<https://doi.org/10.1137/0515019>

A chain of $n + 1$ weakly coupled oscillators with a linear gradient in natural frequencies is shown to exhibit “frequency plateaus,” or sequences of oscillators having the same frequency, with a jump in frequency from one plateau to another. We first show that the equations for the coupled oscillators admit an invariant $(n + 1)$ -torus on which the equations have a special form, one in which an n -dimensional subsystem is approximately invariant. We then show that when the linear gradient becomes too steep to allow phaselocking, there emerges a large-scale invariant circle in this n -dimensional system which corresponds to the existence of a pair of plateaus, and whose homotopy class within the n -torus corresponds to the position of the frequency jump. Also discussed are the effects of anisotropic and nonuniform coupling.

Study the phenomena of how interactions among neurons with random frequencies can lead to clusters of locked oscillators.

8 - Propagating Neuronal Discharges

Golomb and Amitai 1997

Propagating neuronal discharges in neocortical slices: Computational and experimental study

Journal of Neurophysiology 78:1199-211

<https://doi.org/10.1152/jn.1997.78.3.1199>

Propagating neuronal discharges in neocortical slices: computational and experimental study. *J. Neurophysiol.* 78: 1199–1211, 1997. We studied the propagation of paroxysmal discharges in disinhibited neocortical slices by developing and analyzing a model of excitatory regular-spiking neocortical cells with spatially decaying synaptic efficacies and by field potential recording in rat slices. Evoked discharges may propagate both in the model and in the experiment. The model discharge propagates as a traveling pulse with constant velocity and shape. The discharge shape is determined by an interplay between the synaptic driving force and the neuron's intrinsic currents, in particular the slow potassium current. In the model, N -

methyl-d-aspartate (NMDA) conductance contributes much less to the discharge velocity than amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) conductance. Blocking NMDA receptors experimentally with 2-amino-5-phosphonovaleric acid (APV) has no significant effect on the discharge velocity. In both model and experiments, propagation occurs for AMPA synaptic coupling g_{AMPA} above a certain threshold, at which the velocity is finite (non-zero). The discharge velocity grows linearly with the g_{AMPA} for g_{AMPA} much above the threshold. In the experiments, blocking AMPA receptors gradually by increasing concentrations of 6-cyano-7-nitroquinoxaline-2,3-dione (CNQX) in the perfusing solution results in a gradual reduction of the discharge velocity until propagation stops altogether, thus confirming the model prediction. When discharges are terminated in the model by the slow potassium current, a network with the same parameter set may display discharges with several forms, which have different velocities and numbers of spikes; initial conditions select the exhibited pattern. When the discharge is also terminated by strong synaptic depression, there is only one discharge form for a particular parameter set; the velocity grows continuously with increased synaptic conductances. No indication for more than one discharge velocity was observed experimentally. If the AMPA decay rate increases while the maximal excitatory postsynaptic conductance (EPSC) a cell receives is kept fixed, the velocity increases by ~20% until it reaches a saturated value. Therefore the discharge velocity is determined mainly by the cells' integration time of input EPSCs. We conclude, on the basis of both the experiments and the model, that the total amount of excitatory conductance a typical cell receives in a control slice exhibiting paroxysmal discharges is only ~5 times larger than the excitatory conductance needed for raising the potential of a resting cell above its action potential threshold.

Based on the model of neocortical cells present in this study, construct a one-dimensional network with spatially decaying synaptic efficacies, and investigate the dynamic mechanisms that lead to the creation, propagation, and cessation of discharges. Specifically, seek the answers to the following questions:

- 1) What are the spatiotemporal properties of the discharge propagation, and how are they related to the network architecture?
- 2) How do the velocity and shape of the discharge depend on the synaptic parameters, such as the strength of AMPA and NMDA synaptic efficacies, and the level of synaptic depression?
- 3) What is the relationship between the propagation velocity and the synaptic kinetics?

9 - Recurrent Networks

Seung, Lee, Reis and Tank 2000

Stability of the memory of eye position in a recurrent network of conductance-based model neurons

Neuron 26.1: 259-271

[https://doi.org/10.1016/S0896-6273\(00\)81155-1](https://doi.org/10.1016/S0896-6273(00)81155-1)

Studies of the neural correlates of short-term memory in a wide variety of brain areas have found that transient inputs can cause persistent changes in rates of action potential firing, through a mechanism that remains unknown. In a premotor area that is responsible for holding the eyes still during fixation, persistent neural firing encodes the angular position of the eyes in a characteristic manner: below a threshold position the neuron is silent, and above it the firing rate is linearly related to position. Both the threshold and linear slope vary from neuron to neuron. We have reproduced this behavior in a biophysically plausible network model. Persistence depends on precise tuning of the strength of synaptic feedback, and a relatively long synaptic time constant improves the robustness to mistuning.

Develop an oculomotor integrator model as a network of conductance-based neurons which interact with each other by recurring excitatory synapses. The integrator neurons receive feedforward inputs from three neurons. The vestibular neuron which is tonically active at a constant rate, simulating the background activity present in vestibular afferents when the head is stationary. The excitatory and Inhibitory burst neurons that are normally silent, except for occasional brief burst of action potentials that cause saccadic eye movements. These bursts change the firing rates of neurons in the network which is maintained by recurrent excitation after the feedforward input is over. Signals from the integrator neurons lead to the oculomotor plant so that persistent changes in these signals cause persistent changes in the angular position of the eyes.

This model should reproduce the following properties of biological integrator -

- Each integrator neuron in the model should exhibit a linear relationship between firing rate and eye position when it is active. However, there is also a threshold eye position below which it is silent. The linear slope and the threshold vary from neuron to neuron.
- Because of some imperfection in persistence, there is some drift of neural activity with time, which leads to drift in the eye position during fixation. The drift velocity depends systematically on eye position, generally in a nonlinear manner.
- The persistence of neural activity degrades when synaptic strengths are mistuned, neurons are destroyed, or the strength of feedback is otherwise perturbed. Decision making

AND/OR

Wong and Wang 2006

A recurrent network mechanism of time integration in perceptual decisions

Journal of Neuroscience 26:1314-1328

<https://doi.org/10.1523/JNEUROSCI.3733-05.2006>

Recent physiological studies using behaving monkeys revealed that, in a two-alternative forced-choice visual motion discrimination task, reaction time was correlated with ramping of spike activity of lateral intraparietal cortical neurons. The ramping activity appears to reflect temporal accumulation, on a timescale of hundreds of milliseconds, of sensory evidence before a decision is reached. To elucidate the cellular and circuit basis of such integration times, we developed and investigated a simplified two-variable version of a biophysically realistic cortical network model of decision making. In this model, slow time integration can be achieved robustly if excitatory reverberation is primarily mediated by NMDA receptors; our model with only fast AMPA receptors at recurrent synapses produces decision times that are not comparable with experimental observations. Moreover, we found two distinct modes of network behavior, in which decision computation by winner-take-all competition is instantiated with or without attractor states for working memory. Decision process is closely linked to the local dynamics, in the “decision space” of the system, in the vicinity of an unstable saddle steady state that separates the basins of attraction for the two alternative choices. This picture provides a rigorous and quantitative explanation for the dependence of performance and response time on the degree of task difficulty, and the reason for which reaction times are longer in error trials than in correct trials as observed in the monkey experiment. Our reduced two-variable neural model offers a simple yet biophysically plausible framework for studying perceptual decision making in general.

Understand the cortical circuit for decision making (Fig 1), and then re-derive the reduced version decision making model. With stability analysis and numerical simulation, try to investigate the following questions (you may use either eleven-variable or two-variable model for simulation depending on the questions):

- How does the recurrent dynamics give rise to a much longer integration time? Is this slow linear ramping a consequence of a network with slow recurrent excitation?
- Can the model still work when recurrent excitation is solely mediated by the much faster AMPA receptors (AMPA receptors)?
- Is it necessary that neurons subserving integration during stimulation also show persistent activity during working memory?

10 - Dynamic gain control

Borst, Flanagin and Sompolinsky 2005

Adaptation without parameter change: Dynamic gain control in motion detection, , *PNAS* 102: 6172-6176

<https://doi.org/10.1073/pnas.0500491102>

Many sensory systems adapt their input-output relationship to changes in the statistics of the ambient stimulus. Such adaptive behavior has been measured in a motion detection sensitive neuron of the fly visual system, H1. The rapid adaptation of the velocity response gain has been interpreted as evidence of optimal matching of the H1 response to the dynamic range of the stimulus, thereby maximizing its information transmission. Here, we show that correlation-type motion detectors, which are commonly thought to underlie fly motion vision, intrinsically possess adaptive properties. Increasing the amplitude of the velocity fluctuations leads to a decrease of the effective gain and the time constant of the velocity response without any change in the parameters of these detectors. The seemingly complex property of this adaptation turns out to be a straightforward consequence of the multidimensionality of the stimulus and the nonlinear nature of the system.

Develop a motion detection model using Reichardt detectors, which extract the direction of motion by multiplying the brightness signals from neighboring image locations after asymmetric temporal filtering. The model should illustrate the following properties,

- Increasing the amplitude of the velocity fluctuations (variance) suppresses the contribution of the stimulus past, which leads to a marked reduction in the response gain.
- Increasing the stimulus variance shortens the time scale of the motion detection response thereby reducing it to the correlation time of the stimulus fluctuations.

As we know that nervous system is inherently nonlinear and multidimensional, develop a simple neuron model which shows that changing the form of nonlinearity may have significant effects on the magnitude of the resultant adaptive response.

11. Hebbian learning rule

Song, Miller and Abbott 2000

Spike-timing-dependent Hebbian plasticity as temporal difference learning competitive Hebbian learning through spike-timing-dependent synaptic plasticity (STDP)

Nature Neuroscience 3: 919-926

<https://doi.org/10.1038/78829>

Hebbian models of development and learning require both activity-dependent synaptic plasticity and a mechanism that induces competition between different synapses. One form of experimentally observed long-term synaptic plasticity, which we call spike-timing-dependent plasticity (STDP), depends on the relative timing of pre- and postsynaptic action potentials. In modeling studies, we find that this form of synaptic modification can automatically balance synaptic strengths to make postsynaptic firing irregular but more sensitive to presynaptic spike timing. It has been argued that neurons *in vivo* operate in such a balanced regime. Synapses modifiable by STDP compete for control of the timing of postsynaptic action potentials. Inputs that fire the postsynaptic neuron with short latency or that act in correlated groups are able to compete most successfully and develop strong synapses, while synapses of longer-latency or less-effective inputs are weakened.

A spike-timing-dependent Hebbian mechanism governs the plasticity of recurrent excitatory synapses in the neocortex: synapses that are activated a few milliseconds before a postsynaptic spike are potentiated, while those that are activated a few milliseconds after are depressed. Using a biophysical model of a cortical neuron, this paper shows that a temporal difference rule used in conjunction with dendritic backpropagating action potentials reproduces the temporally asymmetric window of Hebbian plasticity observed physiologically. Furthermore, the size and shape of the window vary with the distance of the synapse from the soma. This work shows how a spike-timing-based temporal difference learning rule can allow a network of neocortical neurons to predict an input a few milliseconds before the input's expected arrival.

Reproduce the paper's simulation: model the balanced excitation network with LIF neurons ($n = 1000$). The (excitation) synapses are updated by STDP modification rule.

- 1) Given different mean presynaptic input rates (stochastic spike trains), what is the
 - (a) equilibrium distribution of synaptic strengths arising from STDP
 - (b) the spiking variability, i.e., coefficient of variation (CV) of the postsynaptic spike train.
- 2) Latency reduction: given the presynaptic inputs are correlated in various ways, how does the latency between post- and presynaptic spikes change, before and after the synaptic strengths are learned (and stable) through the STDP rule.

12 - Neuronal spike trains data analysis

***Aljadeff, Lansdell, Fairhall and Kleinfeld 2016
Analysis of Neuronal Spike Trains, Deconstructed
Neuron, 91: 221–259.***

<https://doi.org/10.1016/j.neuron.2016.05.039>

As information flows through the brain, neuronal firing progresses from encoding the world as sensed by the animal to driving the motor output of subsequent behavior. One of the more tractable goals of quantitative neuroscience is to develop predictive models that relate the sensory or motor streams with neuronal firing. Here we review and contrast analytical tools used to accomplish this task. We focus on classes of models in which the external variable is

compared with one or more feature vectors to extract a low-dimensional representation, the history of spiking and other variables are potentially incorporated, and these factors are nonlinearly transformed to predict the occurrences of spikes. We illustrate these techniques in application to datasets of different degrees of complexity. In particular, we address the fitting of models in the presence of strong correlations in the external variable, as occurs in natural sensory stimuli and in movement. Spectral correlation between predicted and measured spike trains is introduced to contrast the relative success of different methods.

Objective goals: Based on the statistics of applied stimuli (uncorrelated or correlated), select the proper analysis approaches to analyze the spike trains and extracting the relevant stimuli features.

- Spike-triggered average (STA)
- Spike-Triggered Covariance (STC)
- Maximum Noise Entropy Method (MNE)
- Generalized Linear Models (GLM)

The group can use the dataset acquired from their research project or search the available dataset online. Dataset and accompanying code presented in the paper is also available (<https://github.com/NeuroInfoPrimer/primer>).

13 - Analysis of Neuronal Morphology

Wen, Stepanyants, Elston, Grosberg and Chklovskii 2009

Maximization of the connectivity repertoire as a statistical principle governing the shapes of dendritic arbors

PNAS, 106: 12536-12541

<https://doi.org/10.1073/pnas.0901530106>

The shapes of dendritic arbors are fascinating and important, yet the principles underlying these complex and diverse structures remain unclear. Here, we analyzed basal dendritic arbors of 2,171 pyramidal neurons sampled from mammalian brains and discovered 3 statistical properties: the dendritic arbor size scales with the total dendritic length, the spatial correlation of dendritic branches within an arbor has a universal functional form, and small parts of an arbor are self-similar. We proposed that these properties result from maximizing the repertoire of possible connectivity patterns between dendrites and surrounding axons while keeping the cost of dendrites low. We solved this optimization problem by drawing an analogy with maximization of the entropy for a given energy in statistical physics. The solution is consistent with the above observations and predicts scaling relations that can be tested experimentally. In addition, our theory explains why dendritic branches of pyramidal cells are distributed more sparsely than those of Purkinje cells. Our results represent a step toward a unifying view of the relationship between neuronal morphology and function.

Wen et al. analyzed basal dendritic arbors of pyramidal neurons sampled from mammalian brains and discovered three statistical properties:

1. the dendritic arbor size scales with the total dendritic length,
2. the spatial correlation of dendritic branches within an arbor has a universal functional form,
3. and small parts of an arbor are self-similar.

Wen et al. proposed that these properties result from *maximizing the repertoire of possible connectivity patterns between dendrites and surrounding axons while keeping the cost of dendrites low*. They solved this optimization problem by drawing an analogy with *maximization of the entropy* for a given energy in statistical physics. The solution is consistent with the above observations and predicts scaling relations that can be tested experimentally.

- Scaling, universality, and self-similarity (Eq. 1) of basal dendritic arbors of pyramidal neurons.
- Derived an analytical expression for connectivity repertoire (Eq. 2), and dendritic cost.
- Maximizing the connectivity repertoire for a given dendritic cost and compared the results with experimental measurements.

14 - Bring up a topic that interests you or is related to your research project

You are welcome to find a topic that you are interested in or relevant to your on-going research.