Abstracts of papers presented
at the 2004 meeting on

COMPUTATIONAL & SYSTEMS NEUROSCIENCE

March 24–March 28, 2004

Arranged by

Carlos Brody, *Cold Spring Harbor Laboratory*
Alexandre Pouget, *University of Rochester*
Michael Shadlen, *University of Washington*
Anthony Zador, *Cold Spring Harbor Laboratory*

Cold Spring Harbor Laboratory
Cold Spring Harbor, New York

Major funding for this meeting is provided by the **Swartz Foundation** and the **Redwood Neuroscience Institute**. Funding also provided in part by the **National Science Foundation;** and the **National Institute of Mental Health,** a branch of the **National Institutes of Health;** with additional support provided by **Advanced Acoustic Concepts, Inc.**

Contributions from the following companies provide core support for the Cold Spring Harbor meetings program.

Corporate Benefactors

Amgen Inc.
Aventis Pharma AG
Bristol-Myers Squibb Company
Eli Lilly and Company
GlaxoSmithKline
Pfizer Inc.

Corporate Sponsors

Applied Biosystems
BioVentures, Inc.
Cogene BioTech Ventures, Ltd.
Diagnostic Products Corporation
Genentech, Inc.
Hoffmann-La Roche Inc.
Kyowa Hakko Kogyo Co., Ltd
Lexicon Genetics Inc.
Merck Research Laboratories
New England BioLabs, Inc.
OSI Pharmaceuticals, Inc.
Pall Corporation
Wyeth Genetics Institute

Plant Corporate Associates

Monsanto Company

Corporate Affiliates

Affymetrix, Inc.

Corporate Contributors

Axxora, LLC
Biogen
Epicentre Technologies
Illumina
Integrated DNA Technologies
IRx Therapeutics, Inc.
KeyGene

Foundations

Albert B. Sabin Vaccine Institute, Inc.

Cover: Illustrated by Marjorie Domenowske.

COMPUTATIONAL & SYSTEMS NEUROSCIENCE

Wednesday, March 24 – Sunday, March 28, 2004

Wednesday	7:30 pm	Opening Lecture
Thursday	9:00 am	Session 1
Thursday	2:00 pm	Poster Session I
Thursday	4:30 pm	Wine and Cheese Party *
Thursday	7:30 pm	Session 3
Friday	9:00 am	Session 4
Friday	2:00 pm	Poster Session II
Friday	7:30 pm	Session 6
Saturday	9:00 am	Session 7
Saturday	2:00 pm	Session 8
Saturday	6:00 pm	Banquet
Sunday	9:00 am	Session 9

Poster sessions are located in *Bush Lecture Hall*

* *Airslie Lawn*, weather permitting

Mealtimes at Blackford Hall are as follows:

Breakfast	7:30 am-9:00 am
Lunch	11:30 am-1:30 pm
Dinner	5:30 pm-7:00 pm

Bar is open from 5:00 pm until late

PROGRAM

WEDNESDAY, March 24—7:30 PM

OPENING LECTURE

Abbott, L.F., Volen Center for Complex Systems and Dept. of Biology, Brandeis University, Waltham, Massachusetts: Cascades and multi-timescale plasticity. 1

THURSDAY, March 25—9:00 AM

SESSION 1

Meister, M., Harvard University, Cambridge, Massachusetts: Adaptation in the neural code of the retina. 2

Brainard, D.H.,[1] Longère, P.,[2] Kraft, J.M.,[3] Delahunt, P.B.,[4] Freeman, W.T.,[5] Xiao, B.,[1] [1]University of Pennsylvania, Philadelphia; [2]24 Chemin de la Cavalerie, Grasse, France; [3]Institute of Science and Technology, University of Manchester, United Kingdom; [4]University of California, Davis; [5]Massachusetts Institute of Technology, Cambridge: Computational models of human color constancy. 3

Deneve, S., Gatsby Computational Neuroscience Unit, University College London, United Kingdom: Bayesian inference with recurrent spiking networks. 4

Newsome, W.T., Sugrue, L.P., Corrado, G.S., Howard Hughes Medical Institute and Dept. of Neurobiology, Stanford University School of Medicine, California: Decision-making and the neural representation of "experienced value". 5

Niv, Y.,[1] Duff, M.,[2] Dayan, P.,[2] [1]Interdisciplinary Center for Neural Computation, Hebrew University, Jerusalem, Israel; [2]Gatsby Computational Neuroscience Unit, University College London, United Kingdom: The effects of uncertainty TD learning. 6

THURSDAY, March 25—2:00 PM

SESSION 2 POSTER SESSION I

THURSDAY, March 25—4:30 PM

Wine and Cheese Party

THURSDAY, March 25—7:30 PM

SESSION 3

FRIDAY, March 26—9:00 AM

SESSION 4

FRIDAY, March 26—2:00 PM

SESSION 5 POSTER SESSION II

FRIDAY, March 26—7:30 PM

SESSION 6

SATURDAY, March 27—9:00 AM

SESSION 7

SATURDAY, March 27—2:00 PM

SESSION 8

SATURDAY, March 27

BANQUET

Cocktails 6:00 PM Dinner 6:45 PM

SUNDAY, March 28—9:00 AM

SESSION 9

Bonhoeffer, T., Max-Planck Institute of Neurobiology, München, Germany: Competitive interactions between potentiated synapses.

Fiete, I.R.,[1,2] Seung, H.S.,[2,3] [1]Dept. of Physics, Harvard University, [2]Dept. of Brain and Cognitive Science,[3]Howard Hughes Medical Institute, Massachusetts Institute of Technology, Cambridge: Rapid learning by reinforcement of variation due to injected noise in a spiking neural network. 234

Johnston, D., Baylor College of Medicine, Houston, Texas: Dendritic K+ channels in hippocampal pyramidal neurons—Biophysics, diversity and function.

AUTHOR INDEX

CASCADES AND MULTI-TIMESCALE PLASTICITY

L.F. Abbott
Volen Center for Complex Systems and Department of Biology
Brandeis University, Waltham, MA 02454-9110

Neurons, synapses, and neural circuits exhibit temporal dynamics on time scales ranging from milliseconds to years. Furthermore, these dynamics appear to be scale-invariant rather than being characterized by fixed time constants. I will discuss ways of modeling and thinking about such systems, and explore the functional implications of scale-invariant dynamics for adaptation and memory.

ADAPTATION IN THE NEURAL CODE OF THE RETINA

Markus Meister

The optic nerve is a bottleneck for visual information. Correspondingly, the retina can transmit to the brain only a tiny fraction of the raw visual information that reaches the photoreceptors, and must discard the rest. Many aspects of retinal processing can be understood as clever strategies for extracting those visual features that are unexpected and not already known to the brain. Recent work shows that the retina can adapt these functions dynamically to the spatial and temporal statistics of the visual environment, leading to a general strategy for novelty detection in the visual scene.

COMPUTATIONAL MODELS OF HUMAN COLOR CONSTANCY

David H.Brainard[1], Philippe Longère[2], James M. Kraft[3], Peter B. Delahunt[4], William T. Freeman[5], and Bei Xiao[1] 1) University of Pennsylvania 2) 24 Chemin de la Cavalerie, Grasse, France 3) University of Manchester Institute of Science and Technology, 4) UC Davis 5) Massachusetts Institute of Technology

The challenge of perception arises because the retinal image provides ambiguous information about the physical properties of a viewed scene. In the case of object color, the ambiguity arises because the object's surface reflectance is confounded with the illuminant's spectral power distribution. Many modern accounts articulate the principle that human vision resolves the ambiguity by matching its processing to the statistical structure of natural images. There are few examples, however, where this principle has been used to develop a quantitative account of psychophysical measurements. This paper provides such an account. First, we develop a Bayesian algorithm that estimates illuminant and surface properties from the retinal image. The algorithm resolves ambiguity in the image through explicit priors that represent the statistical structure of spectra in natural images. These priors provide a parameterization that allows us to control the algorithm's performance. Second, we measure color appearance by asking observers to adjust the appearance of test objects, embedded in complex three-dimensional scenes, until they appear achromatic. We then develop a model for the psychophysical data by assuming that the algorithm's illuminant estimates provide the information used by the visual system to recover object color. The model accounts for the data well. The parameters of the model that yield a good fit may be interpreted in terms of how well matched the priors used by the visual system are to the statistics of natural scenes.

BAYESIAN INFERENCE WITH RECURRENT SPIKING NETWORKS

Sophie Deneve
Gatsby Computational Neuroscience Unit
University College London, London, UK WC1N 3AR

Most perceptual and motor tasks performed by the central nervous system are essentially probabilistic. They infer a few important but "hidden" properties, such as object shapes, time to collision, or motor commands, from many noisy, local and ambiguous sensory cues. These cues must be integrated via partial probabilistic knowledge of the environment: Prior probabilities, i.e. the sensory and motor events most likely to occur, and conditional probabilities arising from functional or causal relationships between these events. This raises the issue of how uncertainty and probability distributions are represented and computed in biological neural networks.

We show that leaky integrate and fire neurons with adaption can be interpreted as Bayesian integrators accumulating evidence over time about events in the external world or the body, and communicating to other neurons their certainties about these events. In this framework, the membrane potential is interpreted as a log-likelihood ratio for a hidden binary state variable, such as the presence or absence of the preferred stimulus in the cell receptive field. The spikes emitted by the neuron gives a sparse representation for its log-likelihood ratio. Population of neurons, each of which accumulate and report the likelihood ratio that their preferred values are the best interpretation of the sensory input, can represent uncertainties about analog variables such as orientation.

Networks of these neurons perform exact and approximate Bayesian inference, using standard models of synaptic transmission. Thus, we propose that the connectivity in cortical circuits implement graphical models, or Bayesian networks, representing the statistical structures of the tasks they perform. Propagation of spikes in these neural networks implements belief propagation in the corresponding Bayesian networks. Balanced excitation and inhibition is essential to perform Bayesian inference in these recurrent networks just as it is for their stability.

This reinterpretation of the integrate and fire neurons has many implications and makes many predictions for single cell computations, the statistics of neural firing, cortical microcircuits, connectivity between brain areas and modularity. In particular, we will consider dendrite computations, the origin and meaning of irregular inter-spike intervals and neural synchrony, neural integration and decision making, invariance of neural tuning, and modularity as a tractability/accuracy tradeoff.

Decision-making and the neural representation of ‘experienced value’

William T. Newsome, Leo P. Sugrue and Gregory S. Corrado

Howard Hughes Medical Institute and Department of Neurobiology
Stanford University School of Medicine

In the study of decision-making, psychophysicists and sensory physiologists traditionally emphasize the effects of sensory stimuli on the outcome of the decision process. Psychologists and economists, however, have long known that decision-making is influenced not only by the sensory stimulus, but also by an organism’s prior experience or beliefs concerning the “value” of the alternative choices, expressed in terms of likely positive or aversive consequences. Brain circuitry that mediates decision-making must presumably reflect both influences, and we have recently been able to demonstrate both effects at the behavioral and neurophysiological levels. To measure “experienced value” objectively, we have developed a probability matching paradigm for rhesus monkeys in which the animal’s valuation is revealed through the proportion of choices allocated to alternative behaviors. In neurophysiological recordings, we have found that neurons in the lateral intraparietal area (LIP) code experienced value, although LIP is unlikely to be the site where experienced value is originally computed. Our results to date suggest that signals from multiple sources within the brain, including representations of sensory stimuli as well as internal valuation converge in areas such as LIP to render simple behavioral decisions.

THE EFFECTS OF UNCERTAINTY ON TD LEARNING

Yael Niv[1], Michael Duff[2] and Peter Dayan[2]
[1]Interdisciplinary Center for Neural Computation, Hebrew University
[2]Gatsby Computational Neuroscience Unit, University College London

Substantial evidence suggests that the phasic activities of dopamine (DA) neurons in the primate midbrain represent a temporal difference (TD) error in the predictions of future reward. TD offers a computationally compelling account of a role for DA in appetitive classical and instrumental conditioning, and a precise and parsimonious computational theory of the generation of DA firing patterns. Recent DA recordings (Fiorillo *et al Science,* 2003) in an experiment involving inherently stochastic reward delivery, show activity *ramping* towards the time of the reward, with ramp height related to the degree of reward uncertainty. *Prima facie*, these data present a crucial challenge, as both classical and instrumental facets of the TD would require the reliable activity in the ramp to be *predicted away* by earlier stimuli.

We use analysis and simulations to show that the apparently anomalous ramps are in fact to be expected under a standard TD account, if, as suggested by the low baseline firing rates of the DA cells, positive and negative prediction errors are differentially scaled. Using a simple tapped-delay-line representation of time between the stimulus and the reward (as commonly adopted in TD models), together with a fixed learning rate, a ramp in the DA activity emerges just as in the experimental data.

Analytically deriving the average response at the time of the reward from the TD learning rule, we show that the height of the modelled ramps are indeed proportional to the variance of the rewards, in accordance with the data. There is, however, a key difference between the uncertainty and TD accounts of the ramps. According to the former, ramps are within-trial phenomena, coding uncertainty; by contrast, the latter suggests they arise only through averaging across multiple trials. Under the TD account, the non-stationarity engendered by constant learning from errors makes the PSTH traces potentially misleading, as they average over different trial histories.

Our study suggests that uncertainty need not be playing any explicit part in determining these aspects of DA activity, a conclusion also consistent with various other suggestive results from Fiorillo and others. We also study the effects on TD learning of other sources of representational and learning noise.

PRECISELY REPEATING MOTIFS OF SYNAPTIC ACTIVITY IN NEOCORTEX

Gloster Aaron[1], Yuji Ikegaya[1], Ilan Lampl[2,3], David Ferster[2] and Rafael Yuste[1]
(1) Dept. Biological Sciences, Columbia University, New York, NY 10027,USA, (2) Dept. Neurobiology and Physiology, Northwestern University, Evanston IL 60208, USA, and (3) Dept. of Neurobiology, Weizmann Institute of Science, Rehovot 76100,Israel

The essence of cortical function is the propagation and transformation of neuronal activity by the cortical circuit. At the same time, it is not understood how activity can propagate at all through a network composed of small, unreliable and depressing synapses. As a solution, it has been proposed that cortical activity takes the form of sequences of synchronous activity ("synfire chains"), which propagate through the cortical circuit with high temporal fidelity. Consistent with this, repetitions of temporally precise firing sequences of spikes have been reported, although the existence of these repetitions is controversial. If network activity is precisely repeating, then intracellular recordings from a single neuron should report these repeating patterns in terms of EPSPs and IPSPs. We record intracellularly from layer 5 pyramidal cells in slices of mouse visual cortex (n = 20) and from intact cat visual cortex (n = 21) and, using a supercomputer, find repeated temporal patterns of spontaneous synaptic inputs, consistent with the predictions of the synfire chain hypothesis. These repeated patterns, which we call "motifs", can last for several seconds, are non-oscillatory, and can repeat after many minutes and yet display millisecond precision in the repetition. Repeated motifs occur at higher frequency than chance, can repeat more than once, and can occur simultaneously in pairs of neurons. Concurrent imaging of spontaneous network activity in slices shows repeated activation of the same sequences of cells during the occurrence of repeated synaptic motifs in nearby cells. Repeated motifs and precise sequences of network activation occur during synchronization of neuronal firing but also in their absence. Our data demonstrate that neocortical circuits *in vitro* and *in vivo* are repeatedly activated in precise temporal sequences. Precise sequences most likely arise from exact reactivations of the same neuronal ensembles, indicating a high reliability in the synaptic dynamics of the circuit. These ensembles may represent basic multicellular units of cortical function.

BACKPROPAGATING ACTION POTENTIALS: ROLE IN SYNAPTIC PLASTICITY

Corey D. Acker[1,2], Kamal Sen[1,2,3], John A. White[1,2,4]
1. Department of Biomedical Engineering
2. Center for BioDynamics
3. Hearing Research Center
4. Center for Memory and Brain
Boston University, Boston MA 02215, USA

Advances in experimental techniques have led to an explosion in knowledge of the biophysical properties and behavior of dendrites. One of the most salient features of apical dendrites in hippocampal pyramidal cells is the backpropagation of axonally initiated action potentials into the dendritic tree. Experiments have shown that these backpropagating action potentials (BAPs) play an important role in several aspects of neuronal function. Here, we use computational modeling of pyramidal cell dendrites to first study the biophysical properties that control the efficacy of backpropagation. In both morphologically detailed and simplified models, biophysical details including passive membrane properties, shape, and distribution of voltage-gated ion channels are manipulated to study the effects on backpropagation. Spike time dependent synaptic plasticity (STDP) is believed to depend on backpropagation; the BAP is thought to report back to synapses the timing of output action potentials. In models including NMDA receptor mediated Ca^{2+} dynamics and a standard model of Ca^{2+}-controlled changes in synaptic weight, we see that the predicted STDP learning window is sensitive to BAP waveform shape. In these models, we find it difficult to reproduce some important features of experimentally measured STDP learning windows. Inclusion of a more detailed description of Ca^{2+} dynamics and voltage-gated Ca^{2+} channels may mitigate these discrepancies.

Support Contributed By: NIH R01 NS34425

NEURAL CORRELATES OF VIBRISSA RESONANCE: BAND-PASS AND SOMATOTOPIC REPRESENTATION OF HIGH-FREQUENCY STIMULI

Mark L. Andermann, Jason Ritt, Maria A. Neimark and Christopher I. Moore
Program in Biophysics, Harvard University and
McGovern Institute For Brain Research, MIT,
Cambridge, MA 02139

The array of vibrissae on a rat's face is the first stage in a high-resolution tactile sensing system. Recently, we discovered that vibrissae (whiskers) resonate, generating several-fold increases in motion amplitude when stimulated at specific frequencies. We have investigated the neural correlates of vibrissa resonance in peripheral and central somatosensory neurons by presenting low-amplitude, high-frequency vibrissa stimulation. We found that single neurons in the trigeminal ganglion and primary somatosensory cortex (SI) showed band-pass tuning and enhanced sensitivity to threshold stimuli near the vibrissa fundamental resonance frequency (FRF). Further, a somatotopic map of high-frequency selectivity was observed in SI, with isofrequency columns extending along the barrel-column representations of arcs of vibrissae, in agreement with the gradient in vibrissa FRFs across the vibrissa pad. To extend these findings to more natural stimuli, we presented gratings, sandpaper and punctate taps to the vibrissae. We found that trigeminal ganglion neurons can encode the changes in gain and the temporal signature associated with vibrissa resonance amplification of these stimuli. These findings suggest several parallels between high-frequency processing in the vibrissa sensory system and in the auditory system, and have important implications for the detection and discrimination of high-frequency information.

MAPPING SPATIO-TEMPORAL DYNAMICS

Jörn Anemüller, Terrence J. Sejnowski and Scott Makeig
Swartz Center for Comp. Neuroscience, UCSD, La Jolla, California 92093,
Comp. Neurobiology Lab., The Salk Institute, La Jolla, California 92037

The phenomenon of spatio-temporal dynamics observed in recordings of neuronal systems activity has been noticed for long, using different imaging techniques. Appropriate modeling of the dynamics, in particular for unaveraged single-trial data, should take into account that the recorded activity may be due to several ongoing processes. Independent component analysis (ICA) can, under certain assumptions, separate activity from several mutually independent processes. However, in its original formulation ICA has very limited capabilities of capturing the aforementioned dynamics. We recently incorporated modeling of spatio-temporal dynamics into the ICA model, resulting in a complex-valued ICA algorithm that obtains frequency-domain representations of the separated signals (Anemüller, Sejnwoski, Makeig, *Neural Networks* 16:1311-1323, November 2003).

To visualize the spatio-temporal dynamics corresponding to the obtained frequency-domain components, these have to be projected back into the time-domain, which involves three steps. First, components in different frequency-bands are identified, that originate from the same physiological process. Subsequently, these components are projected to the measurement sensors, resulting for each sensor and frequency-band in the part of the measurement that can be attributed to this process. Finally, the result is projected to the time-domain, giving the picture of the spatio-temporal dynamics that is associated with the process in question.

We have analyzed data obtained from 32-channel EEG recordings during a visual selective attention task. Complex independent components in the spectral range from 4 to 8 Hz were found to be highly specific for a subject response (right hand button press) related process. When projected back to the time-domain, this process was characterized by a central focus of activation with highest signal amplitude near 50 ms after response time. Subsequently, the focus of activation shifted in a posterior and left-hemispherical direction within about 20 ms, a pattern that could reflect the subject's motor response. Using a conventional averaging approach without ICA processing, the shift in the posterior direction was less pronounced, and the shift in the contra-lateral direction was not visible.

Preliminary data from an EEG experiment involving saccades have also resulted in spatio-temporally dynamic components that are related to stimulus presentation and neuronal processes. We are also investigating the application of the method to multi-channel recordings obtained by optical imaging.

STIMULUS COMPETITION AND ATTENTIONAL MODULATION IN V2/V4 COULD BE COMPUTED BY NONLINEAR DENDRITIC SUBUNITS

Kevin A. Archie and Bartlett W. Mel
Neuroscience Graduate Program and Dept. of Biomedical Engineering
University of Southern California, Los Angeles, CA 90089-2520

Neurons in V2 and V4 exhibit competition between simultaneously presented stimuli: given two stimuli A and B, the response to the combined stimulus $A + B$ lies between the responses to A and B presented individually. Attention directed to one of the stimuli drives the response to the combined stimulus towards that for the attended stimulus alone (Luck et al., *J. Neurophysiol.* **77**:24, 1997; Reynolds et al., *J. Neurosci.* **19**:1736, 1998). We propose an extremely compact neural circuit in which the competition and attentional effects result from interactions between synaptic excitation and inhibition in the dendrites of a single postsynaptic neuron. Using a newly developed compartmental model of a layer 5 pyramidal cell, and a software front end which allows us to translate realistic visual stimuli into trains of synaptic events, we show that our model provides a good match to single unit recordings in V2 and V4. Using a simplified algebraic model for the pyramidal neuron, developed in tandem with the biophysical model, we gained a number of insights into the key biophysical properties of the pyramidal neuron needed to account for the competition and attentional data.

INTERACTIONS BETWEEN EXCITATION AND INHIBITION IN BIOPHYSICALLY DETAILED AND SIMPLIFIED MODELS OF NEOCORTICAL PYRAMIDAL CELLS

Kevin A. Archie and Bartlett W. Mel
Neuroscience Graduate Program and Dept. of Biomedical Engineering
University of Southern California, Los Angeles, CA 90089-2520

Finding compact functional descriptions of individual neurons is a core problem of neuroscience. Recent theoretical studies showed that the firing rate of a detailed hippocampal pyramidal cell (model) could be described by a simple mathematical formula that mapped individual thin dendrites onto the nonlinear subunits of a two-layer neural network (Poirazi, Brannon, and Mel, *Neuron*, **37**:977–87, **37**:989–99, 2003). However, this study manipulated only the spatial patterning of synaptic excitation across the dendritic arbor, neglecting spatial interactions between excitation and inhibition that could in principle contribute to many forms of cortical computation. To address this shortcoming, and to facilitate further studies of dendritic computation in neorcortical sensory processing contexts, we developed a new compartmental model of a layer 5 pyramidal cell containing 16 types of ion channels, whose distribution and parameters were constrained in part by recent data on the integrative properties of these cells' basal dendrites (Schiller et al., *Nature* **404**:285, 2000; A. Polsky and J. Schiller, personal communication). In controlled studies in which we stimulated two sets of thin basal dendrites with random high frequency inputs, we found that the two-layer sum-of-sigmoids model predicted the firing rate of the biophysical cell model under the simplest assumption that synaptic excitation and inhibition combine linearly within each branch prior to the sigmoidal modulation stage. We conclude that the combination of biophysically detailed and simple algebraic models provides a useful toolkit for investigating the computing functions of individual neurons.

CIRCUIT ATTRACTORS IN THE NEOCORTEX: DYNAMICS AND MECHANISMS OF NETWORK UP STATES

Dmitriy Aronov, Rosa Cossart, and Rafael Yuste

The cerebral cortex has traditionally been thought of as a feedforward system that processes input through successive stages to reach an appropriate output. Yet, cortical circuits are known to be strongly interconnected and are active even in the absence of sensory input. Such spontaneous firing can be coordinated among specific groups of neurons in a way that probably reflects the underlying functional architecture of the cortex. Moreover, the membrane potential of cortical neurons fluctuates spontaneously between a resting DOWN state and a depolarized UP state. The elevated firing rate in the UP state provides a substrate for persistent activity, which may be crucial for mediating intrinsic cortical functions, such as working memory.

To investigate the dynamics of persistent cortical activity, we used two-photon calcium imaging to reconstruct the spontaneous activity of up to 1,500 neurons in slices of mouse primary visual cortex. We combined this large-scale imaging with targeted recordings of individual neurons to investigate the intracellular correlates of network events. We find that, in the absence of any stimulation, groups of cortical neurons simultaneously enter membrane potential UP states. These 'network UP states' involve small, but highly significant numbers of cells (2-3%). Such coactive ensembles of neurons are often organized into patterns whose anatomical structure may reflect the underlying cortical architecture. Successive synchronizations frequently involve similar ensembles of neurons, and single neurons can be active in multiple coactive assemblies. UP states typically last for several seconds, mediating persistent activity on timescales much longer than those of the feedforward information flow through the cortex. The stereotyped spatiotemporal dynamics of this activity are highly reminiscent of those theoretically predicted for feedback and recurrent neural networks. Thus, we propose that network UP states are circuit attractors, defined as emergent stable states in network dynamics that can be used to implement solutions to computational problems and working memory. We reinforce this idea by demonstrating that network UP states can be initiated by electric stimulation, which serves as an initial state for network dynamics.

Furthermore, we investigated the synaptic and intrinsic mechanisms of network UP states. We find that some UP states persist in the absence of glutamatergic transmission. However, they are abolished in the absence of GABAergic transmission, and pressure application of GABA to the slices triggers simultaneous sustained depolarizations in groups of cortical neurons. Thus, the GABAergic network may be responsible for generating the network synchronizations we observe. We propose that the coordinated activity of specific interneuron subclasses provides a circuit mechanism for triggering the stereotyped occurrence of network UP states.

TOPOGRAPHIC PLASTICITY IN BARN OWLS

Gurinder Atwal
Department of Physics,
Princeton University, Princeton, NJ 08544

The internal mapping of the external environment is carried out using the receptive fields of topographic neurons in the cortex, and in a normal barn owl the aural and visual cortical maps are aligned from early experiences. The visual system is important in instructing the auditory system how to translate audible cues into representations of space in the optic tectum within the midbrain. Experiments on juvenile barn owls involving instantaneous misalignment of the aural and visual stimuli have revealed adaptive behavior, manifested by functional and anatomical changes of the auditory localization system.

Using methods of information theory and statistical mechanics a model of the adaptive dynamics of the aural receptive field is presented and analyzed. The dynamics are determined by maximizing the mutual information between the neural output and the weighted sensory neural inputs, admixed with noise, subject to biophysical constraints. Lagrange parameters are introduced quantifying the costs associated with the gain, spatial extent and dynamics of the receptive fields.

Reduced costs of neural rewiring, presumably as in the case of juvenile barn owls, reveal two qualitatively different types of receptive field adaptation depending on the magnitude of audio-visual misalignment. On the other hand, large neural costs, as in the case of adult barn owls, prohibit learning.

However, by letting the misalignment increase sufficiently slowly with time, it is shown that the ability to adapt can be significantly increased even when neural rewiring costs are high, in agreement with recent experimental reports of increased plasticity of the auditory space map in adult barn owls due to incremental learning. Finally, a critical speed of misalignment is identified, in terms of the cost Lagrange parameters, demarcating the crossover from adaptive to non-adaptive behavior.

PRINCIPAL AND INDEPENDENT COMPONENTS OF MACAQUE VOCALIZATIONS

Bruno B. Averbeck[1] and Lizabeth M. Romanski[2]
[1]Center for Visual Science, Department of Brain and Cognitive Sciences, and [2]Department of Neurobiology & Anatomy, and Center for Navigation and Communication Sciences, University of Rochester, Rochester NY 14627

Neurons in high-level sensory cortical areas respond to complex elements in sensory stimuli. Feature elimination can be used to explore the components of the stimuli that are driving neural responses. In this approach, a complex stimulus that evokes a neuronal response is simplified, and if the cell responds to the reduced stimulus it is considered selective for the remaining features. We have developed a feature elimination technique which utilizes either the principal or the independent components of a stimulus to define a subset of features, to which a neuron might be sensitive. The original stimulus can be filtered using these components, resulting in a stimulus which retains only a fraction of the features present in the original. We have applied this technique to macaque vocalizations, and we have used both spectral and bispectral techniques to explore the original and filtered stimuli, as well as the features extracted by each technique. The results of our analyses show that the features extracted by principal components analysis are closely related to the dominant Fourier components of the stimuli, often called formants in the study of speech perception. In contrast, the features extracted by independent components analysis preserve the relative phase across a set of harmonically related frequencies. We have also carried out preliminary experiments, in which we presented both filtered and unfiltered stimuli to fixating macaques, while recording the activity of single neurons in the ventral lateral prefrontal cortex. Our data indicate that the principal and independent component filtered stimuli evoke robust responses in many prefrontal auditory neurons. This novel approach provides a useful method for exploring the essential features within complex stimuli that activate higher order sensory neurons.

CIRCUIT MECHANISMS OF DIFFERENTIAL MOTION SELECTIVITY IN THE RETINA

Stephen A. Baccus, Bence P. Ölveczky, Markus Meister
Harvard University, Dept. of Molecular & Cellular Biology,
Cambridge, MA.

Certain types of retinal ganglion cells respond very sensitively to differential motion between the receptive field center and surround, as produced by an object moving over the background, but are strongly suppressed by global motion, as produced by fixational eye movements (Olveczky, Baccus & Meister, *Nature*, 2003). We investigated the circuit basis for this object motion selective (OMS) response by recording intracellularly from bipolar, amacrine or OMS ganglion cells while simultaneously recording the spiking output of ~50 retinal ganglion cells. Under a jittering global motion stimulus, polyaxonal amacrine cells depolarized at precisely the right times to inhibit the depolarization and firing of OMS ganglion cells. Amacrine and OMS ganglion cell membrane potential responses conformed to a receptive field model that pools over many small rectified subunits, as described previously for Y-type ganglion cells. Bipolar cells had small receptive fields and responses consistent with the postulated subunits. Finally, current injection into polyaxonal amacrine cells showed that they selectively inhibit the OMS ganglion cells. These results give a circuit explanation of differential motion selectivity: Bipolar cell output is rectified and then summed by polyaxonal, transient Y-type amacrine cells, which then inhibit Y-type ganglion cells.

MULTINEURON RESPONSE DYNAMICS IN CAT VISUAL CORTEX DURING THE PRESENTATION OF TIME-VARYING NATURAL SCENES.

Jonathan L. Baker, Shih Cheng Yen and Charles M. Gray. Center for Computational Biology, Montana State University, Bozeman MT 59717.

The most common paradigm for studying sensory processing in the visual system has been to examine the activity of single neurons in response to highly artificial stimuli such as bars and gratings. While this approach has been highly successful, very little has been learned about how groups of neurons jointly respond to natural scenes. To address this issue, we have recorded spike activity from small groups of 3-10 well isolated, single units in the striate cortex of anesthetized cats and analyzed their responses to the repeated presentation (n=100 trials) of short clips (30 sec) of time-varying natural scenes (movies). Receptive fields of simultaneously recorded neurons were mapped using sparse noise stimuli according to previously published methods (Ohzawa et al., 1996). We were interested in answering the following questions: 1) What are the properties of cortical neuronal responses to time-varying natural images? Can the responses be characterized as sparse or dense? What is the distribution of response durations and firing rates across cells? During significate response episodes, how reliable do individual neurons respond in both spike count and spike time? 2) How often do the responses of simultaneously recorded neurons overlap in time? During periods of response overlap, what are the properties of the joint activity? Is there some form of response covariation in firing rate? 3) Are the individual/joint responses of cortical neurons consistent with the space-time receptive fields mapped using classical techniques?

The results revealed a number of interesting properties not predictable from the responses of single neurons to artificial stimuli. 1) Neuronal responses in cat striate cortex to natural time-varying images are sparse ($\mu = 72.8$) (Rolls & Tovee, 1995) and brief. The probability distributions of firing rate for all cells were highly peaked at 0 with heavy tales. Significant responses occur approximately 10% of the time, and these events are usually less than 200 ms in duration. Spike timing within events can be very precise (SD of spike times ranged from 2 – 400 msec, mode = 9 msec). Periods of high spike count reliability often co-occurred with high spike time reliability. 2) Nearby cells exhibit a high degree of heterogeneity in their responses to natural scenes. Joint responses occur with a wide distribution of probabilities and have an average value of ~0.5. In many instances, cells recorded from the same probe showed a variety of overlapping and non-overlapping responses. Periods of joint activity displayed a broad range of spike-count correlations across trials (ranging from .2 to .8) that could not be accounted for by the stimulus alone. Thus, even cells recorded within the same column often respond in markedly different ways to the same stimuli. It remains to be determined the extent to which this heterogeneity can be explained by differences in the classical spatiotemporal receptive fields. 3) In spite of the heterogeneity, when pairs of cells do respond at the same time, roughly half of these episodes are characterized by large values of spike count correlation that exceed those published previously. This finding suggests that responses to natural scenes involve dynamic patterns of neuronal interaction that are selective to the properties of the image. Thus, periods of joint activity could reflect independent responses to the stimulus as well as interactions within the cortical network that vary dynamically with the stimulus.

Ohzawa, I., DeAngelis, G.C., and Freeman, R.D. (1996) *J. Neurophysiology*., 75:1779-1805.
Rolls E.T. and Tovee (1995) *J. Neurophysiol.,* 73:713-726.

Supported by the National Eye Institute, and a Kopriva Fellowship to JLB

PERSISTENT MODULATION BY TEMPORAL CONTEXT IN AUDITORY CORTEX OF AWAKE PRIMATES

Edward L. Bartlett and Xiaoqin Wang, Dept. of Biomed. Eng., Johns Hopkins University, Baltimore, MD, USA

An important issue in sensory neuroscience is the dynamic representation of a stimulus. This study addresses how auditory cortical responses to a given sound are affected by preceding sounds. Changes in an acoustic stream may be due to the presence of more than one sound source or an acoustic transition within a single source. These transitions may signify behaviorally relevant portions of the stream, such as consonant-vowel combinations in speech or phrases of animal vocalizations. Thus, a sound cannot be unambiguously grouped or identified without reference to its temporal context.

Previous studies have mainly investigated temporal context in relation to forward masking, using short-duration pure tone stimuli and recording mainly from anesthetized animals (Calford and Semple 1995, Brosch and Schreiner 1997). However, psychophysical studies using long-duration stimuli found that sound context affected the formation of perceptual streams and the perception of sound changes (Vliegen et al. 1999, Holt and Lotto 2002). This study used complex stimuli that had long-durations, various modulation properties (e.g., amplitude and frequency modulations), and both tone and noise carriers. We tested the notion that context dependent auditory processing is a general phenomenon that is related to the properties of the preceding stimulus beyond the frequency domain. This context dependent processing may allow for tracking small fluctuations in a single stream while enhancing the contrast between large changes in an acoustic stream that could indicate the introduction of a new auditory object.

To investigate how the temporal acoustic context affected neural responses, we recorded from single units in the auditory cortex of awake primates (marmosets). Using combinations of modulated tone and modulated noise stimuli, we found that the preceding stimulus could suppress or facilitate responses to the succeeding stimulus for durations greater than 1 second. These persistent effects were dependent on the duration, intensity and modulation property of the preceding stimulus, in addition to the carrier frequency and type. They occurred whether or not the two stimuli were separated by a silent interval. The strength of suppressive effects was often tuned such that preceding stimuli whose parameters were similar to succeeding stimuli produced the strongest suppression, but suppressive effects often occurred regardless of the succeeding stimulus used (nonspecific suppression). Facilitatory effects usually occurred for particular two-sound combinations (specific facilitation). The modulatory effects could be observed in the absence or near-absence of discharges evoked by the preceding stimulus, suggesting subthreshold interactions between the neural responses to the two sounds. Our results demonstrate that auditory cortex neurons have lasting sensory memories of temporal context and suggest that facilitatory and suppressive effects operate via different mechanisms.

A SELF-ORGANISING SENSORIMOTOR SPIKING NETWORK.

Anthony J. Bell
Redwood Neuroscience Institute
1010 El Camino Real, Suite 380, Menlo Park, CA 94025

Neurons spike and are embedded in a loop with the world, apparently learning without supervision. We propose an unsupervised learning principle called *self-modeling* for such loopy spiking systems. It is argued that optimal self-modeling, when coupled to the world, also performs world-modeling, and thus incorporates the learning of sensorimotor contingencies.

The objective of the (discrete time) recurrent spiking network is to take the binary neural state vector, $\mathbf{y}$, over a certain time window, and build from it a probability density (p.d.f.) estimate, $\hat{p}(\mathbf{y})$, of *itself*, using only what the 1's and 0's say about other 1's and 0's. This estimate should be as close as possible to the true density, $p(\mathbf{y})$, which is a delta-function at the actual value of $\mathbf{y}$. Minimising the KL distance, $D[p||\hat{p}]$, corresponds to maximising $\log \hat{p}(\mathbf{y})$, the log likelihood of the network's global model of itself.

In this model, each spike (or silence) carries a small piece of information in an estimate, $\hat{p}(\mathbf{y}_j)$, of the *marginal p.d.f.* over the vector of variables that caused the jth spike (the 'local' model). The simplest method for maximising the likelihood of the global model is to (1) formulate the meaning of a spike, ie: $\hat{p}(\mathbf{y}_j)$ (2) combine those meanings to form a spike-based 'population code' $\hat{p}(\mathbf{y})$ (3) alter the local marginals, $\hat{p}(\mathbf{y}_j)$ (and thus also $\hat{p}(\mathbf{y})$), so that they are maximally consistent with each other (a maximum-entropy problem) (4) learn by altering synaptic weights, $\mathbf{W}$, and leakages, $\mathbf{w}$, using Contrastive Divergence approximations to the gradients of $\log \hat{p}(\mathbf{y})$. The gradient estimates required here are shown to be of the 'Boltzmann Machine' type: respectively, contrasts in spike counts (3) and relative spike timings (4) under the data, $\mathbf{y}$, and the model distribution, $\hat{p}(\mathbf{y})$.

When the network is in 'thinking' mode (no sensory or sensorimotor interactions), we conjecture that the optimal self-model is achieved along the lines of phase boundaries, as in 'critical states' in simpler media such as the Ising spin model. This 'Santa Fe' idea draws on the intuition that because the network's spatio-temporal mutual information peaks at critical phase boundaries, so does the self-representation of the system as an 'autopoetic' unit. The coupling of the thinking system to the world through sensory and sensorimotor contingencies creates a 'disturbance' in the self-model which must be *compensated for* to keep $\hat{p}(\mathbf{y})$ optimal. It is yet to be demonstrated to what degree such unsupervised sensorimotor compensations can account for the action-taking properties of natural neural networks. But here, the learning rules are derived for the linear discrete-time integrate-and-fire model.

ENCODING COMMUNICATION SIGNALS WITH TIMESCALE SEPARATION BY SPIKE FREQUENCY ADAPTATION

Jan Benda, André Longtin and Len Maler
Department of Physics and Department of Cellular and Molecular Medicine.
University of Ottawa, Ontario, Canada K1N 6N5

Spike-frequency adaptation is commonly observed in almost every spiking neuron and is involved in various computational tasks. Many of the observed phenomena are network effects mediated by inhibitory feedback, feedforward inhibition, or depressing synapses. Here we present an example where spike-frequency adaptation induced by slow intrinsic currents of the spike generator plays an important role in separating fast communication signals (chirps) from slow background beat patterns in the electrosensory system of the weakly electric fish *Apteronotus leptorhynchus*. These fish actively generate a weak quasi-sinusoidal electric field (EOD: electric organ discharge) whose amplitude modulation (AM) is encoded in the fast adapting electroreceptors (Nelson *et. al.* (1997): J Comp Phys A 181, 532).

The stimuli are sinusoidal amplitude modulations of arbitrary frequency (the beat resulting from the superposition of two fishes EOD) which are disrupted by fast up- or downstrokes generated by the chirps. The firing frequency response of electro-receptor afferents recorded *in vivo* follows the beat. However, the response to the chirps is enhanced by about 40% relative to the beat response, as long as the beat frequency is lower than about 30 Hz.

This observation can be completely explained by the high-pass filter properties induced by spike-frequency adaptation. Applying a universal model for spike-frequency adaptation (Benda & Herz (2003): Neural Comput. 15, 2523) to measured data of f-I curves and adaptation time constants we find a cutoff frequency for the high-pass filter at 23 Hz. Thus, beats with frequencies below 23 Hz are transmitted with a much lower gain than the fast chirp, resulting in a relatively enhanced response to the chirp. Beats with higher frequencies, however, are transmitted with a gain similar to the one for the chirps. Therefore, the firing frequency response to the chirps on fast beats as large as the one to the beat.

In a behavioral study J. Bastian observed that the probability of the fish emitting a chirp decays down to zero in about the same range of beat frequencies as we find in our analysis of electro-receptor responses (Bastian *et al.* (2001): J Exp Biol 204, 1909). We thus conclude that the main bottleneck of the systems performance might be the electro-receptor afferents, since they set the timescale separating slow beats on which chirps are detectable from fast beats.

TEMPORAL COGNITION AND WORKING MEMORY

Tom J Bennett, Boris Gutkin and Peter Dayan
Gatsby Computational Neuroscience Unit, University College London,
Alexandra House, 17 Queen Square, London, WC1N 3AR, UK

Psychological studies of the way in which animals time intervals show a key scalar regularity. Namely, the standard deviation of the estimates of the interval length is proportional to the (typically nearly unbiased) mean of the estimates over multiple trials. This implies that the discriminative stimuli for animal behaviour in a timing task are subject to multiplicative noise. Despite a rich body of psychological models of timing, there is a dearth of physiologically-based accounts which explain this regularity.

We propose a theory of interval timing based on the experimentally observed dynamical behaviour of cortical cells during Delayed Match to Sample working memory tasks. These neurons display a diverse array of repeatable temporal activity patterns (collectively termed persistent activity) in the delay period that follows the presentation of a stimulus in each trial. We treat these patterns as forming a temporal basis function representation of the time elapsed since the stimulus was shown. Recent electro-physiological data from parietal cortex, by Shafi, Bodner, Zhou and Fuster (2003) suggest that the standard deviation in the activity of individual cells at a given point in time across trials scales linearly the with mean activity of the cell at that time.

Our model of sustained activity uses the inherent unreliability of synapses within a recurrent network of spiking cells to generate multiplicative noise internally. This mechanism, initially proposed by Shapiro, Wearden and Barone (2003), at all times renders the process of spike generation dependent on the post-stimulus history of synaptic releases. Persistent activity is generated without the need for bistable neurons by the percolation of spikes around the network. Activity levels are controlled by making the average probability that a neuron emits a spike, given that a spike has arrived at a presynaptic terminal, inversely proportional to the neuron's recurrent connectivity. Storage is initiated by a signal which resets the activity to a well defined and reproducible level. It is terminated either by an equivalent signal or as a consequence of the decay of the memory trace.

Time estimates are made via threshold discrimination of the mean activity of the neural population. Random synaptic failures between the cells give rise to the proportional relationship between the standard deviation of the distribution of time estimates and the distribution mean over any set of trials. This makes our proposal consistent with relevant psychological and physiological results.

HOW TO 'RESET' A NEURAL NETWORK?

Matthias Bethge
Redwood Neuroscience Institute
1010 El Camino Real, Suite 380, Menlo Park, CA 94025

The ongoing activity in the cortex as well as the large amount of feedback connections put strong emphasis on the role of the internal network state in neuronal processing. This suggests that perceptual inference or pattern recognition is performed by the cortex in a recurrent or 'iterative' fashion. In particular, this view has been outlined in terms of Kalman filtering ('predictive coding') as well as in more general terms of Bayesian belief networks.

While recurrent networks appear to be best suited for inference problems dealing with temporal sequences, they are considered to be disadvantages in case of rapid classification tasks. In fact, network memory, which may easily become infinite in case of recurrent neural networks, can appear either as a curse or a blessing: if the task is to classify temporal sequences that are ambiguous at any instance of time but unique over a larger history, then network memory is clearly necessary in order to perform the task successfully. On the other hand, network memory is desirable not without fail, because it imposes an upper bound on the processing speed.

While the limitation in speed has been used as a heuristic argument against recurrent neuronal computations, I present simulation studies, which indicate that the accumulation of memory in neural networks constitutes a more fundamental problem that is not limited to recurrent networks, but applies to multilayer feed-forward networks with spiking neurons and dynamic synapses as well. This finding suggests that particular neuronal 'reset' mechanisms are necessary in order to make the current network state sufficiently independent from the preceding history either steadily or only on demand. For both cases, the potential use of shunting inhibition as a 'reset signal' is emphasized and discussed.

Reducing spike train variance: A computational theory of spike-timing dependent plasticity

Sander M. Bohte and Michael C. Mozer
Department of Computer Science
University of Colorado, Boulder, USA

A phenomenon currently of great interest in neuroscience is spike-timing dependent plasticity (STDP): paired neuron recordings reveal that the precise timing of pre- and post-synaptic action potentials determines the sign and magnitude of the resulting plasticity [3]. These findings have resulted in a flurry of models that focus on various aspects of the phenomenon, from biological models that give rise to the STDP mechanism (i.e. [5]), to models that explore the consequences of using STDP-like rules in large ensembles of spiking neurons [6], to models that propose underlying computational justifications for STDP [7,8,4]. In our work, we present an alternative computational motivation for STDP: If a neuron attempts to reduce the variance in the timing and number of spikes in its output train, given the current input spike train, we show that an STDP learning rule emerges.

To derive the learning rule, we compute the derivative of the probability of post-synaptic events with respect to the input weights via a stochastic version of the Spike Response Model [1,2]. We apply this learning rule to a simulation mimicking the *in vivo* STDP experiment in [9]. The change in efficacy of the weight of the relevant input-spike as a function of the offset of the input-spike relative to post-synaptic spiking shows the typical STDP curve.

The model presents a computational-level account that exploits the dynamics of a realistic neural model. Furthermore, it enables us to make predictions about how the STDP learning kernel varies as a function of the biological parameters that can be measured or derived from measurements; in contrast, in the model proposed by Dayan et. al. [4], predictions are a function of the signal and noise distributions, which are harder to ascertain. Our model explains the major phenomena surrounding STDP from standard learning theory and accepted neuron models, whereas other computational-level accounts have not been as comprehensive.

[1] W. Gerstner (2001) A Framework for Spiking Neuron Models: The Spike Response Model In: The Handbook of Biological Physics, Vol.4 (Ch. 12), pp 469-516 [2] X. Xie and H.S. Seung (2003) Learning in neural networks by reinforcement of irregular spiking, submitted. [3] Bi & Poo, 2001 Annu. Rev. Neurosci. 24:139-166. [4] Dayan, Hausser & London, NIPS2003. [5] Porr, Saudargiene & Woergoetter, NIPS2003. [6] Song, Miller & Abbott, Nature Neurosci. 3:919-926. [7] Rao & Sejnowski, 2001, Neural Comp. 13:2221-2237. [8] Chechik, Neur. Comp. 2003 15(7):1481–1510. [9] Zhang et.al., Nature 395:37–44.

IMPLICATIONS OF NONLINEAR CENTER-SURROUND INTERACTIONS IN THE RAT CORTEX FOR TEXTURE DISCRIMINATION

Alireza S. Boloori and Garrett B. Stanley
Division of Engineering and Applied Sciences
Harvard University, Cambridge, MA 02139

Natural stimuli for the rat vibrissa system engage multiple whiskers and thus elicit center-surround interactions which we have studied here. Previous work indicates that rat cortical cells are primarily responsive to velocity transients, and we therefore model whisker stimuli as square waves with rise and fall times representing inputs. In this context, the simplest form of response integration is the suppression of a deflection response that is caused by a preceding adjacent whisker movement. This is quantified via the Condition-Test Ratio (CTR) as the normalization of the second whisker's response in presence of the first deflection to its response in the absence of the first. These pair-wise effects are important in that the response to an arbitrary stimulus results from their nonlinear combinations. We quantified various pair-wise response interactions (see Fig. 1 top), that when combined, give rise to the observed response attenuation of a deflection preceded by two stimuli (Fig. 1 middle and bottom: solid lines show CTRs). We showed that the three-deflection interaction can be predicted from the previously measured pair-wise CTRs (Fig. 1 bottom, dashed lines). The prediction method was extended to describe cortical responses to periodic and random stimuli. Lastly, we investigated the role of center-surround interactions in the coding of tactile stimuli via simulations.

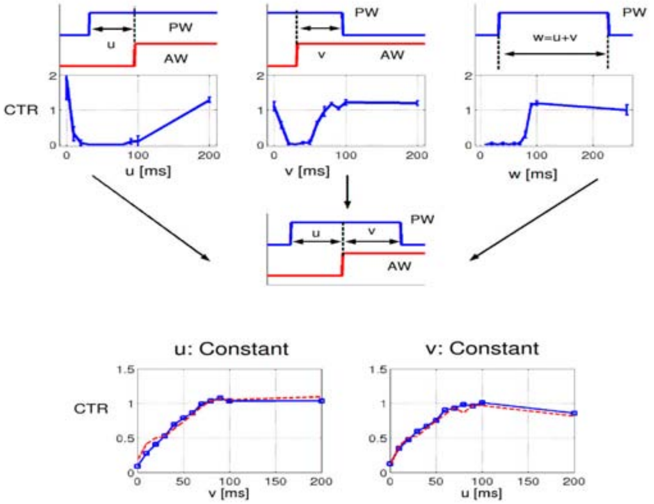

Figure 1: Data are averages over 15 cells. PW: primary whisker, AW: adjacent whisker.

INTEGRATION OF STIMULUS CONTRAST IN THE EARLY VISUAL SYSTEM

Vincent Bonin, Valerio Mante, Séverine Durand and Matteo Carandini
Smith-Kettlewell Eye Research Institute, San Francisco, CA, USA

The responses of neurons in lateral geniculate nucleus (LGN) and in primary visual cortex (V1) depend on stimulus contrast and size. While LGN responses grow almost linearly with contrast, V1 responses grow steeply with contrast at low contrasts, and saturate at high contrasts. Moreover, both LGN and V1 responses grow with stimulus size until an optimal size is reached, and then decrease for larger stimuli. This dependence on size is absent at low contrasts.

We sought to explain these and other observations with a single divisive model, in which the numerator is computed by a *receptive field* (RF) and the denominator by a large overlapping *suppressive field* (SF). In the model, the membrane potential is

$$V = V_{\max} \frac{O_R}{c_{50} + O_S},$$

where V_{max} is *gain*, c_{50} is *semisaturation contrast*, and O_R and O_S are the outputs of RF and SF. A rectification stage transforms V into a firing rate:

$$R = \left[V - V_{thresh} + N(\sigma)\right]_+,$$

where V_{thresh} is a threshold, $N(\sigma)$ is Gaussian noise, and $[\,]_+$ indicates rectification.

To test the model we recorded firing rate responses in LGN and V1 of anesthetized, paralyzed cats. Stimuli were drifting gratings of contrast c in circular windows of radius r. Contrast and radius were varied to cover a matrix of values.

Assuming SF and RF are Gaussians, their responses to these stimuli are

$$O_R = c\left(1 - \exp\left(-r^2 / 2w_R^2\right)\right),$$

$$O_S = c\left(1 - \exp\left(-r^2 / 2w_S^2\right)\right)$$

where w_R, w_S denote the field sizes.

The model provides excellent fits to the data. It explains > 90 % of the variance in the mean data in a majority of cells (33/54 in LGN, 54/100 in V1; >70% 44/54 in LGN, 92/100 in V1).

Because the model describes the responses of both LGN and V1 neurons, the differences between these neurons can be studied from parameter values.

V1 neurons have low c_{50} (median=16%), leading to strong response saturation at high contrast. They have positive V_{thresh} (median=6.0 spikes/s), which combined with the noise $N(\sigma)$ results in a power-law growth of the response with contrast at low contrast

LGN neurons have higher c_{50} (median=54%), leading to weak saturation at high contrast. They have negative V_{thresh} (median= -3.7 spikes/s), i.e. nonzero spontaneous firing rates, resulting in linear increase of the response with contrast at low contrasts. As predicted, LGN responses show frank saturation only for large stimuli, for which O_S approaches c_{50}.

The model also captures the dependence of responses on size. For large, high contrast stimuli O_S approaches c_{50} and effectively suppresses the response. This explains why size tuning is strong at high contrast and absent at low contrast.

Thanks to the noise $N(\sigma)$ the model also predicts trial-to-trial variability. Differences in V_{thresh} explain why response variance increases with mean in V1 while being largely constant in LGN. Differences in V_{thresh} and in response amplitude (median maximum response = 51 spikes/s in LGN vs. 30 spikes/s in V1) together explain why responses are more variable in V1 than in LGN, given noise of similar amplitudes (median σ = 5.1 spikes/s in LGN vs. 6.9 spikes/s in V1).

We conclude that a single model can explain dependence of V1 and LGN responses on stimulus contrast and size. In addition, the model explains response variance and its dependence on the mean.

Support: James S McDonnell Foundation

AN INTERACTIVE RACE MODEL OF COUNTERMANDING SACCADES

L. Boucher, G.D. Logan, T.J. Palmeri, & J.D. Schall

In a saccade countermanding task, a subject is instructed to make a saccade to an eccentric target that appears coincident with a disappearance of a central fixation point. Infrequently, the fixation point reappears (stop signal) indicating that the saccade to the target must be inhibited. The probability that a subject will inhibit the saccade increases as a function of stop signal delay (SSD). The latencies of the saccades that escape inhibition also increase with SSD but with a common minimum value. This behavior can be explained as the outcome of a race between GO and STOP processes with stochastically independent finish times (Logan & Cowan, 1984; Hanes & Schall, 1995).

The control of eye movements is achieved through a mutually inhibitory and excitatory network of gaze-shifting and gaze-holding neurons. To determine if single cells within the frontal eye fields generated signals sufficient enough to control saccadic initiation, Hanes et al. (1998) recorded from movement and fixation neurons while monkeys performed in the countermanding paradigm. They found that both classes of neurons modulated their activity early enough to control saccade initiation. Thus, it seems plausible to assert that the GO and STOP processes may be instantiated by the gaze-shifting and gaze-holding neurons, respectively. Given the above, we have a paradox: how can interacting neural units produce behavior that appears to be the outcome of a race between independent processes?

The race model of Logan and Cowan (1984) was extended by simulating various architectures of the GO and STOP processes. Here we present a class of models which vary the onset of and the amount to which the GO and STOP processes interact with one another. By requiring replication of both the behavioral and the neural data, the network architecture of the simulation is sufficiently constrained to reveal a particular, and more likely correct, account of how the control of saccadic initiation is accomplished. The best fit model is one in which the GO and STOP processes proceed independently throughout most of their duration until some time after the stop signal when the STOP process strongly inhibits the GO process. Thus, it is possible to observe behavior which is the result of an underlying competitive architecture that appears to be the result of independent processes.

GRANGER CAUSALITY ANALYSIS OF SENSORIMOTOR CORTEX IN MOTOR MAINTENANCE BEHAVIOR

Steven L. Bressler[1], Mingzhou Ding[1], Anders Ledberg[1], Andrea Brovelli[1], Yonghong Chen[1], Richard Nakamura[2]; [1]Center for Complex Systems & Brain Sciences, Florida Atlantic University, Boca Raton, FL 33431, USA; [2]Laboratory of Neuropsychology, National Institute of Mental Health, Bethesda, MD 20892, USA

The role of postcentral cortical areas in motor maintenance behavior is poorly understood. To address this issue, we investigated the functional relations of local field potentials (LFPs) from neuronal assemblies in pre- and postcentral areas of two macaque monkeys as they pressed a hand lever during the wait period of a visual discrimination task. It is well established that oscillatory activity in the beta frequency range is synchronized at the various stages of efferent outflow from the motor cortex to the musculature during this type of behavior. Using power and coherence spectral analysis of LFPs from the wait period, we first determined that a beta-synchronized large-scale network also links postcentral areas with motor cortex. We then employed Granger causality spectra as a measure of causal influences among recording sites. Granger causal influence is defined in terms of stochastic processes: one stochastic process is causal to a second if the autoregressive predictability of the second process at a given time point is improved by including measurements from the immediate past of the first. In both monkeys strong causal influences were observed from primary somatosensory cortex to both motor cortex and inferior posterior parietal cortex (area 7), with the latter area also exerting causal influences on motor cortex. Causal influences from motor cortex to postcentral sites, however, were weak in one monkey and not observed in the other. These results are the first to demonstrate in awake behaving monkeys that directional influences are carried by beta-synchronized oscillations from primary somatosensory cortex and area 7 to motor cortex during motor maintenance behavior. The findings are consistent with the idea of a cortico-peripheral-cortical sensorimotor loop in which the primary somatosensory cortex modulates steady motor cortical output, with additional higher-order error control by area 7. Supported by NSF grant IBN0090717, NIMH grants MH64204 and MH42900, and ONR grant N00014-99-1.

IDENTIFICATION OF DECISION MAKING NEURONS IN THE LEECH CNS.

KL Briggman (1), H Abarbanel (2), WB Kristan Jr*(1). Neurobiology(1), Physics(2), UC San Diego, La Jolla CA 92093.

We are interested in the neuronal basis of decision-making. We are using the central nervous system of the leech because it appears to be non-deterministic: it can choose to swim or crawl in response to the same stimulus. The decision to swim or to crawl appears to be made sequentially beginning with an initial decision to locomote, followed by a decision to swim or crawl. We use a FRET-based voltage sensitive dye to record from approximately 150 neurons in a midbody segmental ganglion. We monitored the activity after a stimulus that evokes swimming or crawling. We are interested in neurons that are activated differently leading up to either a swim or a crawl because these neurons are candidate decision-making neurons. We use Principal Component Analysis (PCA) to reduce the dimensionality of these large data sets. PCA highlights which neurons have the largest differences in activity before a swim or crawl. We then impale these candidate neurons and attempt to bias the behavior towards swimming or crawling with current injection. A neuron’s ability to bias the decision is our criteria for labeling a cell a decision-maker. We will extend this analysis to the subesophageal ganglion (“brain”) because several potential decision-making neurons have their somata in this structure.

THE PERCEPTION OF HEADING DEPENDS ON MULTIPLE CORTICAL AREAS

Kenneth H. Britten, Tao Zhang and James B. Maciokas
Center for Neuroscience, UC Davis

Subjects can accurately perceive their direction of self-motion based on solely visual cues, even in the presence of smooth pursuit eye movements, which greatly distort the pattern of retinal motion vectors ("optic flow") upon which the perception depends. While many areas in extrastriate visual cortex and in parietal cortex respond selectively to a variety of optic flow stimuli, most work on heading has focused on the medial superior temporal area (MST). Indeed, previous work from our lab has demonstrated that MST is directly involved in heading perception, since microstimulation of MST led to systematic biases in the perception of heading on a two-alternative discrimination (Britten and van Wezel, 2002). Related experiments (Zhang and Britten, 2003) indicate that the ventral intraparietal area (VIP) contributes to heading perception as well. Microstimulation in VIP produced very similar biases as it did in MST, suggesting a similar functional role.

While microstimulation can test the role of an area in perception or behavior, it is less revealing about the nature of that role than single-cell physiology. We have therefore quantitatively measured signals relating to heading in both areas VIP and MST. We simulated linear trajectories through a three-dimensional cloud of points, which creates an expanding flow pattern on the screen. We varied the angle of heading horizontally, and measured the tuning of single neurons, whose visual receptive fields were centered near the middle of the range of headings employed. Neurons in both areas were well tuned for this dimension, showing either open-ended or band-pass tuning. However, there were quantitative differences in the tuning – MST neurons tended to have lower evoked rates and lower slopes to their tuning. Therefore, it appears that MST contains less information, overall, than does area VIP.

We were also interested in how the representations of heading in MST and VIP were affected by the presence of horizontal smooth pursuit eye movements. Perception is quite tolerant of such eye movements, and this appears to largely depend on non-visual information about eye movements, such as efference copy or proprioception. We therefore also measured neuronal tuning for horizontal heading in the presence of 10 deg/sec horizontal pursuit. In VIP, tuning was remarkably stable in the presence of pursuit, with most tuning functions shifting less than 10 degrees, and the population model shifting less than 2 degrees. Amplitudes and bandwidths also changed relatively little on average. While data collection in MST is ongoing, it is clear that MST cells are less pursuit-tolerant than are VIP cells. Consistent with work from other labs, which implied different stimuli, MST cells frequently exhibited large gain changes under pursuit, and also frequently shifted their tuning functions dramatically. Therefore, it appears that the representation of heading derived from optic flow in VIP is more isomorphic with perception than is the representation in MST.

References:
Britten, K.H. and van Wezel, R.J.A. (2002) "Area MST and heading perception in macaque monkeys", *Cerebral Cortex,* 12:692-701.
Zhang, T. and Britten, K.H. (2003) "Microstimulation of area VIP biases heading perception in monkeys" Program No. 339.9. *2003 Abstract Viewer/Itinerary Planner.* Washington, DC: Society for Neuroscience.

Acknowledgements:
Supported by the National Eye Institute and the Human Frontier Science Program.

AGGREGATION PHEROMONE PROCESSING IN THE LOCUST ANTENNAL LOBE

Bede Broome, Gilles Laurent
Division of Biology, California Institute of Technology, Pasadena, CA, 91125

Pheromones are behaviorally relevant odor compounds used for communication and survival by many insects. Much is known about the processing of general odors in the locust antennal lobe; however, almost all of the detailed information we possess about the sensing of pheromones has been obtained in other model systems. We examined how aggregation pheromones are processed in the locust, which lacks anatomically separate pheromone responsive glomeruli. Using intracellular single unit and extracellular multi-single-unit tetrode recordings, we explored the manner in which pheromones are processed in the locust antennal lobe. We describe the various ways in which processing of these odors is similar or dissimilar to that of general odors. The response characteristics of the antennal lobe appear to be comparable for pheromones and general odors. Projection neurons are broadly tuned with respect to aggregation pheromones and respond to those pheromones similarly to general odors. In addition, pheromones stimulate synchronous and periodic firing among projection neurons in a manner similar to that observed with general odors.

FROM SPIKES TO SPEED-ACCURACY VIA THE LOCUS COERULEUS

Eric Brown[1], Jeff Moehlis[1] , Philip Holmes[1,2], Ed Clayton[3], Janusz Rajkowski[3]
and Gary Aston-Jones[3]
[1] Prog. Appl., Comp. Mathematics, [2] Dept. Mech., Aero. Engineering
Princeton University, Princeton, NJ 08544, U.S.A.
[3] Dept. Psychiatry and Laboratory of Neuromodulation and Behavior,
University of Pennsylvania, Philadelphia, PA 19104, U.S.A.

Recent brain recordings suggest a link between different firing patterns in the brainstem nucleus locus coeruleus (LC) and different levels of performance in simple cognitive tasks. Starting from the dynamics of single spiking neurons, I will describe mathematical models for these firing patterns and possible mechanisms for the observed transitions between them. Then, in an extension of previous work, I will discuss a possible role for the LC in optimizing speed and accuracy in decision tasks, via release of neuromodulators which dynamically adjust the sensitivity (i.e. gain) of neural populations.

Neurons in the alert monkey LC exhibit two distinct modes: *phasic* and *tonic*. In the latter, associated with labile behavior and poor performance on tasks requiring focused attention, cells fire at relatively high rates with little synchrony; in the former, associated with good performance, firing rates are lower but display greater synchrony. The phasic mode also produces greater response to stimuli. Recent data indicates that LC responses differ not only between LC modes, as just discussed, but also among different psychological tasks. We present empirical results demonstrating this difference for the target identification vs. Eriksen flanker tasks.

From membrane voltage and ion channel equations, we derive a phase oscillator model for LC neurons. Average spiking probabilities of a pool of cells over many trials are then computed via a probability density formulation. These show that: 1) Post-stimulus response is elevated in populations with lower spike rates; 2) Responses decay exponentially due to noise and variable pre-stimulus spike rates; and 3) Shorter stimuli preferentially cause depressed post-activation spiking. These results allow us to propose mechanisms for the different LC responses observed across behavioral and task conditions, and to make explicit the role of baseline firing rates and the duration of task-related inputs in determining LC response.

Finally, we identify the distinct time-dependent gain values which implement optimal processing strategies for several different models of neural populations in decision tasks. These provide predictions for the gain manipulations that diverse neural mechanisms, including possibly the LC firing patterns just described, may be expected to implement.

MODELS OF COORDINATE TRANSFORMATIONS FOR REACHING TOWARD AUDITORY TARGETS

Marina Brozović[1], Larry Abbott[2] and Richard Andersen[1]
[1]Division of Biology,
Caltech, Pasadena, California 91125
[2]Volen Center for Complex systems,
Brandeis University, Waltham, Massachusetts 02454

We are interested in how the parietal cortex processes auditory signals and converts this information into a motoric action such as a reach. The experimental study by Y. Cohen and R. Andersen (2000) showed that although the sound is originally coded in the head reference frame, 42% of the parietal reach region (PRR) neurons code the position of the auditory targets in the eye-centered reference frame and an additional 13% code in the intermediate eye-head reference frame. The changes in the receptive fields found in the data can be correlated to the underlying mathematics of coordinate transformations.

Our goal is to build a biologically plausible model, which reproduces the complex structure of the receptive fields observed in data. We would like to test if the coordinate transformations are better modeled using a direct multiplication in the hidden layer, or using a simple summation followed by a non-linear transfer function. This study will provide a deeper understanding of the cellular mechanisms that lead to the formation of gain fields.

The first network is based on the Zipser-Andersen model where the weighted sum of the inputs is passed through the sigmoid transfer function. The method used to train the network is standard backpropagation. The second model is a modified version of Zipser-Andersen network. Here, a complex transfer function is proposed to simulate multiplicative behavior. The input in the hidden layer includes the weighted sum of the pre-synaptic firing rates as well as the variance of the input signal. The resulting transfer function scales with the noise. The backpropagation algorithm is also adjusted for this network architecture.

We used a quantitative method based on correlation analysis to estimate the similarity between the receptive fields in data and model. The neurons in the first model develop dominantly planar, simple looking receptive fields that do not conform to more complicated structures observed in the data. The second model, which simulates a multiplicative behavior of the neurons, shows some encouraging preliminary results. The neurons in the hidden layer developed receptive fields similar to the ones in the data.

SILENT SYNAPSES ARE REQUIRED FOR OPTIMAL INFORMATION STORAGE AT EXCITATORY CONNECTIONS

Nicolas Brunel[1], Vicent Hakim[2], Jean-Pierre Nadal[2], Philippe Isope[3] and Boris Barbour[4]. [1] NPSM, CNRS/Université de Paris 5, Paris, France; [2] Statistical Physics, [4] Neurobiology, Ecole Normale Supérieure/CNRS, Paris, France; [3] Kinsmen Lab, University of British Columbia, Vancouver, Canada.

It is widely believed that synaptic modifications underlie learning and memory. Here, we address the question of what can be deduced about the learning process from experimental measurements of synaptic weight distributions. We consider the case of excitatory synapses, taking the specific example of the cerebellar granule cell--Purkinje cell synapse, whose analogy with the prototypical `perceptron' learning machine is well established. We calculate the weights distribution required for optimal information storage in the perceptron and find it contains a majority of silent synapses, showing that these are necessary to maximise information storage. The optimal distribution fits closely the experimentally-determined distribution for granule cell-Purkinje cell synapses and provides information about the reliability of storage. For the best-fit parameters obtained, we calculate that a Purkinje cell may store up to 0.27 bits of information per synapse, equivalent to 5 Kilobytes per cell. Our optimal distribution also accounts well for weights distributions at hippocampal and neocortical synapses and leads to the hypothesis that many of those synapses have also been silenced in order to optimise information storage.

A BURST-TIME DEPENDENT LEARNING RULE AT THE RETINOGENICULATE SYNAPSE AND IMPLICATIONS FOR ACTIVITY-DEPENDENT DEVELOPMENT.

DA Butts, PO Kanold, CJ Shatz *Department of Neurobiology, Harvard Medical School*

Synaptic refinement in the mammalian lateral geniculate nucleus (LGN) requires neural activity. Learning rules that specify how combinations of pre- and postsynaptic activity lead to changes in synaptic efficacy are thought to underlie this activity-dependent development. However, it is not known how synaptic learning rules function *in vivo* in response to natural activity patterns. The activity-dependent refinement of retinal inputs to the lateral geniculate nucleus (LGN) during development provides a unique opportunity to investigate this because both the input activity from the retina and its importance in development are well studied. In fact, multi-electrode and calcium imaging have given a clear picture of the spatiotemporal patterns in the retina during development. These *retinal waves* provide the only input to LGN neurons for much of development, and in both normal and transgenic animals, studies have shown that altering the spatiotemporal properties of retinal wave activity affects the refinement of their arbors in the LGN. Previously, we used the information contained in the retinal waves to predict the existence of a coincidence-based learning rule based on retinal ganglion cell bursts with time scales on the order of a second.

To see if such a learning rule exists *in vivo*, we use perforated patch recording in slice preparations of rat LGN during the eye-specific segregation of retinal ganglion cell (RGC) axons in the LGN (P6-10). We record from identified LGN neurons and measure the size of postsynaptic currents evoked by optic tract stimulation, both before and after a stimulation paradigm designed to emulate *in vivo* activity by applying bursts of appropriate duration and frequency. We find that the degree of synaptic plasticity is dependent on the latency between pre- and postsynaptic bursts, in a way that is consistent with our theoretical predictions.

Finally, we investigate the ramifications of this "burst-time dependent" learning rule at the retinogeniculate synapse. Using the measured learning rule, we simulate a simple model of retinogeniculate development and demonstrate how such a "burst-time dependent rule" could drive refinement of the inputs to an LGN neuron. Because our measured learning rule has windows for potentiation and depression, for some sets of parameters this refinement does not require additional rules that govern competition, but rather self-regulates the total input to the LGN neuron as observed in theoretical studies of spike-time dependent learning rules.

By measuring a learning rule based *in vivo* activity patterns, we are thus able to give a plausible description of activity-dependent development at the retinogeniculate synapse as well as provide a context for studies of disrupted retinogeniculate development resulting from blocking or modifying retinal wave activity.

Amplification of trial-to-trial response variability by neurons in visual cortex

Matteo Carandini

Smith-Kettlewell Eye Research Institute, San Francisco

The firing rate of neurons in visual cortex is highly variable across trials of stimulus presentation. Indeed, its variance grows with, and is often larger than, its mean. Because the generation of spikes is remarkably noiseless, the source of this variability is thought to lie in the underlying synaptic input. Here this hypothesis is tested through in vivo measurements of membrane potential. Surprisingly, trial-to-trial variability of potential is low, and variance barely grows with the mean. A simple model explains how this low variability of potential is greatly amplified by the firing threshold, resulting in the large variability of firing rate.

ASSOCIATIVE MEMORY AND OLFACTORY REPRESENTATION

S. Cassenaer and G. Laurent
Division of Biology
California Institute of Technology

An identified honeybee neuron, VUMmx1, appears to encode the reward stimulus in conditioning experiments (Hammer 1993). VUMmx1 endogenously responds to sucrose, and activating the cell during odorant presentation subsequently gives rise to proboscis extension in response to the odorant alone. VUMmx1 projects to the antennal lobe (AL), the mushroom body (MB) and the lateral protocerebral lobe, and injection of the VUMmx1 transmitter octopamine in the AL and MB can also replace sucrose as the unconditioned stimulus (Hammer and Menzel 1998). One interpretation of these behavioral experiments is that olfactory representations in the AL and MB could be changed by newly formed associations between odors and sucrose stimuli mediated by VUMmx1.

Neurons analogous to VUMmx1 that project to the AL and to the MB have been identified in the locust as well (Braunig 1991). We characterize how olfactory representations in the locust are affected by pairing odorant presentation with direct activation of these cells or with injection of octopamine.

QUANTAL ANALYSIS OF SENSORY-EVOKED RESPONSES IN CEREBELLAR GRANULE CELLS

Paul Chadderton, Troy Margrie, Michael Häusser

Wolfson Institute for Biomedical Research and Department of Physiology, University College London, Gower Street, London WC1E 6BT

Determining the relationship between sensory-evoked synaptic input and the resulting pattern of output spikes is essential if we are to understand the computations performed by the input layers of cortical structures. Granule cells constitute the input layer of the cerebellar cortex, translating mossy fibre (MF) signals into parallel fibre (PF) input to Purkinje cells. Remarkably little is known about the transformations performed by cerebellar granule cells on sensory-driven MF input in the intact brain, largely because the small size and dense packing of granule cells has hitherto precluded direct recordings from individual granule cells in vivo. We have investigated the relationship between MF synaptic currents evoked by tactile stimulation and the resulting granule cell output patterns using whole-cell patch-clamp recordings. We found that the ongoing firing rate of granule cells is very low, due in part to dampening of excitability by a tonic $GABA_A$ receptor-mediated conductance. Sensory stimulation produced bursts of MF EPSCs that summate to trigger bursts of granule cell spikes. By comparing the amplitude of sensory-evoked EPSCs with those of spontaneous EPSCs in the presence of TTX, we demonstrated that sensory-evoked burst responses were evoked by only a few quantal EPSCs. Nevertheless, spontaneous unitary MF inputs only triggered spike output when inhibition was reduced. Our results show that the input layer of the cerebellum balances exquisite sensitivity with a high signal-to-noise ratio. Granule cell bursts are optimally suited to trigger glutamate receptor activation (Dittman et al., *J. Neurosci.* 20:1374; Takechi et al., *Nature* 396:757) and plasticity (Brown et al., *Nat. Neurosci.* 6: 1048; Casado et al., *Neuron* 33: 123; Wang et al., *Nat. Neurosci.* 3: 1266) at PF synapses, providing a link between input representation and memory storage in the cerebellum.

Acknowledgements: This work was supported by the Wellcome Trust and the Gatsby Foundation.

WHY ARE MOST NEURONS IN THE HEAD OF THE ANIMAL?

Beth L. Chen and Dmitri B. Chklovskii
Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724

Why do neurons follow a particular organization in an animal? Can the placement or misplacement of neurons help us understand their function in the network? We address these questions by computing the optimal position of neurons through minimization of network wiring.

In the nematode *C. Elegans*, the number of neurons, the location of their cell bodies, as well as the wiring diagram have been documented and found to be reproducible from animal to animal. Using a quadratic cost function to minimize wiring, we found the placement of 280 (non-pharyngeal) neurons in the hermaphrodite worm to agree well with actual neuron locations on three levels 1. Neurons in the same gangalia cluster together. 2. The relative ordering of neurons follows actual anterior to posterior order, especially along the body and tail of the worm. 3. The neurons are well positioned relative to the worm’s sensory organs and muscles.

Our work also highlighted interesting and biologically relevant exceptions to the wiring minimization rule. Two classes of neurons, AVA and PVC, are command interneurons responsible for backward and forward locomotion in the worm, respectively. Instead of positioning at a location that minimizes total wire length, the cell bodies of these neurons are preferentially located closer to their post-synaptic terminals.

The results of our work not only provide insights for understanding the organizational constraints of neurons in an animal but also offer a new way to spotlight functionally important neurons in the network.

READING POPULATION CODES: RATE ESTIMATION VERSUS COINCIDENCE DETECTION

Denis Chigirev
Lewis-Sigler Institute for Integrative Genomics
Princeton University, Princeton, NJ 08544

There is a growing amount of evidence that information is represented in the brain in the form of a population response. How to read this information is still an open question. One commonly made assumption is that signals are differentiated by variable firing rate response to the stimuli. Then pooling firing rates together we can infer the stimulus. Such strategy is often justified by the noisiness of individual responses. However, the rates themselves still need to be estimated. Here we show that for biologically plausible response times the error introduced through rate estimation is of the order of 30 percent of the firing rate itself. This may explain the apprarent "noisiness" of individual neurons – the rate estimation error trumps trial to trial rate variability and makes further improvements in the precision of rate coding biologically meaningless.

After developing the theory of Bayesian rate estimation and showing its limitations we investigate an alternative coding strategy: coincidence detection. We consider wether such a system improves correct stimulus identification for quick time response. We show that a coincidence detector is inferior to rate estimator for independent neurons, but once the neurons are coupled coincidence detector outperforms a rate estimator and becomes the proper strategy for correct signal identification.

SYNAPTIC CONNECTIVITY AND NEURONAL MORPHOLOGY: TWO SIDES OF THE SAME COIN

Dmitri B. Chklovskii, Cold Spring Harbor Laboratory mitya@cshl.edu

Brain functionality relies on large numbers of neurons and synaptic connections between them. A cubic millimeter of the mouse neocortex, for example, contains about 10^5 neurons and 10^9 synapses. Building such a network is a formidable challenge because connections must be implemented in the physical world using biological wiring, i.e. axons and dendrites, which come at a considerable cost. Axons and dendrites take up valuable space, introduce delays and attenuation, require material and metabolic energy, and rely on genetic information for guidance in development. Although we do not know the exact origin of the wiring cost, it may be approximated by the wiring volume. Then the assumption that evolution minimized the wiring cost, while maximizing the network functionality, leads to the following optimal design problem. For a fixed functionality of the network, as specified by the synaptic connectivity, find the wiring design that minimizes the wiring volume.

Here, I consider a network with all-to-all connectivity, where each neuron makes a synaptic connection onto every other neuron. I calculate the minimal physical size of the all-to-all connected network in terms of the number of neurons, N, and the wire diameter, d. The network size depends on the chosen wiring design. Starting with the simplest possible wiring design, non-branching (or point-to-point) axons, I add features of neuronal morphology, such as branching axons, branching dendrites, and dendritic spines. Inclusion of each feature reduces the size of the network, implying that neuronal morphology makes wiring more efficient. Moreover, only the final wiring design, including all the salient morphology features, yields the correct size of a cortical column. This calculation suggests that the neuronal morphology has evolved to implement high synaptic connectivity, necessary for brain function. This teleological argument should help in understanding brain evolution and in inferring synaptic connectivity from the arbor morphology.

PREDICTIVE ENCODING OF TIME-VARYING SIGNALS USING SPIKE-TIME-DEPENDENT PLASTICITY

Erik Cook and Markus Meister, Harvard University

We used computational modeling to examine how spike trains efficiently encode time-varying signals. We assumed that continuous signals are converted into spike trains for effective transmission over long distances. Given this hypothesized role for spike trains, what is the best way for single neurons to encode time-varying signals using the fewest spikes?

Predictive encoding has been widely proposed as a model for neuronal communication. This encoding method was invented in 1952 by de Jager for converting analog signals into digital representations. We used predictive encoding as a model of single neuron spike generation for efficient transmission of a time-varying band limited signal, $X(t)$ (Fig. 1). The key component of this encoder is the prediction filter which converts a sequence of spikes into the best estimate of the encoded waveform $X'(t)$. The same prediction filter resides both in the feedback loop at the spike generator and in the decoder (which corresponds to the synaptic response). The prediction filter gets its name from the fact that it predicts future values of $X(t)$ based on past spike activity.

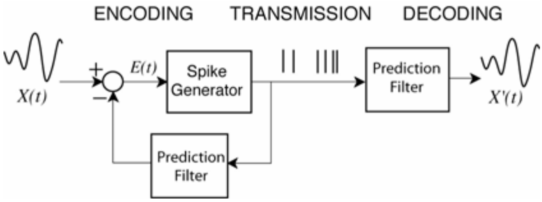

Fig 1. Predictive encoder and decoder

The prediction filters in Fig. 1 can be optimized to transmit $X(t)$ with minimal spikes, which maximizes encoding efficiency. This optimization uses the correlation between spikes and the encoding error, $E(t)$, to adaptively adjust the weights of the prediction filter. The major limitation of this model is that the prediction filter residing at the decoder must also be optimized using the error signal that is only available at the spike generator. How can the prediction filter at the decoder be optimized when no error signal is available?

To solve this problem, we introduced two new components to the model. First, inspired by the retina, we add a second spike channel that encodes an inverted copy of the input. Thus, the model contains the familiar retinal on- and off-pathways. Next we used spike-time-dependent plasticity to optimize the prediction filters at both the spike-generator and decoder. In this case, the relative timing of the spikes in the on- and off-pathways is used to optimize the prediction filters. Thus, the error signal used to drive the optimization is implicitly represented in the interaction between spikes. We show that this model efficiently transmits the time-varying signal $X(t)$ and that the resulting on- and off-pathways demonstrate transient spike responses that are similar to retinal ganglion cell activity. Thus, this model may account for efficient encoding of time-varying information in the retina and other areas of the brain. Importantly, spike-time-dependent plasticity may represent a mechanism that the nervous system uses to transmit error signals and maintain a match between the encoding and decoding of information in spike trains.

LINEAR-NONLINEAR-PROBABILISTIC MODELS OF ANIMAL FORAGING: BEHAVIORAL METRICS OF SUBJECTIVE VALUE

Greg S. Corrado, Leo P. Sugrue and William T. Newsome
Department of Neurobiology.
Stanford University, Stanford, CA 94305

Animal foraging behavior has received much attention of late from researchers studying free choice and the attribution of value. As the simplest natural context engaging these mechanisms, ethological foraging seems the ideal setting for these experiments. To understand how the brain directs free choice in this context, we must localize neural signals carrying information related to covert decision variables like subjective value. To do this, we need a model of how perceived value is sculpted by experience, and how it impacts choice.

Linear-Nonlinear-Probabilistic (LNP) models have been very successful in describing neural computation in early sensory systems through hidden variables like membrane potential. We apply this same framework to the behavior of whole animals, foraging in a nonstationary environment. Our goal is to construct an adequate yet compact model of animal choice, that allows the estimation of hidden decision variables correlated with neural activity.

We find LNP models to be effective descriptors of choice behavior in humans and other animals. These models allow realtime estimation of covert decision variables that guide animal behavior. The resulting operationalized metric of subjective value is also successful in extracting putative neural correlates of value in area LIP. These neural correlates have the desired properties of being largely independent of immediate sensory signals and motor responses, but highly depenedent on recent reward experience.

In constructing these models, linear filters of reward history resembling exponential kernels are extracted directly from animal behavior using Choice Triggered Averaging (CTA). The integration time constant (τ) of these kernels is typically on the order of 9 choices. In simulation, LNP models using similar kernels produce behavior indistinguishable from real data. Moreover, these models harvest rewards more efficiently than models with either shorter or longer τ, gathering over 90% of the rewards earned by an ideal harvester.

This correspondence between τs recovered by CTA and those of the best performing LNP model appears to be largely independent of the choice of nonlinearity. Both divisively-normalized and subtractive-opponent formulations of the nonlinear stage produce adequate descriptions of animal choice, matched performance of ideal and observed harvesting behavior, and generate hidden variables which regress well on the neural signals.

We conclude that LNP models adequately describe simple animal foraging behavior, and successfully extract neural correlates of subjective value.

EXPECTATION-MAXIMIZATION (EM) ESTIMATION OF AN INTEGRATE-AND-FIRE MODEL WITH SPIKE-FREQUENCY ADAPTATION

Gabriela Czanner
Department of Statistics
University of Pittsburgh, Pittsburgh, PA 15260

We are studying stochastic integrate-and-fire (I&F) models of membrane potential which incorporate the spike-frequency adaptation. Our aim is to estimate parameters of the model from voltage trace data using expectation-maximization (EM) methodology.

We assume that spike-frequency adaptation is caused by an adaptation current. Hence the I&F model we are using consists of two stochastic differential equations: one describes the membrane potential, the second describes the adaptation current which is hidden. The adaptation current brings the membrane potential down, and so makes it harder to reach a threshold and to fire. On the other hand, the membrane potential affects the hidden process through spikes. When the membrane potential reaches the threshold the adaptation current increases by a constant amount.

Since the adaptation current process is hidden we use a state-space model approach (Jazwinski, 1970). In traditional models (Shumway and Stoffer, 1982; Smith and Brown, 2003) a latent process affects the observed process, but not *vice versa*. In our model, both processes affect one another.

We develop an EM algorithm to estimate the parameters of unobserved adaptation process. The algorithm maximizes the complete data log likelihood and it combines a recursive linear filter algorithm, fixed interval smoothing algorithm and state-space covariance algorithm. A key step in our approach is in a discretization of the I&F model, so that the resulting model retains the behavior of the model and theoretical derivations are possible.

We illustrate the EM algorithm with simulated data. We simulated the membrane potential and the adaptation current of cortical regular spiking pyramidal neurons (Liu and Wang, 2001) and used the algorithm to estimate the hidden adaptation process and its parameters.

SHORT PATHS AND SIGNAL PROPAGATION IN A MODEL CORTEX

Rava A. da Silveira and Daniel S. Fisher
Department of Physics, Harvard University, Cambridge, Massachusetts 02138
Laurence F. Abbott
Volen Center for Complex Systems and Department of Biology, Brandeis University, Waltham, Massachusetts 02454

We present a set of qualitative and approximate quantitative arguments about simple issues of connectivity and dynamics, motivated by experiments.

Anatomical measurements in the cortex indicate that any neuron is connected to any other neuron through paths of up to about six intermediary synapses, rarely more. Also, electrophysiological experiments on both human and animal subjects show that visual information is processed rapidly, with, *e.g.*, discriminatory activity in frontal areas of the cortex less than 150 milliseconds after stimulus onset, a performance that indeed requires '*short*' neural paths (*i.e.*, with few intermediary synapses) spanning the whole cortex. At the same time, cortical connectivity is known to be sparse and distributed, and it appears fairly random, leaving aside large targeted white matter tracts and averaging over columns and selectivities. Furthermore, the histogram of axon lengths is not flat, but falls off as a power law, so that there is only a small supply of very long axons among the *non-targeted* connections.

We ask, first, how we can reconcile these two sets of observations or, more specifically, whether it is possible to have many *short paths* even in a (model) cortex where connections are random and long axons are in a minority. We then turn to the problem of *propagation* of a front of activity (a '*signal*') in such a model cortex: how does propagation depend on basic parameters (such as the density of excitatory and inhibitory neurons, the sparseness of the connectivity, synaptic efficacies, *etc.*)? And what constraints do general requirements, such as stability, impose on signal propagation?

To answer questions on connectivity, we define a simple model, in which connections are random with a probability that decreases with distance. We show that, nevertheless, the number of synapses needed to connect two neurons grows *very slowly* with the distance separating them. Typically, about six synapses should indeed suffice to span the whole cortex, in the absence of any 'design.' As a first attempt to answering the more difficult questions on dynamics, we pose a very crude model of neural activity, in which neurons have a fixed efficacy (probability of spiking upon receipt of an excitatory spike) and are otherwise linear. In spite of its simplicity, the model displays interesting sensitivity to the values of its parameters (such as, *e.g.*, the strength of inhibition), and basic requirements such as stability impose stringent constraints on propagation. A preliminary analysis of the model yields general properties of propagation (relating, *e.g.*, propagation speed and input strength) and establishes a context for extensions that include finer effects (of, *e.g.*, non-linearity and fluctuation).

A COMPUTATIONAL MODEL FOR BOTTOM-UP AND TOP-DOWN INFORMATION FLOW IN THE VISUAL CORTEX

Akira Date
Communications Research Laboratory, Kyoto 619-0289 Japan

We have developed a computational model which performs image interpretation in degraded environments; Given a degraded image, the model try to find the best interpretation of bounding contours. The interpretation is represented by the distributed pattern of activity in a tree-structured graph consisting of a large number of units. The bottom level of the graph is the pixel process, corresponding to the actual digitized image. Successively higher levels correspond to the images in coarser scales, each concerned with a small piece of information. The model is generative one based on stochastic grammars, which can accommodate *a priori* information at multiple levels of a knowledge hierarchy. Given a degraded picture, we seek the image that maximizes the posterior distribution where maximization is performed via dynamic programming. There is no local minima for this computation. This is a powerful theoretical framework provided by S. Geman & K. Manbeck (1993) where they demonstrated potential advantages of the model-based hierarchical approach over purely bottom-up or top-down alternatives in recognition. Since they imposed sufficient contextual constrains, some nodes need to have more than 100,000 states when experimented with a 256×256 image. We have developed a simplified version of this model with weak constraints in which each node has two states. By this, the computational operations are greatly reduced. The capability of the model to perform image interpretation has been investigated. We demonstrate that 1) incomplete generative models still work for this recognition problem, 2) damages of top-down signals affect boundary-finding, which intuitively depicts a role of the top-down signals.

Our model shares many properties with others: a) it keeps multiple hypothesis and propagates ambiguity in the bottom-up processes, b) top-down signals from higher level in the hierarchy serve to disambiguate states in lower levels. Our model is not aimed at accounting for biological data such as end-stopping and other extra-classical receptive-field effects. We, rather, focus on the problem of part-whole hierarchy, how a large set of locally-connected units represent a class of shapes which has global-structure from the engineering point of view. Obviously it does not directly address the question of whether analogous processes operate in visual cortex, and the model is too simple to explain. However, general features of more realistic and sophisticated models are probably reflected in this simple example.

NATURAL STIMULUS STATISTICS MODULATE RECEPTIVE FIELD STRUCTURE IN PRIMARY VISUAL CORTEX

Stephen V. David[1], William E. Vinje[2,3] and Jack L. Gallant[3,4]
[1] Graduate Group in Bioengineering, University of California, Berkeley
[2] Department of Molecular and Cellular Biology, University of California, Berkeley
[3] Program in Neuroscience, University of California, Berkeley
[4] Department of Psychology, University of California, Berkeley

Studies of primary visual cortex (V1) have produced models that account for neuronal responses to synthetic stimuli such as sinusoidal gratings. How well do these models generalize to activity during natural vision? To answer this question, we recorded responses in V1 of awake macaques under conditions simulating natural vision and during stimulation by synthetic grating sequences. We compared how well a nonlinear receptive field model fit using each of these data sets predicted time-varying neural responses to a novel natural stimulus. In more than half of the neurons studied (59%), the model fit using the natural stimulus predicted natural visual responses better than the model fit to the synthetic stimulus.

The superior predictive power of the model fit using the natural stimulus suggests that during natural vision V1 response properties are modulated in a way that cannot be predicted from responses to the synthetic stimulus. In order to characterize this modulation, we compared temporal and spatial response properties of the models using the different stimuli. We validated differences in response properties using a hybrid model that combined response properties from both stimulus classes. For example, we fit one hybrid model that fit temporal response properties using natural stimuli and spatial response properties using grating stimuli. By comparing predictive power of the hybrid model to that of the model fit using only the natural stimulus, we were able to measure the influence of natural stimuli on spatial response properties alone.

During natural stimulation, temporal responses often showed a stronger late inhibitory component, indicating an increase in short term adaptation. In addition, spatial tuning underwent complex shifts from the structure observed with synthetic stimuli. The majority of these shifts lay in suppressive, rather than excitatory, elements of the spatial response profile.

CONTRASTING NEURONAL CORRELATES BETWEEN DORSAL AND VENTRAL STRIATUM IN THE RAT

Nathaniel D. Daw,[1,3,4] David S. Touretzky,[1,3] and William E. Skaggs[2,3,5]
[1]Computer Science Dept., Carnegie Mellon University
[2]Dept. of Neuroscience, University of Pittsburgh
[3]Center for the Neural Basis of Cognition
[4]Present affiliation: Gatsby Computational Neuroscience Unit, University College London
[5]Present affiliation: University of Arizona

Cortico-striatal interactions are famously organized in a set of “loops” joining associated areas of cortex and striatum. These anatomical distinctions are assumed to reflect a functional breakdown; e.g., the ventral part of the striatum (innervated by limbic structures such as orbitofrontal cortex) is thought to have more of a motivational role than the dorsal striatum, which receives input from sensorimotor cortices. While such a separation of function is supported by targeted lesions and drug infusion studies, it has not clearly been reflected in neuronal recordings. Here, dorsal and ventral striatal units are compared in rats performing a task designed to tease apart responses related to the substructures’ putative functions. The distinction our experiment probed, suggested by computational models, is between firing related to an animal’s *decisions* versus firing related to those decisions’ expected *outcomes*.

Rats performed a series of cued decisions (left or right turns) on a T-maze for food and chocolate milk reward. During a block of trials, correctly executing a cued turn direction was rewarded with a consistent outcome (e.g., right → milk; left → food). In a subsequent block of trials, the association between direction and reward was reversed (e.g., right → food; left → milk), to dissociate neuronal responses to the actions from their outcomes.

Here we report on 187 units recorded from the dorsal striatum of two rats, and 247 units from the ventral striatum of a third rat. (Because they come from a single animal, the ventral data are preliminary; we are now analyzing ventral data from another rat.) About half of the units were unresponsive; most of the rest responded transiently at some phase of the task.

Of the responsive neurons, 39 dorsal and 49 ventral neurons responded differentially during left- versus right-turn trajectories. After contingencies were reversed, dorsal units tended to follow the turn direction (33/39) rather than the outcome type (6/39), while ventral units showed the opposite pattern: all but one neuron followed the rewards rather than the actions. These results are consistent with the functional anatomy in that they demonstrate more motoric correlates dorsally and more reward-related correlates ventrally.

Supported by National Science Foundation grant DGE-9987588.

DIFFERENTIAL EFFECTS OF SINGLE AMACRINE CELLS ON DIFFERENT GANGLION CELL TYPES

Saskia E.J. de Vries and Markus Meister
Program in Neuroscience, Department of Molecular and Cellular Biology, Harvard University, Cambridge, MA

Retinal ganglion cells differ in the kinetics of their light response, the prominent distinction being between transient and sustained responses. These differences may arise early in the retinal network with distinct bipolar cell types, or later on through dedicated amacrine circuits in the inner retina. We directly examined the involvement of amacrine cells in these responses by stimulating a single amacrine cell intracellularly while recording extracellularly from 50 ganglion cells in the salamander retina. Such stimulation affected both the spontaneous and light-evoked activity of ganglion cells within 400 m of the amacrine cell. Injecting hyperpolarizing current into sustained OFF amacrine cells, in the absence of visual input, elicited spiking activity from both transient and sustained ganglion cells. The time course of this response, however, differed depending on the ganglion cell type. Amacrine cell hyperpolarization caused sustained firing in ganglion cells with a sustained light response, but only transient firing in ganglion cells with a transient light response. We conclude that sustained amacrine cells influence both sustained and transient ganglion cell responses, but these interactions have different kinetics. This suggests that the circuit element that confers the distinction between sustained and transient light responses lies between the amacrine and ganglion cells.

POPULATION DYNAMICS OF SPIKING NEURONS: QUANTIFYING THE CONTRIBUTION OF SPIKE TIMING TO CODING.

V Del Prete and ACC Coolen
Department of Mathematics
King's College London, Strand WC2R2LS London UK

The question whether time is a relevant dimension for information processing and transmission in the brain is controversial: information estimates from real data recorded in vivo have shown that in many cases most information is contained in the time averaged activity, suggesting a dominance of a rate code versus a temporal code. On the other hand synchrony and oscillatory behaviour have been detected in many brain regions, but the role of such phenomena remains unclear. Spike timing has also been shown to play a relevant role in synaptic plasticity. From a theoretical point of view the inclusion of a spike generating mechanism makes the mathematical analysis of the dynamics much more complicated, so that the common choices are to restrict oneself to rate models or to resort to purely numerical simulations. We propose here an effective spiking neuron model which aims at bridging the gap between the realistic parameterization of conductance based models (whose complexity does not allow for an analytical treatment of the population dynamics) and the extremely simplified description of the integrate and fire models. We study the dynamics of two interconnected populations of excitatory and inhibitory neurons, each endowed with a recurrent structure of the connectivities, where in both populations the single neuron dynamics is given by the effective model. In particular, a separation of degrees of freedom between supra and sub threshold activity and an analytical mechanism for the reset after the spike, allow us to go beyond the mean field approximation on the synaptic current fed into the post-synaptic neuron, which is often considered as a sum of independent inputs and treated as a signal term plus gaussian noise. We can actually define new global functions quantifying the amount of **synchrony between pairs of neurons** and study their dynamical evolution in time. Varying the different neuronal time constants, the level of noise and the synaptic strengths across populations we study the effect of inhibition on the activity of the excitatory cells in driving oscillations and synchrony at different time scales. We are currently introducing external currents into the network, mimicking the response of previous stages of processing to external stimuli. We are computing the mutual information between the response of the neurons and the currents, estimating separately the contribution of average rates and spikes to coding.

Partially funded by EU contract QLGA-CT-2001-51056.

SHARED AND PRIVATE VARIABILITY IN THE AUDITORY CORTEX

Michael R. DeWeese & Anthony M. Zador

Cold Spring Harbor Laboratory, 1 Bungtown Rd., Cold Spring Harbor, New York 11724, USA

The high variability of cortical sensory responses is often assumed to impose a major constraint on efficient computation. In the auditory cortex, however, response variability can be very low. We have used *in vivo* whole-cell patch clamp methods to study the trial-to-trial variability of the subthreshold fluctuations in membrane potential underlying tone-evoked responses. We partitioned this subthreshold variability into a *private* component (which includes synaptic, thermal and other sources local to the recorded cell) and a *shared* component arising from network interactions. We found that the private component was remarkably small, usually about 1-3 mV as quantified by variance over mean PSP height. The shared component was often much larger, and showed more heterogeneity across the population, ranging from about 1-10 mV. Our observations indicate that neurons in the auditory cortex may operate in the low noise regime.

SUBTHRESHOLD VOLTAGE NOISE OF CORTICAL PYRAMIDAL NEURONS

Kamran Diba[1], Gilad Jacobson[2,3], Yasmin Oz[2], Anat Yaron[2,3], Christof Koch[1], Idan Segev[2,3], Yosef Yarom[2,3]
[1] Computation and Neural Systems Program, Division of Biology 139-74, California Institute of Technology, Pasadena 91125, USA.
[2] Department of Neurobiology and [3] The Interdisciplinary Center for Neural Computation, The Hebrew University, Jerusalem 91904, Israel.

Noise affects the input-output characteristics of neurons. Nevertheless, intrinsic voltage noise has not previously been studied systematically in cortical pyramidal neurons. We measure the voltage power spectral density and impedance in layer IV-V pyramidal neurons of the rat somatosensory cortex at different holding potentials and with different pharmacological treatments, in vitro. We dissect the relative contributions of holding potential, synaptic activity and Na^+ conductance to the voltage noise, experimentally, and compare this with a ``ball-and-stick'' model of a neuron with stochastic Markovian ion channels and spontaneous low-rate synaptic input. Voltage noise increases significantly with holding potential. We show that Na^+ conductance is responsible for noise amplification with depolarization and that this amplification is a consequence of the voltage-dependent effect of Na^+ channels on the impedance properties of the membrane. On the other hand, the stochasticity of Na^+ channels in itself does not contribute significantly to the measured noise. Thus, voltage dependence of one conductance may dramatically alter the manifestation of noise from a different current noise source as holding potential varies. Spontaneous synaptic activity, which varies considerably across cells in the slice, is a minor source of noise in our preparation. Nevertheless, it is still the significant factor in shaping the noise spectrum in the frequency range above 10 Hz. Here, synaptic input acts as a current noise source without contributing significantly to the impedance of cells. We compare this data with output from our model neuron to elucidate the impedance behaviour and spectral properties of the voltage noise. Furthermore, we make predictions about the intrinsic noise properties of cortical pyramidal neurons under realistic in vivo conditions.

STOCHASTIC MODELS OF DECISION MAKING:
BEHAVIOR AND PHYSIOLOGY IN A MOTION DISCRIMINATION TASK

Jochen Ditterich[1,2] & Michael N. Shadlen[2]
Center for Neuroscience, Section of Neurobiology, Physiology & Behavior,
University of California, Davis[1] &
HHMI, WaNPRC, Dept. of Physiology & Biophysics, University of Washington, Seattle[2]

Computational models based on diffusion processes or random walks have been proposed to account for human decision making behavior in a variety of tasks (e.g., Luce, 1986; Link, 1992; Ratcliff et al., 1999; Smith, 2000; Usher and McClelland, 2001). Here we explore whether such models account for the speed and accuracy of perceptual decisions in a reaction-time motion-discrimination task and whether such models explain the decision related activity of neurons recorded from the parietal cortex.

We analyzed published behavioral and neurophysiological data from two monkeys (Roitman and Shadlen, 2001). The monkeys were trained on a variant of the 2-alternative forced-choice direction-discrimination task used by Newsome and colleagues in their studies of motion perception in the macaque (e.g., Britten et al., 1992). The difficulty of the discrimination was controlled by varying the fraction of coherently moving random dots in the display. The monkeys communicated their direction choice by making an eye movement from the fixation point to one of two choice targets. They were free to make this eye movement response as soon as they were ready, thus providing a measure of both response time and direction choice.

The diffusion models we analyzed share the following features. The response of direction-selective neurons constitutes evidence in favor of one direction and against the opposite direction. The evidence is expressed as a simple difference and is integrated as a function of time. When the accumulated evidence reaches a criterion level — a positive or negative barrier — the process terminates. The sign of the termination dictates the choice. The time of termination dictates decision time, to which a random non-decision time is added to yield response time. The simplest diffusion model, described fully by these features, can explain the psychometric function and the mean response times, but it fails to account for the response time distributions. The model predicts more skew than is seen in the observed distributions.

This deficiency was ameliorated by allowing the terminating barriers to decline as a function of time. This “urgency mechanism” instantiates the reasonable assumption that the monkeys base their decisions on less compelling evidence as time passes. This “time-variant” diffusion model explains the psychometric function and mean response times, and it accounts for the shape of the response time distributions. Several possible implementations of the time-variant diffusion model are compatible with the observed behavior. For example, a declining decision criterion and an increasing gain applied to the sensory evidence make nearly identical predictions for response time and choice. However, each of these mechanisms make distinct predictions about the time course of the evolving decision variables in the brain. We compared simulations of the expected time course of the decision variable to LIP responses obtained from the monkeys during the discrimination task. The best match was achieved by increasing the gain of the sensory evidence as a function of time.

Thus, a straightforward extension of the diffusion-to-barrier model can explain the dependence of accuracy and mean response time on motion strength, the shapes of the response time distributions, and the pattern of neural responses recorded from the parietal cortex of the monkey while performing a perceptual decision. Interestingly, such a “time-variant” decision process allows the monkey to maximize reward rate in light of the risk of aborting a trial by breaking fixation before a choice can be reported.

(Supported by DFG (DI 819/1-1), HHMI, and EY11378.)

MODELLING CORTICAL REORGANIZATION IN TINNITUS

Melissa Dominguez[1] and Suzanna Becker[1] and Ian C. Bruce[2]
Department of Psychology[1] and Department of Electrical and Computer Engineering [2] McMasterUniversity, Hamilton, Ontario, Canada

Hearing loss due to peripheral damage is associated with cochlear hair cell damage or loss and some retrograde degeneration of auditory nerve fibers. Surviving auditory nerve fibers in the impaired region exhibit elevated and broadened frequency tuning and the tonotopic representation of broadband stimuli such as speech is distorted. In impaired cortical regions, increased tuning to frequencies near the edge of the hearing loss coupled with increased spontaneous and synchronous firing is observed. Tinnitus, an auditory percept in the absence of sensory input, may arise under these circumstances as a result of plastic reorganization in the auditory cortex, particularly on recurrent cortical connectons that tend to reinforce Hebbian cell assemblies having correlated firing.

We present a model of auditory cortex that captures several key features of auditory frequency tuning, and of cortical organization in general. This spiking neuron model compensates for the loss of afferent excitation in the damaged region by changing the balance of excitation and inhibition on the lateral connections. These changes allow the model to capture some of the cortical correlates of tinnitus, including increased synchrony of firing. This model will be used to investigate the role of neural plasticity in cortical reorganization. Further, it may also be useful for evaluating procedures designed to segregate synchronous activity underlying tinnitus and for evaluating adaptive hearing devices that compensate for selective hearing loss.

SCALE-INVARIANT ADAPTATION IN NEURAL SYSTEMS

Patrick J. Drew & L.F. Abbott
Volen Center for Complex Systems and Department of Biology
Brandeis University, Waltham, MA 02454-9110

Virtually all neural responses show adaptation, and this adaptation frequently takes a scale-invariant form, meaning that it has no single characteristic time scale (Thorson & Biederman-Thorson, 1974; Xu et al., 1996; Fairhall et al., 2001). Examples of scale-invariant adaption range from single channels to human psychophysics. Nevertheless, most models of adaptation are based on dynamics operating at only a few different time scales. Such models do not capture the distinctive features of scale-invariant adaptation in which decays are power-law or logarithmic, rather than exponential, and no intrinsic parameters set the relevant scales for the system.

We base our model of scale-invariant adaptation on a general description of perfect or complete adaptation used, for example, in studies of bacterial chemotaxis (Yi et al., 2000). Perfect adaptation refers to a system that ultimately stops responding to a constant stimulus. In the standard version of this model, a stimulus drives a response that is fed into an integrator. The output of the integrator is then fed back to inhibit the response being integrated. The circuit produces a response that is a low-pass filtered version of the derivative of the stimulus. This reproduces the well-known connection between adaption and differentiation but fails to match the scale-invariant feature because the low-pass filtering occurs with a specific time constant.

We construct a scale-invariant model of adaptation by modifying the integral in the feedback loop to a dimensionless form. Specifically, the time integral of previous responses is weighted by a factor inversely proportional to how far back in time the response occurred. This modification allows the model to match data showing scale-invariant adaptation in responses to electric fields in electric fish over 4 orders of magnitude (Xu et al., 1996). The response generated by this form of adaptation does not correspond to differentiation but, instead, approximates a Hilbert transformation of the stimulus. For example, responses to sinusoidal signals are phase shifted with a phase and amplitude that are roughly independent of frequency. This allows such systems to extract the envelope of a complex oscillatory signal. We explore these and other aspects of this type of adaptation.

ENCODING OF BEHAVIORALLY GENERATED OPTIC FLOW BY BLOWFLY NEURONS THOUGHT TO BE INVOLVED IN OPTOMOTOR COURSE CONTROL

M. Egelhaaf[1], R. Kern[1], J.H. van Hateren [2] and J.P. Lindemann[1]
[1]Department of Neurobiology, Bielefeld University, 33501 Bielefeld, Germany
[2] Dept. of Neurobiophysics, University of Groningen, 9747 AG Groningen, The Netherlands.
e-mail: martin.egelhaaf@uni-bielefeld.de

Knowing about the wiring of a neuronal circuit and the general response characteristics of its elements does not allow us to infer safely what visual features are encoded during natural behavioral situations, because the mechanisms underlying neuronal information processing are often non-linear. Therefore, results obtained with relatively simple stimuli as are conventionally used in visual sciences cannot be generalized to natural situations which usually are characterized by much more complex stimuli. Although sensing is often implicitly assumed to be the acquisition of information by a passive observer, part of the visually information is generated actively when animals move. For instance, in free-flight several insects shift their gaze by saccadic turns of the entire body and keep their gaze basically fixed between saccades, a viewing strategy reminiscent of primate eye movements. Thanks to recent technological developments, the optic flow experienced on the eyes of free-flying blowflies can be replayed to the visual system in electrophysiological experiments. We show that visual motion computation in blowflies is tuned to make efficient use of the characteristic dynamics of retinal image flow resulting from their saccadic flight style which enables them to obtain information about the spatial layout of its environment. Between saccades the retinal image flow resulting from sideslip movements contains information about object distances. This information is confined to low frequencies, whereas intersaccadic rotational optic flow is coded separately in an adjacent frequency band. Our conclusions based on neuronal responses to natural optic flow are in stark contrast to earlier conclusions based on conventional visual motion stimuli. We stress that taking into account the constraints encountered in behavioral situations is essential for understanding the functional role of neuronal computations.

Glomerular input results in distinct modes of Ca^{2+} signaling in olfactory bulb granule cell dendrites

Veronica Egger, Karel Svoboda & Zachary F. Mainen, CSHL

In the mammalian olfactory bulb (OB), axonless granule cells (GCs) mediate self- and lateral inhibitory interactions between mitral/tufted cells via a dendrodendritic reciprocal synapse. While GCs have been reported to spike during odorant stimulation, inhibitory interactions do not necessarily require GC action potentials (APs), thus implying the possibility of more than one GC output mode. Do these modes exist and if so, do they involve spatially differential activation of the GC dendritic tree? We imaged calcium transients in GC apical dendrites, using two-photon microscopy in rat brain slices, and mimicked sensory input by activating GCs via extracellular stimulation of a glomerulus.

We observed a sequence of distinct GC activation patterns with increasing strength of stimulation that rely on different mechanisms of dendritic integration:

1. The most local activation corresponds to synaptic input to an isolated spine, which results in robust Ca^{2+} transients (n = 31 spines), mostly restricted to the spine head. Half of this synaptic Ca^{2+} influx is due to NMDAR activation (n = 7), while the remainder is blocked entirely by 100 µM Ni^{2+} and thus due to activation of voltage-dependent Ca^{2+} channels and/or internal stores (n = 5).

2. Large EPSPs above a certain threshold (~ 30 mV) cause Ca^{2+} influx throughout the GC dendrite ('global EPSP'; n = 44 cells). This Ca^{2+} signal is inhibited by T-type Ca^{2+} channel blockers (n = 4) and most likely corresponds to the previously unknown substrate for AP-independent lateral inhibition in the OB. So far, no evidence for regional or branchlet activation has been observed.

3. When the GC ultimately spikes, Ca^{2+} influx is also seen throughout the dendrite (n = 34 cells), on average with faster rise times than for the global EPSP. Glomerularly-evoked APs are usually followed by a long-lasting depolarization with concomitant additional Ca^{2+} influx.

4. Moreover, in a given spine, Ca^{2+} evoked by APs or global EPSPs can summate with Ca^{2+} evoked by direct synaptic activation, resulting in increased mean Ca^{2+} transient amplitudes. Thus these overlaps correspond to additive operations that might render release from this spine more likely.

Thus there is a set of distinct GC output modes, differentiated by the spatial extent of the Ca^{2+} signal, its amplitude and kinetics. These observations also show that the contribution of APs is not to increase the degree of GC activation, suggesting that APs do not enhance the extent of lateral inhibition. However, AP-mediated Ca^{2+} signals appear to provide greater reliability and precision of temporal onset.

POSSIBLE FUNCTIONAL PURPOSE FOR THE SPIKE-TIMING AND FREQUENCY DEPENDENCE OF NEOCORTICAL SYNAPTIC PLASTICITY

Michael Eisele and Ken D. Miller
Keck Center for Integrative Neuroscience
San Francisco, CA 94143-0444
eisele@phy.ucsf.edu

Neocortical synaptic plasticity depends not only on the relative timing between pre- and postsynaptic spikes, but also on the timing between presynaptic spikes, on the timing between postsynaptic spikes, and on the frequency of pre- and postsynaptic spike trains. Here a simple learning rule is presented, which can reproduce the spike-timing dependence that was seen in superficial layers of cat visual cortex. It is similar to the spike-efficacy rule that was proposed by Froemke and Dan. But unlike this rule, our model can also qualitatively reproduce the frequency dependence that was seen in several other cortical areas. It potentiates the synapse if the postsynaptic neuron responds to input with a large firing rate and depresses the synapse if the firing rate is low.

This rule is based on an underlying principle: Before pre- and postsynaptic spikes are associated, they are grouped to firing episodes, and the weight change depends only on the timing of these episodes, but as little as possible on the timing of individual spikes. Because of this principle, the rule has only 5 free parameters, and its functional purpose becomes easier to understand.

Several functional interpretations are possible, but one fits particularly well. We call this interpretation a ”mending” rule. Two randomly chosen neurons in the neocortex are not very likely to be connected, even if they are close to each other, but if such neurons form a synaptic connection, they often form multiple ones. By keeping such a strong connections strong, and weak connections weak, synaptic plasticity may mend and maintain connections that have been previously learned.

We have constructed such a mending rule. It is based on several assumptions about how the presynaptic input and other inputs through other synapses excite the postsynaptic neuron. Under the right set of assumptions, it becomes the same as our 5-parameter rule, which fits the neocortical data. This does not mean that mending is the main goal of synaptic plasticity, as the same rule could also be used for learning. But it does mean that both the spike-timing and frequency dependence of neocortical synaptic plasticity may already be needed for a rather simple task: the mending of existing connections.

TRADEOFF IN SPECTROTEMPORAL FEATURE SELECTIVITY, RESPONSE SPARSITY, AND MUTUAL INFORMATION AND THE EFFECTS OF INTRACELLULAR THRESHOLDING.

Monty A. Escabi[1.2], Reza Nassiri[2], Heather L. Read[3], Christoph E. Schreiner[4]
[1] Electrical and Computer Engineering, [2] Biomedical Engineering, and [3] Department of Psychology, University of Connecticut, Storrs, CT 06109 and [4] Keck Center for Integrative Neuroscience, University of California, San Francisco, CA 94143.

Contrasting hypotheses of sensory coding range from the notion of the pinnacle "grandmother cell" or "feature detector" to linear coding strategies that use firing rate to encode sensory information. Few studies, however, address the reliability of the encoding on a single action potential-stimulus feature basis.

In the central nucleus of the inferior colliculus (ICC), neurons can respond selectively to spectro-temporal features in complex auditory stimuli and can be selectively activated by the instantaneous statistics of complex sound stimuli (Escabí and Schreiner 2002). Spectro-temporal receptive field (STRF) analysis can be used to assess the structural feature preferences of auditory neurons and to approximate their linear and nonlinear contribution of the response (Escabi & Read, 2003; Fritz et al., 2003). STRFs provide an initial glance at the stimulus features that evoke neural responses, however, such an analysis only conveys the average stimulus pattern that evokes an action potential; and provides no information on how sound patterns preceding each action potential contribute to the average sound waveform. We consider whether stimulus patterns are highly stereotypical from one action potential to another or alternatively whether they differ significantly from spike to spike. Thus, the same STRF could be built up conceptually by averaging stimulus-response ensembles with strikingly different statistical properties and exhibiting different degrees of stimulus-response variability. This variability, in turn, determines the precision of the neural response.

We develop a second-order spectro-temporal reverse correlation technique to measure the fidelity of the encoding. We do this by quantifying the variability of the spectro-temporal stimulus patterns that evoke action potentials. We find that reliability of spectro-temporal feature encoding in the ICC is inversely related to neuronal spike rates. Neurons with the highest encoding reliability respond to complex sound patterns with exceptionally low spike rates (down to 0.05 spikes/sec), but do so with great fidelity, while at the alternate extreme, neurons can exhibit substantially high spike rates (up to ~ 100 Hz) but do so with low reliability. A similar tradeoff in the response reliability is also observed between neuronal spike rates and spike-normalized mutual information.

We next study this behavior in a modified spectro-temporal integrate-and-fire neuron model and find that we can accurately replicate this behavior by simply changing the model neuron's intracellular threshold. We exploit this model to identify the specific tradeoffs between neuronal spike rates, spectrotemporal feature selectivity, and mutual information. This approach allows us to characterize how the neuronal representation changes from a highly accurate single spike per stimulus "sparse code" to an imprecise rate representation. We propose that similar threshold-dependent tradeoffs in spike rates and information encoding ability should be evident in other sensory systems.

DEVELOPING A STATISTICAL MECHANICS FOR NEURAL POPULATIONS

Michael A. Farries Depts. of Biology and Otolaryngology, Univ. of Washington

Many behavioral functions are executed by large neural populations, and the functions of individual neurons themselves emerge from the collective action of millions of ion channels and signaling molecules. Thus, neuroscience needs a means to identify those aspects of overwhelmingly detailed biophysical states that suffice to predict behavior and use them to define simplified, "higher level" state descriptions that are more directly related to psychological or behavioral variables. In physics, statistical mechanics is used to connect massively detailed "microstates" to directly measurable parameters that define the "macrostate"; neuroscience needs a statistical mechanics that can be applied to neural populations. Unfortunately, any attempt to apply statistical mechanics to neural systems in a fully general way encounters grave difficulties. Statistical mechanics generally presupposes a system that is at or near equilibrium and relies on constants of motion (e.g., total energy, particle number, angular momentum) to define macrostates. Neural systems, on the other hand, function far from equilibrium, and most behavioral variables will be dynamic and differ dramatically from the conserved quantities that are so useful in physical applications. An alternative formulation of statistical mechanics avoids these issues. This approach, proposed by E.T. Jaynes, treats statistical mechanics as an exercise in statistical inference. This formulation, which I call "subjective statistical mechanics," drops the distinction between "physical entropy" and "information theoretic entropy"; entropy is in all cases treated as a measure of uncertainty in our knowledge of a system. The computation of maximum entropy microstate ensembles from knowledge of the macrostate is now treated as a statistical inference rather than justified by physical postulates. This interpretational shift permits the use of any well-defined macroscopic variables, and lifts the restriction to ergodic systems at equilibrium. A fully general nonequilibrium statistical mechanics can be developed by carefully treating the question as a problem in statistical inference. This generalized statistical mechanics can be directly applied to neural systems.

NEURAL CORRELATES OF OLFACTORY-GUIDED BEHAVIOR IN PIRIFORM AND ORBITOFRONTAL CORTEX

C.E. Feierstein, M.C. Quirk, Z.F.Mainen.
Watson School of Biological Sciences and Cold Spring Harbor Laboratory, Cold Spring Harbor, New York, USA.

Piriform and orbitofrontal cortices are reciprocally connected structures that have been implicated in high level olfactory processing. Both olfactory and non-olfactory (e.g. reward-related) responses have been described in single unit recordings from both regions, and thus their respective contribution to olfactory-guided behavior remains unclear.

To characterize the roles of piriform and orbitofrontal cortices in olfactory discrimination, multiple tetrode recordings were made from Long Evans rats during the performance of a two-alternative olfactory discrimination task.

Both orbitofrontal and piriform cortex neurons displayed a variety of responses during task events; even nearby cells recorded on the same tetrode showed heterogeneous correlates.

Neurons in both areas showed odor-selectivity during the odor sampling period, but odor-selectivity was more common in piriform cortex. Furthermore, odor-selectivity developed earlier in piriform cortex: In odor-selective piriform neurons, selectivity typically occurred during a rapid and transient peak in firing rate following odor onset. In contrast, in orbitofrontal cortex, selectivity usually appeared not earlier than withdrawal from the odor port, but could sometimes begin during late odor sampling.

Neurons in both piriform and orbitofrontal cortex showed choice-selectivity during the response period (moving from the odor port to one of the water ports). Analysis of error trials confirmed that activity during this period reflects the response of the rat rather than the odor identity.

Finally, a large fraction of neurons in both areas modulated their activity during the delivery of water reward. Orbitofrontal cortex showed a larger fraction of reward-responsive neurons, and orbitofrontal but not piriform neurons showed selectivity for the port in which the water was delivered.

In summary, we find evidence for more rapid and prominent development of olfactory information in piriform cortex and for richer representation of reward-related information in orbitofrontal cortex. These results suggest that piriform cortex may contain representations of all the information required to actively guide olfactory decisions in well-trained rats. This is consistent with the proposal that orbitofrontal cortex may play a supervisory rather than executive role in well-learned goal-directed behavior.

RESPONSES OF COMPLEX CELLS TO VISUAL FEATURES IN NATURAL SCENES

Gidon Felsen, Jon Touryan, and Yang Dan
Division of Neurobiology, Department of Molecular and Cell Biology; Group in Vision Science; Helen Wills Neuroscience Institute.
University of California, Berkeley, CA 94720-3200

Since sensory systems have evolved and developed in the natural world, they should be well suited to process stimuli with naturalistic characteristics. To explore the idea that cortical cells are particularly adapted to encode information in natural visual stimuli, we compared their responses to natural images and to spatially unstructured random noise.

We first estimated the receptive field (RF) of a complex cell using a modified spike-triggered correlation analysis in which the spatial structure of the RF is represented by two orthogonal eigenvectors. These eigenvectors are oriented, localized, bandpass, and in approximate spatial quadrature. We then calculated the response of the cell to every image in a natural stimulus ensemble and a random stimulus ensemble as a function of the ”projection” of the image onto the eigenvectors (dot product of the image and the eigenvector). The gain (slope) of such a projection-response function represents the sensitivity of the cell to the relevant features (those represented by the RF) embedded in the stimulus ensemble. We found that the function had a higher gain for natural scenes than for random noise.

In order to determine whether the power spectrum of natural scenes was sufficient to account for this difference in response gain, we also compared the responses to natural-power noise (a stimulus with a natural spatial power spectrum but a random phase spectrum) and random noise. We found that the response as a function of projection had similar gain for these two types of stimuli, which suggests that the phase spectrum of natural scenes plays a critical role in the responses of complex cells.

CORRECTIVE SUBMOVEMENTS IN THE MONKEY: PREDICTIVE CONTROL UNDER UNCERTAINTY

A. Fishbach, S.A. Roy, C. Bastianen, L.E. Miller, J.C. Houk
Department of Physiology, Northwestern University,
Chicago, Illinois 60611

Two macaque monkeys were trained to turn a knob rapidly and accurately to align a cursor to a target on a computer screen for juice reward. The position and velocity traces were recorded while the monkey was performing the behavioral task. Visual perturbations of the target were randomly introduced at movement onset. Target perturbations toward the cursor ("near perturbation"), target perturbations away from the cursor ("far perturbation"), and no-perturbation trials were randomly presented in equal proportion.

We introduce an algorithm for automatic decomposition of the movements into an initial primary plus one or more corrective submovements even when the latter overlap the former. The algorithm is based on the principles laid down by Novak et al. (*Exp. Brain Res.* **132**: 419-433, 2000). Our main results are:

1. The amplitude of primary and corrective submovements tend to have a constant coefficient of variation (SD/Mean). Similar linear scaling of the noise SD with the signal mean has been reported for human isometric force production (Jones et al., *J. Neurophysiol.* **88**:1533–44, 2002).
2. Primary movements and corrective submovements tend to undershoot the target. The undershoot bias was noise dependent; the undershoot size for the unperturbed trials significantly increased after target perturbations were introduced to the experimental procedure. A similar phenomenon has been demonstrated in human subjects and is consistent with stochastic optimal control theory (Engelbrecht et al., *Psychol. Sci.* **14**(3):257-61, 2003).
3. The primary movement relative error is defined as the ratio of the error at the end of the primary movement over the initial distance to the target. We show for the unperturbed trials that the onset of error correction is correlated with the expected value of the primary movement relative error. The greater the expected relative error of the primary movement the shorter the latency of the ensuing corrective submovement. Since in most of the cases the correction starts well before the primary movement ends, these findings support a stochastic predictive control mechanism.
4. We show that visual feedback from the perturbed target had a significant impact on the corrections the monkeys made only at latencies greater than 176 ms in one monkey and 220 ms for the second monkey.
5. Our findings set the stage for single unit studies of optimal predictive control in the monkey.

ENCODING OF VISUAL INFORMATION BY SIMULTANEOUSLY RECORDED NEURONS IN INFERIOR TEMPORAL CORTEX IN NATURAL SCENES

Leonardo Franco, Edmund T. Rolls and Nikolaos C. Aggelopoulos
Dept. of Experimental Psychology
University of Oxford
South Parks Road – OX1 3UD
Oxford – UK

A fundamental issue in understanding brain function is how information is encoded by populations of neurons under natural conditions. Information theoretic techniques were used to analyse recordings from simultaneously recorded neurons from the macaque inferior temporal cortex. The information was measured about which of two objects was present in a complex natural background during a visual discrimination task in which the monkey had to touch one of the objects so that we could measure whether stimulus-dependent cross-correlations between neurons were present in a natural vision task requiring segmentation against a complex background and object separation.

We found for all cases that most of the information was conveyed by the neuronal firing rates; that stimulus-dependent cross-correlations contributed only 9%; and that redundancy between the neurons was 4%. Comparison with the same stimuli presented on a plain background indicated little information loss was produced when the complex background was used, consistent with the evidence that IT neuron receptive fields become small in complex backgrounds (Rolls et al, 2003, J Neurosci 23:339-348). Indeed, it was found in the complex scene that many IT neurons responded only when the monkey looked at the effective one of the two stimuli, which were separated by 11 deg.

It is concluded that under natural vision conditions, the encoding in the inferior temporal visual cortex is quantitatively mainly in the number of spikes, with stimulus-dependent synchrony not being a major part of the code.

RECEPTIVE FIELD STRUCTURE AND SHAPE PROCESSING IN MACAQUE AREA V4

Winrich A.Freiwald[1,2,3], Doris Y.Tsao[1,2], Roger B.H.Tootell[2], Margaret S.Livingstone[1]
1. Dept. NeuroBiol., Harvard Med. Sch., Boston, MA, USA, 2. MGH, Athinoula A. Martinos Ctr. for Biomed. Imaging, Charlestown, MA, USA
3. Brain Res. Inst. and Ctr. for Advanced Imaging, Univ. of Bremen, FRG

Area V4 is an important intermediate stage in the ventral shape processing pathway. Receptive field structure of cells in early areas (V1 and V2) within this pathway can be characterized by reverse correlation techniques, while shape selectivity in higher areas (TEO and TE) has mainly been characterized by response profiles to sets of complex stimuli. Can coding in V4 be understood within the same conceptual framework as that in early cortical areas? To what extent can response selectivity for complex shapes be explained by RF substructure?
Here we used a novel reverse correlation technique (Livingstone et al., Neuron 30, 2001, 781ff.) to analyze RF substructure of neurons in cortical area V4d in the fixating macaque monkey. Cells (n=133) were stimulated with a sparse noise stimulus consisting of pairs of small squares of same or opposite contrast flashed at 60 Hz at random positions within the RF. Reverse correlation of spikes to the absolute and relative position of the squares allowed the construction of a temporal sequence of first and second order RF maps. This technique revealed RF substructure for more than 90% of V4 neurons with varying degrees of orientation tuning (determined in separate experiments). Significant 1st and 2nd order structure, both contrast-dependent and -independent, was repeatedly observed. Some of these interaction patterns and their temporal dynamics are indicative of new coding principles beyond what has so far been reported for V1 and V2. In separate experiments on a subset of cells, complex shape selectivity was characterized by another reverse correlation technique invoking the presentation of a rapid sequence of Cartesian and non-Cartesian stimuli (Gallant et al., Science 259, 1993, 100ff.). A comparison of RF substructure with shape selectivity revealed systematic relationships between the two, but indicates the necessity for inclusion of higher order terms as well.
Support Contributed By: NIH, R01 MH67529 A01, R01 EB00790 A01, EY13135, and DFG, SFB 517

MODEL REDUCTION AND OPTIMIZATION IN A DECISION MAKING TASK

J. Gao[1] and P. Holmes[1,2]
[1] Department of Mechanical and Aerospace Engineering,
[2] Program in Applied and Computational Mathematics,
Princeton University
Princeton, New Jersey 08544

The relationship among three different models of the two alternative forced choice (TAFC) task is explored. Specifically, we consider the firing rate model, the connectionist model and the drift diffusion model. The first two are abstracted models of the firing rates and *activations* of neural populations trained to respond to different stimuli. Noise is usually added to these models to account for unmodeled inputs. The third is a *formal* model of decision making suggested by Laming, Ratcliff and others to reproduce behavioral data.

We show that the connectionist and firing rate models, without noise, are identical up to linear affine variable and parameter changes. We reduce both models to a two-dimensional piecewise linear model, in which the phase plane is divided into nine tiles, in each of which the system is governed by different linear functions. We address two cases, taking into account the time-dependence of the the signal input. Based on this reduction, the dynamics described by two-dimensional sigmoidal functions collapses, via piecewise linearization and restriction to an attracting line, to a one dimensional drift diffusion problem. For stimuli with constant signal to noise ratio, this drift diffusion model is equivalent to the sequential probability ratio test (SPRT), which is optimal in the sense that the decision time needed for fixed error rate is smaller than in any other tests on noisy accumulating data. Moreover, analytical expressions for error rates, reaction time distributions and reward rates are easily found for the drift diffusion model.

We compute error rates and reaction times for the nonlinear models, the piecewise linear and pure linear models, and for the reduced drift diffusion system. We show that the results compare well for appropriate parameter choices. Our analysis strengthens the link among commonly-used neural network and decision-making models.

A NETWORK MODEL FOR THETA OSCILLATIONS IN THE SEPTO-HIPPOCAMPAL SYSTEM

Caroline Geisler and Xiao-Jing Wang
Volen Center for Complex Systems, Brandeis University, Waltham, MA 02454

During exploratory movement and REM sleep the rodent hippocampus exhibits a prominent coherent theta-frequency (3-10 Hz) rhythm, which is often temporarily nested with faster gamma frequency ($\sim$40 Hz) oscillations. Lesion studies have shown that the hippocampus theta rhythm depends critically on the integrity of the afferent inputs from the medial septum and the diagonal band of Brocca (MS) (Steward and Fox, Trends in Neurosci., 1990). The MS contains GABAergic neurons which are thought to play distinct roles of pacemakers in the generation of hippocampal theta rhythm. Specifically, it has been hypothesized that MS GABAergic afferents impose the theta rhythm in the hippocampus that in turn phasically paces the firing of pyramidal neurons. However, intrinsic mechanisms within the hippocampus may also contribute to the synchronized theta rhythmogenesis. Moreover, recent work revealed that discharges of distinct classes of hippocampal interneurons are phase locked to different phases of the theta field potential (Csicsvari J. et all. 1999; Klausberger et al., Nature, 2003). This is a particularly well defined phenomenon in which to study the interaction of different interneuron subtypes. It remains to be determined how the relative phase preferences of these interneurons are produced, and exactly how these interneurons and pyramidal cells are coordinated in the synchronous theta oscillations.

In the present study we focus on the theta and gamma rhythmogenesis and relative phase relationship between the discharge of different cell types. We present a four population model that incorporates distinct classes of conductance-based GABAergic neurons and pyramidal cells: MS GABAergic cells and hippocampal low-threshold spiking interneurons form a reciprocal loop that constitutes a pacemaker at theta frequency (Wang, J. Neurophysiol. 2002) for hippocampal pyramidal cells. A second class of hippocampal interneurons, fast-spiking interneurons, is in a reciprocal loop with pyramidal cells that generates gamma oscillations. We report modeling results on (1) the interplay between intrinsic cell properties and synaptic interactions, that lead to theta and gamma oscillations, (2) the phase relationship between the discharge probabilities of different classes of interneurons and pyramidal cells during theta oscillations and (3) how the activation of distinct neuron populations promotes oscillations at distinct frequencies. Our four population model provides a framework to understand the mechanisms which lead to synchronous activity of the pyramidal cell population in the theta and gamma frequency range.

BAYESIAN COMPOSITION SYSTEMS FOR VISION

Stuart Geman and Elie Bienenstock
Division of Applied Mathematics and Department of Neuroscience.
Brown University, Providence, RI 02912

High-level vision has proved to be remarkably resilient to computational approaches. What makes unconstrained vision so challenging is the simultaneous presence of two types of difficulties: (1) any single object can present itself in widely different conditions, giving rise to highly variable images, and (2) natural scenes are typically cluttered with objects and/or structured background. Various existing computational approaches deal successfully with each of these two difficulties taken separately, yet to date no artificial vision system is able to emulate, even remotely, the ability of brains to interpret natural visual scenes. The presence of clutter renders bottom-up segmentation ineffective, and the variability of objects makes top-down search impractical.

Consider a Complex Cell (CC) in V1. The firing of a CC signals the presence of one of several alternative stimuli (e.g. elongated bars) in the cells receptive field. Consider constructing a hierarchical model consisting of layers of **CC-like** processing units that generalize across a variety of stimulus dimensions in addition to position. In building such a system, one faces the following dilemma. If one lets the CC discard all information about the position (or other features) of the stimulus, there will be no way for neurons in an upper layer to know the relative positions (or other features) of stimuli that elicited the activity of CCs in a lower layer. Yet this information is precisely what is required to interpret scenes where parsing cannot be done on the basis of local information alone. If on the other hand one requires that information about all stimulus dimensions get transmitted to upper layers, one is led to consider cells that generalize across some dimensions but carry detailed information across others; this results in the necessity to include prohibitively large numbers of cells, and defeats the purpose of having CCs in the first place.

This paper explores the hypothesis that brains use mechanisms that allow a CC to *commit itself* to the one stimulus that elicited its activity, and propagate *some* information about this commitment to cells in upper layers. The commitment mechanism, which may use a combination of reversible STDP and coincidence detection, is not discussed here in detail. We focus on the computational advantages of using Committed Complex Cells (CCCs) in a hierarchical system for vision.

We propose a model–which we call a Bayesian Composition System–whose states are Directed Acyclic Graphs (DAGs) defined on an array of spatially and hierarchically arranged CCCs. The DAG is, essentially, the collection of commitments of all active cells in a given state of the system. We define a probability distribution on the space of all admissible states (DAGs) using a recursion on the depth of the DAG.

There are two main components in this definition: a *context-free* one, and a *content-dependent* one. The context-free component is closely related (in fact almost identical) to the usual notion of a context-free grammar. Equivalently, it results from endowing the space of activity states of a hierarchical network of classical (i.e., non-committed) complex cells with a suitable probability distribution.

Content dependency refers to the fact that the probability that two DAGs will bind with each other is allowed to depend on the "contents" of each DAG in addition to the identity (and activity level) of the cells through which binding is effected. Using the example of complex cells in V1, the binding rule for two such cells is made dependent on the respective commitment values of these two cells. In particular, it is allowed to depend on the *relative* retinal locations of the two bars that elicit the activity of the two complex cells. This makes it possible to construct invariant representations of higher-level objects, e.g., right angles.

We illustrate the use of composition systems on simple examples, and briefly discuss learning (parameter estimation) for such systems.

ACTIVITY-DEPENDENT RECOVERY FROM INACTIVATION IN A SINGLE NEURON MODEL

Gail Gilboa[1] and Naama Brenner[2]
Depts. Of Mathematics and Chemical Engineering
Technion, Israel Institute of Technology, Haifa, ISRAEL

Inactivation following high levels of neural activity is a universal property of neurons. While traditionally described as “desensitization”, implying a limitation of neural performance, recent experimental results are suggesting that this property may be an important ingredient in the adaptive behavior of neural systems. The additional time scales involved in neural inactivation and recovery are usually longer than those associated with spike-related excitability. In some cases, these time scales span a wide range, and can be themselves activity dependent. Observation of multiple time scales in recovery from inactivation in isolated patches indicate that they are an intrinsic property of ion channel kinetics; the related observation of multiple time scales in adaptation in sensory systems indicate this phenomenon may have an important functional role in signal processing.

This situation calls for a theoretical investigation of both microscopic mechanism and functional role of multiple time scales in neural systems. The theoretical study of inactivation processes has focused on mainly on the level of ion channel kinetics. Here we construct an effective model for a neuron, which is based on diffusion among a large set of inactive states of the ion channels. Neural activity is described at a coarse-grained level, emphasizing slow changes in excitability and mean activity rather than details of the action potentials themselves. The excitability level of the neuron is described by an “availability” variable, indicating the global properties of the ensemble of ion channels. This variable interacts nonlinearly with signals coming into the cell from the environment. We study the response of this effective neuron to different signals, following the dynamics of both neural activity and internal excitability.

This model exhibits several interesting features that are qualitatively consistent with recent experimental results. The most important of these is that recovery from inactivation is characterized by a time scale that depends on the history of activity. A scaling relation between stimulation duration and recovery time is found: the longer the stimulation duration, the longer is the time scale for recovery of neural excitability. This relation has the functional form of a power law, in agreement with experimental findings.

Under some stimulation protocols, the model neuron can set into a “steady state” of response, which involves some degree of inactivation with intermittent responses. This results in a locking of the response to a lower harmonic of a periodic stimulus, similar to results observed in experiments. The availability variable in this situation hovers around the threshold value for activity.

The underlying mechanism suggested here for adaptive response times is quite general: memory of past activity is effectively registered in the distribution of channels among their multiple inactivation states. The clustering of many microscopic states of a protein to form one functional (macroscopic) state is well known in ion channels as well as in receptors and other proteins. This degeneracy can provide additional functionality over longer time scales and add to the adaptive properties of the system, as shown here for the case of ion channels.

ANALYSIS OF NON STATIONARY NEURAL RESPONSES

Amir Globerson, Eran Stark, Ron Paz, Eilon Vaadia and Naftali Tishby
The Interdisciplinary Center for Neural Computation
Hebrew University, Jerusalem, Israel

Neuronal responses are dynamic in nature. Responses to a constant stimulus may have complex dynamic structure, that may take different forms: the firing rate of single neurons could change over time, correlations between groups of neurons may evolve, or their relative precise timing. All these non-stationary response properties may be stimulus dependent and be used to enhance the efficiency with which neurons encode behaviorally relevant events. For example, a neuron may fire the same number of spikes in response to two different stimuli, but do so using different rate functions, with identical integrals. Indeed spike counts were often shown to convey only a fraction of the information in spike trains.

Information in temporal structures was considered by various methods. For example, Strong *et al* (1998) used the distribution of spike patterns at various temporal resolutions and Panzeri and Treves (1998) developed a second order Taylor expansion of the information. One possible shortcoming of the *direct* method is that the relevant temporal window is too large to allow estimation of the required response distributions.

In the current work, we give an information theoretic, non parametric account of such coding. Denote the neural response over time by $R_1, \ldots, R_T$ and a behavioral events by S. Then one would like to estimate the mutual information $I(R_1, \ldots, R_T; S)$. This requires the estimation of $p(R_1, \ldots, R_T|S)$ which has an exponential number of parameters. We approximate it by assuming conditional independence of the responses given the stimulus $p(R_1, \ldots, R_T|S) = \prod_{i=1}^{T} p(R_i|S)$. This approximation is obeyed for example by inhomogeneous Poisson processes, a commonly used spike train model. Importantly, this model does not imply unconditional independence $p(R_1, \ldots, R_T) \neq \prod_{i=1}^{T} p(R_i)$, and does take into account the different temporal structure between stimuli that can be used for coding.

We applied our method to study populations of motor cortical neurons, recorded in primates during behavioral tasks. The method succeeds to identify cells with significantly high mutual information, which are not revealed using standard analysis. We find a significant increase in the number of significantly informative cells once the non stationary characteristics of the responses are taken into account. This also allows us to estimate the time scale in which non stationarity is used in encoding, which is typically on the order of tens of milliseconds.

SEQUENTIAL FLUCTUATIONS IN CHOICE AND REACTION TIME FOR A DIRECTION-DISCRIMINATION TASK.

Joshua I. Gold[1], Mark E. Mazurek[2] and Jamie D. Roitman[3]. [1]University of Pennsylvania Department of Neuroscience, Philadelphia, PA 19104. [2]University of Washington PBIO and RPRC, Seattle, WA 98195. [3]Duke University Department of Neurobiology, Durham, NC 27710.

Perceptual tasks typically involve distinct trials presented sequentially. The trials are distinct because stimulus information presented on an individual trial is sufficient to respond correctly on that trial. However, the computations used by the brain to interpret the stimulus might also use information from other trials. For example, information about the direction of visual motion appears to be encoded as a likelihood, or conditional probability of a given neural response given a particular stimulus. According to Bayes' rule, this likelihood must be combined with the *a priori* probability of that stimulus to discern accurately what was presented. We studied whether an *a priori* probability computed from the recent history of trials is used.

We trained monkeys on a reaction-time version of a two-alternative direction-discrimination task. Motion strength and direction were varied randomly trial-to-trial. In some experiments, the relative probabilities of the two directions were unequal. We examined the behavioral data for sequential correlations in choice and reaction time. Choice was analyzed by calculating the optimal linear filter between the recent history of stimulus directions and the current choice. In many experiments, a weighted average in the direction of recent correct trials and opposite to recent incorrect trials were correlated slightly with the monkey's choice. Reaction time was analyzed by subtracting out the effects of motion strength and accuracy for each session and computing the autocorrelation function for the residuals. Fluctuations in residual reaction time were small or non-existent. The results suggest that the decision variable underlying performance can take into account information from recent trials. This mechanism appears to play only a small role in the monkey's decisions but is consistent with a form of matching law, in which *a priori* probabilities are related to the rates of reward associated with each decision. Funded by NEI, NCRR, HHMI, Burroughs-Wellcome and the McKnight Foundation.

STRENGTH AND TIMING OF INHIBITION CAN EXPLAIN A VISUAL MASKING ILLUSION

Mark S. Goldman[1,2], Itsaso Olasagasti[1,2], Virginia Hafer[1], Susana Martinez-Conde[3], Stephen L. Macknik[3]
[1]Department of Physics and [2]Program in Neuroscience, Wellesley College, Wellesley, MA [3]Barrow Neurological Institute, Neurosurgery and Neurobiology Division, Phoenix, AZ

Visual masking is a phenomenon in which a target can be rendered imperceptible by a previously (forward masking) or subsequently (backward masking) presented masking object. Here we model a visual masking illusion, the standing wave of invisibility, in which a target bar of light is presented in alternation with a neighboring masking bar of light. For appropriate durations of presentation of the target and masking bars, with the mask being presented for a longer time than the target, the target bar can be rendered psychophysically invisible by the presence of the masking bar. Neuronal recordings demonstrate that the psychophysical perception of invisibility is mirrored by significantly weakened neuronal responses to the target bar in LGN neurons and, to an even greater extent, V1 neurons. Modeling of the neural basis of this illusion enables us to relate the response properties of early visual neurons to the perception of visibility. We find that a simple linear filter plus threshold nonlinearity model of the LGN neuron response fails to capture the illusion because it overly constrains the timing and strength of inhibition relative to excitation. Modification of this model to represent the anatomically separate pathways through which excitation and inhibition arrive can explain the illusion. Motivated by experiment, we propose two means by which inhibition could give rise to the observed illusory invisibility of the target. First, if inhibition is delayed relative to excitation by a time consistent with experimental recordings, then the inhibition from the presentation of the mask can cancel the excitatory drive due to the target. Second, if the amount of inhibition depends upon the duration of the mask (as found in experiments), then the longer-duration mask will inhibit the target more than the target inhibits the mask. These results hold generally for receptive fields with inhibitory and excitatory inputs that are antagonistic in time, and do not depend explicitly on spatially antagonistic mechanisms such as lateral inhibition. In combination with further experiments on dichoptic masking, we believe this work illustrates general properties arising from the timing of push-pull excitation and inhibition.

CLASSIFICATION AND MEMORY BEHAVIOUR OF MAN REVISITED BY MACHINE

Arnulf B. A. Graf, Felix A. Wichmann, Heinrich H. Bülthoff and Bernhard Schölkopf
Max Planck Institute for Biological Cybernetics
Spemannstraße 38, D-72076 Tübingen, Germany

The main aim of this study is to unravel the mechanisms responsible for classification and memorization of visual stimuli in the human brain using methods from machine learning. We perform a visual gender discrimination task of images from a human face database. Man and machine are studied using an embedded feedback loop methodology as explained below.

In a *first classification experiment* the human subjects are asked to classify faces from a subset of the database according to their gender. Their responses, i.e. estimated gender, reaction time and confidence rating are recorded. The gender estimated by the subject together with a low-dimensional representation obtained using Principal Component Analysis of the presented stimuli then form the subject's personal dataset. On this dataset, various linear separating hyperplane algorithms obtained from machine learning, Support Vector Machines, Relevance Vector Machines, mean-of-class prototype learners and K-means clustering combined with a nearest-neigbour classifier, are trained. The separating hyperplane (SH) is then computed for each of these algorithms using a possibly sparse subset of the subject's dataset: the representations. We observe that stimuli far from the SH are classified more accurately, faster and with higher confidence than those near to the SH.

A new dataset, consisting of stimuli drawn from the representations, from the dataset seen by the subject and from a subset of the database which has not been presented to the subject is created on-line. In a subsequent *memory experiment* the subject is asked to classify the stimuli from this dataset as seen or unseen in the first classification experiment. This step actually constitutes the feedback loop: the results of the subject are used by the machine to create a novel set of stimuli which are then presented to the subject. Here we assess the usefulness of the representations given by the learning algorithms by studying whether a given representation was actually useful for man during classification. In this case we assume this representation to be activated during classification and thus possibly memorized better and hence recognized when presented a second time.

Finally, in the *second classification experiment*, the subject is asked to perform a gender classification experiment identical to the first one, except that all the stimuli from the database and all the representations are presented. This experiment allows to verify the essential assumption that the internal representation of the faces in the subjects is stable and, if not, to see whether the "unstable" ones lie close to the SH of one of the algorithms.

EVIDENCE OF A REINFORCEMENT-LEARNING SIGNAL ENCODED BY SIMPLE SPIKE FIRING IN THE MACAQUE POSTERIOR LATERAL CEREBELLUM

Bradley Greger[1], Scott A. Norris[2] and W. Thomas Thach[2]
[1]Division of Biology, Caltech, Pasadena, CA, 91125
[2]DBBS, Washington Univ. Medical School, St. Louis, MO, 63110

Since David Marr originally proposed that the cerebellum subserves motor learning, several theorists have hypothesized that motor-learning signal are encoded by complex spikes while motor-control signals are encoded by simple spikes. Here we present evidence that simple spike firing was correlated with success or failure in the performance of a reaching task, and thus may also have been encoding a motor-learning signal.

Cells were recorded in the posterior-lateral region of the cerebellar cortex. In 27/89 cells there was a significant difference in the success versus failure mean firing rates during the 500ms after the reach was completed. In 8/27 cells the mean firing rate on success versus failure was also significantly different during the 500ms prior to the beginning of the subsequent trial. To differentiate whether this was a motor-error signal or a reinforcement-learning signal, reaches were parsed into three groups by distance: one group of successful trials, and two groups of failure trials. 22 cells had significant differences in firing rate across the hit-miss boundary while only 4 cells exhibited significantly different firing rates by distance. Receiver operator characteristic (ROC) analysis showed that discriminability rose rapidly around, and in some cases even before, the time that the screen was touched and could persist at a relatively high level into the subsequent trial. While recording from five cells both small and large targets were presented. This allowed for a comparison of successful and failed reaches while reaches were made to similar touch-target distances. While the monkey was reaching to similar touch-target distances, 4/5 cells exhibited significant differences in firing rate on success versus failure after the reach was completed. 2/5 cells also exhibited significant differences in mean firing rate on success versus failure rapidly after the presentation of the target.

These findings suggest that simple spikes in the posterior lateral cerebellum encoded success versus failure in the binary nature characteristic of reinforcement learning systems, and that this putative reinforcement-learning signal may have been both predictive and temporally extended. In addition to teaching signals encoded in complex spikes, the cerebellum may be making use of a reinforcement-learning signal encoded in simple spike firing to subserve a role in motor learning.

GAIN MODULATION IN FLY MOTION-SENSITIVE NEURONS BY VISUAL BACKGROUND INPUT

Jan Grewe, Anne-Kathrin Warzecha and Martin Egelhaaf
Department of Neurobiology, Bielefeld University, 33501 Bielefeld, Germany
e-mail: jan.grewe@uni-bielefeld.de

The limited operating range of sensory neurons given by their biophysical properties may not suffice to encode the range of stimulus strengths that occur under natural conditions. It is a widespread phenomenon among nervous systems that the neuronal operating range adjusts to the actual stimulus conditions. Modulation of the input-output relationship of a neuron, i.e. its gain, plays an important role in this adjustment and, thus, provides a valuable tool to increase coding efficiency. Recently, it has been proposed that gain modulation in cortical neurons might be a result of synaptic background input which, on its own, does not drive the cell in either direction but decreases the cell's input resistance and thereby reduces the sensitivity for exciting input.

Here we investigate the impact of background input on the tuning properties of wide-field direction-selective motion-sensitive neurons of blowflies that ore thought to act optic flow sensors in visual course control. The responses of these neurons can be analyzed in unanaesthetized animals at the level of integrated postsynaptic potentials and of spike trains. By stimulating the neurons with their natural sensory input, i.e. visual motion it is ensured that we are working in a behaviourally relevant activity range. In our experiments the motion input is combined with different levels of visual background noise that, on its own, does not drive the neuron, but affects several aspects of neuronal motion tuning in a pronounced way. The characteristic dependences of the neuronal input-output relationships on background activity are interpreted in terms of the opponent excitatory and inhibitory inputs of the neuron. These opponent inputs are driven in varying proportions, depending on the characteristics of the visual stimulus. According to this scheme, visual background noise, as is quite common in natural environments, activates excitatory and inhibitory inputs in an approximately balanced way, thereby increasing the conductance of the neuron and, thus, considerably affects its tuning to motion stimuli. Under natural behavioral condition, background input might, thus, shape the selectivity of motion sensitive neurons for optic flow stimuli.

Supported by the deutsche Forschungsgemeinschaft (DFG)

Object Categorization: Greater Reliance on Preferred than Nonpreferred Responses

Kalanit Grill-Spector and Nancy Kanwisher

Human object recognition is stunningly fast and accurate. What is the nature of the cortical representation of objects that underlies this ability? Haxby et al. (2001) used fMRI to demonstrate that information about object category is distributed over a broad region of the human ventral visual pathway. They further argued that non-maximal responses carry important information about object category. Specifically, even if maximal responses were excluded from the analysis (e.g., the fusiform face area: FFA), the pattern of activation across other ventral regions could be used to determine the category (in this case a face). Conversely, Haxby et al. also argued that the pattern of activation across regions that are selective for a category could be used to distinguish between non-preferred categories (e.g., the pattern across the FFA could be used to distinguish between bottles and chairs). However, this finding was not replicated by Spiridon & Kanwisher (2002) in humans, or by Tsao et al (2003) in monkeys.

The critical question, ignored in these studies, concerns not what information is *present* in the ventral visual pathway, but which of this information is *used* in object recognition. Here we addressed this question directly by imaging brain activity while subjects performed object recognition tasks at threshold, and then sorting the fMRI data by behavioral outcome. We presented stimuli briefly which were immediately masked in a rapid event-related design. In a run subjects viewed images from one category (e.g., flowers) and textures. 50% of the objects contained different images of a pre-assigned within-category target (e.g., rose) and the rest were different objects from the same category (e.g., other flower species). Subjects were instructed to indicate whether the stimulus contained: (1) the within-category target, (2) an object, or (3) a not an object. The experiment was performed using faces, birds, flowers, houses, cars and guitars.

This design enabled us to compare the magnitude of the fMRI response when subjects identified the target (an identification hit), when they detected an object but did not identify it (a detection hit), and when they did not detect the object (a detection miss). Importantly, it allowed us to examine: (1) Which regions in the ventral stream were correlated with detection or identification of these categories? (2) For a given region of interest (ROI), e.g., the FFA, do low signals for non-preferred categories correlate with success at object detection or identification?

In all 5 subjects we found regions in lateral and ventral occipito-temporal cortex that were correlated with object identification and detection. However, different patterns of activation across the ventral stream corresponded to identification (or detection) of different categories. For faces we found extensive overlap between regions involved in face detection (and identification) and the FFA. In addition, we found voxels outside the FFA that were correlated with face perception. These voxels were located in ventral regions that were either selective to faces, or responsive to both faces and objects. However, we failed to find any evidence that non-preferred regions form part of the representations used for face detection (or identification). For example, the PPA, which responds more strongly to houses and scenes than faces did not show either a positive or negative correlation with face detection or identification.

We also examined whether low signals for non-preferred categories were correlated with subjects' perceptual performance. For example, ROI analysis of the FFA (which responds maximally to faces) revealed correlation between FFA activation and face detection and identification. FFA activation was lower for birds, but correlated with bird detection and identification. Perhaps because birds have faces. However, the FFA response was not correlated (either positively or negatively) with detection or identification of other non-faces tested here. In contrast other regions in the ventral stream were correlated to detection and identification of these categories. These findings suggest that the low response to stimuli from non-preferred categories does not form part of the neural code used by subjects in detecting (or identifying) those categories.

NEURAL CORRELATES OF DECISION CRITERIA

Jack Grinband, Joy Hirsch, & Vincent P. Ferrera,
Center for Neurobiology and Behavior, Columbia University

One of the most important requirements for making accurate decisions is setting an appropriate decision criterion. To study behavioral and neural correlates of the decision criteria, we have trained subjects to categorize a set of stimuli that can vary continuously along one dimension (length of lines or speed of moving dots). One of two cues instructed the subjects on how to set their decision criterion for categorizing the stimuli as "short/long" or "slow/fast" responses. Behavioral analyses indicate that this task minimizes the effects of previous trials and stimulus specific biases, and allows us to dissociate neural activity related to categorization vs. response selection. Subjects performed the task during event-related fMRI to determine which brain areas correlate with different decision variables. There were two versions of the task: one in which the criterion was cued explicitly and one in which it was cued implicitly. We identify neural networks involved in specifying the criterion, categorizing the stimulus, and selecting the response.

ENHANCED DETECTION OF TEMPORALLY-CORRELATED SIGNALS THROUGH DOPAMINERGIC MODULATION IN A MODEL OF MEDIUM SPINY NEURONS

Aaron J. Gruber[1,2] and Sara A. Solla[1,3]
Departments of Physiol.[1], Biomed. Eng.[2], and Phys. and Astronomy[3].
Northwestern University, Chicago, Illinois 60611

The modulation of neuronal response properties by monoamines may enhance the ability of the brain to detect signals embedded in a noisy background. Many models and theoretical investigations of monoaminergic effects on neuronal processing have relied on the assumption that increased levels of monoamines result in an increased steepness of sigmoidaly shaped neuronal response functions. This sharpening of the sigmoid can increase the signal to noise ratio measure of signal detection in the output; however, this metric may not be an appropriate measure of performance since the error rate of an optimal classifier operating on the transformed output is identical to the error rate for the input. An existing theorem shows that any monotonic mapping of a signal will not alter the Bayesian error, which is the minimum classification error achievable for a temporally uncorrelated input signal.

We have developed a simple yet biophysically grounded rate model of medium spiny neurons, the predominant input neuron type of the basal ganglia. The function that relates the steady state output of the model to synaptic input has a sigmoidal shape with a steep slope in conditions of low dopamine; this monotonic relationship becomes bistable in high dopamine conditions.

The classification error of a noisy binary signal is compared to the classification error of the same signal after being mapped through either the monotonic function in low dopamine conditions or the bistable function in high dopamine conditions. The Bayesian error of the output is not altered by the induction of bistability in high dopamine conditions; this is an extension of the original theorem for a non-monotonic mapping function. A biologically plausible classification scheme based on the spike threshold of medium spiny neurons produces higher errors in conditions of high dopamine as compared to the error in low dopamine, indicating a degradation of the input signal. However, resonance effects of bistability cause a reduction in classification error as temporal correlations are introduced to the input. This reduction is not observed for the monotonic mapping in low dopamine. For moderate temporal correlations, the classification error of the output is lower in high dopamine conditions, and can even become less than the Bayesian error of the input. An analytic treatment of the model shows that it is only the width of the bistable region in relation to the statistics of the input that controls the enhancement of signal detection in conditions of high dopamine.

ACTIVITY-BASED, LOCAL UNSUPERVISED LEARNING RULES FOR COMPUTING TWO MOTION SELECTIVITIES IN POPULATION OF OPPOSING NEURONS

Ming Gu, Chunhui Mo and Christof Koch
Department of Computation and Neural Systems
California Institute of Technology, Pasadena, California 91125

Abstract

An activity-based, learning model is proposed to account for direction selectivity in neurons in visual cortex. This model is based on a local veto operation between excitation and inhibition of the shunting type, and keeps track of the calcium concentration change at excitatory synapses. Biophysical simulations demonstrate that a model cell with an unsupervised learning algorithm, based on whether or not the cell generated one or more somatic action potentials, will develop direction selectivity organically after unsupervised training. Here we investigate if the same learning principle, together with mutual inhibition between two such model cells, can give rise to selectivities to opposing directions of motion. We included an initial inhibition on each cell to account for inhibition induced by other areas of V1 that have receptive fields close to these two cells. Noise is introduced to reproduce stochastic effects during training. The time of initial inhibition is randomly selected within the first 20msec of every trial for each cell. Each cell inhibits the other through a delayed inhibitory synapse from its soma onto the other cell's soma. Our biophysical simulations show that using mutual inhibition, it is possible to generate two direction-selective cells tuned to motion in opposing directions. Nonlinear shunting inhibition acts during the initial stage of visual cortical processing to tune cells' direction selectivity of motion and to set the balance between the opposing directions.

RAPID VISUAL PROCESSING AND STDP

Guyonneau Rudy[1,2], Rufin VanRullen[1] and SimonJ.Thorpe[1,2]
[1] CerCo, Faculte de Medecine, Toulouse, France
[2] Spikenet Technology, Revel, France

The speed with which images can be analysed by the visual system constitutes a major challenge for computational models. In particular, our ability to initiate rapid behavioural responses to complex visual stimuli implies that at each stage of the visual system, processing needs to be performed in 10 milliseconds or less. How might selective neuronal responses be generated so rapidly? Here we argue that Spike-Time Dependent Plasticity (STDP) may provide a key mechanism by concentrating high synaptic weights on inputs that systematically tend to fire early.

Interestingly, in many brains areas temporal aspects of spike trains have been found to be highly reproducible. It means that multiple exposures to the same stimuli in the course of a life-time or, on a shorter timescale, a single sequence of similar processing waves induced by the rythmic activity of cortical oscillations, would thus lead to the repetition of similar input patterns. If the same input pattern is repeatedly presented to a neuron, how will STDP affect its behaviour? Here we show that the post-synaptic latency decreases and stabilizes at a minimal value while synapses transmitting the earliest reproductible afferents become selected and later ones neglected. This trend is remarkably robust: realistic conditions such as jitter and random spontaneous activity do not prevent the same trend from emerging. Moreover, not only does the neuron react faster to the input it was exposed to, it also becomes selective to it. Indeed, other input patterns either cause the neuron to fire later or not at all.

These theoretical results have critical implications for our understanding of neural information encoding and processing: that STDP naturally leads neurons to respond rapidly and selectively on the basis of the first few spikes in their afferents lends support for the idea that even complex visual recognition tasks can be performed using a single wave of spikes. To illustrate this claim, we show how neurons in the visual system can be made to be selective to particular aspects of the visual scene and how these selectivities can be used for visual representation and recognition.

ODOR PROCESSING BY DYNAMIC NEURAL CIRCUITS IN THE ZEBRAFISH OLFACTORY BULB

Christopher J. Habermann[1], Gilles Laurent[2], Rainer W. Friedrich[1]
[1]Max-Planck-Institute for medical Research, Department of biomedical Optics, Jahnstr. 29, D-69120 Heidelberg, Germany; [2]Caltech, Division of Biology, Pasadena, CA 91125, USA.

Odor information is encoded combinatorially in patterns of activity throughout the early olfactory pathway. In the olfactory bulb (OB), the principal neurons, the mitral cells (MC), receive sensory input through olfactory glomeruli and are interconnected by inhibitory interneurons, the granule cells (GCs). MCs respond to odors with electrical activity that is patterned on two time scales: (1) over hundreds of milliseconds, firing rates change in an odor-dependent manner, and (2) odor-specific subsets of MCs transiently synchronize their action potentials.

Firing rate patterns across the MC population change in an odor-dependent manner during the first few hundred milliseconds of an odor response. This process decorrelates initially similar activity patterns evoked by related odors and enhances their discriminability. Synchronization of MCs results from reciprocal interactions with GCs. In GCs, synchronized activity propagates across the OB in waves. The direction and velocity of wave propagation are variable from cycle-to-cycle; however, the average direction and velocity is reproducible and odor-invariant. MCs that phase-lock their action potentials to the GC oscillation, however, synchronize without phase lags. Hence, the complex spatio-temporal pattern of synchronized GC activity is not transmitted from the OB to higher brain regions.

In response to natural odorants at an intermediate concentration, only a subset of MCs transiently synchronized their spikes. When analyzed separately, synchronized spike patterns convey information about odor category, while the residual spikes convey information about odor identity. Hence, the neuronal circuitry in the OB extracts different stimulus features and conveys them simultaneously through separate channels, defined by the temporal structure of spiking activity. These results indicate that different messages can be conveyed simultaneously by an activity patterns when analyzed on different time scales, suggesting that encoding strategies operating on different time scales may coexist and multiplex information about complementary stimulus features in the same population of neurons.

DISINTEGRATION OF CELL ASSEMBLIES UNDER GENERAL ANESTHESIA

K.D.Harris; P.Bartho; M.B.Zugaro; L.Monconduit; S.Marguet; G.Buzsaki

Neocortical activity in vivo falls into one of two states. The “synchronized” or “inactivated” state occurs during slow-wave sleep and deep anesthesia, characterized by large amplitude slow EEG activity. The “desynchronized” or “activated” state occurs during waking and REM sleep, characterized by less low-frequency activity, but increased power in the beta and gamma bands (20-100Hz).

Under urethane anesthesia, cortical EEG is usually inactivated. However, activated EEG may be induced by stimulus such as a tail pinch, even though consciousness is not regained. EEG activation can therefore occur in three different conditions: wakefulness, REM, and anesthesia. Nevertheless, one might expect different operating modes of cortical circuits in these conditions. We contrasted patterns of population activity in these activated states, using multi-site recordings in rat neocortex.

During waking or REM sleep, individual spike trains showed an irregular structure. At the population level, a “peer prediction” analysis (Harris et al, *Nature* 424:552-6, 2003) revealed repeatedly co-activated groups of neurons, suggesting that the irregular spike trains resulted from an organization of neurons into transiently synchronous cell assemblies.

During activated periods under anesthesia, a strikingly different pattern was observed. Spike trains were more regular and often oscillatory (5-20 Hz). However, the frequencies and spike times of simultaneously recorded cells showed no discernable global coordination, suggesting that the oscillations seen in each spike train were generated by cell-specific mechanisms, rather than an ongoing global oscillation. We suggest that the lack of neuronal coordination under anesthesia is related to the absence of mental activity in the anesthetized animal.

RECURRENT INHIBITION CONTRIBUTES STRONGLY TO MEMBRANE POTENTIAL VARIABILITY DURING CORTICAL NETWORK ACTIVITY

Andrea Hasenstaub, Udo Kraushaar, Yousheng Shu, Bilal Haider, Alvaro Duque, David McCormick

Synaptic activity, both spontaneous and sensory-evoked, in cortical networks contains a wide range of frequency components from near 0 to over 100 Hz. While low-frequency components are important for determining the general amplitude and timing of the action potential response, high-frequency components are critically involved in determining the precise timing of action potential generation. We have previously reported that ferret visual and prefrontal cortical slices maintained *in vitro* in a medium that closely resembles CSF in vivo generate spontaneous UP states through recurrent excitation and inhibition within the local network. These UP states closely resemble the slow oscillation generated *in vivo* in the cortex of ferrets anesthetized with ketamine-xylazine, or during natural sleep. Here we used these recurrent network activities *in vivo* and *in vitro* to investigate the contribution of excitatory and inhibitory synaptic potentials to the different frequency components of network synaptic activity, as measured in layer 5 pyramidal cells and fast-spiking interneurons. We measured the power spectra of somatically recorded inhibitory synaptic potentials, currents, and conductances arriving during UP states *in vivo* and *in vitro*. Comparing these to the power spectra of EPSPs or EPSCs revealed that inhibitory potentials, currents, and conductances were significantly more variable and exhibited significantly more high-frequency power than excitatory potentials, currents, or conductances. Similarly, the local block of inhibition with the application of picrotoxin resulted in the loss of high-frequency components in synaptic barrages. The importance of inhibitory neurons to high frequency generation could be partly explained by the intrinsic differences between fast-spiking interneurons and pyramidal cells; interneurons, compared to pyramidal cells, had briefer absolute and relative refractory periods and better sub- and supra- threshold frequency transfer. Models indicated that the positioning of interneurons in the cortical network also contributed to their ability to generate high frequencies in recurrent activity. These results indicate that inhibitory inputs are critically involved in the generation of high frequency components of synaptic activity in cortical networks. Given that the high-frequency components are necessary for the generation of precise spike timing, these results indicate that an important role for inhibition in cortical networks is the determination of spike times.

CONFABULATION
THE BUILDING BLOCK OF THOUGHT

Robert Hecht-Nielsen
Computational Neurobiology Program
Institute for Neural Computation
ECE Department
University of California, San Diego

This talk will sketch the author's cortical theory (Hecht-Nielsen, R. **A Theory of Cerebral Cortex** (2003) Technical Report #0301, Institute for Neural Computation, University of California, San Diego (http://inc2.ucsd.edu/inc_videos/index.html), and Hecht-Nielsen, R. and McKenna, T. [Eds] (2003) **Computational Models for Neuroscience**, Springer-Verlag). First, a simple information processing architecture will be introduced to illustrate the theory's explanation for knowledge acquisition and storage. The mathematical form of this knowledge is then defined. Next, the hypothesized basic building block of thought (*confabulation*) is described, again using this architecture. The results of computer experiments with confabulation will then be presented to illustrate its efficacy, novelty tolerance, and apt characteristics. Finally, the three cortical neuronal circuits proposed by the theory will be described. Two of these are hypothesized to implement the functions of the contrived architecture. The third is the key cortical component involved in launching the *action commands* that implement the moment-by-moment control of movements and thoughts.

NEURO-KINEMATIC MODELING OF ZEBRAFISH LOCOMOTOR CONTROL SYSTEMS

Scott A. Hill, Xiao-Ping Liu, Melissa A. Borla, Donald M. O'Malley and Jorge V. José.
CIRCS & the Depts. of Physics and Biology, Northeastern University, Boston, MA, USA.

The 7-day old larval zebrafish has an extensive locomotive repertoire, including a sophisticated prey capture behavior (Borla et al., 2002). These behaviors are controlled by the approximately 300 neurons that project from brain into spinal cord (O'Malley et al., 2003). Based on models of Xenopus spinal cord (reviewed in Tunstall et al., 2002; Dale, 2003), we constructed a simple network model that uses reciprocal inhibition to produce the alternating (left-right) patterns of neural activity that underlie undulatory swimming. The first larval-zebrafish specific locomotive behaviors that we investigated were the distinct *burst* and *slow* swim patterns which differ greatly in terms of tail-beat frequency (TBF), bend amplitude and yaw (Budick and O'Malley, 2002). Burst swim TBFs range between 45 and 75 Hz, while slow swims fall in the 25 to 40 Hz range. How TBF is rapidly modulated is unknown, so we attempted to model this using a 2-cell Hodgkin-Huxley model implemented in the *NEURON* modeling program, using crossed-inhibitory and self-excitatory synapses. By varying the strengths of excitatory and inhibitory synaptic inputs, we were able to produce a range of TBFs that extended across the *slow swim-burst swim* frequency range. While some sets of synaptic strengths produced systematic changes in frequency, other values failed to produce regular rhythms. One model behavior observed was an acceleration of oscillator frequency early in the swim bout that gradually tapered off leaving the model cycling at a steady frequency. The cause of this is unclear, but extending the model to include additional features of spinal circuitry, such as greater numbers and types of spinal interneurons, and additional conductances, are being tested to determine if the observed outputs of zebrafish spinal networks can be more faithfully replicated.

The one segment model was extended to a chain of 25 segmental oscillators, representing the 25 segments of the larval spinal cord; these were coupled by descending connections. This model successfully propagated neural activity, representing the undulatory waves of swimming. It also yielded quite complex behaviors, such as the appearance of “humps” in TBF where downstream segments accelerated to frequencies above the plateau level of the lead oscillator, and then fell back to join that plateau. This behavior developed several segments caudal to the lead oscillator, with more caudal spinal segments showing successively delayed humps. Eventually, all segments joined the plateau frequency of the lead oscillator. Changes in the coupling strength between segments was used to create a more biologically relevant propagation of neural activity down the spinal model and this resulted in yet more complex patterns of neural activity. The potential impact of such neural outputs on larval locomotive behaviors was visualized by creating a 25-segment kinematic model which could receive the output of the neural model and use it to drive the bending of individual axial segments. This yielded relatively normal kinematic patterns during the stable time-epochs of the neural model. The kinematic model could alternatively be driven by a set of parameters reflecting descending control signals that might (e.g.) regulate trunk stiffness or the spatial distribution of outputs to spinal cord. Manipulating such parameters produced trunk bending patterns that mimicked patterns exhibited by larval zebrafish during swimming and during a complex turning maneuver called a J-turn (J-turns are used by zebrafish larvae to orient towards prey without hydrodynamically disturbing them). Such control parameters do not have correlates to known descending control signals, but they do provide insight into the possible set of controls that might be used to generate the larva's intricate locomotive repertoire. These models are simple in relation to the actual biomechanical and neural architectures, but as they grow in complexity and fidelity, they should develop into a set of useful tools to visualize the operation of locomotor control systems. This will constitute, in effect, an artificial spinal cord which can be used to test different descending control strategies. Ultimately we hope these models will facilitate the discovery of fundamental principles of descending motor control.

Support Contributed By: *NINDS-NS37789* (DMO) and the *Center for Interdisciplinary Research in Complex Systems (CIRCS)* at Northeastern University.

INTERACTIONS BETWEEN BLUE-YELLOW AND LUMINANCE SIGNALS IN V1

Greg Horwitz

To elucidate the principles underlying color representation in V1, we studied individual neurons in awake macaque monkeys using a reverse correlation technique. The monkeys maintained visual fixation while a computer-generated stimulus was displayed at the neuron's receptive field. The stimulus was an 8 x 8 checkerboard, each element of which changed color randomly and independently 100 times each second. Color changes were implemented by independently modulating the intensity of each of three cathode ray tube phosphors according to Gaussian distributions.

We analyzed data from this experiment in several different ways. First, we extracted the set of 20 stimulus frames immediately preceding each spike. We then averaged these spike-triggered stimuli together to derive the spike-triggered average stimulus (STA), which is the average sequence of frames that preceded a spike. This analysis revealed spatial, temporal, and spectral features of the average excitatory stimulus.

If the neuron can be well-modeled as a cascade of a linear filter and a static nonlinearity, the STA converges to the linear kernel asymptotically. Thus, if V1 neurons combine their cone inputs linearly, STAs can be used to estimate the weights with which input from each cone type is integrated. When we analyzed the data this way, we found that cone weights differed widely and unsystematically across cells, consistent with previous reports of color tuning in V1. This result may indicate that color tuning in V1 is essentially random, or that an assumption of the analysis is flawed.

Under the linear-nonlinear cascade model, the probability that a neuron will fire a spike depends only on the projection of the stimulus onto the cell's linear kernel: the neuron is blind to any orthogonal stimulus projection. We tested this assumption directly by taking the part of each spike-triggered stimulus that was orthogonal to the STA analyzing at their covariance structure. Specifically, we examined the largest eigenvector of this matrix, which the first principal component (PC1).

Many neurons in our data set had statistically significant, structured PC1s in violation of the linear-nonlinear cascade model. One group of neurons in particular, which had STAs indicating blue/yellow opponency, had PC1s consistent with a luminance modulation. This was true whether we considered PC1s in the temporal or the spatial domain. These neurons thus appear to respond to the blueness (or yellowness) of a stimulus as well as spatio-temporal luminance contrast.

To understand how these neurons combined blue/yellow and luminance signals, we considered a generalized linear-nonlinear cascade model. In this model, firing rates depended on the stimulus projection onto the plane spanned by the STA and PC1. Consideration of this model revealed that firing rates depended jointly on the blue/yellowness of the stimulus as well as its luminance contrast, irrespective of polarity. The interaction between these signals was facilitatory and roughly consistent with gain modulation.

This result has two important implications. First, many neurons in V1, including the majority of blue/yellow neurons, are poorly described by the linear-nonlinear cascade model. Linear estimates of cone weights are therefore misleading for these cells. Secondly, V1 cells combine signals relating to color an luminance in a computationally sensible manner: color signals are boosted at luminance discontinuities. This may provide an efficient means of estimating surface reflectance.

RETINAL ADAPTATION TO SPACE-TIME CORRELATIONS.

Toshihiko Hosoya[1,2] and Markus Meister[1]
[1]Harvard University, [2]Present address: RIKEN Brain Science Institute

Image processing in the retina serves for efficient transmission of strongly correlated natural visual scenes, by removing redundant components and signaling deviation from the statistical regularities. For example, in images drawn from natural environments, the light intensities at nearby points are very similar on average. Through lateral inhibition, retinal circuits amplify the difference between nearby regions and suppress components common to many neurons. Thus the responses of ganglion cells are less redundant than those of the overlying photoreceptors. Similarly in the temporal domain, natural images are mostly static. The ganglion cells have biphasic temporal receptive fields, which make them sensitive to changes in the light intensities and avoid repeated transmission of the same signal. However, in the course of daily vision, animals encounter various image environments with distinct correlations. Are the lateral interactions and the temporal processing in the retina hardwired for the average image structure, or do they adapt flexibly to the environment? Here we show that when presented with visual stimuli having various space-time correlations, e.g. orientations, the vertebrate retina modifies its stimulus selectivity by changing the spatial and temporal profiles of the ganglion cell receptive fields (Figure). These changes invariably had the effect of suppressing the redundant components and enhancing the deviation from the regularities. We suggest that the retina adapts flexibly to the statistical structure of the visual input, in a manner that detects statistical regularity and emphasizes deviations from the expected structure.

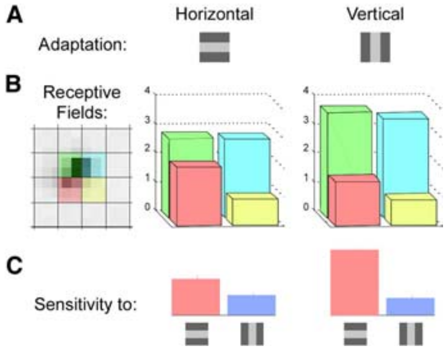

Adaptation to orientations. The salamander retina was adapted to horizontal or vertical flickering bars (A). Four regions (left in B) in the receptive field center (left in B, grayscale) showed modifications in the sensitivity under the adaptation (center and right in B). The adaptation suppressed the cell's sensitivity to the dominant orientations (C).

IN VIVO WHOLE-CELL AND CELL-ATTACHED RECORDINGS IN PRIMARY AUDITORY CORTEX IN AWAKE RESTRAINED RATS

Tomas Hromadka[1], Michael DeWeese[2] , Anthony Zador[2]
[1] Watson School of Biological Sciences, Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, NY, USA, [2] Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, NY, USA

Much of what is known about responses in auditory cortex is based on recordings obtained in anesthetized animals. However, it has recently become increasingly clear that response properties of neurons in awake animals can differ considerably. We have therefore recorded stimulus-evoked responses from well-isolated single units in the primary auditory cortex of head-restrained rats. We used cell-attached and whole-cell recording techniques in order to minimize both of the main sources of error in spike detection: failure to detect a spike (false negatives) and contamination by spikes from nearby neurons (false positives). These techniques also differ from conventional extracellular tungsten-electrode recording in their selection bias: With cell-attached recording, neurons are selected solely on the ability of the experimenter to form a gigaohm seal, rather than on the basis of spontaneous neuronal activity or responsiveness to particular stimuli.

We sampled neuronal responses across a range of cortical layers and probed the neurons with various types of stimuli, including pure tones, FM sweeps, and white-noise bursts. We find that single neurons in primary auditory cortex of awake rats display a rich repertoire of responses, ranging from binary onset responses to sustained high-frequency onset or offset responses; sometimes nearby neurons responded in very different ways.

INFORMATION ANALYSIS SHOWS GRADUAL SELECTIVITY FOR NATURAL SOUNDS IN AUDITORY PROCESSING OF SONGBIRDS

Anne Hsu, Sarah M. N. Woolley, Thane Fremouw, and Frédéric E. Theunissen

Department of Physics, Psychology, University of California Berkeley, Berkeley, CA 94720

Animals are highly capable of processing complex natural sounds. In particular, the ability to discriminate behaviorally relevant sounds, such as vocalizations in the auditory system, is necessary for social interaction and reproduction. Animals have been shown to select for natural sounds over synthetic ones. Furthermore, higher information rates have been found in the neuronal responses of frogs and insects when the spectral content and the amplitude distribution of the song is matched to that of natural vocalizations.

We tested the hypothesis that the joint temporal and spectral modulations in natural sounds lead to higher information rates in neurons from the auditory system of songbirds. To test our hypothesis, we calculated information rates obtained in single neurons of the zebra finch auditory midbrain and forebrain in response to conspecific birdsong and synthetic stimuli. The synthetic stimuli consisted of modulation-limited noise and synthetic song. The modulation-limited noise stimuli consist of all the modulations found in zebra finch song, sampled uniformly. The synthetic songs have the same spectral and temporal modulation spectra as zebra finch song, but have random phase. To calculate the information rates, we modeled each neuronal response with its estimated time-varying mean firing rate and a gamma point process. We then used this gamma process and time-varying mean rate to generate enough model spike trains to obtain estimates of the information using the direct entropy calculation methodology. The mutual information depends on spiking precision and response dynamics. We assessed the relative contribution of these factors by calculating one metric for spiking precision and two metrics for response dynamics: the entropy of the mean rate distribution and bandwidth.

Mean rates only showed selectivity for song-ripples over birdsong in the midbrain. Information analysis, however, show a gradual selectivity for features in the natural song with ascension in the auditory processing stream: Natural modulation power spectrum leads to higher information values at all stages of auditory processing. Natural phase leads to higher information at the highest level of auditory processing. The increases in information in response to the natural sounds were the result of larger bandwidth and response distributions with higher entropy. On the other hand, neural reliability was similar for all three stimuli ensembles.

A NEUROMORPHIC MODEL OF THE THALAMIC RELAY CELL

Kai Hynna and Kwabena Boahen
Department of Bioengineering.
University of Pennsylvania, Philadelphia, PA 19104

Neuromorphic engineering involves the design of silicon circuits that morph circuits within the brain. This "morphing" occurs at all levels: from intrinsic neuronal currents to synaptic circuitry to population architecture. Here we demonstrate a neuromorphic model of the thalamic relay cell, complete with a silicon low threshold calcium current. Previous models of silicon layers have focused more on the synaptic circuitry usually involving simple integrate and fire neurons or adaptive neurons. Models of more complicated intrinsic currents are few, and none attempt to model the low threshold calcium current of the thalamic cell. Our T current is inspired from common computational models: a pair of Hodgkin-Huxley variables define channel activation and inactivation dynamics. We demonstrate the voltage dependence of the neuron's firing mode as a result of the low threshold calcium current, and extract the channel dynamics from the transistor equations. By designing a biologically-realistic relay cell, our goal is to understand how the T current affects information processing in the awake thalamus, both at the single cell level and in a large thalamo-reticular network.

HOW GANGLION CELLS BECOME TUNED TO DETECT PATTERNS AT DIFFERENT SPATIAL SCALES: A TARGETED CELL-CLASS ABLATION STUDY

Adam L. Jacobs, John R. Sinclair and Sheila Nirenberg
Department of Neurobiology.
UCLA, Los Angeles, CA 90095

One of the most basic functions of retinal ganglion cells is to extract information about spatial patterns. This process takes place on multiple spatial scales, presumably because the visual scene contains biologically relevant information at different scales. For example, some cells appear to be built to extract information on a large scale, as would be necessary to pull out a large object from a busy background (e.g., a tree on a grassy hillside). Other cells appear to be built to extract information on a small scale, as would be necessary to pull out fine details (e.g., the branches on the tree).

The mechanisms that tune cells to different scales are only partially understood. Here, we tapped into a part of the circuitry involved in tuning cells specifically to large spatial patterns. Using a method for targeted cell class ablation, we eliminated a class of interneurons from the second processing layer of the retina. The ablation caused a shift in tuning away from large spatial patterns, and the shift was caused by change in the balance between the ganglion cell's receptive field center and surround. Specifically, the ablation caused a shrinkage in the size of the surround without affecting the center.

Our results show that the tuning of ganglion cells to different spatial scales is a two-step process - as opposed to the single-step process classically assumed (Marr 1982, Dowling 1987). In the first step, a basic center/surround receptive field is set up to allow ganglion cells to detect spatial patterns in general. In the second step, the receptive fields are modified so that different cells become tuned to detect different spatial scales.

Why might the retina use such an approach, with the extra complexity it entails? By using separate subcircuits to further adjust the fundamental center-surround organization common to ganglion cells, the retina may be able to ensure that each point in visual space is covered by a wide range of spatial frequency tuning. A further, intriguing possibility is that the additional circuitry gives the retina flexibility to adjust spatial tuning ”on the fly” to adapt to the changing visual scene.

References

Marr, D. *Vision.* W.H. Freeman and Company, 1982.

Dowling, J. *The retina, an approachable part of the brain.* Harvard University Press, 1987.

EXPLORING THE ATTENTIONAL MODULATION OF NEURAL CODES

Santiago Jaramillo and Barak A. Pearlmutter
NUI Maynooth, Ireland.

Several experiments have shown that neuronal activity can be modulated by attention even while the sensory stimulus is held constant. This modulation implies changes on the tuning curve (or receptive field) of the neurons involved in sensory processing. We propose an information-theoretic hypothesis for the purpose of this modulation, and show using computer simulation that the same modulation appears in a system that is optimally encoding a sensory stimulus as the system is informed about the changing relevance of different features of the input. Specifically, we present a simple model that learns a covert attention mechanism, given input patterns and tradeoff requirements. We show that, after optimization, the system gains the ability to reorganize its computational resources (or coding strategy) depending on the incoming covert attentional signal, and that the modulation of activity of the encoding units for different attentional states qualitatively matches that observed in animal selective attention experiments.

The model we propose is not confined to visual attention, or to spatial attention. Instead, due to its generality, the model can be applied to any modality, and to any attentional goal.

A spiking neural network for recognizing spatiotemporal sequences of spikes

Dezhe Z. Jin
Howard Hughes Medical Institute and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139

Sensory neurons in many brain areas spike with precise timing to stimuli with temporal structures, and encode temporally complex stimuli into spatiotemporal spikes. How the downstream neurons read out such neural code is an important unsolved problem. In this paper, we describe a decoding scheme using a spiking recurrent neural network. The network consists of excitatory neurons that form a synfire chain, and two globally inhibitory interneurons of different types that provide delayed feedforward and fast feedback inhibition, respectively. The network signals recognition of a specific spatiotemporal sequence when the last excitatory neuron down the synfire chain spikes, which happens if and only if that sequence was present in the input spike stream. The recognition scheme is invariant to variations in the intervals between input spikes within some range. The computation of the network can be mapped into that of a finite state machine. Our network provides a simple way to decode spatiotemporal spikes with diverse types of neurons.

Synchronization of inhibition as a mechanism for attentional gain modulation

Jorge V. José, P.H.E. Tiesinga, J.-M. Fellous, E. Salinas, and T.J. Sejnowski.
Center for the Interdisciplinary Research on Complex Systems & the Depts. of Physics, Northeastern University, Boston, MA, Department of Physics & Astronomy, University of North Carolina, Chapel Hill, NC, Computational Neurobiology Lab, Salk Institute, La Jolla, CA, HHMI, Salk Institute, La Jolla, CA, Department of Neurobiology and Anatomy,Wake Forest University School of Medicine, Winston-Salem, NC

Naturally occurring visual scenes contain large amounts of spatial and temporal information that are transduced into spike trains along the visual sensory pathway. Human psychophysics indicates that only a small part of that information is attended. For instance, differences between two images shown in rapid succession are not visible unless the subject pays attention to the particular spatial location of the change (Simons et al. 2000, Rensink et al. 2000.) Attention allows the selection of certain items from the stream of information processing. When monkeys focus attention on visual stimuli within the receptive field of a cortical neuron, two distinct changes are observed: there is an increase in the neuron's firing rate, or there is an increase in coherence in the gamma-frequency range (30-50 Hz) between the neuron's spikes and the local field potential (LFP), with only a small increases in firing rate (Fries et al, 2001). Attentional modulation of the synchrony of local interneuron networks could potentially account for these observations. The temporally-patterned inhibitory output of these networks modulates the response of cortical neurons to visual stimulation. To explore this hypothesis, we perform computer simulations of a Hodgkin-Huxley type neuron driven by a constant depolarizing current, I, representing visual stimulation and a modulatory inhibitory input representing the effects of attention. We determine how the neuron's firing rate and the coherence of its output spike train with the synaptic inputs modulated by the degree of synchrony of inhibitory inputs. When inhibitory synchrony was increased, the coherence of the neuron spike train with the synaptic input increased, but the firing rate either increased or remained the same. The mean number a_{IV} of inhibitory inputs is a key determinant of the shape of the firing rate versus current (f-I) curves. For large a_{IV}~ 50, f-I curves saturated for large I and an increase in input synchrony resulted in a shift of sensitivity -- the model neuron responded to weaker inputs. For small a_{IV}~ 10, the f-I curves were non saturating and an increase in input synchrony led to an increase in the gain of the response -- the firing rate in response to the same input was multiplied by an approximately constant factor. The firing rate modulation with inhibitory synchrony was highest when the input network oscillated in the gamma frequency range. The observed changes in firing rate and coherence of the model neuron are consistent with experimental observations and lend support to the hypothesis that attentional modulation results from the activity of an inhibitory network spiking at gamma frequencies.

MODELING COMPLEX CELLS IN V1 OF ALERT MONKEYS

Igor Kagan, Moshe Gur, and Max Snodderly
Division of Biology, Caltech, Pasadena, CA 91125
Department of Biomedical Engineering, Technion, Haifa 32000,Israel
Medical College of Georgia, Augusta, GA 30912

Cortical complex cells are usually described as nonlinear energy operators that sum squared outputs of quadrature pairs of linear subunits, responding to drifting sinusoidal gratings with unmodulated elevation of firing rate (F0 harmonic). However, several lines of evidence suggest that the view of complex cells as a uniform class is over-simplified, since energy models do not capture many complex cell behaviors. In alert monkeys complex cells with strongly overlapping increment and decrement regions exhibit a considerable F1 modulation, and a subset of these cells have a *relative modulation* (RM=F1/F0) >1. We have also found that most complex cells show profound dependence of *the response form* (harmonic content), and not only the amplitude, on grating parameters such as spatial and temporal frequency and size, displaying a variety of behaviors ranging from nonlinear unmodulated firing (F0) and frequency doubling (F2) to pseudolinear modulation (F1).

One of the parsimonious explanations could be that at least some of these behaviors ,e.g. F1 modulation, result from the imbalance of increment and decrement mechanisms such as incomplete spatial overlap and/or difference in amplitudes of the two regions. We tested this hypothesis using a model that approximates an *apparent* structure of complex receptive fields in our data by pooling two linear (increment and decrement) inputs with Gaussian spatial profile and same biphasic temporal response function. Model cells with various overlaps and amplitude ratios were stimulated with drifting gratings of different spatial frequencies. To quantify the measure of spatial (im)balance we computed a product of overlap index and amplitude ratio. In the model, maximal modulation increased with spatial imbalance, and the correlation for the two measures was high ($r=-0.86$, $p<0.01$). However, the model consistently yielded *lower* RM values than those in the data. Moreover, weak insignificant correlation between RM and spatial imbalance in the data for complex cells ($r=-0.15$, $p>0.01$) was inconsistent with model predictions. Thus, a static spatial imbalance of increment and decrement mechanisms cannot fully predict the presence of strong F1 harmonic in responses of complex cells.

These results and effects of temporal frequency suggest that temporal properties of input channels and possibly the dynamics of interaction between them play an important role in shaping the responses of complex cells. To account for the response diversity exhibited by complex cells, we are developing more realistic models that also include influences of the surround.

Order effects on neural representation of multisyllabic call sequences in a frontal auditory field.

Jagmeet S. Kanwal, Palle D. Prasada Rao, Zhicheng Lai and Matthew Clement

Dept. of Physiology & Biophysics, Georgetown University, 3900 Reservoir Rd, Washington, DC 20057

For auditory communication, temporal integration is an important neural mechanism for the representation and perception of complex, multisyllabic calls. To test for this function in a frontal auditory field (FAf), we recorded naturally emitted syllable sequences from a captive colony of free-flying mustached bats. Syllable sequences were digitized at a sampling rate of 250 kHz and prepared as stimuli. Using tungsten microelectrodes (~10 µm tip diameter), we recorded extracellular responses of single neurons (n=36), multiunit activity as well as evoked local field potentials (LFPs) from 27 sites in the FAf of 6 awake animals. Stimuli consisted of 3-4 second long multisyllabic sequences presented repeatedly, once every 5 seconds, at ~90 dB SPL. Syllables within sequences were presented as normal, in reversed order and as time-reversed.

Of the 32 different syllables present within 5 separate sequences, at least 16 produced a distinct response at one or more recording sites. In 61% of the cases (n=65), the response latency, amplitude and/or shape of the averaged (over 50 trials) LFP elicited by a normal syllable was modified because of order-reversal; in 64% of the cases, it was modified because of time-reversal and in 41% of the cases, it was modified because of both order- and time-reversal of the syllable sequence. In nearly 30% of the cases, this affect was specific to the re-ordering of syllables in a sequence. Single unit responses frequently matched the LFP data, but were less robust. These data suggest that the response of FAf neurons to a particular syllable in a sequence can be influenced by the syllables preceding it. FAf neurons, however, did not show a nonspecific, end-of-sequence response, suggesting that in untrained bats, temporal integration of neural activity up to the level of the FAf is most likely limited to the subsecond range.

Supported by NIH/NIDCD research grant DC02054 to JK.

DYNAMIC REPSONSES OF CA1 PYRAMIDAL NEURONS TO DENDRITIC AND SOMATIC INPUTS.

Tara Keck, Theoden I. Netoff, and John A. White
Boston University: Department of Biomedical Engineering, Center for Biodynamics, Center for Memory and Brain, Boston, MA 02215

Understanding single cell neuronal dynamics is a necessary first step in understanding dynamic spatiotemporal patterns of activity in neural networks. Neuronal dynamics can be characterized using spike time response curves (STRCs), which measure how much an artificial or controlled biological synaptic input, delivered at a particular time, advances or delays postsynaptic spikes. Under testable assumptions, STRCs can then be used to predict the degree of phase-locking in network activity of multiple, mutually coupled neurons.

Past implementations of this technique have relied upon artificial synaptic stimuli delivered at the soma. We hypothesize that the features of STRCs depend on whether the synaptic input is given at the dendrite or the soma. Using dynamic clamp techniques, we have tested this prediction in hippocampal brain slices using a picospritzer to give glutamatergic inputs to CA1 pyramidal cells.

In initial experiments, we have found that dendritic inputs lead to temporal differences in the STRC relative to the STRC generated from the somatic inputs. This difference has implications for the interactions of these cells in a small network. When a biological cell is coupled with a vitrual cell (based on experimental data) via somatic inputs, they synchronize; however, when coupled via the dendrites, the cells fire anti-synchronously. This behavior was previously predicted with simulations of a two compartment Pinsky-Rinzel neuron (Keck et. al., Society for Neuroscience Abstracts, 2003) and more complex models (Crook et. al., J Comput Neurosci, 1998).

Preliminary results indicate that while STRCs differ for dendritic and somatic inputs, they are not statistically different for inputs at varying distances up the dendrite. Based on previous studies that show inputs to the dendrites reach the soma with the same magnitude independent of input location (Magee and Cook, Nat Neurosci, 2000), this result is not surprising. Further studies will be done to examine the effects of input location with inhibitory inputs. Once these effects on the STRC have been fully characterized, we plan to investigate the mechanisms behind the differences between the dendritic and somatic STRCs.

BUMPS MOVE TO REMEMBER:
ACCOUNTING FOR THE DYNAMICS OF PERSISTENT ACTIVITY

Adam Kepecs

Cold Spring Harbor Laboratory
Cold Spring Harbor, NY 11724

Recordings from animals during working memory tasks have revealed that single neurons can fire persistently without external cues and their firing rate can be persistent at graded levels to encode stimuli. Such persistent neuronal activity is thought to be responsible for holding stimuli “on-line” available for working memory processes. However, persistent activity is rarely static and most recordings show strongly, yet systematically time-varying firing rates. Here I present a recurrent network model for graded persistent activity that shows dynamic variation in firing rates. Specifically, the model accounts for the dynamics of firing rates in the best described case of time-varying persistent activity: prefrontal recordings during a parametric working memory task reported by Romo and colleagues (1999).

A standard bump attractor is shown to be able to encode a scalar parameter with neurons that are either positively or negatively tuned to the stimuli. This requires an intensity-to-space conversion and a wide bump state. Asymmetries at the recurrent connections or in the external inputs cause systemic drift in the bump state, which can account for the temporal behavior of neurons. The model can account for both persistent as well as ramping up or ramping down activity that is either positively or negatively monotonic in the stimulus. Changing the drift speed results in a time-rescaling of ramping activity as occurs in prefrontal neurons recorded at different delay durations (Brody *et al*, 2003). This suggests that ramping activity could be used to encode time on top of stimulus parameters. Alternatively, the activity waves that occur during the delay period could serve as a slow delay-line, transferring information from one cortical column to another.

While the model is agnostic as to why persistent activity is time-varying, it makes a number of specific predictions about the noise-correlation between different neuron classes, the existence of non-stimulus coding but persistent neurons and the temporal features of firing. Importantly, the model requires few specialized mechanisms and instead builds on a wide literature supporting the notion of bump attractors and waves of activity in the recurrent networks of cortex.

DECODING A TEMPORAL POPULATION CODE

Philipp Knüsel, Reto Wyss, Peter König, Paul F.M.J. Verschure
Institute of Neuroinformatics, University/ETH Zürich
Winterthurerstr. 190, CH-8057 Zürich, Switzerland

The neuronal encoding of sensory information requires decoding by other neural structures. Encoding can involve both the spatial distribution of neuronal activity and its temporal dynamics. Decoding, in turn, has to be adapted to the encoding format. Here we investigate the encoding-decoding process in the context of the processing of visual information.

As an encoding stage we use a recently proposed model of the visual cortex that generates a, so called, Temporal Population Code (TPC) [1]. Here a cortical network, by virtue of its lateral connections, transforms a stimulus into a spatio-temporal response pattern. The temporal trace of the average population response has been shown to render an invariant and robust representation of the stimulus classes and provides a substrate for the formation of place fields.

The TPC can be read out via one continuously varying signal. Hence, for its decoding we investigate the so called liquid state machine [2]. This model of cortical microcircuits is proposed for computations on continuous temporal codes. A combination of these encoding- and decoding-networks serves to assess the properties of an optimal decoder in the context of a temporal code.

The results show that a decoding-network that continuously processes the sequence of TPC renders 35% correct classifications. However, a stimulus locked clocking of the decoding-network by resetting its state at stimulus onset leads to 66% correct classifications. A quantitative analysis of the internal states of the decoding-network shows that this noticeable difference is caused by overlaps due to the mixing of information from past and present stimuli in the internal state of the decoding-network. This raises the question how the brain solves this fundamental trade-off: While retaining a trace of the present stimulus, the mixing of information across stimuli should be avoided. Our study suggests that this trade-off is solved by resetting the decoder at stimulus onset. To test this hypothesis, we have started to relate these results to the olfactory system of insects.

[1] R. Wyss et al., Proc Natl Acad Sci USA, 100:324–9

[2] W. Maass et al., Neural Comput, 14:2531–60

LATERAL-INHIBITORY-NETWORK MODELS OF THE CENTRAL PROCESSING OF AUDITORY NERVE ACTIVITY

Jennifer Ko and Ian C. Bruce
Department of Electrical and Computer Engineering.
McMaster University, Hamilton, ON, Canada

Lateral inhibitory networks (LINs) of neurons are thought to be widely used throughout the nervous system and are known to enhance spatial edges and peaks in their input excitation patterns. It is postulated, based on experimental findings (Suga, 1995; Ramachandran et al., 1999), that lateral inhibition contributes to the central, sub-cortical, auditory processing. Previous LIN models of the central processing of auditory nerve activity were based on highly simplified phenomenological models of neural activity (Gerken, 1996; Kral and Majernik, 1996; Shamma, 1985). A more biologically realistic model of an auditory LIN has thus been developed to investigate the plausibility of such networks in these central pathways. The effects of peripheral hearing impairment on central auditory activity are also investigated. In particular, the effects on speech input and abnormal spontaneous input edges are presented.

A single layer, uniform, recurrent LIN structure was modeled. Each neuron in the LIN is described by a conductance-based, leaky integrate-and-fire model. Input spike instances were obtained from Bruce and colleagues' (2003) model of the auditory periphery for speech sounds and by using the Bernoulli approximation of a Poisson process to represent spontaneous activity from the auditory nerve.

The characteristic edge-effect of lateral inhibition that can be seen in the mean spike rates results from a spatial edge at a regional reduction of spontaneous input. The effect of neuronal parameters on the prominence of the edge-effect peak was measured. It was found that the values of the neuronal parameters must fall within a very specific and narrow range for the edge-effect to be seen. These results may have implications in the revelation of the mechanisms of tinnitus, the phantom perception of sounds.

Synthetic speech processed by Bruce and colleagues' model of the normal and impaired ear was also presented as input to this LIN. To ascertain how formants (vocal tract resonances) in the speech-driven input are affected by this LIN, spatio-temporal response patterns were observed and power ratios (Miller et al., 1997) were used to measure the synchrony of the output of the LIN to the format frequencies. It was found that the LIN significantly reduces synchrony to the formant frequencies. Abnormal speech input from the auditory periphery alters the spatio-temporal response pattern such that the neural representation of speech is degraded.

EFFECT OF PROLONGED AND RAPID ADAPTATION ON NEURONAL CORRELATION IN V1

Adam Kohn, Matthew A Smith, and J Anthony Movshon
Center for Neural Science, New York University

The responses of cortical neurons depend both on the current stimulus and on stimulus history. We have evaluated whether noise correlation and spike time synchronization between pairs of V1 neurons can also be altered by adaptation. In general, neurons with similar tuning properties tend to be more strongly correlated than those with different preferences. Since adaptation in V1 alters stimulus selectivity, we expected that adaptation at an orientation which would shift the tuning curves of two neurons apart might also reduce the strength of correlation.

We recorded from pairs of single units in the primary visual cortex of anesthetized, paralyzed macaque monkeys with multiple microelectrodes. Neurons were separated by less than 300 microns, had overlapping receptive fields, and similar orientation and spatiotemporal frequency preferences. We characterized noise correlation by calculating spike count correlation, rsc, and cross-correlograms (CCGs) for a stimulus–typically of an orientation between the preferred orientations of the two cells–that evoked a moderate response in both cells. We then adapted the cells by presenting that same grating for 40s, which reduced responsivity by roughly 50 per cent on average. Adaptation did not lead to decorrelation. rsc increased slightly but not significantly after adaptation, and the width of the CCG increased substantially with little decrease in peak height. The effect of adaptation on both noise correlation and synchronization was much like the effect of lowering stimulus contrast, suggesting that adaptation reduces the effective contrast of a test stimulus but has no other effect on neuronal correlation.

In a second set of experiments, we studied the effect of rapid adaptation by recording responses to brief (0.3 s) stimuli presented in a continuous pseudorandom sequence. The response to each stimulus in the sequence was affected by recent stimulus history. For neurons with similar tuning, the modulation of response strength due to this rapid adaptation caused a small but statistically significant increase in rsc, similar in magnitude to the effect observed after more prolonged adaptation.

Our results show that correlation between cells is not as easily altered by stimulus history as are responsivity and neuronal tuning. This suggests that stimulus history affects the gain and dynamics of individual neuron responses, but has little effect on the functional connectivity of the circuits within which the neurons live.

DYNAMICAL DIMENSION OF NEURAL TISSUE WITHIN A BRAIN-COMPUTER INTERFACE

Michael Kositsky, Ferdinando A. Mussa-Ivaldi
Department of Physiology, Northwestern University, Chicago, IL 60611

We developed a research paradigm in which the brainstem of a lamprey is interconnected in a closed-loop manner with a mobile robot or a dynamical system simulated in real time. In any brain-machine interface setting, the internal dynamics of the neural component is of critical importance. We have studied the dynamics of the neural tissue based on the interaction of this tissue with the robot. In particular, we report findings concerning the dimension of the dynamics introduced by the neural element in the behavior of the hybrid system. This information can be used for design of elaborate brain-machine interfaces. It also provides a new approach to studying various properties of neural tissue and mechanisms of neural information processing.

To assess the dynamical dimension of the neural tissue, we represent the phase-space of the hybrid system by constructing a surrogate state-vector composed of external device configurations sampled with a fixed time-step. The embedding theorem states that the dimension of a dynamical system can be deduced from the minimal number of components in the surrogate state-vector leading to trajectories without self-intersections.

The experiment consists of a series of episodes. The initial position of the external device for each episode is chosen at random. Then the entire closed-loop system is run for a certain time. Each episode produces a sequence of observations (trajectory). We used the “delta-epsilon” method (Kaplan, Physica D 73:38-48, 1994) to detect intersections in the trajectories due to an insufficient dimension for estimating the dimension of the entire hybrid system. Using this estimate, we computed the dynamical dimension of the neural element.

While in most cases it is assumed that the preparation functions as an input-to-output mapping box, our experiment yielded a very different result: four or five (for different preparations) independent dynamical components were picked up by a single electrode. Possible explanations of this finding might involve the fact that the extracellular electrode picks up spikes from four or five different axons, various mechanisms like various channel and autoreceptors dynamics or multi-synaptic connections between the stimulated and recorded populations.

HOW TO FIND DECISION-MAKERS IN NEURAL CIRCUITS?

Alexei Koulakov

Neural circuits often face the problem of classifying stimuli into discrete groups and making decisions based on such classifications. Neurons of such circuits can be distinguished according to their correlations with different features of stimulus or response. In this study we define the third class of neurons, which is responsible for making perceptual decisions. We suggest two descriptions for contribution of units to decision-making: first, as a spatial derivative of correlations between neural activity and the response; second, as an impact of variability in a given neuron on the response. These two definitions are shown to be equivalent, when they can be compared. We also suggest an experimental strategy for determining contributions to decision-making, which uses electric stimulation with time-varying random current.

BASAL GANGLIA-PROJECTING HVC NEURONS DO NOT RESPOND TO DISTORTED AUDITORY FEEDBACK IN A SINGING BIRD.

Alexay A. Kozhevnikov and Michale S. Fee
Department of Brain and Cognitive Science and McGovern Institute,
MIT, Cambridge, MA 02139

The songbird has emerged as an important model system for studying the generation and learning of neural sequences. The bird's extremely stereotyped vocal behavior is implemented by a discrete and anatomically well-characterized neural circuit. Brain area HVC (used as a proper name) is uniquely situated in the song control circuit. It is a part of descending motor pathway, which is necessary for song production and, at the same time, is the main source of neural input to the basal ganglia circuit necessary for song learning. The connectivity of HVC implies that it might play a crucial role in both song production and song learning. We study the activity of identified neurons in the nucleus HVC of zebra finches during singing. We identify HVC neurons projecting to areas RA and X and also putative HVC interneurons by in vivo antidromic stimulation. The 3 types of HVC neurons have remarkably distinct firing patterns during singing. HVC interneurons (n=49, 6 birds) are active throughout the song and display tonic patterns of activity. Projection neurons display highly phasic stereotypical firing patterns: X-projecting neurons burst on average 1.6 times per motive (n=103, 6 birds) while RA-projecting neurons (n=18, 4 birds) burst extremely sparsely - at most once per motif. The bursts of projection HVC neurons are tightly locked to the song and frequently have the rms jitter of less than 1 ms.

It is necessary for the bird to evaluate its song performance in order to learn the song as a juvenile and to maintain the song through adulthood. One of the important and, as yet, unanswered questions is how the bird evaluates its vocal output. It has been proposed that the activity of X-projecting neurons during singing might encode auditory information used for evaluation of the song during learning. We recorded the activity of X-projecting HVC neurons in the presence of distorted auditory feedback. Exposure of the birds to distorted auditory feedback does not cause immediate changes of the song, but causes a gradual degradation of the song. Both in young birds (¡90 days old, 3 birds) and in adult birds (2 birds), the firing patterns of X-projecting neurons did not change in the presence of distorted auditory feedback. Our results are inconsistent with the idea that X-projecting neurons provide an auditory-related signal to the basal ganglia pathway during singing. The activity of these neurons during singing is purely premotor-related and may provide timing information to basal ganglia circuit similar to that transmitted to premotor pathway by RA-projecting neurons.

SPEED ADAPTATION IN MACAQUE MT

Bart Krekelberg[1]. Richard. J.A. van Wezel[2], Thomas. D. Albright[1]
[1] The Salk Institute, Howard Hughes Medical Institute, La Jolla,USA.
2. Helmholtz Institute, Utrecht University, The Netherlands.

Prolonged exposure to a moving stimulus changes speed perception in two salient ways. First, the speed of a subsequently presented stimulus is underestimated. Second, the ability to discriminate speeds is improved. This can be interpreted as a sacrifice in accuracy of absolute speed perception in favor of the ability to discriminate speeds. We investigated this perceptual phenomenon and its neural correlates in the macaque middle-temporal area.

We used patches of dots both as adaptation and test stimuli. One patch with dots moving in the cell's preferred direction was positioned in a cell's receptive field for 2s. After this adaptation and a brief (0.5s) blank period, a test stimulus was shown (5 test speeds: 80%, 40% or 0% slower or faster than the adaptation speed). In a control condition, the dots were stationary during the adaptation period. We recorded from 83 cells in two monkeys who fixated a central dot.

Adaptation reduced the firing rate of nearly all cells. Assuming a code in which higher firing rates correspond to higher perceived speeds, this could explain the underestimation of perceived speed after adaptation. An ideal observer analysis showed that - on average - the adapted stimuli should appear 10% slower than the non-adapted stimuli.

The reduction in firing rate was not constant for all test speeds; it was strongest when the test speed was most different from the cell's preferred speed. This implies that the population speed-tuning curve became narrower. Such a narrower tuning curve may explain why motion discrimination improves with adaptation.

Taken together these data show that activity in area MT is the neural correlate of the perceptual phenomenon of speed adaptation. More generally, it shows how adaptation can lead to changes in the neural responses that are detrimental (mismatched absolute speeds) and beneficial (improved discriminability) at the same time.

TRANSMITTING MORE THAN SPIKES

Jutta Kretzberg[1,2], Martin Egelhaaf[3], William Kristan[2], Terrence Sejnowski[1,2]
1: The Salk Institute, La Jolla, California 92037
2: University of California San Diego, La Jolla, California 92037
3: Bielefeld University, Bielefeld, Germany

The discussion about neural coding and synaptic transmission is usually focused on spikes. However, some neurons do not spike at all. The amount of continuously released transmitter is determined by their membrane potential. Depending on, e.g., time constants and noise levels the analog signal of a graded synapse could transmit more information than a binary spike-triggered synapse. Moreover, some neurons use a combination of spikes and graded voltage. In these neurons both modes of information transmission appear to code for different aspects of the neuronal signal.

In the visual system of the fly a population of large field motion sensitive neurons consists of three types of neurons: spiking, graded and “mixed”. Even though all of these neurons are well characterized, the properties of their input cells remain unknown. We developed a computational model for synaptic transmission that can simulate the continuum between spiking, mixed and non-spiking transmission. The kind of transmission is charactirized by only two parameters of a sigmoidal transfer function, the slope and the midpoint. Applying this model to experimental data from the fly, we tried to predict which kind of synaptic transmission the input elements of the fly’s motion-sensitive neurons are using. The transfer function which reproduced experimental data best has a steep slope like it is typical for spike-triggered synapses but it is centered around the presynaptic resting potential like in the case of graded transmission.

The second system we analyzed is the nervous system of the leech, consisting of several ganglia. Some well-characterized neurons connect to cells within the same ganglion as well as to neurons in adjacent ganglia. For these cells synaptic transmission within one ganglion does not require the production of spikes. However, spikes are needed to transmit information to the next ganglion. Simultaneous intracellular recordings reveal the transfer functions of graded and spike-triggered synaptic transmission by a single presynaptic neuron. We investigate whether both signal types are employed for different aspects of the information contained in the neuron’s membrane potential and how the transfer functions change activity-dependent over time.

Supported by the Sloan foundation and the Swartz foundation.

REPRESENTATION AND PROCESSING OF NOVEL STIMULI BY AN ASSOCIATIVE NEURAL NETWORK

Reimer Kühn and V Del Prete
Department of Mathematics
King's College London, Strand WC2R2LS London UK

Cortical areas are constantly processing thousands of new stimuli arriving from the external world, and their representation by the activity of neurons is thereby continuously modified. In some cases previous associations learnt between two stimuli are disrupted or replaced by new associations (e.g. stimulus-reinforcer association learning and reversal, in experiments of classical conditioning). It is therefore of interest to estimate the amount of information which a population of neurons, trained to represent a set of pre-learnt stimuli, can carry about a set of novel stimuli, which are presented to the net. In order to investigate this issue, we consider a network of graded response neurons storing a given number of patterns, which represent the learnt stimuli. They are embedded via a Hebbian covariance learning rule. Novel stimuli are modeled as external inputs to the network, and are endowed with a varying degree of correlation with the learnt patterns. For such a setup, we can calculate the distribution of the firing rates across the population of neurons, as the novel stimuli are presented as input to the network, and hence the mutual information between the firing rate distribution, and the novel stimuli. We examine the dependence of the mutual information (i) on the degree of correlation between novel and previously learnt stimuli as well on other parameters of the network, such as (ii) the loading level, (iii) the sparseness of the representation of previously learnt stimuli, and (iv) the average activity of the network. Although the graded response neuron model is admittedly a somewhat poor approximation to the real dynamics of neurons in the brain, our approach does have the advantage that it allows us to derive most of our results analytically. We are currently examining extensions of this analysis to more realistic models and to continuous patterns as representing the learnt stimuli.

Partially funded by EU contract QLGA-CT-2001-51056.

A Biophysical Model of Graded Persistent Activity in a Single Neuron

Mauktik Kulkarni and Kechen Zhang

Departments of Neuroscience and Biomedical Engineering
Johns Hopkins University School of Medicine
720 Rutland Ave, Traylor Building 426
Baltimore, Maryland 21205
Emails: mkulkar1@jhmi.edu, kechen@bme.jhu.edu

Neurons in various brain regions demonstrate graded persistent activity: the ability to sustain firing activity at multiple firing rates for seconds or, in some cases, even several minutes after transient external input has disappeared. Recurrent feedback loop in a neural network is one mechanism that could give rise to this phenomenon. Alternative mechanisms may involve only a single neuron, as suggested by a recent electrophysiological study of entorhinal cortical neurons that exhibited graded persistent activity after blocking all synaptic transmissions. This persistent activity is dependent on a particular type of calcium-dependent non-specific cationic channels. Models of persistent activity in single neurons have been studied by assuming variable stationary positions of calcium wavefronts in the dendrites or hysteresis in the activation of the cationic current. We propose a Hodgkin-Huxley type model of a single neuron that exhibits sustained firing activity using only fast conductances, incorporating high-threshold calcium channels and calcium-dependent non-specific channels. With proper choice of parameters, the dynamics of the influx and decay of intracellular calcium allow its concentration to act as a short-term memory that approximately integrates the external input. The model may provide a possible biophysical mechanism for a single neuron to hold sensory information for up to several seconds during working memory tasks.

NEURAL MICROCIRCUITS WITH ADAPTING SYNAPSES

Eduard Kuriscak and Petr Marsalek
Department of Patological Physiology,
Charles University of Prague, First Medical Faculty.
U nemocnice 5, Praha 2, CZ 128 53, Czech Republic

Basic building blocks of mammalian neocortex consist of the pyramidal excitatory cells connected to the inhibitory interneurons. We modeled a neuronal network of these elements. We used one point neurons with the Hodgkin-Huxley equations adapted for neocortical neurons. The f-I curve of these equations is close to the linear function. It is linearized by the presence of the K^+ AHP current. The Na^+, K^+, and Ca^{++} currents are chosen with the aim to reach the compromise between the biological complexity and numerical tractability of the network. Cells are randomly connected. We study, what is the effect of replacing the static, conventional synapses by dynamic, adaptive synapses, described in last years. Differential equations describing both neurons and synapses are solved by a custom written Runge-Kutta method differential equation solver. One property which cannot be modeled solely by differential equations are delays. Therefore we also implemented delays as one of the features toward modeling more realistic network. As output variables we observe: mean firing rate in the network, synchronous firing of groups of neurons and the degree of their synchronization and also the development of state variables in dynamic synapses. The aim of the simulation is to keep local and global parameters of the network close to the experimentally reported values.

We ask, how do networks with static and dynamic synapses differ. We found that dynamic synapses keep network firing with plausible outputs for larger range of sensory-like inputs in comparison to static synapses. We found that bursting is a network property, not dependent on the type of synapse.

In the current state of the project we investigate following questions:
1) Dynamical synapse can reproduce the LTP/LTD, shown originally in slices. We also demonstrate the LTP/LTD in one neuron. We go beyond the slices experiment and show, what parameters of dynamic synapses are important for inducing the LTP/LTD. We also show, what is a typical distribution of potentiated and depressed synapses during the network activity close to the *in vivo* situation. 2) Nonlinear synapses and delays enable more storage capacity than classical synapses and no delays in network. We again compare setups with and without these properties. We show, how selected parameters of the network, which represent biophysical quantities, can be tracked back into fewer parameters describing the memory in the network.

MAXIMUM LIKELIHOOD ESTIMATION OF A STOCHASTIC INTEGRATE-AND-FIRE NEURAL ENCODING MODEL

Jonathan W. Pillow, Liam Paninski, and Eero P. Simoncelli
Howard Hughes Medical Institute, New York University, New York

We examine a cascade encoding model for neural response in which a linear filtering stage is followed by a noisy, leaky, integrate-and-fire spike generation mechanism.This model provides a biophysically more realistic alternative to models based on Poisson (memoryless) spike generation, and can effectively reproduce a variety of spiking behaviors seen *in vivo*.We construct the maximum likelihood estimator for the model parameters, given only extracellular spike train responses (not intracellular voltage data).Specifically, we prove that the log likelihood function is concave and thus has an essentially unique global maximum that can be found using gradient ascent techniques. We develop an efficient algorithm for computing the maximum likelihood solution, demonstrate the effectiveness of the resulting estimator with numerical simulations, and discuss a method of testing the model's validity using time-rescaling and density evolution techniques.Finally, we will show preliminary applications of the analysis to both in vitro and in vivo data.

UNEXPECTED SPATIO-TEMPORAL STRUCTURE IN V1 SIMPLE AND COMPLEX CELLS REVEALED BY SPIKE-TRIGGERED COVARIANCE

Nicole C Rust, Odelia Schwartz*, Eero P Simoncelli, and J Anthony Movshon
Center for Neural Science, New York University and *Salk Institute

V1 neurons are commonly classed as simple and complex based upon their responses to drifting sinusoidal gratings. The periodic response of a simple cell is typically modeled using a single linear filter followed by a half-squaring nonlinearity. Phase-insensitivity in complex cells arises in the motion-energy model by squaring the output of two linear filters in quadrature. In fact, one observes a continuum of response patterns to sinusoidal gratings in V1 cells. We performed a spike-triggered characterization of macaque V1 cells with a range of phase-sensitivities to estimate the number and type of linear filters involved in the generation of these cell's responses.

We stimulated neurons with a dense, random, binary bar stimulus confined to the classical receptive field. For each cell, we recovered a set of linear filters describing the generation of the neuron's response by calculating the spike-triggered average (STA) and significant axes revealed by applying a principal components analysis to the spike-triggered covariance matrix (STC). Assuming a linear-nonlinear-poisson (LNP) model, the number of recovered filters in this analysis sets a lower bound on the number of filters the cell uses in performing its computations.

The results revealed by this analysis predict the continuum of modulation in response to gratings across the population of V1 neurons. In simple cells we recovered an STA with clear spatio-temporal structure in addition to at least one additional filter. The additional filter tended to differ from the STA by a phase shift and decreased the modulation of the cell's response. However, in the more modulated simple cells, the weight of this filter was weak relative to the STA and thus had only a small effect on response. For less modulated simple cells, the weight of this additional filter increased. For complex cells, the STA weakened and two full-rectified, quaderature phase filters were revealed by the STC, as predicted by the energy model.

In complex cells, our analysis often recovered more than one filter pair. While the first pair of filters always had clear spatio-temporal structure in the middle of the receptive field, additional filter pairs tended to have more structure near the receptive field's edges. These additional pairs thus describe the spatio-temporal tuning along the receptive field fringe. These filters were not the product of eye movements, deviations from Poisson spiking, or the particular stimulus used for the characterization. In the context of the LNP model, the presence of additional filter pairs beyond the first suggests the existence of multiple spatially-shifted subunits in complex cells.

EXCITATORY AND INHIBITORY SYNAPTIC INPUTS TO DIRECTION-SELECTIVE SIMPLE CELLS OF CAT VISUAL CORTEX

Nicholas J. Priebe and David Ferster
Department of Neurobiology and Physiology, Northwestern University, Evanston, IL, USA

The receptive field properties of neurons in the visual pathway change as visual information moves from the retina through the lateral geniculate nucleus (LGN) to primary visual cortex (V1). Whereas the spatial receptive fields of LGN neurons are circular, containing opposing center and surround fields, the spatial receptive fields in V1 contain elongations that endow them with sensitivity to the orientation of edges. As with the processing of spatial information, the processing of temporal information differs between the LGN and V1. While LGN neurons in the cat are not selective for object direction (Cleland and Levick, 1974), many V1 neurons respond selectively to the direction of moving objects (Movshon, 1975).

Two fundamentally different models have emerged to explain cortical direction selectivity, the basis of which is a comparison across both space and time of luminance or contrast cues. A comparison over time and space could emerge in the V1 responses from the specificity of LGN input. Because of the heterogeneity of response latencies of LGN cells, grouped into lagged and non-lagged cells, a specific arrangement of LGN inputs onto a cortical neuron could endow the cortical neuron with direction selectivity (Saul and Humphrey, 1990; Saul and Humphrey, 1992; Jagadeesh et al., 1993). Alternatively, the emergence of direction selectivity could be a strictly cortical phenomenon based upon inhibition between neurons of different spatial receptive fields (Torre and Poggio, 1978). Regardless of the different principles by which may be computed, most of the models of direction selectivity require a nonlinear step to sharpen selectivity. For direction selectivity based solely on the LGN input, this nonlinearity is generally thought to be the spike threshold, since it appears that the voltage modulations of directional cells are largely linear. However, the presence of shunting inhibition (Grzywacz and Koch, 1987) has raised the possibility that it may indeed be cortical inhibition which endows neurons with direction selectivity (BorgGraham et al., 1998; Hirsch et al., 1998; Anderson et al., 1999). Cortical inhibition would suppress a neuron from firing when a stimulus moves in the non-preferred direction, but would allow responses to a stimulus in the preferred direction (Torre and Poggio, 1978).

We have explored the role of cortical inhibition by extracting the excitatory and inhibitory components of synaptic inputs from intracellular records obtained *in vivo*. We find that both excitation and inhibition are tuned for the same direction, but differ in their relative timing. Using whole-cell patch recordings, we recorded the voltage response of simple cells to gratings moving in both the preferred and null directions as well as to bright and dark bars flashed randomly within the neuron's receptive field (1-D noise). Excitatory and inhibitory inputs were distinguished by repeating the stimulus while injecting hyperpolarizing current of different amplitudes into the cell. Preferred-direction motion of the grating stimulus elicited both the largest excitatory and largest inhibitory input onto simple cells. While the inhibition and excitation onto simple cells were greatest to grating motion in the preferred direction, these inputs were temporally out of phase. To our flashing stimuli, simple cells exhibited spatiotemporal orientation in their receptive field maps derived from the voltage responses. Estimates of the excitatory and inhibitory inputs derived from the receptive field maps were tuned for the same direction, although the timings of the peak excitation and inhibition did not overlap. Therefore, just as the spatial push-pull mechanism operates for orientation selectivity, our inhibitory and excitatory maps reveal a similar push-pull mechanism, extended into time and space, for direction selectivity.

INTRACELLULAR ANALYSES OF RECEPTIVE FIELD CLASS AND LAMINAR POSITION IN THE CAT'S VISUAL CORTEX.

Luis M. Martinez[1], Friedrich T. Sommer[2], Qingbo Wang[3], Judith A. Hirsch[3]; [1] Dept. Medicina, Univ. A Coruña, 15006 Spain, [2]RNI, Menlo Park, CA 90025, [3]Biol.Sci. USC, Los Angeles, CA 90089.

We are interested in understanding the synaptic basis of neural response in the primary visual cortex. Thus, we make whole-cell recordings *in vivo* with dye-filled electrodes from the different cortical layers and map responses to sparse noise, individually flashed bright and dark spots (Jones & Palmer, 1987, *J Neurophys* **58**:1187). The immediate goal of the current study was twofold. First, we aimed to determine the intracellular signature of traditional receptive field types in the visual cortex. Second, we wished to correlate response pattern with anatomical station in the cortical microcircuit.

Cells in the primary visual cortex fall into two main physiological classes, simple and complex. Simple receptive fields, as first described, are built of parallel and adjacent subregions in which stimuli of the opposite contrast evoke responses of the opposite sign—*push-pull*. Complex receptive fields, originally defined by exclusion, lack segregated subregions. We evaluated the receptive fields of cells in layers 2+3, 4, 5 and 6 for three parameters that provide indices of simpleness and complexness. First, we used a push-pull index to compare the relative weight and distribution of excitation, push, and inhibition, pull, within each subregion of the receptive field; the index ranges from 0, for perfect push—pull, to 2 for perfect push—push. Second, we used an overlap index (modified from Schiller et al., *J Neurophys* **39**:1288) to gauge the spatial segregation of component subregions; values of the index range from ≤ 0 for separated subregions to 1 for concentric subregions. Last, we measured the number of segregated *on* and *off* subregions in a given cell's receptive field.

Cells whose receptive fields had scores that indicated simpleness—i.e. small push-pull indices, multiple subregions and small overlap indices—were only found in thalamorecipient zones, layer 4 and upper layer 6 or bordering regions. This clustering held true for excitatory and inhibitory cells alike. Cells with scores that indicated complexness were found in all layers, though the specific profiles varied with laminar location. For example, complex cells in thalamorecpient zones responded well to both bright and dark spots throughout the receptive field. Thus, their push-pull and overlap indices were near the maximum. Cells in regions further removed from thalamus, layers 2+3 and 5, rarely responded to both bright and dark stimuli and so had push-pull index values near 1.

In sum, our finding is that receptive field class varies systematically with cortical depth. Simple cells are restricted to the thalamocortical stage of processing while the response profile of complex cells varies as a function of position in the microcircuit. Supported by NIH EY09593.

NEURAL CODING OF TIME-VARYING SIGNALS IN AUDITORY CORTEX

Xiaoqin Wang, Thomas Lu and Li Liang
Department of Biomedical Engineering, Johns Hopkins University School of Medicine, Baltimore, MD 21205

The neural representation of time-varying signals in the auditory cortex is of special interest to our understanding of mechanisms underlying speech processing. Time-varying signals are fundamental components of communication sounds such as human speech and animal vocalizations, as well as musical sounds. Low-frequency modulations are important for speech perception and melody recognition, while higher-frequency modulations produce other types of sensations such as pitch and roughness. Both humans and animals are capable of perceiving the information contained in temporally modulated sounds across a wide range of time scales from millisecond to tens and hundreds of milliseconds. Compared with the visual or somatosensory system, the auditory system has much longer ascending pathway leading to the cortex. A consequence of such a long serial processing pathway is that the temporal precision of stimulus-synchronized spikes is progressively reduced due to synaptic integrations at each stage. It has long been noticed that neurons in auditory cortex do not faithfully follow rapidly changing stimulus components. Studies in both anesthetized and awake animals have shown that cortical neurons can only be synchronized to temporal modulations at a rate less than ~100 Hz, compared with a limit of >1 kHz at the periphery auditory system. The lack of stimulus-synchronized cortical responses to rapid, but perceivable temporal modulations has been puzzling.

Our recent studies showed that auditory cortex solves this problem by performing temporal-to-rate transformations. We have showed that auditory cortical neurons could implicitly code rapid temporal modulations of sounds by the magnitude of their average discharge rate, whereas slow temporal modulations of sounds were explicitly coded by stimulus-synchronized temporal discharge patterns of neurons in auditory cortex. The firing rate-based cortical representations were also found to be capable of encoding fine temporal variations in complex sounds. Moreover, cortical neurons appear to extract temporal profiles embedded in complex sounds, regardless how they are generated (in amplitude or frequency domain). We further showed that auditory cortical neurons mark sparse acoustic events (including stimulus onset time) with precise spike timing and transform rapidly occurring acoustic events into firing rate-based representations. The temporal integration window varies between cortical neurons, suggesting multiple time scales in temporal processing. We suggest that the "slow-down" of temporal response rates in auditory cortex is for the purpose of multi-sensory integration. Other sensory systems (visual, tactile) are much slower at the periphery, but discharge synchrony rates are similar across sensory cortical areas. The functional importance of the firing rate-based cortical representations is that they are the results of temporal-to-rate transformations and therefore represent processed instead of preserved information at this level of the auditory system.

SIMULTANEOUS MULTIPLE RECALL WITH SPIKING NEURONS

Sebastian A. Wills and David J.C. Mackay
Cavendish Laboratory, University of Cambridge,
Madingley Rd, Cambridge, CB3 0HE, UK
`{saw27,mackay}@mrao.cam.ac.uk`

A spiking neural network architecture capable of storing temporal memories and concurrently recalling multiple memories is presented. The network consists of identical spiking neurons. Inputs to each neuron are grouped; a group becomes activated if s out of the S inputs in that group simultaneously receive a spike. The neuron fires when a certain number d of groups become active. With suitable choice of parameters, the neuron therefore approximates a binary sigma-pi function. The connections between neurons have a range of (possibly noisy) propagation delays, which are set during training of the network.

A memory is a temporal pattern of spikes, each spike defined by the identity of the neuron generating it and the time at which it occurs. The network is able to recall memories in two ways. Firstly, it can 'fill in' the missing spikes when an incomplete memory is played to the network. Secondly, if the network is trained on periodic patterns, valid memories are repeatedly replayed in a stable fashion. Simulations show that a 1000-neuron network (with $s=S=3$, $d=1$), using memories consisting of 50 spikes over a 100ms period, is capable of retrieving any one of more than 80000 stored memories when stimulated with only the first six spikes in the recall pattern.

The network is also capable of retrieving multiple superposed memories. For example, the same 1000-neuron network trained on 1000 memories can simultaneously recall any desired 7 of the 1000 memories.

LINEAR MODELS OF AUDITORY TUNING IN SONGBIRD NEURONS REVEAL EFFICIENT CODING OF NATURAL SOUNDS

Sarah M. N. Woolley, Thane Fremouw and Frédéric E. Theunissen

Department of Psychology, University of California Berkeley, Berkeley, CA 94720

Songbirds are advantageous model systems in which to study auditory processing because: 1) normal auditory processing leads to the learning, discrimination, and identification of acoustically complex vocalizations; and 2) songbird brains contain well defined auditory processing regions. Tuning to natural sounds has been proposed as a mechanism for the selective attention to and acquisition of meaningful communication sounds such as conspecific song. We have hypothesized that songbird auditory neurons may be tuned to the spectral and temporal modulations characteristic of their own species' songs.

We have used linear models called Spectro-temporal Receptive Fields (STRFs) and the Fourier transforms of those STRFs, Modulation Transfer Functions (MTFs), to characterize the tuning properties of single neurons in zebra finch auditory system. The STRF gives the linear time/frequency tuning of a neuron and the MTF shows the spectral and temporal modulations to which a neuron responds best. The modulation tuning of neurons can then be compared to the modulations that are represented in classes of sounds such as zebra finch song by computing a modulation power spectrum for each sound class and comparing them with the MTFs obtained from neural responses.

The linear auditory tuning properties of single neurons in the zebra finch midbrain region, the Mesencephalicus Lateralis, dorsalis (MLd), primary auditory forebrain region, Field L, and the secondary forebrain region, the Caudal Mesopallium (CM) during the processing of song and other synthetic stimuli designed to match some but not all of the acoustic features of song have now been characterized. Analysis of STRFs and MTFs has revealed that populations of auditory neurons whiten the modulation spectrum of natural sounds and in particular of vocalizations. Neural tuning for temporal modulations emphasizes the modulations with high frequency and de-emphasizes the very low frequency modulations. The low modulation frequencies are common in all natural sounds and have large power. Therefore, they provide little information in terms of discriminating and identifying sounds. Finally, an information theoretic analysis shows that this coding for natural sounds leads to higher information rates to natural sounds than to synthetic sounds with flat modulation power spectra.

COMPUTATIONAL NEUROINFORMATICS: ALGORITHMS FOR INFORMATION TRANSMISSION IN CORTICAL NEURONAL NETWORKS

Daniel Gardner[1] and Jonathan D. Victor[2], [1]Laboratory of Neuroinformatics, Department of Physiology & Biophysics, and [2]Department of Neurology and Neuroscience, Weill Medical College of Cornell University, New York, NY 10021

We announce a new collaborative thrust, funded by the NIH's Human Brain Project, to develop computational neuroinformatics techniques and bring them to bear on the fundamental problem of neural coding, the information content of spike trains and other neuronal signals. Current analyses and models are often tested on datasets from one preparation or protocol, limiting the ability to test theories and distinguish general from domain-specific. To remove limits due to restricted exchange of algorithms or datasets, we are enabling a wide range of project-developed and user-submitted parallelized computational algorithms to be applied to a correspondingly broad assemblage of neural datasets collected from multiple cortical areas, many protocols, and several preparations. The project synthesizes:

- a suite of parallelized data-driven information-based analytical algorithms, with
- a powerful multiprocessor computational array, each linked to
- open, public databases that acquire, aggregate, and archive single- and multi-unit neuronal data from visual and somatosensory cortex, and other vertebrate and invertebrate networks collected by multiple labs, techniques, data formats and protocols.

Together, these methods will aid elucidation of general or universal methods used by nervous systems for encoding and processing of information. Information content will be characterized using a range of complementary approaches including direct information-theoretic methods, metric space methods, embedding methods, spectro-temporal analysis, and other methods, with hypothesis-testing facilitated by procedures for surrogate data generation.

Accessible via neurodatabase.org, our neurodatabase system currently allows users to download actual datasets and descriptive metadata; using Java-based tools, XML-wrapped data can be viewed on expanded or contracted time scales, or acquired for further analysis. Community participation is encouraged, with open access to datasets, to algorithm utilization and design, to data submission, and to computational resources. To support the new NIH Data Sharing initiative, and encourage submission of both data and algorithms, intellectual property rights are retained by submitters, and database access requires users to acknowledge these rights. We invite additional collaborations—in neuroinformatics, experimental neuroscience, computational neuroscience, and computer science—to advance our understanding of neural coding. Such collaborations may also advance the design and execution of protocols for the study of visual and somatic sensation that produce neural data —including multiunit data—that focus on hypothesis testing using these evolving algorithms.

The Human Brain Project supports research and development of advanced technologies through cooperative efforts among neuroscientists and information scientists. A coordinated effort of fifteen federal institutes across four federal agencies (NIH/NSF/NASA/DOE), it currently funds 36 grants, including two to our Lab. A major goal is a federated World Wide Web based information management system for neuroscience, linking interoperable databases and data management tools to provide a better understanding of the structure and function of the brain.

Supported by Human Brain Project/Neuroinformatics grants MH068012 from NIMH, NINDS, NIA, NIBIB, and NSF, by MH057153 from NIMH and NINDS, and by SBIR MH060538.

INTRINSIC REPRESENTATION: BOOTSTRAPPING SYMBOLS FROM EXPERIENCE

Stephen David Larson
MIT Computer Science and Artificial Intelligence Laboratory

One of the important pursuits of computational neuroscience is the development of an understanding of the neural correlates of mental activity. In order to reach this understanding, we find it useful to develop mathematical and computational models that attempt to explain how the brain represents information. To that end, I have developed a model of an adaptive knowledge representation scheme, inspired by the function of the cerebral cortex, which provides a new perspective on how we may explain some of the brain's informational properties.

My model consists of a hierarchical system composed of three main layers. On the bottom layer of the hierarchy, a self-organizing map capturing sensory data distributes images learned by the system according to their statistical similarity. One layer above, a clustering algorithm divides the map space into clusters of statistical similarity. Finally, above that, these clusters are considered to be units that can undergo a form of Hebbian association with other clusters. Additionally, separate instances of this hierarchy can be set in concert to gather input from multiple sensory information spaces, enabling a form of cross-modal learning.

Using this architecture, my implemented system builds descriptions up from low-level perceptual information and, without supervision, discovers regularities in that information. Then, my system, with supervision, can associate the regularity with "symbolic tags". Experiments conducted with the implementation shows that it successfully learns symbols corresponding to blocks in a simple 2D blocks world, and learns to associate the position of its eye with the position of its arm.

I have implemented a computer program for my model that serves as an existence proof that the model can be instantiated, demonstrated that the program is capable of learning symbols for blocks in a blocks world, and demonstrated that the program is capable of associating the movements of its eye with the movements of its arm.

MODELLING HOW MONKEYS LEARN FROM THEIR DECISIONS

Brian Lau and Paul W. Glimcher
Center for Neural Science
New York University, New York, NY 10003

We trained monkeys to perform a free-choice task based on Herrnstein's Matching Law to determine how choices vary according to the probability and magnitude of reward. Offset of a central target cued the monkeys to choose by making a saccadic eye movement to one of two targets for which rewards were scheduled at different independent rates. If a reward was scheduled for the target the monkey chose, fluid reward was given. If a reward had been scheduled for the target not chosen, the reward remained available until that target was next chosen. In a daily session we varied the probability or magnitude of the rewards associated with each target over 3-8 blocks of 100-250 trials. The log of the ratio of choices was linearly related to the log of the ratio of reward probabilities and magnitudes, with slopes of $\approx$0.6 and $\approx$1 for probability and magnitude respectively.

We modelled trial-by-trial behavior with reinforcement learning algorithms. The basic model had two stages. The first stage used a reward prediction error based on dopamine neuron responses to determine how obtained rewards were weighted to estimate the current utilities or preferences for each choice (e.g., Rescorla & Wagner 1972). The second stage used a noisy winner-take-all rule to map differences in utility into probabilities of choosing the targets. This basic model fit the choice data well with two free parameters, capturing steady-state behavior and fluctuations that occurred at block transitions. We elaborated this model to determine whether: 1) the exponential shape of the reward weighting predicted by the Rescorla-Wagner rule was sufficient, 2) utilities were updated for choices *not* taken, 3) differences or ratios of utility better characterized choice, 4) monkeys' choices exhibited structure independent of that predicted by reward history. We fit the models by maximum likelihood, and used Akaike's Information Criterion to select the most parsimonious model. We found that: 1) reward weighting was close to exponential, but that the most recent reward had extra weight, 2) utilities were updated for choices not taken (treated as though not rewarded), 3) differences characterized choices slightly better than ratios and 4) the only structure (in addition to that predicted by reward history) was a tendency to alternate.

We assessed the goodness-of-fit of these models by measuring their predictive accuracy using the receiver-operating characteristic. We found that the best-fitting model predicted $> 80\%$ of the monkeys' choices, suggesting that simple reinforcement learning algorithms accurately characterize behavior in our free-choice task.

Supported by: EY-010536, DOD NDSEG Fellowship

BAYESIAN MODELLING OF VESTIBULAR ILLUSIONS

Jean Laurens and Jacques Droulez
Laboratoire de la Physiologie de la Perception et de l'Action.
11 place Marcelin Berthelot, 75005 Paris, France

Complex self-motion stimulations in the dark can be powerfully disorienting and create illusory motion percepts. Modelling the underlying cognitive processes provide a tool to explore the bases of vestibular perception.

In the absence of visual cues, the CNS has to use rotational and linear acceleration information provided by the vestibular canals and the otoliths respectively. These sensory signals are imprecise because of sensors and neural noise, a critical issue when the acceleration signals have to be integrated to estimate head velocity and position. Moreover, linear acceleration sensors, such as the otoliths, report the sum of acceleration and gravity, the gravito-inertial force (GIF). Consequently head tilt as well as acceleration stimulate the otoliths.

Consistent with these observations, experimental evidences have revealed biases and inaccuracies of motion perception in the dark. During sustained rotation, the percept of velocities fades away, due to a low pass filtering by both the vestibular canals and the neural system. Similarly, sustained linear acceleration are misinterpreted as head tilt. These phenomenons are commonly explained by the rare occurence of sustained rotations or accelerations in ''natural'' situations. However, in other conditions, canal and otoliths informations can be merged to generate accurate estimates of head tilt.

To understand how vestibular cues are combined and integrated, we designed a Bayesian model of vestibular acceleration processing. This model transcribes in a probabilistic framework the following assumptions : 1) A geometric internal model of motion in the space is used, allowing canal-otolith signal fusion. 2) Canal's responses to acceleration are noisy. 3) Low rotational velocities and translational acceleration are more likely to occur. 4) Gravity magnitude is known with low precision (i.e. somewhere around 9.8 $m.s^{-2}$). 5) When the motion estimate is too implausible, it can be randomly reset.

By maximizing the posterior probability of self-motions given the sensory inputs, we successfully predicted human subject's motion estimates in a variety of experimental conditions, such as those described above, rotations at the end of a centrifuge arm (OAR) and rotations around a tilted axis (OVAR). Thus we show that a wide range of psychophysical results can be explained by humans acting as ideal observer of a noisy vestibular signals, with an imprecise knowledge of gravity magnitude, in a context where low velocities and accelerations are more likely.

EFFECTS OF NOISE ON THE TUNING OF RESPONSES IN V1

Thomas Lauritzen (1,2) and Ken Miller (1)
(1) Graduate Group in Biophysics, Department of Physiology, Sloan-Swartz Center, W.M. Keck Center, University of California, San Francisco, CA 94143.
(2) Redwood Neuroscience Institute, Menlo Park, CA 94025.

The orientation selectivity of the spiking of cells in V1 remains constant with contrast, a phenomenon known as contrast-invariant orientation tuning. Recent results from the Ferster group reveal that the combination of voltage noise and contrast-invariant tuning of the voltage yields contrast-invariant spike responses (Anderson et al., 2000). We have previously shown that the voltage fluctuations in the Ferster group's raw data are reasonably described by an Ornstein-Uhlenbeck process and, in a single-cell model, that the contrast-invariant spike tuning is only achieved when the inhibition received by a cell balances or dominates the excitatory input to it (SFN2002, 456,1). This inhibition is required to cancel the untuned mean (DC) component of the LGN input and thus allow voltage tuning to be contrast-invariant. However, in a single-cell model the voltage DC did not show the orientation tuning seen in Ferster's data. We speculated that intracortical excitation in a network model would yield the orientation-tuned, contrast-invariant DC voltage response.

Here we include the in vivo noise level in a large-scale network model previously introduced (Lauritzen and Miller, 2003). We show that the network connectivity can provide the contrast-invariant voltage tuning, including an orientation-tuned voltage DC, and verify that this yields contrast-invariant spike tuning. Further, the circular variance of the cells decreases with contrast as found by Shapley et al. (SFN2002, 720.6).

We then address the effect of correlation of noise in the network and find that the contrast-invariance of orientation tuning can be stable even for large correlations between the membrane potential fluctuations of individual cells. Further, increased noise correlations decrease the spontaneous activity of the cells, while they increase the gain as well as shift the temporal-frequency tuning of the cells in the presence of visual stimuli.

SPATIO-TEMPORAL MODELS FOR TIME ENCODING AND STIMULUS RECOVERY

Aurel A. Lazar
Department of Electrical Engineering
Columbia University, New York, NY 10027

A key question arising in sensory systems is how to represent an arbitrary stimulus as a sequence of action potentials. The temporal requirements imposed on this representation might dependent on the information presented to the sensory neurons. For example, the temporal precision of auditory processing involves measurements of interaural time delays with sub millisecond accuracy. Rapid intensity transients appear to be a key stimulus feature for triggering precisely timed spikes. All sensory systems use ensembles of neurons to encode information but direct experimental insights into the operation of biological neural networks is scarce.

In [7] the question of stimulus (signal) representation was formulated as one of time encoding, i.e., as one of encoding amplitude information into a time sequence. Formally, a time encoding of a bandlimited function $x(t), t \in \mathbb{R}$, is a representation of $x(t)$ as a sequence of strictly increasing times $(t_k), k \in \mathbb{Z}$.

There are two natural requirements that a time encoding mechanism should satisfy. The first is that the encoding should be implemented as a real-time *asynchronous circuit*. Secondly, the (ideal) encoding mechanism should be *invertible*, that is, the amplitude information can be recovered from the time sequence with arbitrary accuracy. A Time Encoding Machine (TEM) is the realization of such an encoding mechanism.

The first example of a Time Encoding Machine satisfying the above requirements was given [7]. It consists of a feedback loop that contains an adder, a linear filter and a noninverting Schmitt trigger. The invertibility property of the TEM is due to a representation of the bandlimited function $x(t)$, $t \in \mathbb{R}$, as a discrete set of integral values $\int_{t_k}^{t_{k+1}} x(u)du$ evaluated on time intervals that satisfy in average the Nyquist rate. Hence, under simple conditions, bandlimited stimuli encoded with the Time Encoding Machine can be recovered loss-free from the time sequence at its output. A Time Decoding Machine (TDM) is the realization of an algorithm for stimulus recovery with arbitrary accuracy.

The average rate of $(t_k)_{k\in\mathbb{Z}}$ in [7] is proportional with the bandwidth of the stimulus. Clearly the output of a neuron can not support large spike averages and a natural physical limit has to be imposed in modelling of sensory neurons. A number of sensory systems, including the retina and the cochlea have been modelled as a bank of filters, with each of the filters feeding a signal into an integrate-and-fire neuron. Such a model is a possible realization of time encoding using filter banks and integrate-and-fire neurons. This model raises a number of theoretical questions. The ones addressed in this paper investigate whether (i) the ideal encoding mechanism is invertible and if so, what algorithm achieves perfect stimulus recovery, and (ii) what is the best estimate of the stimulus when the TEM is noisy.

Using the theory of frames [3], we shall derive a channelization of the bandwidth of the stimulus that leads to a multidimensional time encoding representation of the stimulus $(t_k^m)_{k\in\mathbb{Z}}, 1 \leq m \leq M$, where M is the number of channels. By choosing M, the transfer function of the filters describing the filter bank, and the parameters of the integrate-and-fire neuron model, the average spike rate at the output of each integrate-and-fire neurons can be controlled. We shall demonstrate that time encoding based on filter banks and integrate-and-fire neurons provides, under certain natural conditions, an equivalent representation of information, i.e., under ideal conditions the stimulus $(x(t))$, $t \in \mathbb{R}$, can be recovered loss-free from its multidimensional spike train representation (t_k^m), $k \in \mathbb{Z}$ and $m = 1, 2, ..., M$. If the m'th integrate-and-fire neuron is characterized through an arbitrary, possibly time varying threshold $\delta^m(t), k \in \mathbb{Z}$ and $m = 1, 2..., M$ that satisfies a simple upper bound condition the stimulus can be perfectly recovered if $(t_k, \delta^m(t_k^m))$ is known for all $k \in \mathbb{Z}$ and $m = 1, 2..., M$.

Finally, we shall provide estimation algorithms for stimulus under the following constraints: (i) noisy input, (ii) non-zero refractory period, (iii) random threshold, (iv) noisy spike train.

TRACKING RECEPTIVE FIELD MODULATION DURING NATURAL STIMULATION

Nicholas A. Lesica and Garrett B. Stanley
Division of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138

Traditional approaches to characterizing the stimulus/response mapping in sensory systems make a number of simplifying assumptions: 1) the stimulus is stationary and uncorrelated, 2) the mapping does not change over time, and 3) the response of the neuron depends only on the stimulus and is independent from one interval to the next. However, recent studies of sensory function under natural stimulus conditions have demonstrated important features of neural encoding that are in violation of these assumptions. Characterizing the stimulus/response mapping in a natural setting demands a more realistic model of sensory encoding in which stimuli of arbitrary complexity are adaptively filtered into a history dependent neural response. To identify the stimulus/response mapping in this context, a new analytic technique must be developed. In this poster, we detail a point process extended recursive least-squares (ERLS) approach to receptive field estimation. Using simulated and experimental neural responses, we demonstrate the ability of the ERLS technique to: 1) estimate RFs from responses to complex natural stimuli, 2) track adaptation of receptive field properties during a single trial, and 3) capture the behavior of a neuron more accurately by including history dependence in a point process response model. The ERLS technique lends tremendous flexibility to experimental design, which is essential for the investigation of sensory function in the natural environment. The retinal and LGN data used in this poster have been published in previous studies [2, 1].

References:

[1] S. A. Baccus and M. Meister. Fast and slow contrast adaptation in retinal circuitry. *Neuro*n, 36:909–919, 2002.
[2] G. B. Stanley, F. F. Li, and Y. Dan. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *J. Neurosci*., 19(18):8036–8042, 1999.

VIRTUAL CONDITIONING OF RAT PRIMARY MOTOR CORTICAL ENSEMBLE ACTIVITY

Lei. Li; Shaohua Xu; Joseph T. Francis; J.K. Chapin

Dept Physiology & Pharmacology, SUNY Health Science Center,
Brooklyn, NY, 11203-2098 USA

Previous studies have shown that primary motor cortical single units can be dissociated from movement by conditioning. More recently, studies have indicated that neural population (NP) functions extracted from neural ensemble recordings can replace overt movement as the operant behavior. In this study, we are exploring explicitly whether a reward (water or electrical stimulation in the medial forebrain bundle, MFB) can shape ensemble motor cortical activity, and further, whether arbitrary neural activity patterns can be conditioned. Rats were trained to press a lever for MFB stimulation or a water drop. During this schedule, neural activity was recorded through chronically implanted multi-electrode arrays in the motor cortex. We then developed a computer program to transform these multineuron spike recordings into neural population functions and use these NP functions directly as rewarding criterion instead of bar-pressing behavior. Preliminary data shows that rewards can shape motor cortical activity. In addition, under NP function rewarding scheme, neurons appear to fire synchronously and generate higher NP functions. Our study may help resolve the long-standing question about whether the motor cortex is a low-level executive that is obligatorily tied to specific overt movements, or alternatively can operate at a more abstracted level. On a more practical level, such motor cortical adaptability would enable paralyzed patients to control neuroprosthetic devices using their brain activity alone.

CROSSNETS: POSSIBLE HIGH-PERFORMANCE NEUROMORPHIC NETWORKS

Konstantin Likharev, Jung Hoon Lee, Xiaolong Ma, and Özgür Türel
Stony Brook University, Stony Brook, NY 11794-3800

The past two years evidenced a spectacular breakthrough in molecular electronics – see, e.g., Ref. 1. These exciting developments give a hope for the transfer of electronic industry, within the next 10 to 20 years, from the currently dominating semiconductor-transistor ("CMOS") technology to "CMOL" integrated circuits combining an advanced CMOS sub-system with a layer of molecular devices self-assembled on nanowire fabric.[2] Such circuits may eventually combine unprecedented areal density (above 10^{12} cm^{-2}) of active devices with high operation speed, at acceptable power consumption. Thus for the first time there is a prospect of creating integrated circuits with density exceeding that of the cerebral cortex, at much higher speed and (hopefully) comparable functionality.

Our group is working on a family of neuromorphic architectures based on localized crossbar-type networks ("CrossNets")[3, 4] that allow to reach very high cell connectivity (e.g., 10^4 typical for the neocortex) in quasi-2D CMOL circuits. In CrossNets, a CMOS sub-system is used for the implementation of cell bodies (somas) and possible rare and fast long-range communications, while simple molecular ("latching switches") are used for synaptic functions. Previously we had shown[4] that such networks can operate as Hopfield networks. The objective of this talk would be to review our work and report three major new developments.

First, we have demonstrated that CrossNets operating as Hopfield networks may be very defect-tolerant: in some cases the 99% pattern recognition fidelity could be reached with almost 90% of bad molecular devices. Second, we have shown how to use such essentially binary devices as molecular latching switches to reproduce very closely the performance of neural-network pattern classifiers with continuous weights. Finally (probably most importantly), we have understood how to implement these operations using two-terminal rather than three-terminal latching switches. This means a considerable simplification of the possible structure of the molecules and a dramatic increase of the possible yield of their chemically-directed self-assembly.

We will also discuss in brief more distant prospects of CrossNet development, that may eventually allow to put an artificial analog of the human cerebral cortex, capable of advanced information processing and (hopefully) self-development at a speed some four orders of magnitude higher than its biological prototype, on a ~30x30 cm^2 silicon wafer.

1. J. R. Heath and M. Ratner, Phys. Today **56**, 43 (2003).
2. K. Likharev, in: J. Greer *et al.* (eds.), *Nano and Giga Challenges in Microelectronics* (Elsevier, Amsterdam, 2003), pp. 27-68.
3. Ö. Türel and K. Likharev, Int. J. of Circuit Theory and Applications **31**, 37 (2003).
4. Ö. Türel, I. Muckra, and K. Likharev, in: *Proc. of the Int. Joint Conf. on Neural Networks* (Portland, OR, July 2003), pp. 365-370.

GLOBAL FOREBRAIN DYNAMICS PREDICT ANIMAL BEHAVIORAL STATES AND THEIR TRANSITIONS

Shih-Chieh Lin*[1], Damien Gervasoni*[1], Sidarta Ribeiro*[1], Ernesto S. Soares[1], Janaina Pantoja[1], Miguel A.L. Nicolelis[1, 2, 3, 4]

[1] Department of Neurobiology, Duke University Medical Center, Durham, NC, USA
[2] Department of Biomedical Engineering, Duke University, Durham, NC, USA
[3] Department of Psychological Brain Sciences, Duke University, Durham, NC, USA
[4] Duke University Center for Neuro-Engineering, Duke University, Durham, NC, USA
* These authors contributed equally to the work.

Despite the ubiquity of behavioral state alternation among higher vertebrates, the physiological mechanisms underlying the spontaneous succession of these states remain poorly understood. As a first step to address this issue, we investigated the relationship between the different behavioral states and the concurrent local field potentials (LFP) recorded in multiple forebrain areas of rats. A 3D state-space, defined by two amplitude ratios and a coherence measurement, allowed all global rat brain states to be clearly identified in all subjects as distinct clusters within a dynamic structure formed by the combined activity of cortical, hippocampal, striatal, and thalamic LFP. This simple spectral structure, solely based on neural signal, provides an objective classification of all major brain states (93-100% accuracy). Within this structure, state transitions follow specific spectral trajectories with characteristic duration, spectral path and coherence bandwidth. Most state transitions involve striking changes in coherence spectra, with synchronization peaks that coincide with the dominant frequencies of underlying large-amplitude LFP oscillations. Synchronization within neuronal ensembles and coherence among interconnected brain areas have been proposed to underlie the integration of behavior. Given the strong phase-locking between neuronal firing and LFP oscillations, transient increases in synaptic input coherence may functionally connect distributed neuronal populations at the boundaries of different state-specific neural processing modes. We propose that such increased neural coherence may function as a “handshake protocol” within and across brain areas, synchronizing neuronal populations to facilitate the global exchange of neural information during successive states of the wake-sleep cycle.

HOW LOCAL IS THE “LOCAL FIELD POTENTIAL”?

Jing Liu and William T. Newsome, HHMI and Dept. of Neurobiology, Fairchild D209, Stanford University, Stanford, CA 94305-5125

The local field potential (LFP), generally defined as a bandpass filtered (2-200 Hz) version of the analog voltage recorded from a microelectrode, is being used increasingly by neurophysiologists to assess neural activity within a local volume of cortical tissue. Because low frequency electrical signals are capable of traveling substantial distances through the brain (up to several centimeters in the case of the EEG), the LFP may lead to erroneous conclusions unless the spatial characteristics of the signal are well understood.

To address this issue, we recorded LFP, multiunit (MU), and in some cases single unit signals from cortical area MT while alert monkeys performed a fixation task and a speed discrimination task. Our LFP signals were initially bandpass filtered (3-200 Hz) using hardware filters (Krohn-Hite) with a typical sampling frequency of 1 kHz. We used software filtering *post hoc* to analyze the LFP signal within specific frequency bands and to remove excessive 60 Hz noise when necessary. The amplitude of the LFP was considered to be the root-mean-square noise of the filtered analog trace during the stimulus presentation interval.

Our first goal was to assess the spatial selectivity of the LFP signal by comparing quantitatively the direction, speed and disparity tuning curves derived simultaneously from the LFP and MU signals. It is well known that MU signals (bandpass filtered, 100-6kHz) recorded from MT are well tuned for direction, disparity and speed (DeAngelis & Newsome, ’99; Liu & Newsome, ’03). MU tuning curves are only slightly less selective than those of single units recorded simultaneously, presumably reflecting the spatially clustered (or columnar) organization of MT neurons tuned for these properties. We analyzed LFP data from 3 to 100Hz, and found that LFP signals are also well tuned for direction, disparity and speed, *but only for frequencies >30 Hz*. The preferred direction, speed and disparity were highly correlated between simultaneously recorded LFP and MU signals. Tuning was substantially weaker or nonexistent for signals <30 Hz. We conclude that the LFP, if filtered properly, fairly represents local activity on the spatial scale of cortical columns.

We then asked whether LFP signals, filtered from 30-200 Hz, reflect local fluctuations in sensory signals that are correlated with trial-to-trial variation in perceptual decisions made near psychophysical threshold. We measured “choice probabilities” with psychophysical and physiological data obtained while monkeys performed a 2AFC speed discrimination task. Choice probabilities were statistically significant and equal in magnitude for MU and LFP (3 to 100Hz) signals. Similar to the tuning curves described above, however, choice probabilities were absent from the LFP signal at frequencies <30 Hz. We conclude that the LFP, appropriately filtered, is indeed reasonably local and can be used to probe the relationship between local neural signals and perceptual decisions.

(Supported by HHMI and by NEI (05603)).

COMPUTATIONAL NEUROETHOLOGY IN MOUSE AUDITION

Robert C. Liu and Christoph E. Schreiner
Sloan-Swartz Center for Theoretical Neurobiology
University of California, San Francisco 94143-0444

Ecological theories hypothesize that an organism's neural coding strategy reflects the statistics and regularities of the natural world that it experiences. Evidence for this view is building in both vision (Simoncelli and Olshausen, 2001) and audition (Lewicki, 2002). Further progress could be made by developing new models to explore the neural representation of natural stimuli.

We have developed a computational neuroethological framework to address this issue in the mouse auditory system. Neural coding is studied using a stimulus that naturally elicits an observable behavioral response. After characterizing the statistical properties of this natural stimulus (etho- component), neural areas are examined for their ability to represent these characteristics (neuro-). Coding algorithms are then evaluated (compu-), using behavior as a constraint. This provides a systematic approach to the question of how the neural code represents the statistics of behaviorally-relevant sensory stimuli.

We are carrying out this paradigm in the context of mouse ultrasound (> 25 kHz) communication. These vocalizations are produced in at least two behavioral situations. First, the calls of isolated mouse pups elicit retrievals to the nest by their mothers. Second, ultrasounds by adult males encountering females seem to attract the latter and reduce its aggressiveness. We have previously shown that the calls from these two contexts fall into two acoustically distinct spectral and temporal categories, laying a possible foundation for their categorical identification and discrimination (Liu *et al.*, 2003). Recently, we have demonstrated that auditory cortical neurons in mothers appear particularly well matched to a specific temporal statistic of pup calls: the repetition period between successive vocalizations. This is reflected first by the ability of the neural population to follow pup calls up to (and not far beyond) the dominant vocal repetition rate; and second, by oscillations in neural firing that have periods matching the distribution of natural vocalization periods.

Most recently, we have begun to explore the representation of artificial pup calls synthesized from bandpass filtered ultrasound noise. Mothers categorically perceive these as pup-like for only a limited range of narrow bandwidths (< 22.5 kHz) (Ehret and Haack, 1982). Beyond this sharp categorical bandwidth border, they do not exhibit a natural preference for larger bandwidth model ultrasounds. Our preliminary data indicates that auditory cortical neurons abruptly increase their activity with a small change in bandwidth near this border, possibly establishing a representation that facilitates the discrimination of ultrasound categories downstream of auditory cortex.

RHYTHMS SYNCHRONIZED BY ELECTRICAL SYNAPSES IN THE THALAMIC RETICULAR NUCLEUS

Michael A. Long, Carole E. Landisman, and Barry W. Connors.
Dept. of Neuroscience, Brown University, Providence, RI, USA.

Neurons of the thalamic reticular nucleus (TRN) generate synchronized rhythms under certain conditions, using mechanisms that are only partly understood. We recently found that some TRN neurons are interconnected by electrical synapses; here we tested the possibility that electrical synapses mediate rhythmic synchrony in rat TRN slices. We measured the incidence and strength of electrical coupling in pairs of TRN neurons and found that both were a steep function of intersomatic distance. Electrical coupling was robust when cells were closely spaced, but coupling was absent at distances >40 mm. When metabotropic glutamate receptors (mGluRs) were activated with a bath-applied agonist or by an endogenous ligand released during tetanic stimulation, vigorous spiking occured in all neurons tested, similar to the response seen in neocortical LTS interneurons. About 40% of TRN neurons displayed robust mGluR-induced rhythms of subthreshold membrane potential, with a mean frequency of about 10 Hz. Such rhythmic TRN neurons had intrinsic membrane properties that differed from nonrhythmic cells. In the absence of fast chemical synaptic transmission, subthreshold rhythms were only well synchronized between closely spaced, electrically coupled pairs; rhythms in noncoupled cells were not synchronized. The results suggest that electrical coupling can coordinate spindle-frequency rhythms among small, localized clusters of TRN cells.

Present address for M. Long: Dept. of Brain and Cognitive Sciences, MIT, Cambridge, MA, 02139, USA.

Bayesian inference in auditory-visual segregation and integration

Wei Ji Ma, Ladan Shams

Like most animals, humans are endowed with multiple sensory modalities. While this attribute tremendously enriches our perception of the environment, it also poses the nervous system, at any instant, with the eminent problem of whether and how to integrate the signals from the different modalities. Temporally or spatially coincident signals in different sensory modalities may or may not originate from the same source in the world. The nervous system is constantly faced, therefore, with the problem of estimating which sensory signals have been caused by the same source and should be integrated, and which have been caused by different sources and should be segregated. Here we show that human auditory- visual perception is surprisingly consistent with an ideal observer (derived using Bayes' rule), indicating that combining and segregating auditory and visual signals is statistically optimal. These results provide the first unifying account for the entire spectrum of cue combination, ranging from no integration, to partial interactions, to complete fusion. Our findings also show that a previously reported auditory-visual illusion known as sound-induced flash illusion, is not a perceptual aberration or oddity, and is in fact consistent with a general, statistically optimal strategy used by the brain for combining auditory and visual information.

A PUSH-PULL MODEL OF PREFRONTAL CORTEX DURING A SEQUENTIAL DISCRIMINATION TASK

Christian Machens and Carlos Brody
Cold Spring Harbor Laboratory
1 Bungtown Rd., Freeman Bldg.,
Cold Spring Harbor, NY 11724
machens@cshl.edu, brody@cshl.edu

Sequential discrimination tasks are widely used in psychophysical studies. In a typical such task, a subject is presented with a first stimulus (f1), and then, after a delay of a few seconds, with a second stimulus (f2), after which the subject must make a decision based on a comparison of the two (f2 > f1?). Sequential discrimination thus requires at least three components: loading a stimulus with a particular value (f1) into the working memory system, storing that value over a few seconds, and then comparing the memory of f1 to f2 when the second stimulus f2 is presented.

For a somatosensory variant of this task, the neurophysiological pathway has been largely identified. In particular, neurons that participate in all three components of the task have been found in the prefrontal cortex of macaques [Romo et al., 1999; Brody et al., unpublished observations]. Neurons that exhibit activity as a function of the first stimulus (f1) fall into two classes: those with firing rate proportional to f1 (positive tuning) and those with firing rate proportional to -f1 (negative tuning).

Based on these findings, we set up a simple neural network model for the prefrontal cortex. The network consists of oppositely tuned neurons that form a line attractor and are therefore able to store the value of a single stimulus (such as f1). We show how the input of neurons from somatosensory cortex S2 manipulates the attractor dynamics of the network model so that stimulus f1 can be loaded or that stimulus f2 can be compared to the memory of f1. The model agrees with key aspects of the electrophysiological evidence and is able to carry out all three components of the task (i.e., loading, storing, comparing) within a single, integrated, framework.

INFLUENCE OF THALAMIC STIMULATION ON NETWORK ATTRACTORS IN NEOCORTEX: HUBEL & WIESEL VS. HOPFIELD

Jason N. MacLean*, Brendon O. Watson, Rochelle Urban, and Rafael Yuste, Department of Biological Sciences, Columbia University, New York, New York 10027.

How much of our perception of the world is determined by our internal mental states? In sensory cortex it remains unclear whether receptive fields are the product of thalamic recruitment, or of pre-existing cortical circuit attractors. Network UP states are spatially organized and have stereotypical dynamics which imply that they are network attractor states[1]. However, it remains to be elucidated whether thalamic input pre-determines or influences this spatial organization. Utilizing mouse thalamocortical slices, bulk loaded with the calcium indicator fura-2AM, we have simultaneously examined the spatiotemporal dynamics of several hundred neurons during spontaneously occurring UP states and activations triggered by thalamic stimulation. Brief (4 to 8 pulses) stimulation of the ventrobasal (VB) nucleus at 40 Hz produces activation of neuronal ensembles, recruiting UP states. Networks activated by thalamic stimulation overlap greatly with those active during spontaneous UP states and are characterized by a clustered organization. Further, voltage clamp recordings of neurons involved indicate similar current amplitude and reversal potential during both evoked and spontaneous UP states suggestive of the same synaptic drive. A response by a neuron to thalamic stimulation is a very good predictor of spontaneous UP states occurring in that neuron. However, field potential recordings, intracellular recordings of neurons receiving direct thalamic input, and experiments in slices without intact thalamic input indicate that spontaneous UP states do not require and are not preceded by thalamic activity. We conclude that thalamic activation evokes UP states and that these UP states are very similar to those occuring spontaneously. Our data are consistent with the hypothesis that preexisting cortical circuit attractors determine and dominate the effect of thalamic inputs.

Supported by NEI and NINDS.

1 Cossart et al (2003) Nature, 423: 483-488

PHASE LOCKING STATES PROMOTED BY ELECTRICAL OR SYNAPTIC CONNECTIONS BETWEEN PAIRS OF INHIBITORY INTERNEURONS.

J. Mancilla[1], T. Lewis[2], D.J. Pinto[1], J. Rinzel[2], and B.W. Connors.[1]
[1]Department of Neuroscience, Brown University, Providence, RI 02912.
[2]Center for Neural Science, New York University, New York, NY 10003.

The barrel cortex of rodents has two networks of inhibitory interneurons, fast-spiking (FS) and low-threshold spiking (LTS). While FS cells can have both electrical and synaptic connections amongst each other, LTS cells only have electrical connections. Our theoretical studies suggest that the ability of electrical coupling to synchronize activity between interneurons depends on a number of factors, such as coupling strength and firing frequency. In fact, the theory predicts that under some conditions electrical synapses alone can promote anti-phasic activity between coupled pairs of neurons. We examined the frequency-dependence of synchrony for connected pairs of FS or LTS cells with electrical coupling, chemical inhibitory connections, or a combination of both. In some unconnected pairs, connections were simulated using a dynamic clamp protocol to test a range of connection strengths. The dependence of synchrony on firing frequency was tested by either increasing the frequency of both cells continually, with simultaneously injected current ramps, or stepwise with simultaneous current steps. In both cases, strong electrical coupling promoted synchrony at all frequencies tested and the correlation between coupled pairs increased with coupling strength. In addition, strong electrical coupling led to robust frequency-locking (± 0.5Hz) in cell pairs with different intrinsic firing rates. Consistent with modeling studies, cell pairs with either weak electrical coupling or mixed electrical and inhibitory coupling exhibited both synchronous and anti-phasic activity at low frequencies, i.e. bistability. Cell pairs connected only by inhibitory synapses displayed relatively weak phase-locking and electrical synapses dominated the effects on phase-locking in pairs with mixed coupling. However, strongly coupled pairs could only be forced into anti-phase when inhibition was present, suggesting that, in pairs with mixed coupling, inhibition does affect phase-locking under certain conditions.

Supported by NIH and the Epilepsy Foundation through the support of the American Epilepsy Society and Milken Family Foundation.

INTERACTION OF GAIN CONTROL MECHANISMS IN EARLY VISUAL SYSTEM

Valerio Mante, Vincent Bonin and Matteo Carandini
Smith-Kettlewell Eye Research Institute, San Francisco, CA, USA

In the early visual system, visual signals are shaped by two fast, nonlinear gain-control mechanisms: *light adaptation* and *contrast gain control*. Light adaptation takes place in retina, and controls responsiveness based on local light level. Contrast gain control takes place both in retina and in subsequent stages, and controls responsiveness based on local stimulus contrast (Shapley & Enroth-Cugell, 1984).

Are these two mechanisms independent of each other? Light level and local contrast are likely to vary independently in a natural environment (one depends on the illuminant, the other on surface reflectance), so there would seem to be an advantage to enforcing independence. On the other hand, one might not expect independence, because light adaptation precedes contrast gain control, and is likely to distort the signals before the computation of local contrast can take place.

To test independence, we measured the effects of light adaptation and contrast gain control as seen by neurons of lateral geniculate nucleus (LGN), which read out the output of the retina. We recorded from LGN of anesthetized and paralyzed cats while displaying sinusoidal gratings of various light levels (6-58 cd/m^2) and contrast (3-100%). Drift rate varied logarithmically in time from 0.5 to 40 Hz and back to 0.5 Hz, in 10 s. From the responses we estimated the impulse response for a given combination of light level and contrast.

We modeled the effects of light level and contrast with two resistor-capacitor (RC) circuits in series, one for light adaptation and one for contrast gain control. The conductances g_L and g_C of the two RC circuits depend one on light level L and the other on contrast C, and are the only parameters allowed to vary across stimuli.

We fitted the model to impulse responses measured at various light levels at fixed contrast (97 ± 3% of variance explained, N = 23) and at various contrast values at fixed light level (96 ± 4%, N = 23). To test independence, we used the fitted parameters to predict impulse responses obtained at novel combinations of light level and contrast: quality of the predictions (71 ± 14%) was lower and not correlated with quality of the fits (91 ± 5%, N = 9). In a few cells, quality of predictions was comparable with quality of fits, indicating independence in the effects of light level and contrast.

Fits yielded values of g_L approximately linear in light level L, growing as $g_L = L$. At low temporal frequencies response gain is inversely proportional to g, and therefore light adaptation is perfect (i.e. follows Weber's law) over the measured range of light levels.

Fits yielded values of g_C that were approximately proportional to contrast, growing as $g=(C_{50}+C)$, with C_{50} a constant. Therefore, at low temporal frequencies responses grow with contrast as the well-known expression $R = C / (C_{50}+C)$. At higher temporal frequencies both light adaptation and contrast gain control were less effective.

At high contrast, changes in g were more pronounced for ON-center cells than for OFF-center cells. These differences might be an effect of light adaptation: ON-center cells respond to increases in light level, which rapidly reduce the light adaptation gain, while OFF-center cells respond to decreases in light level, which increase the light adaptation gain.

In summary, light adaptation and contrast gain control can each be modeled by an RC-circuit, and are described by changes in conductance alone; however, they are independent of each other only on a subset of cells.

Supported by the James S. McDonnell Foundation

TRANSIENT INFORMATION FLOW IN A CORTICAL-LIKE NETWORK

Julien Mayor and Wulfram Gerstner
Laboratory of Computational Neuroscience, Brain and Mind Institute.
Swiss Federal Institute of Technology, CH-1015 Lausanne, Switzerland

The ultra-short-term memory properties of recurrent neural networks of spiking neurons with random connectivity are investigated. We demonstrate that the network can process in a distributed fashion many superimposed streams of input. A simple output structure is powerful enough to extract the information distributed across the network. The dimensional blow up, which is crucial in kernel methods, is efficiently achieved by the dynamics of the network itself.

The autocorrelation of the input plays a crucial role in the way the information is extracted from the network. We find a non-monotonic dependence upon the intrinsic noise level, an effect similar to the well-known phenomenon of stochastic resonance.

Analytical results based upon an effective mean-field model of the network show good agreement with the simulations. Our results could help to clarify the relation between ”liquid computing” and the theory of randomly connected integrate-and-fire networks.

Evoked gamma activity is sensitive to the spectrotemporal structure of communication sounds

Andrei V. Medvedev, Jagmeet S. Kanwal
Dept. of Physiology & Biophysics, Georgetown University, 3900 Reservoir Rd, Washington, DC 20057

Neural computations underlying perception of communication sounds or “calls” are poorly understood both in terms of spiking activity and local field potentials (LFPs). To examine the emergent computation that results from synaptic activity within neuronal networks, we recorded LFPs from the posterior region of the primary auditory cortex of awake mustached bats in response to the presentation of 14 species-specific simple syllabic calls.

Significant differences between the empirical distributions of spectral estimates of the LFP at low (<20 Hz) versus high (>20 Hz) frequencies indicated that gamma activity 20-100 Hz is not part of a broadband complex signal but an oscillatory response that is relatively independent of the low frequencies. Call-evoked gamma activity was attenuated with call reversal showing a correlation with the spectrographic parameter of frequency modulation (FM) characterizing asymmetry in the temporal structure of a call. Another spectral parameter ‘Harmonic complexity’ of a call was quantified through calculation of the ‘effective’ frequency within the normalized power spectrum of that call. The relative amount of gamma activity in the power spectrum of a call-evoked LFP significantly correlated with the ‘effective’ frequency of the corresponding call. These data show that the call-evoked gamma activity is also sensitive to the harmonic structure of calls. Amplitude modulation, even when asymmetric across call reversals, had no significant effect on gamma activity.

Our results on sensitivity of the evoked gamma activity to the spectrotemporal structure of communication sounds provide the first experimental evidence that stimulus-locked gamma oscillations may be an important prerequisite for recognition of complex sounds. The observed effect may be based on transient synchrony manifested by gamma oscillations within neural clusters as has been suggested in a recent modeling study (Hopfield & Brody 2000, 2001).

Supported by NIH/NIDCD research grant DC02054 to J.K.

FREQUENCY DISCRIMINATION BY INTEGRAL FEEDBACK CONTROL.

Paul Miller and Xiao-Jing Wang
Volen Center for Complex Systems,
Brandeis University, 415 South St, Waltham, MA 02454

In a delayed comparison task [Romo *et al.* Nature **399**:470 (1999)] macaque monkeys indicate which of two temporally separated vibrotactile stimuli is of greater frequency. Labeling the initial stimulus, f1, and the second stimulus, f2, Brody et al. [Phil. Trans. R. Soc. Lond. B **357**:1843 (2002)] observed neurons encode first f1, and later in the task the difference, f2-f1. The particular dynamics of the coding of neurons varies with location through somatosensory, prefrontal and premotor cortices.

If two concurrent inputs enter the decision-making area through separate pathways, cross-inhibition allows a response to the stronger stimulus alone, through a winner-takes-all mechanism [Wang, Neuron **36**:955 (2002)]. However, in the discrimination task, the two inputs are not concurrent, and enter the cortex through the same pathway.

In this paper we propose a network that can solve the delayed comparison task, using the method of integral feedback control. The key concept is that neurons which fire monotonically with the vibrotactile stimulus excite an integrator circuit, which also serves as the working memory buffer. The integrator then inhibits the stimulus-dependent input neurons. Once the input neurons are silenced, the integrator's activity persists at a rate that depends only on the stimulus frequency, so long as the duration of the stimulus is above a critical value. The inhibiton of the input neurons remains throughout the delay period. When the second stimulus is presented, the input neurons are unable to fire, unless the second stimulus strength is greater than the first, so that the feedback inhibiton is overcome. Hence a response to the second stimulus occurs only if it is greater than the first, and the basis of a discriminatory circuit is present. We point out that the discrimination behavior remains robust, even with a large variation of connection strengths to and from the integrator.

We produce more populations of neurons with some of the experimentally observed properties (ramping neurons, and other neurons that are tuned negatively as well as positively monotonically) to generate a network that contains two output populations. One output is active to indicate a response to f2 $>$ f1. Activity at the other output signifies the alternate response, to f1 $>$ f2. In this sense the circuitry translates a single, temporally separated sensory input, into one of two possible choices for a motor response.

A Model of Low Frequency Oscillatory Visual Responses in Macaque Inferotemporal Cortex

Samat Moldakarimov, Julianne E. Rollenhagen, Carl R. Olson, and Carson C. Chow

Some neurons in inferotemporal cortex (IT) of the macaque monkey respond to visual stimuli by firing action potentials in a series of bursts at a frequency of around 5 Hz. We speculated that the oscillatory responses depend on competitive interactions among neurons selective for different stimuli. To test this hypothesis, we recorded neuronal responses to a preferred foveal stimulus (the 'object') presented either in isolation or against the backdrop of an already present peripheral non-preferred stimulus (the 'flanker'). The presence of the flanker enhanced the oscillatory component of the response to the object. We have now constructed a model that can account for the experimental data. The model consists of two pools of neurons, one representing the object and the other the flanker, with reciprocal inhibition and adaptation. We have found that depending on the strength of the inputs, synapses, adaptation, there is either oscillatory activity, both pools are active (on-on), both pools are quiescent (off-off), or one pool is active and the other is silent (on-off). Additionally, our model shows a normalization phenomenon in the on-on region. When an object and flanker are presented simultaneously, the resulting activity lies in between the activity of the object and flanker presented alone. In the oscillatory and on-off regions, normalization was not observed. We explain mathematically why these behaviors exist in the network and make predictions for future experiments.

SUBTHRESHOLD MICROSTIMULATION OF SUPERIOR COLLICULUS (SC) MEDIATES ATTENTION: BEHAVIOR AND PHYSIOLOGY

James R. Müller, Marios G. Philiastides*, and William T. Newsome
HHMI and Dept. Neurobiology, *Dept. Electrical Engineering
Stanford University, Stanford, CA 94305-5125

When attention is directed to a localized region within the visual field, discrimination of stimuli in that region is enhanced. At the same time, attention modulates responses of visual neurons in specific ways. Recent evidence suggests these effects may be mediated via feedback to visual cortex from oculomotor structures, for example the SC (Gattass & Desimone, SFN 1992). To test this, we examined the effects of microstimulation of the SC (current $<$ eye-movement threshold) on the direction discrimination performance of monkeys, and on neural responses in the middle temporal visual area (MT).

Monkeys discriminated the direction of coherent motion of a localized patch of random dots positioned at one of two locations within a large noise field. Performance (i.e. coherence threshold) was greatly improved by a spatial cue indicating the location of the patch of coherent dots. To test the SC's role in improving performance, we stimulated the SC while measuring the monkeys' coherence thresholds. If SC microstimulation mediates attention, it should improve performance when the coherent patch is within the movement field (MF) of the stimulated SC site, but not when the patch is remote from the SC MF. Both predictions were confirmed by the experimental results, consistent with SC playing a role in mediating visual attention.

At the physiological level, attention affects MT neurons in two ways (Treue & Maunsell, 1999). When two competing stimuli, one preferred and one nonpreferred, are presented simultaneously within a receptive field (RF), attention to the preferred stimulus increases the response (effect ♯1). But attention to a non-preferred stimulus decreases the response, whether the nonpreferred stimulus is within the RF (effect ♯2A) or remote from it (effect ♯2B). To test the SC's role in causing these effects, we measured visual responses in MT while microstimulating SC. A preferred stimulus was always placed within the MT RF. Non-preferred stimuli were placed within the RF in some experiments and remote from the RF in others. Consistent with the effects of visual attention described by Treue and Maunsell, we found that electrical stimulation of the SC 1) increased responses in MT when the SC MF contained the preferred stimulus (♯1 above), and 2) decreased MT responses when the SC MF contained the non-preferred stimulus, whether it fell within (♯2A) or outside (♯2B) the MT RF. Thus both lines of evidence – behavioral and electrophysiological – are consistent with SC playing a role in mediating effects of attention on visual responses in MT. (Supported by HHMI, NEI (05603)).

STRUCTURAL PLASTICITY AND AUDITORY LOCALIZATION IN BARN OWLS - A COMPUTATIONAL MODEL

Shreesh P. Mysore and Steven R. Quartz
Control and Dynamical Systems Program
California Institute of Technology
1200 E. California Blvd., MC 107-81
Pasadena CA 91125 USA

Traditionally, learning is viewed as a change in the connection strengths between neurons in a fixed-architecture network in response to experience. Recent evidence from neuroscience suggests that structural plasticity is more central to learning than previously assumed. A model system in which behavioral plasticity is implemented primarily through structural change is the auditory localization system in the mid brain of barn owls.

It is known that the auditory and visual space maps in barn owls are integrated in the optic tectum (OT), where auditory cue values are associated with locations in visual space and are subsequently used to drive orienting behavioral responses. When fitted with prismatic spectacles that shift the visual field in the horizontal direction, young owls ($<$180 days old) display behavioral adaptation over the course of several weeks. This has been shown to be neurally implemented through axo- and synaptogenesis between the central and external nuclei of the inferior colliculus (the latter, referred to as ICX, is one level upstream of the OT). These anatomical changes alter the auditory tuning curves of the ICX neurons and result in a modification of the existing topographic map of space in a behaviorally consistent manner.

We present a computational model that attempts to unify experimental evidence on localization plasticity in young owls. We apply a Hebbian-like scheme to a recurrent network of sigmoidal units and show that it allows for a gated, topographic, visual error signal to drive the intricate spatiotemporal changes that occur in the network and in behavior following prismatic visual shifts. Previous models exist that capture post-developmental localization plasticity in the barn owl OT. However, they make assumptions (that have since been shown to be incorrect) about the nature (foveation basis) of the instructive signal driving plasticity, and/or rely heavily on reward-based processes (which are not well-studied in owls). We show how computational mechanisms and experimental data interact to place constraints on (a) the potential site of an important (inhibitory) component of plasticity, and on (b) possible biological signals that can drive the observed structural changes, and make two predictions. We discuss the temporal sensitivity of the signals involved, and also briefly discuss ties between structural plasticity and “complexity” of networks, and the future use of spiking units in the model.

DISTRIBUTED CODING OF NEURONAL ENSEMBLES

INFORMATION THEORETIC AND RANDOM FOREST ANALYSIS OF NEURONAL ENSEMBLES REVEALS DISTRIBUTED CODING OF INFORMATION

Kumar Narayanan, Eyal Kimchi, Artur Luczak, and Mark Laubach
John B. Pierce Laboratory
Yale University, New Haven, CT, 06519

Distributed processing (DP) is the sharing of processing load among elements of a network. DP has been proposed as a model for the function of neuronal ensembles. A hallmark of DP is redundancy, i.e., continued performance despite the degradation of information or processing capacity. To examine the distributive properties of neuronal ensembles in the cerebral cortex, we recorded neuronal ensembles using chronic microwire arrays in the motor cortex of rats that performed a reaction time task and in somatsensory cortex during sensory mapping experiments of the rat forepaw under pentobarbital anesthesia. This methodology records pseudo-random samples of cortical activity within a given cortical area.

We analyzed the recorded data using a statistical pattern recognition approach. We identified features in firing patterns associated with different experimental conditions (reaction time in motor cortex, location in sensorimotor cortex) using a wavelet based method called discriminant pursuit (Buckheit and Donoho 1995; Laubach 2004). Features were classified (i.e., short vs. long reaction times, different locations on the paw) using the Learning Vector Quantization [LVQ] (Kohonen 2001) or Random Forest [RF] (Breiman 2001; Oh et al. 2003) algorithms. Additionally, we performed a neuron-dropping analysis (Nicolelis et al. 1999) to quantify the contributions of individual neurons to neural coding a the ensemble level. The information theoretic measure of redundancy was used to define the information lost by the removal of individual neurons.

We found that information about task performance or stimulus location was not represented by a few strongly driven neurons in each ensemble. Instead, these parameters were encoded by the activity of many, weakly responsive neurons. Moreover, the importance of neurons depended on whether the neurons were tested alone or as part of an ensemble. These results suggest that a single neuron's response properties may not be the primary determinant of the neuron's contribution to coding at the population level. We are currently examining this issue using computational simulations and are also examining the spatial determinants for this type of distributed coding.

THE CONTRIBUTION OF INTRACELLULAR THRESHOLDING TO SPIKE-TIMING JITTER, TRIAL-TO-TRIAL RELIABILITY AND INFORMATION CONTENT IN A DYNAMICALLY ACTIVATED NEURON.

Reza Nassiri[1], Monty A. Escabí [1,2]
[1]Biomedical Engineering and [2]Electrical and Computer Engineering,
University of Connecticut, Storrs, CT 06269

Individual neurons do not produce identical spike trains in response to identical stimuli. While the reliability of the neuronal responses varies widely across trials for different neurons, this variability may be quantified in two basic dimensions, namely trial-to-trial reproducibility and spike-timing jitter. Inspection of raster plots showing spike timings of neuronal outputs to identical stimuli shows that the temporal alignment of spikes often do not lineup perfectly, thus contributing to spike-timing jitter. Neurons, however, can also exhibit failures in spike generation: these errors are typically characterized as the trial-to-trial reproducibility or reliability of the neuron. From an information theoretic perspective, a hypothetical noiseless neuron would convey the maximal *information* for a given spike train sequence. Neuronal *noise* in the form of reproducibility errors and spike-timing jitter, however, limit its ability to encode viable information.

We use a simple dynamically driven integrate-and-fire neuron to study the relationship between intracellular noise fluctuations, the characteristics of the extracellular spiking noise, and the information that can be communicated between two or more neurons. Specifically, we use a combination of neuronal thresholding and random intracellular voltage fluctuations (intracellular noise) to quantitatively asses the information content which a cell communicates. We develop a spike crosscorrelation procedure for simultaneously extracting spike-timing jitter, trial-to-trial reliability, and information content of a single neuron from two independent response trials. Varying the relative threshold voltage of our computational integrate-and-fire neuron model with the addition of random noise provides systematic changes in spike-timing jitter and trial-to-trial reproducibility which provide physiologically plausible results as to the source of jitter and reproducibility noise in neuronal outputs. We investigate the details of the relationships between intracellular thresholding, spike-timing precision, and trial-to-trial reliability in this simplified model neuron. The results are analyzed in an information theoretical context, with focus on identifying how spike-timing precision and trial-to-trial reproducibility individually limit information transmission in real neuronal spike trains.

RETINAL GANGLION CELLS REVERT FROM OFF TYPE TO ON TYPE DURING VISUAL SACCADES

Maria A. Neimark and Markus Meister
Program in Biophysics and MCB Department
Harvard University, Cambridge, MA

Saccades are large and rapid gaze shifting eye movements that produce large image shifts on the retina. Our visual sensitivity is greatly diminished in the period surrounding a saccade. This suppression is attributed partly to input from the retina, and partly to signals originating in the eye movement system. We explored the effects of saccades on retinal processing, by recording spike trains from retinal ganglion cells (GCs) in the tiger salamander, and probing the changes in their visual responses due to large image shifts.

We have found that visual saccades have a pronounced effect on retinal processing. While the firing of OFF GCs is inhibited by a saccade, the firing of ON GCs is enhanced. Further, we discovered a class of OFF ganglion cells whose linear kernel reversed from OFF type to ON type 100-200 ms after a saccade. This effect could be suppressed by the application of APB, an ON bipolar pathway blocker, suggesting that kernel reversal is due to unmasking of ON bipolar inputs during selective suppression of OFF inputs. The kernel reversal effect was localized to subregions of the GC's receptive field, indicating that the inhibition of the OFF inputs may be presynaptic, and is likely to be carried out by wide-field transient amacrine cells, as revealed by the application of synaptic blockers.

This study illustrates that the identity of the ganglion cells as ”ON” or ”OFF” may not be hard-wired in retinal circuitry, but result from a dynamic balance between their ON and OFF inputs, regulated by amacrine cells.

POSITIVE SENSORY FEEDBACK IN A RAT BRAINSTEM SENSORIMOTOR LOOP

Quoc Thang NGUYEN and David Kleinfeld
Department of Physics, UCSD, La Jolla CA USA

The rat vibrissa sensorimotor system has become a prominent experimental model to study active tactile sensation in mammals. This system consists of nested, sensorimotor feedback loops. In the lowest-order feedback loop, the vibrissa-trigemino-facial (VTF) loop in the brainstem, sensory signals from the vibrissae on their way to thalamic and cortical areas are fed back to vibrissa motoneurons in the facial nucleus via oligosynaptic connections.

To investigate the VTF loop, we developed a novel slice preparation termed "slab". Slabs, which are 1000-1200 µm thick horizontal slices from rats that begin to whisk (P12-P15), contain the nuclei of the VTF loop and the input nerves, allowing us to study the VTF circuit in the open-loop configuration.

We recorded intracellularly from vibrissa motoneurons ($E_m = -62 \pm 7$ mV, $R_m = 30 \pm 12$ MΩ, mean ± SD, n = 75). Stimulation of the input nerve ipsilateral to the motoneurons elicited delayed, subthreshold EPSPs in 47 % of the neurons. The evoked EPSPs, referred to as the VTF response, had a stimulation threshold between 0.5 and 20 V, were variable from trial to trial and could induce one or several action potentials in 15 % of the cells. Upon averaging, the earliest EPSPs summed and became gradually larger as the stimulation was increased. Stimulation of the nerve contralateral to a responsive motoneuron failed to elicit a VTF response. The averaged EPSPs had maximal amplitudes between 0.5 to 12 mV that did not change noticeably when motoneurons were hyperpolarized. The latencies Δt of the initial EPSPs were distributed in a short latency group ($\Delta t = 8 \pm 2$ ms) and a longer latency cluster ($\Delta t = 25 \pm 4$ ms). There was no relation between latencies and the stimulation threshold for EPSPs and also no correlation between the latencies of the initial EPSPs and the maximal amplitude of the corresponding VTF response. The response latencies fell into two groups which were compatible with positive sensory feedback during whisking. We speculate that positive sensory feedback may boost the amplitude of vibrissa movements subsequent to the contact of the vibrissae with an object.

We are currently investigating the dynamics of the VTF loop by determining the phase-response curve of vibrissa motoneurons in response to intracellular current injection or nerve stimulation.

A LOCAL SYNAPTOTROPIC GROWTH MECHANISM CAN RESULT IN EFFICIENT GENERATION OF SYNAPTIC CONNECTIVITY

Cristopher Niell and Stephen Smith
Neurosciences Graduate Program, Stanford University

It is generally believed that the genome can not encode explicit instructions for all synaptic connections, but may provide general neurite growth mechanisms which will result in proper connectivity. Our recent in vivo imaging has provided evidence for a synaptotropic mechanism, wherein synapses could influence dendrite growth by selectively stabilizing filopodia upon which they form. Stable filopodia then serve as branches from which further growth can occur. We undertook a theoretical investigation into the consequences of such a growth process, and find that the positive feedback inherent in this local rule results in an efficient large-scale search for regions of high density pre-synaptic partners. Specifically, mathematical analysis shows that growth by selective stabilization of filopodia can be approximated as a reaction-diffusion system, with a spatially varying diffusion constant that is proportional to the probability of synapse formation. Thus, growth will occur faster in regions of appropriate synaptic connections, and the net growth can be shown to climb a gradient, thereby guiding a neuron to regions dense in synaptic partners. Furthermore, computational simulations provide a direct comparison of search efficiency for a synaptotropic mechanism versus growth without feedback, and can recapitulate specific patterns of dendrite growth and connectivity. Coupling synaptotropic growth with competitive Hebbian plasticity is shown to create an efficient search for pre-synaptic partners with correlated activity. The synaptotropic mechanism can thus provide an explanation for the development of connectivity based on local properties of the circuit elements, rather than explicit dependence on global guidance molecules or innate predetermined growth.

HOW INHIBITORY OSCILLATIONS CAN TRAIN NEURAL NETWORKS AND PUNISH COMPETITORS

Kenneth A. Norman, Ehren L. Newman, and Sean M. Polyn
Dept. of Psychology and Center for the Study of Brain, Mind, and Behavior.
Princeton University, Princeton, NJ 08544

Several theorists have argued that the function of neocortex is to model the external world, so we can fill in the properties of stimuli based on their similarity to previously experienced events. A major challenge for cognitive neuroscience is to characterize the mechanisms that neocortex uses to acquire this predictive model. In the research presented here, we show (using rate-coded neural network simulations) how oscillations in the strength of inhibition, coupled with a learning rule that is sensitive to the phase of the oscillation, can train neocortical networks. The learning rule has two main parts: First, it increases weights to units that turn off during the high inhibition phase of the oscillation. Increasing inhibition constitutes a ”stress test” of the network’s representations: Weak parts of the representation drop out when inhibition is increased, and the learning rule strengthens these weak parts. The second part of the learning rule reduces weights to units that turn on during the low inhibition phase of the oscillation. Reducing inhibition allows competitors to become active, and the learning rule acts to punish these competitors. Formally, the rule that we use to update weights is equivalent to the Contrastive Hebbian Learning (CHL) rule, except the sign of the rule depends on the phase of the inhibitory oscillation. The rule is applied to snapshots of network activity taken at successive time steps. When the inhibitory oscillation is moving towards its midpoint, the sign of the learning rule is positive (i.e., successive time steps are moving the network closer to the ”correct” level of activity). When the inhibitory oscillation is moving away from its midpoint, the sign of the learning rule is negative.

We show that our learning algorithm can be used to train large numbers of overlapping patterns into a network, and we compare its efficacy to other unsupervised and supervised learning algorithms for rate-coded networks. We also use the rule to address puzzles in the behavioral literature. Findings from several different domains (e.g., cognitive dissonance reduction, episodic memory retrieval) all show the following regularity: When representations are initially activated by a cue, but subsequently lose the competition to become active, these representations become less accessible. We show how our rule can account for this basic principle as well as nuanced patterns of data within these domains. Finally, we discuss how this research relates to neurobiological data on theta oscillations and learning (e.g., Huerta and Lisman’s finding that the direction of synaptic plasticity depends on the phase of theta).

FUNCTIONAL DISSECTION OF A LEARNING RULE AT THE HIPPOCAMPAL CA3-CA1 SYNAPSE.

Daniel H. O'Connor[1,3], Gayle M. Wittenberg[2,3] and Samuel S.-H. Wang[2,3].
Departments of Psychology[1] and Molecular Biology[2] and Program in Neuroscience[3], Princeton University, Princeton, NJ 08544.

Long-term potentiation (LTP) and long-term depression (LTD) result from distinct biochemical cascades, suggesting that plasticity might be separable into learning rule components that are potentiation-only and depression-only. In extracellular recordings from P14-21 rats we have taken the bidirectional, frequency-dependent learning rule observed by Dudek and Bear (1992 PNAS 89:4363; also known as the Bienenstock-Cooper-Munro, or BCM rule) and separated it pharmacologically into two components: (a) a phosphatase-dependent component that gives depression at stimulus frequencies from 1 to 100 Hz; and (b) a kinase-dependent component that gives potentiation at stimulus frequencies from 10 to 100 Hz. In computational simulations, these two processes can be put back together to reconstitute the original BCM rule observed in the entire synaptic ensemble under a range of assumptions, including the possibility that at single synapses, plasticity is an all-or-none event (Petersen et al. 1998 PNAS 95:4732).

An all-or-none model would suggest that CA3-CA1 synapses are a heterogeneous ensemble, with some synapses in a low state that can only potentiate; and others in a high state that can only depress or depotentiate. To generate more uniform ensembles of synapses, we have saturated LTP or LTD. Fully depressed populations of synapses showed a magnitude and activity-dependence of potentiation that suggests that all synapses in a low state express the same potentiation-only learning rule observed in the phosphatase inhibitor experiments. Conversely, fully potentiated populations of synapses expressed a decrease-only learning rule. However, the ability of some synapses to depress/depotentiate disappeared over a period of minutes. Thus potentiated synapses can enter a ”locked-in” state from which decreases in strength are not possible. Our observations can be captured in a model in which synapses can be in one of three states: low, high, or locked-in.

A WORKING MODEL OF SPATIALLY DISTRIBUTED ACTIVITY INCREASES THE EFFECTIVE RESOLUTION OF FUNCTIONAL IMAGING

Cheryl A. Olman, Department of Neuroscience and Center for Magnetic Resonance Research, University of Minnesota

Stimulus or behavior differences are coded by changes in neural activity on columnar or sub-columnar scales, but the hemodynamic response limits the resolution of standard Gradient Echo (GE) BOLD fMRI to a few millimeters, at best. One way of circumventing this limitation in spatial resolution is to work with a model of neural activity that connects microscopic changes in local neural activity (energy consumption) with the macroscopic BOLD signal.

One example of the utility of such a model is the study of the role of spatial phase structure in scene perception, where the spatial distribution of image features and magnitude of neural responses vary simultaneously and on different spatial scales. To address problems such as this, I have developed a computational model of neural and BOLD fMRI activity in primary visual cortex. Energy consumption resulting from neural activity is predicted with 0.5 mm resolution, and then convolved with an appropriate point spread function (3 mm for GE BOLD imaging; 1 mm for spin echo BOLD imaging), generating a prediction for the spatially distributed BOLD activity. For a first test of the model, the visual stimuli used were patterns of Gabor patches, randomly placed at different densities on a mean gray background, and contrasted against uniformly distributed noise with the same spatial frequency content. Good agreement between modeled and measured distributions of neural activity demonstrates the utility of this working model. In particular, the results show that the interaction between feature density and BOLD response can be understood and controlled, permitting the study of spatially distributed neural activity in the context of complex stimuli.

These results demonstrate that simplistic data analysis based on regional means can overlook important information contained in the BOLD signal. The success of the model hinges on accurate predictions of the spatial distribution of energy consumption. Further work on the modeling effort is focusing on an improved description of the microcircuitry in V1, in order to better predict the effects of both excitatory and inhibitory synaptic activity.

RETINAL RESPONSE TO THE ONSET OF DIFFERENTIAL MOTION

Bence P. Ölveczky, Stephen A. Baccus and Markus Meister
Department of Molecular and Cellular Biology, Harvard University, Cambridge, MA

The task of detecting moving objects is complicated by the presence of eye movements that continually scan the image across the retina, even during fixation. When an object moves within the scene, the corresponding image patch on the retina experiences a different motion trajectory from that in the surrounding region. We recently reported that certain ganglion cells respond selectively to this condition, when the motion in the cell's receptive field center is different from that in its surround [1]. These object-motion-sensitive (OMS) ganglion cells might serve to flag regions of object motion for subsequent attentive processing. Here we examine the dynamics of this response at the very onset of differential motion.

We recorded extracellularly from ganglion cells in the isolated retinas of salamander and rabbit, and also intracellularly from various cell types in the salamander. The visual stimulus consisted of two gratings: one presented in a circular patch overlying the cell's receptive field center ("Object"), the other covering the rest of the retina ("Background"). To simulate a stationary scene, both gratings were moved in unison with a random walk trajectory resembling the animal's fixational eye movements. Object motion was simulated by having the two gratings jitter with different trajectories having the same statistics.

After a period of retina-wide correlated motion the onset of differential motion between a cell's center and its wider surround caused an increase in the firing rate of all ganglion cells. For the OMS cells this increase was, on average, more than 30-fold. The firing rate of all ganglion cells adapted to differential motion with a time constant of ~7 sec. Sensitivity to object motion was stronger and more robust the longer the preceding period of global motion. This recovery from adaptation to differential motion had a time constant of ~30 sec.

The only aspect of our visual stimulus that changed at the onset of differential motion was the correlation between the motion trajectories in the object and background regions. Thus the adaptation to differential motion must occur after the signals from these two regions are combined and therefore likely happens in the inner retina where long-range connections allow for a comparison between signals from different parts of the retina. Considering the object motion sensitive circuitry [1], possible sites of adaptation include the bipolar cell terminals, the polyaxonal amacrine cells, the OMS ganglion cells, and the synapses between these cells.

[1] Ölveczky B.P., Baccus S.A. & Meister M. Segregation of object and background motion in the retina. *Nature* **423** (6938): 401-8 (2003).

RECORDINGS IN THE RAT AUDITORY CORTEX DURING A TWO ALTERNATIVE FORCED CHOICE AUDITORY TASK

Gonzalo H. Otazu and Anthony M. Zador, Cold Spring Harbor Laboratory,
1 Bungtown Road, Cold Spring Harbor, New York 11724, USA

Neural activity in the auditory cortex is modulated both by the auditory stimulus and by behavioral context. We trained rats to perform a two alternative forced choice auditory discrimination task. Subjects poked left or right depending on the frequency (1 or 10 KHz, respectively) of a pure tone imbedded in background noise. Task difficulty was controlled by the varying the signal to noise ratio.

We used chronic tetrodes to record single unit, multi-unit and local field potential responses from the auditory cortex of freely moving rats performing this task. Neuronal activity was modulated both by the stimulus and by the subject's behavior. While some neurons showed robust, short-latency stimulus-locked responses to the stimuli, others did not. In some cases, neurons with stimulus-locked responses were recorded simultaneously on the same tetrode as other nearby units that failed to show clear stimulus modulation, suggesting that tone-responsiveness during this task is not a simple function of cortical location. In some neurons, firing probability was correlated with the subject's left/right decision. These experiments may elucidate the role of the auditory cortex in decision making during a demanding auditory task.

ORIGINS OF RESPONSE TIME VARIABILITY IN A RAPID SERIAL VISUAL PRESENTATION TASK

Lucas Parra[1], Adam Gerson[2] and Paul Sajda[2]
1 Department of Biomedical Engineering, City College of New York
2 Department of Biomedical Engineering, Columbia University

Background: Several researchers have used electroencephalography (EEG) and a rapid serial visual presentation (RSVP) paradigm to demonstrate that object recognition in human's requires as little as 150ms of processing time (Thorpe et al., *Nature* 1996; Johnson and Olshausen, *Journal of Vision* 2003) The earliest motor response however is not observed until 350 ms after stimulus presentation and may be as late as 600 ms. An open question is what is the cortical origin(s) of this large response time (RT) variability. **Methods:** Data was collected for 7 right handed subjects ages ranging from 23 to 37. Images of natural scenes were presented using a "barrage" style RSVP paradigm (Blocks consisting of 100 images presented at 10Hz). For each block there was a 50% chance the sequence contained a target image of a person or man-made structure. Subjects were instructed to respond rapidly, with a button release, to the appearance of a target image. EEG was recorded at 60 scalp locations and spatially filtered to reduce activity due to eye-movement and eye blinks. A zero-phase high pass filter was used to remove DC drifts. The activity in the multiple sensors was combined via a spatially weighted integration to optimally discriminate between target and non-target presentations. The time course of this discriminant activity was analyzed on a single-trial basis. **Results:** The first reliable indication of a single-trial difference between target and non-target images arises between 200-250ms after stimulus onset. The corresponding activity is stimulus locked. It has a bilateral frontal distribution which is consistent with activation of lateral prefrontal cortex (LPFC). The activity remains frontal for about 100 ms. A posterior activity over parietal areas (e.g. posterior parietal cortex, PPC) arises 50-100 ms prior to button release. Its onset time is delayed as compared to stimulus onset. **Discussion:** The "barrage" RSVP task requires high vigilance and likely involves no explicit memory retrieval. Presumably processing time is largely required for decision making and formation and execution of the motor response. This task appears to successively engage two cortical areas: LPFC presumably involved in executive control followed by PPC often implicated in visual-motor integration and intention/attention. We do not find a single cortical source of RT variability, rather half of the variability can be attributed to the processing from frontal activity, beginning stimulus locked, transitioning to parietal activity. The remaining variability/delay appears to be introduced after the activity has arrived in parietal and motor areas.

BUILDING REPRESENTATIONS SPIKE BY SPIKE

Klaus R. Pawelzik, David Rotermund, and Udo A. Ernst

Brains enable animals and humans to rapidly recognize and classify sensory scenes using spikes as signals, that come from the receptors and from within the brain. It appears that on average one spike per neuron and processing step is sufficient to achieve a very good performance. This is particularly astonishing when considering the apparent randomness of cortical spike trains.

Here we focus on the question how reliable representations can be built from noisy input spike patterns containig only a small number of spikes per input neuron.

We present a novel framework for neural computation which utilizes synaptic nonlinearities to extract most of the information contained in each stochastic spike. This is demonstrated in a network that iteratively builds a sparse and noise-robust representation spike per spike. The weights in this model can also be adapted on-line using an unsupervised Hebbian rule that optimizes the representation with respect to reconstruction error. Since our framework corresponds to a linear mixture model it decomposes the rate vectors underlying the inputs into independent causes thereby solving spatial deconvolution problems (like e.g. blind source separation) on the basis of stochastic input.

We tested the performance of our network, using different learning rules and reconstruction paradigms, by building representations for stochastic variants of Boolean functions (like e.g. X-OR) and find that PERFECT performance (e.g. zero error) is on average achieved from less than two spikes per neuron. When applied to hand written digits data we obtain a performance that comes close to the benchmarks despite the stochastic nature of the input used. Furthermore, the performance is particularly robust with respect to pixel noise and partial coverage.

We interpret our rather general ansatz as a model for the cortex. When trained with natural stimuli, our framework reproduces many properties of visual cortical cells including simple cell receptive fields, and non-classical receptive fields. We also investigate, how the structure of long-range connections, as measured within the visual cortex, may emerge during learning by extracting correlations between sequences of inputs.

Our results indicate many neuronal structures in the brain may realize iterative estimators that are optimally adapted for rapidly building representations from sensory scenes or input patterns in other areas from very few spikes per input neuron. The model predicts characteristic temporal evolutions of intracellular responses in visual cortex, that depend on the spatial structure of the input, which can be tested experimentally.

MONAURAL SOURCE SEPARATION USING SPECTRAL CUES

Barak A. Pearlmutter and Anthony M. Zador
NUI Maynooth, Co. Kildare, Ireland & Cold Spring Harbor Laboratory, USA

Animals use a variety of binaural and monaural cues to perform sound localization. Important cues include intra-aural intensity and phase disparity, and monaural spectral filtering by the head and pinnae (via the head-related transfer function, or HRTF). Such cues are also important in helping to separate individual acoustic sources in a noisy environment, such as the voices at a cocktail party.

Here we propose a model for how the auditory cortex could use monaural spatial filtering (the head-related transfer function, HRTF) arising from the head and external ears for improved source separation. Each acoustic source is assumed to have a different known position in space, and is therefore subject to preprocessing by a different linear filter before being combined with the other sources. We illustrate the utility of spectral cues with an idealized example in which the HRTF allows two sources to be recovered perfectly with a single ear. We then provide a more general framework in which sources are represented in an overcomplete basis, which we posit to exist in the primary auditory cortex. The model makes specific testable predictions about cortical response properties.

INFLUENCE OF RHYTHMIC GROUPING ON TIME PERCEPTION

Amandine Penel
Cold Spring Harbor Laboratory, Cold Spring Harbor, NY

Events in temporal sequences are not perceived independently from each other: Perception organizes auditory input. For example, in music or speech listening, two perceptual organization processes can be proposed: (1) the extraction of a temporal regularity, and (2) sequential grouping (following Gestalt laws of temporal proximity, similarity, etc.). These processes influence time perception. The extraction of a temporal regularity results in higher accuracy of temporal and pitch judgments for on-time tones (vs. early and late, in relation to a regular beat). And several results also suggest that rhythmic grouping (grouping by temporal proximity) may influence time perception, with the last short IOI (Inter-Onset Interval) of rhythmic groups (preceding the longer between-groups interval) being perceived as shorter than the previous one.

First, almost all studies of rhythm production or reproduction and music performance show that this last short IOI is produced longer than the previous one (where musical notation indicates that they should have the same duration). This group-final lengthening has been attributed to the communication of the grouping structure to the listeners, but it could compensate for a bias in time perception, and restore perceptual regularity. Second, perceptual studies have shown that a lengthening of this last short IOI (from an isochronous baseline) is difficult to detect, while a shortening is easy to detect. This too is consistent with the last short IOI of rhythmic groups being perceived as shorter than the previous one.

We confirmed this in a previous study, using an adjustment task and musical stimuli (Penel & Drake, in press, Perception & Psychophysics). Professional pianists had to adjust all IOIs of sequences presented initially as temporally irregular, that is they had to increase or decrease all IOIs, so that the sequences would sound as regular as in the musical score. Results indicated that the last short IOI of rhythmic groups was systematically adjusted longer than the previous one, suggesting that it is perceived as shorter, at a fast tempo (300 msec for the short IOI), but not at a slow one (600 msec). These effects were attributed to obligatory rhythmic grouping, which is more likely to occur at fast than at slow tempi.

They needed to be replicated and further explored however, using simpler sequences of identical tones and less musically trained participants. Indeed, one explanation for this perceptual bias is that it results from “musical” expectations about what is usually produced. If that were the case, it should not be observed in another perceptual modality. Thus, in addition to the auditory modality, we also investigated the visual modality, in which tones were replaced by flashes.

The results with the auditory modality (two tempi, sequences of two rhythmic groups, 14 participants) confirmed previous data (Penel & Drake, in press), and also showed a systematic lengthening of the already long between-groups interval, with no difference between non-musicians (6) and amateur musicians (8). The visual modality is currently under investigation. It will provide the critical test for an explanation in terms of musical expectations or in terms of a purely temporal phenomenon resulting from organization of perceptual input.

MODELLING SIMPLE OSCILLATIONS IN PARKINSON’S DISEASE

M.S. Penney, Modelling and Systems Biology, Novartis Pharmaceuticals AG, Basel, Switzerland.
N.F. Britton, Department of Mathematical Sciences, University of Bath, Bath, UK.

Parkinson's disease is manifest in the progressive degeneration of the dopamine cells of the nigrostriatal pathway. In this work, we have modelled the affected dopaminergic innervation of the striatum and its feedback in order to investigate the changes the system undergoes following progressive degeneration.

There is little quantitative data available with which to adequately define and parameterise a proper neuronal network model. Instead we have created a qualitative model which considers a sheet of dopamine neurons innervating the striatum, with positive feedback mediated by the inhibition of inhibitory input by the striatal GABAergic neurons. By further imposing a series of simplifications and assumptions we have reduced this network of thousands of neurons to an analytically tractable system that essentially describes a loop of two neurons.

Despite these simplifications, the model output displays many of the characteristics of Parkinson’s: the system is robust to initial large losses in the dopamine input before becoming symptomatic. The first divergence of the system from its initial stability is the development of an oscillatory component analogous to the resting tremor, which would ultimately progress to rigidity. Analysis of the model demonstrates that this behaviour arises from the non-linear nature of the striatal feedback.

The symptomatic output of the model may be reversed when the modelled effects of common Parkinson's treatments are imposed. Interestingly, the modelled effect of the D2 receptor agonist ropinirole removes the oscillatory component, leading us to suggest that this may be involved with the known side effect of dyskinesia in high dose L-DOPA patients.

CHARACTERIZATION OF MACAQUE RETINAL GANGLION CELL RESPONSES USING SPIKE-TRIGGERED COVARIANCE

Jonathan W. Pillow[1], Eero P. Simoncelli[1], and E. J. Chichilnisky[2]
[1]HHMI and New York University. [2]The Salk Institute, La Jolla CA

Light responses of retinal ganglion cells (RGCs) exhibit several types of nonlinear behavior, including an accelerating contrast-response function, nonlinear pooling of spatial information, and dynamic gain adjustment. Such behaviors have been separately documented in a number of different experiments, each based on a different type of restricted stimuli. We describe a spike-triggered covariance (STC) analysis that aims to provide a complete characterization of RGC responses.

Multi-electrode recordings from macaque RGCs were obtained from isolated retinas stimulated with one-dimensional spatiotemporal white noise (i.e. flickering bars). Spike-triggered average (STA) analysis revealed center-surround spatial organization and biphasic temporal integration expected from RGCs. Eigenvector analysis of the STC revealed components with spatial and temporal structure distinct from the STA. Excitatory components (i.e. those associated with large eigenvalues) exhibited temporal structure similar to the STA, but finer spatial structure, consistent with receptive field subunits. Suppressive components exhibited spatial structure similar to the STA or to excitatory components, but temporal structure that was time-delayed relative to the STA, consistent with spike generation or contrast gain control nonlinearities.

The detailed spatial and temporal structure of these components was probed with a subspace-conditional analysis. STC components were examined in a subspace corresponding to either temporal or spatial stimulus features, obtained by filtering the stimulus with the spatial or temporal profile of the STA. The reduced stimulus dimensionality increased statistical power and revealed a larger number of significant STC components: multiple excitatory spatial components and one or two suppressive temporal components in most cells.

The STC components were then used to constrain a model of RGC light response. The model consisted of identical spatially shifted subunits whose outputs were combined nonlinearly, and a divisive temporal feedback signal. Subunit profiles were constrained to lie within the subspace spanned by spatial and temporal STC components, and the parameters of the nonlinear combination rule and feedback signal were fit using maximum likelihood. This model accurately reproduced the observed STA, contrast-response function, and detailed spatial and temporal structure of STC components. It also provided a substantially more accurate prediction of novel spike trains than a single-filter model characterized using only the STA.

VISUAL ACUITY IN THE PRESENCE OF FIXATIONAL EYE MOVEMENTS

Xaq Pitkow[1], Dan Fisher[2], Haim Sompolinsky[3], Markus Meister[4]

[1]Program in Biophysics, Harvard University
[2]Department of Physics, Harvard University
[3]Department of Physics and Center for Neural Computation, Hebrew University
[4]Department of Molecular and Cellular Biology, Harvard University

Fixational eye movements pose a challenge for the visual system. These incessant movements blur the visual scene across the retina and thus complicate the coding of small features. To explore the computational limitations that eye movements impose on the visual system, we model its performance in discriminating between stimuli that differ only in fine spatial detail. We apply an ideal observer analysis to the responses of a model retina and the outputs of various hypothetical downstream circuits. By comparing the performance of ideal observers to that of human observers, we find that some plausible cortical circuits are inconsistent with psychophysical experiments. For example, previous studies have assumed that the visual system effectively ignores the precise stimulus trajectory induced by eye movements. In contrast, our preliminary psychophysical results indicate that human performance is better when the stimulus trajectory is continuous than when it is random with the same spatial distribution.

ROUTING OF SPIKE SERIES BY DYNAMIC HIPPOCAMPAL CIRCUITS

Frédéric Pouille & Massimo Scanziani
Neurobiology Section, Division of Biology, University of California San Diego,
9500 Gilman Drive, La Jolla, California 92093-0357, U.S.A

Rate and timing of spikes are used to transmit information in the brain. Neuronal circuits that can detect these two properties are, hence, necessary. Among cortical areas, the hippocampus is one of the most accessible to neuronal circuit analysis. We investigated whether distinct temporal features of spike series in CA1 pyramidal cells, the principal hippocampal output, are selectively captured by local neuronal circuits.

Feedback inhibitory neurons represent the main local target of CA1 pyramidal cells. We found that onset and average rate of spike series in hippocampal pyramidal cells are separately detected by two distinct feedback inhibitory interneurons. Feedback inhibitory interneurons that detect the onset of spike series operate as coincidence detectors. Feedback inhibitory interneurons that detect the average rate of spike series operate as integrators. These two modes, coincidence detection and integration, result from a precise synergism between pre and postsynaptic properties of the circuit. Furthermore, feedback inhibitory neurons that detect the onset of spike series send their axons to the soma of CA1 pyramidal cells. Feedback inhibitory neurons that detect the rate of spike series, on the other hand, selectively project to the distal apical dendrites of CA1 pyramidal cells. Thus a simple and stereotypical cortical circuit transforms different temporal properties of spike series in different spatial pattern of synaptic inhibition.

USING A DATABASE OF 20 MILLION MODEL NETWORKS TO STUDY A PACEMAKER CIRCUIT

Astrid A. Prinz and Eve Marder, Volen Center and Biology Department, Brandeis University, Waltham, MA

How well do synaptic and cellular properties need to be tuned to produce appropriate network dynamics? What network configurations will lead to functional network output in different animals, or within one animal at different times in its life? To address these questions, we created a database of 20 million model networks by varying cellular properties and synapse strengths in a circuit modelled after the pyloric pattern-generating network in the crustacean stomatogastric ganglion (STG). For each network in the database, a pacemaker (AB/PD, representing the biological pacemaker kernel), a lateral pyloric (LP) and a pyloric neuron (PY) were selected from three pools of single-compartment model neurons. These pools had previously been chosen from a database of model STG neurons based on their activity and response to inputs. The three model neurons were connected by inhibitory synapses with slow cholinergic (AB/PD to LP and PY) or faster glutamatergic dynamics (AB/PD to LP and PY, LP to AB/PD and PY, PY to LP). Each network's activity was simulated and output features such as period, burst starts, ends and durations, duty cycles, and phase relationships were determined online and stored.

20% of all networks had tri-phasic output with the same burst order as in the biological rhythm. We then applied more stringent criteria derived from 67 control rhythms recorded from lobsters, including ranges for period, burst durations and delays, gap durations, phases, and duty cycles – altogether 15 features of the biological rhythm. 2.2% of the networks fell in the experimental range for all 15 criteria. This "solution space" of the model circuit reproduces features of the biological circuit. For example, the fast component of the AB/PD to LP synapse in these functioning circuits is stronger than the slow component, and vice versa for the AB/PD to PY synapse, mimicking the relative contributions of the fast and slow components seen in the corresponding biological synapses.

The correlations between cellular and synaptic properties in the solution space indicate how synapse strengths and neuronal properties need to be regulated to maintain a functioning circuit. Our results show that rhythm generation can rely on cellular properties in one part of the solution space, on synaptic properties in another part, and on interactions between cellular and synaptic components in the majority of possible solutions.

PARALLEL VISUAL PATHWAYS AND REDUNDANCY IN THE RETINA

J.L. Puchalla; R.A. Harris; E. Schneidman; M.J. Berry
Department of Molecular Biology, Princeton University, Princeton, NJ, USA

Visual images are completely represented in the retina as a pattern of spiking responses of ganglion cells. Despite having highly overlapping receptive fields, retinal ganglion cells are traditionally thought to fall into several distinct categories (ON-type, OFF-type, direction selective, etc.), each of which tiles visual space. As a result, the population of retinal ganglion cells comprise several parallel channels of visual information. In this view, ganglion cells of the same functional type are independent encoders because their receptive fields do not overlap and center-surround antagonism reduces the redundancy in natural visual stimuli; cells of different functional types are also independent, because they encode different visual features.

We have tested this paradigm by measuring the redundancy of the information encoded by pairs of ganglion cells in the tiger salamander, exploring the dependence on both cell type and stimulus. Since natural images exhibit a wide variety of complex statistics, we have chosen a set of natural stimuli to explore a range of possible scenes and motions encountered by the visual system. In addition, two limiting cases of artificial stimuli are used: (1) randomly flickering checkerboards, where the stimulus is uncorrelated in space and time, which reveal redundant coding that is intrinsic to the circuitry of the retina; (2) spatially uniform flicker, where each ganglion cell receives identical stimulation, which reveals the maximum redundancy possible given the functional diversity and noise of ganglion cells.

Contrary to the traditional view, we find that significant redundancy exists both within and between different functional types of ganglion cells over a wide range of cell pair distances. This redundancy allows for robust coding of visual stimuli, such that the communication between eye and brain is tolerant to the misfiring of individual ganglion cells.

Support Contributed By: Lewis Thomas/Bristol-Myers Squibb Fellowship (JP) and the E. Mathilda Ziegler Foundation (MB)

PROPERTIES OF QUANTAL TRANSMISSION AT CA1 SYNAPSES

Sridhar Raghavachari and John E. Lisman
Volen Center for Complex Systems
Brandeis University, Waltham MA 02454

Rapid progress has been made in understanding the mechanisms of fast transmission mediated by the AMPA type of glutamate channel at central synapses. However, important properties of quantal transmission at hippocampal synapses synapses still remain unclear. We have constructed a model of AMPA channel activation and glutamate diffusion in the synaptic cleft which predicts the rise-time, amplitude and variability of quantal responses without no adjustable parameters. The simulations show that initial channel opening is driven by a glutamate spike near the site of vesicle fusion, producing a ”hotspot” of channel activation much smaller than many synapses. Quantal size therefore depends on local channel density but not on total number, a finding with important implications for measuring synaptic strength. Quantal size is stereotyped ($CV<0.2$) because channels approach saturation within the hotspot. Recent work shows that evoked responses can be multiquantal. We find that summation of multiple quanta is nearly linear even when vesicles are released very close to each other. Thus it appears that synapses have a much larger dynamic range than previously thought.

Whereas data on evoked release indicates that quanta are stereotyped (CV $\sim$ 0.2), measured mEPSCs have highly variable amplitudes (CV $\sim$ 0.6). According to one theory (Bekkers and Stevens, 1990, Franks et al. 2003), this arises from variation in vesicle glutamate content. However, both our modeling results and experiments (Ishikawa et al. 2003) indicate that this theory cannot account for the observed rise-time/amplitude correlation of mEPSCs. In contrast, this correlation and the high mEPSC variability can be accounted for if some mEPSCs are generated by two or more vesicles released with temporal jitter. Multivesicular mEPSCs have been demonstrated at other excitatory and inhibitory central synapses. In particular, amplitude histograms of mEPSCs at cerebellar mossy cells as well as thalamic relay cells show two or more narrow peaks that coincide with the peaks in the histograms of evoked responses.

We conclude that a broad range of results can be accounted for by simple principles: 1) quantal amplitude ($\sim$ 10 pA) is stereotyped, 2) some mEPSCs are multivesicular, 3) at moderate and large synapses, evoked responses are generated by quasi-linear summation of multiple quanta.

ACTIVITY-DEPENDENT REGULATION OF CONDUCTANCE THRESHOLDS

Kanaka Rajan E. Marder and L.F. Abbott
Volen Center & Department of Biology, Brandeis University, Waltham, MA

Homeostatic regulation of the intrinsic properties of neurons is likely to play an important role in the maintenance of stable network function. Activity-dependent changes in the ionic conductances of neurons is an essential component of homeostatic maintenance. Previous work studied the effects of activity on the maximal conductances of membrane currents and explored its homeostatic impact. Here we consider another aspect of intrinsic current maintenance, the activity-dependent regulation of activation curves.

Certain models based on neurons of the stomatogastric ganglion will only burst if the activation threshold of their hyperpolarization-activated cation conductance I_H lies in an appropriate range that extends over only a few mV. Furthermore, within this range, burst frequency is quite sensitive to the threshold value. Dynamic clamp experiments have revealed a similar sensitivity in real neurons (Sharp, Skinner & Marder, 1996). We use such a model to study the impact of allowing activity to regulate the threshold of I_H. In the regulated model, the threshold potential at which I_H is half activated is modulated at a rate proportional to the difference between the intracellular Ca^{2+} concentration and a target value for this concentration.

The model bursting neuron with activity-dependent I_H threshold exhibits robust stability with respect to shifts in its resting potential. The model automatically adjust the threshold so that the voltage operating range of the neuron falls on the sloped part of the I_H activation curve. In addition to keeping I_H active, this stabilizes the bursting activity of the neuron. The narrow operating range of the original, unregulated model is broadened when threshold regulation is included, and the burst frequency remains much more constant when the resting potential of the cell is shifted.

The unregulated model neuron tolerates a ± 5mV change in its resting membrane potential before it stops bursting. With threshold regulation, the bursting model neuron broadens its operating range over ± 10mV around a default value of -50mV. Over this voltage range, the change in burst frequency with voltage is about twice as sensitive to resting potential in the unregulated model as in the regulated model.

We are currently testing the idea that the activation threshold of I_H can be shifted by activity in real neurons. Using intracellular voltage-clamp recordings under synaptic block in dorsal gastric neurons of the crab, *Cancer Borealis*, we measure the I_H threshold before and after prolonged current-clamp induced shifts of activity.

NATURALISTIC AUDITORY CONTRAST IMPROVES SPECTROTEMPORAL CODING IN THE CAT INFERIOR COLLICULUS.

Heather L. Read[1], Monty A. Escabí [2], Lee M. Miller[3], and Christoph E. Schreiner[4]

[1]Department of Psychology, [2]Electrical and Computer Engineering, University of Connecticut, Storrs, CT 06269-2157; [3]Helen Wills Neuroscience Institute, University of California, Berkeley, CA 94720; [4]W.M. Keck Center for Integrative Neuroscience, University of California, San Francisco, CA 94143

Abstract: A central hypothesis of sensory coding suggests that sensory systems efficiently make use of statistical structure inherent in natural environments. The possibility that sensory systems are adapted for encoding natural signals has been a topic of discussion since the early work of Barlow (Barlow, 1953, 1961). Natural visual (Ruderman and Bialek, 1994; Dong and Atick, 1995; Ruderman, 1997) and acoustic signals (Voss and Clarke, 1975; Attias and Schreiner, 1998a; Nelken et al., 1999; Lewicki, 2002) show robust statistical properties such as scale invariant contrast statistics and 1/f modulation spectrum. Although numerous studies have looked at these statistical characteristics of natural signals, only a few studies have addressed how such statistics can be used for efficient sensory coding (Rieke et al., 1995; Dan et al., 1996; Attias and Schreiner, 1998b; Nelken et al., 1999; Stanley et al., 1999; Vinje and Gallant, 2002). Direct application of information theoretic approaches has revealed that sensory neurons respond most efficiently to sensory signals with natural statistics, although the exact mechanisms enabling such efficient processing have not been established.

We perform statistical analysis of natural sounds and speech and show that these have logarithmically-distributed (i.e., decibel or dB) spectrotemporal modulations that can cover several orders of magnitude. The dynamic range of spectro-temporal amplitude fluctuations in natural signals was closely matched to the average operating range of peripheral and central auditory neurons to pure-tones (~30-50 dB). By contrast, most artificial stimuli used to probe auditory function, including pure-tones and white noise, have linearly-distributed amplitude fluctuations with a limited average dynamic range.

MODELING THE NEURONAL BASIS OF THE PULFRICH EFFECT

Jenny C. A. Read[1], Bruce G. Cumming[1]
1. Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health

The Pulfrich effect is an illusion of stereo vision, in which delays between the signals reaching the brain from the two eyes lead to a perception of stereoscopic depth. In the classic demonstration, with one eye delayed, a pendulum which is in fact swinging from left to right in a plane in front of the observer is perceived as describing an ellipse.

The traditional explanation was that spatial and temporal disparities are geometrically equivalent. A pendulum moving in a plane would, with interocular delay, create exactly the same stimulus on the retina as a pendulum swinging on an elliptical path when no such interocular delay exists. However, in the 1970s a modified version of the Pulfrich experiment cast doubt on this interpretation. Here, the pendulum was illuminated only intermittently, as in a stroboscope. Yet the Pulfrich illusion still occurred. This presented a challenge to the accepted explanation. Stereopsis depends on correctly pairing the image of an object in one eye with its corresponding image in the other eye. In the classic stimulus, the brain can pair images of the pendulum which reach the eyes at the same time. If the right eye's image is delayed by Δt, then the brain ends up pairing the left eye's image of the pendulum at time t with the right eye's image of the pendulum as it was at some earlier time (t-Δt). Since the pendulum is moving, its position at time t is different from at time (t-Δt). There is therefore a spatial disparity between the paired images, leading to a perception of depth. The problem is that, in the stroboscopic experiment, there are no images of the pendulum which reach the eyes at the same time. If the strobe flashes at time t, both eyes see an image of the pendulum as it was at time t, but the left eye sees this image immediately, while the right eye does not see this image until time (t+Δt). The brain must therefore "remember" the left eye's image in order to pair it with the image which occurs later in the right eye. But since there is no spatial disparity between the paired images, it is unclear why depth should be perceived.

In order to explain this puzzle, modelers have invoked sensors which jointly encode both disparity and motion. These would be activated even by the stroboscopic stimulus, so it has been suggested that these could underlie the Pulfrich illusion. These explanations received a boost when electrophysiological experiments in the cat striate cortex revealed evidence of such joint encoding. However, recent experiments in our lab have shown that such cells are rare in the monkey. This made us wonder whether it was really necessary to postulate joint encoding of motion and disparity in order to explain the stroboscopic Pulfrich effect.

Previous analyses of the stroboscopic experiment have assumed that the stereo system matches the left- and right-eye images of the pendulum which occur closest together in time. These have zero spatial disparity. However, matches between images more widely separated in time do contain spatial disparity. Since cortical cells have finite integration times, temporally disparate matches should not be ignored. We show that such matches can explain both the stroboscopic and the classical Pulfrich effect. Under this scheme, if the weight given to a potential match falls off as a Gaussian function of temporal disparity, the dependence of the perceived depth on temporal delay is quite different from that predicted by joint motion-disparity encoding. We therefore carried out psychophysical experiments in which we measured perceived depth as a function of temporal delay. We found that the data followed the pattern predicted by our model. We conclude that sensors which encode only motion, or only disparity, can contribute to the perception of both the classic and stroboscopic Pulfrich effects.

SELF-SUSTAINED ACTIVITY AND FAILURE IN A SMALL-WORLD NETWORK OF EXCITABLE NEURONS

Hermann Riecke[1], Alex Roxin[1,3], and Sara A. Solla[1,2]
[1]Engineering Sciences and Applied Mathematics, [2]Physiology and Physics & Astronomy, Northwestern University, Evanston, Il 60208
[3]CNRS Neurophysics and Physiology of the Motor System, Université René Déscartes, 75273 Paris Cedex 06

While great strides have been made in unraveling the electrophysiological properties of neurons during the last half century, much less can be said of the networks through which these neurons interact. Despite a large body of theoretical work dedicated to analyzing the dynamics of networks of model neurons, their connection to networks in the brain remains tenuous. Often the actual functional topology of the brain region of interest is known only approximately at best. Thus, the uncertainty inherent in choosing an apt topology has led most researchers to assume such simplified models as all-to-all, locally-coupled, or randomly coupled networks. It is becoming increasingly clear, however, that many networks in the brain have a more complex topology. Recent interest in complex networks has lead to an array of analytical and numerical results that may be useful for researchers in neuroscience.

Inspired by the topology of the cortex, which exhibits both a sense of neighborhood (i.e. local coupling) as well as long-range excitatory connections, we consider a network of excitable integrate-and-fire neurons with a so-called 'small-world' topology. It is characterized by a bidirectional coupling between neighboring neurons and the addition of a fraction p of uni-directional short-cuts connecting randomly chosen neurons.

We find that for non-zero p the short-cuts allow the emergence of persistent activity in which pulses of excitation propagate through the network in a complex spatio-temporal pattern. As p increases such persistent activity undergoes a transition to over-excitation and hence extinction. We characterize the transition analytically using existing mean-field results for small-world networks. While for high propagation velocities failure occurs after a single burst of the population activity, for lower velocities the network can undergo 10,000 and more population bursts before it fails suddenly. In these long transients the activity is chaotic.

In the low-p regime the network is bistable and can be turned on and off between the rest state and the active state by a sufficient stimulus. Despite the complex topology the activity is typically periodic, but many spatio-temporal attractors co-exist. Possibly, this may allow complex 'memories' to be encoded in equally complex spatio-temporal patterns.

OLFACTORY INTERNEURONS IN THE BEHAVING ODOR-CONDITIONED MOUSE

Dima Rinberg[1], Michale Fee[2], Fred Ollinger[1] and Alan Gelperin[1]
[1] Monell Chemical Senses Center
[2] Brain and Cognitive Science, MIT

Recording neural activity in a sensory system in a behaving animal has three important aspects: 1. In an awake animal feedback from higher brain areas is fully active, which can change sensory information processing even at early processing stages; 2. The changes in the neural representation of an odor due to learning and plasticity can be recorded in real time; and 3. The animal's behavior can report the result of sensory information processing. Motivated by these considerations, we developed a system for studying olfactory information processing in the behaving mouse. A microdrive with three motors moving three electrodes optimized for single unit recording was implanted to record from the olfactory bulb of a mouse previously trained in an odor discrimination task. Simultaneous recordings of neural activity and mouse behavior showed a wide spectrum of different responses of mitral/tufted and granule cells in the olfactory bulb.

The water-deprived mouse was trained to discriminate two odors on a go/no-go task, such that one odor signaled water availability at an adjacent water spout while the other odor signaled the lack of water. After initial discrimination training, we then train the mouse to make a new association between an odor and water in one training session lasting 2 hours. We record nose pokes into the odor sampling port and the water reinforcement port, stimulus delivery timing along with the single-unit responses from olfactory bulb cells.

Temporal patterns of presumed mitral/tufted and granule cells were analyzed relative to the timing of nose pokes signaling onset of odor sampling or results of olfactory computations (responses at the water port). As an example of unexpected behavioral modulation of units in olfactory bulb, we found cells that reduced their firing prior to execution of a nose-poke into the odor sampling port.

As our sample of units recorded during odor-guided behavior increases, our analysis of single and multiunit recordings will be directed to understand the neural code for odor, the way it is read by the animal, and the way in which the code may be modified by odor learning.

HAPHAZARD WIRING OF SIMPLE RECEPTIVE FIELDS AND ORIENTATION COLUMNS IN VISUAL CORTEX

Dario L. Ringach
Department of Neurobiology & Psychology.
Jules Stein Eye Institute & Brain Research Institute
University of California, Los Angeles
Los Angeles, CA 90095-1563

The receptive fields of simple cells in visual cortex are composed of elongated ’on’ and ’off’ subregions, responsible for the generation of orientation selectivity. Neurons with similar orientation preferences cluster in ’columns’ that tile the cortical surface forming a map of orientation selectivity. Hubel and Wiesel proposed that simple-cell receptive fields are constructed by the selective pooling of geniculate receptive fields aligned in space. A recent analysis of monosynaptic connections between the lateral geniculate and cortical neurons revealed the existence of ’wiring rules’ that appears to support the classical model. The precise origin of the orientation map is unknown, but both genetic and activity dependent processes are thought to contribute. Here, we put forward a haphazard wiring model of simple receptive fields that is consistent with data on the probability of monosynaptic connections and generates orientation maps in the cortex. We suggest that developmental mechanisms may take advantage of the statistics of random wiring to generate a blueprint of simple-cell receptive fields and orientation tuning in visual cortex that is further refined by activity dependent processes. An experiment to test this hypothesis is proposed.

TACTILE NATURAL SCENES IN THE RAT VIBRISSA SYSTEM

Jason Ritt, Mark L. Andermann and Christopher I. Moore
McGovern Institute for Brain Research
Massachusetts Institute of Technology, Cambridge MA 02139.

The rodent vibrissa (whisker) system is an important model for sensorimotor processing. During vibrissa-based tactile behaviors, rats typically employ a combination of “whisks” (rhythmic vibrissa sweeps) and head and body movements. The mechanical signals transduced via vibrissae during these behaviors are largely unknown, making it difficult to relate neural activity, even at the periphery, to the external environment. This is especially the case for textural judgments, in which the rat needs knowledge not only of times of object contact, but also of presumed high-frequency, small amplitude “micro-deflections” resulting from interactions with the surface texture.

We examined transduction properties of ex vivo (isolated) vibrissae driven by a precision motor, to simulate whisking, over a variety of textured surfaces. Our preparation permitted precise control of velocities and contact points. A high-speed (4kHz) digital camera, with automated image processing, provided high temporal resolution time series of vibrissa deformations resulting from surface interactions.

We found that vibrissa mechanical responses ranged from reliable transductions of surface features to seemingly chaotic motions, depending on kinetic parameters and surface type. For a given vibrissa, many conditions could elicit oscillations at a common high frequency ($>$ 80 Hz), implicating mechanical resonance. Moreover, frictional effects were evident in “stick-slip” motions along the surface, sometimes producing dramatic high-frequency (e.g., 1 kHz) transients. The details of the response were strongly determined by the driven velocity and the degree of vibrissa bending, parameters likely controlled by behaving animals.

These results are a preliminary step toward characterizing vibrissal “tactile natural scenes”, in analogy with previous visual and auditory studies, and will guide future electrophysiological and behavioral investigations.

CANNABINOIDS AND ENDOCANNABINOIDS ON THE HIPPOCAMPAL NETWORK OF BEHAVING RATS

David Robbe and Gyorgy Buzsaki
Center for Molecular and Behavioral Neuroscience, Rutgers, The State University of New Jersey and Neurological Institute of New Jersey, Newark, NJ 07102 USA

Cannabinoids have an impact at various levels of neuronal functions. Cannabinoid G-coupled receptors (CB1) are expressed selectively on the GABAergic terminal of CCK positive interneurons and endocannabinoids are release by pyramidal cells following sustained excitation. In vitro, the endocannabinoid system has been shown to contribute to synaptic plasticity, a candidate model for memory storage. At the behavioral level cannabinoids can affect memory formation. In contrast to recent progress on the effects of cannabinoids at the single cell and behavior levels, their in vivo impact of large-scale neuronal activity is poorly understood. We approached this issue by combining tetrode recordings in the hippocampal CA1 region of freely moving rats before and after systemic injection of agonist and antagonist of the CB1 receptor. Preliminary results indicate that cannabinoids 1) decrease “ripple” (200 Hz transient oscillatory epoch) incidence and at the same time increase the intra-ripple frequency of the remaining ripples; 2) decrease power of theta and gamma oscillation during running. We found that these latter effects are independent of the running speed changes brought about by cannabinoids. Importantly, all these effects were reversed by subsequent antagonist injection, indicating the involvement of CB1 receptor-dependant mechanisms. At the unit level, both agonist and CB1 antagonist appear to regulate the mean firing rate and interspike interval distribution, tested during slow-wave sleep episodes, where behavioral modulation of spike patterns are at minimum. These findings may help bridge the observations in vivo and the behavioral alterations brought about the cannabinoids. They also give insight on how endocannabinoids regulate neuronal firing in vivo.

THE ROLE OF AREA LIP IN A DIRECTION DISCRIMINATION TASK WITH MULTIPLE REWARD CONTINGENCIES.

AE Rorie and WT Newsome. HHMI and Dept of Neurobiology, Stanford University, Stanford, CA 94305.

In deciding between potential actions, animals must consider the sensory information supporting each action, and the associated potential gains. Previous research from our lab has investigated how sensory and reward information influence decisions independently. Here we report behavioral results that provide insight into how these two sources of information interact to influence perceptual decisions.

A rhesus monkey discriminated the direction of motion in a family of variable-coherence random dot stimuli employed in many prior studies from this lab (2AFC task). Coherence levels were chosen to span psychophysical threshold. Additionally, the monkey was cued on each trial that one of three reward contingencies was in effect: 1) a correct decision in either direction resulted in a large reward (two drops of fruit juice), 2) a correct decision in either direction resulted in a small reward (one drop of fruit juice), or 3) a correct decision in one direction resulted in a large reward while a correct decision in the other direction resulted in a small reward. All trial types were randomly interleaved and spatially balanced when the potential rewards differed.

Psychophysical data were virtually identical (both threshold and bias) when decisions favoring the alternative directions were rewarded equally; the absolute size of the reward exerted no effect on performance. Reward information exerted a substantial effect, however, when the rewards for the two directions differed. Decisions were biased toward the direction associated with the larger reward without affecting overall psychophysical sensitivity, resulting in a simple lateral shift of the “unfolded” psychometric function. This implies that information concerning reward magnitude exerts a simple additive bias on the neural mechanisms responsible for integrating motion. The magnitude of the shifts, expressed in units of the visual stimulus, can be as high as 20-30% coherent dots, although smaller effects are more common.

Additionally, we have developed a quantitative method of determining the amount of bias optimal for, given the monkey’s psychometric function, harvesting rewards in each of the three reward contingencies. Interestingly, in the third reward contingency, when the rewards are unequal, the monkey bias’s consistently exceeded the optimal. We hypothesize this is the result of a discrepancy between the targets objective and subjective value. It is well known from behavioral economics and foraging theory that subjective, not objective, value is factored into decisions of this kind. Accordingly, we have developed a method to calculate this subjective value, which in keeping with previously mentioned fields we call utility. While the objective value of the large reward is consistently two drops of fruit juice, the monkey’s utility, or his calculated subjective value, is often greater. While his utility can be as large as ten drops of juice it is typically five drops.

Finally, neurophysiological data were obtained from the lateral intraparietal (LIP) area of the monkey’s brain while performing this task. It has been proposed that area LIP represents the value of potential eye movement targets in terms of their relative expected utility. Expected utility is the product of a reward’s utility and the probability of achieving that reward. By combining our calculated utility and the monkey’s probability of being rewarded on a given trial (as described by the psychometric function) we can calculate expected utility. This enables us to ask if systematic changes in the firing rate of single neurons in LIP correlate with changes in expected utility.

DETECTING LEARNING-INDUCED SYNAPTIC PLASTICITY IN THE BRAIN.

Simon Rumpel[1], Joseph LeDoux[2], Roberto Malinow[1] and Anthony Zador[1]
[1] Cold Spring Harbor Laboratory, Cold Spring Harbor, NY, USA
[2] W. M. Keck Foundation Laboratory of Neurobiology, Center for Neural Science, New York University, New York, NY, USA

How does the brain store information for long periods of time? We have tested the hypothesis that long-term changes in the efficacy of synaptic transmission are a mechanism of memory formation. To detect synaptic plasticity in the living brain we have utilized the recent finding that potentiated synapses show an increased amount of GluR1 subunit containing AMPA receptors. This GluR1 containing type of AMPA receptors is additionally incorporated into synapses during long-term potentiation and accounts at least in part for the increase in synaptic efficacy.

We used a herpes-simplex virus based system to overexpress AMPA receptor GluR1 subunits in neurons of the lateral amygdala, a structure known to be essential for memory formation during cued classical fear conditioning. Overexpression of GluR1 subunits leads to the formation of homomeric AMPA receptors, which can be electrophysiologically identified when contributing to synaptic transmission due to altered rectification properties as compared to endogenous AMPA receptors. In a first series of experiments we have conditioned infected adolescent rats in a classical fear conditioning paradigm. Subsequently, AMPA receptor mediated transmission at projections from the auditory thalamus to the amygdala was analyzed in brain slices. Rats that received paired training (tones and footshocks) showed significantly increased rectification (indicating synaptic incorporation of recombinant GluR1) compared to animals from the unpaired control group. In a second experimental series we blocked synaptic incorporation of GluR1 containing AMPA receptors by viral overexpression of a crucial fragment of the receptor. Fear memory was severely impaired in those animals as compared to animals injected with a control virus.

We conclude that long-term synaptic plasticity underlies memory formation in the lateral amygdala during fear conditioning. Additionally, we conclude that the use of tagged AMPA receptors is a powerful tool to detect synaptic plasticity in the living brain and to functionally label the elements of neuronal circuits that undergo changes during memory formation.

EQUALIZATION OF SYNAPTIC EFFICACIES ON PASSIVE, ACTIVE, AND BRANCHED CABLES BY ANTI-STDP

Clifton C. Rumsey and L.F. Abbott, Volen Center for Complex Systems and Department of Biology, Brandeis University, Waltham, MA 02454

The impact of an individual synapse on the firing pattern of a neuron depends not only on its intrinsic strength, but also on its dendritic location in relation to the overall morphology and conductance profile of the neuron. Because of attenuation along dendritic cables, a distal synapse may have to be intrinsically stronger (i.e., have a larger synaptic conductance) than a proximal synapse to have an equal effect on postsynaptic firing, and hence an equal synaptic efficacy. Some cell types, including CA1 pyramidal neurons, have been shown previously to scale up synaptic strengths with increasing distance from the soma such that excitatory postsynaptic potentials are equalized at the soma. How might such a configuration be achieved? Using a multi-compartment model neuron that receives a number of excitatory and inhibitory inputs, we show that a form of spike-timing dependent plasticity (STDP), called anti-STDP, that has a timing dependence opposite to conventional STDP, provides a mechanism for producing such an equalization of synaptic efficacy. This result requires that anti-STDP be combined with a non-Hebbian form of plasticity, such as nonassociative potentiation, for stability. Excitatory synapses subject to both anti-STDP and nonassociative potentiation are modified such that distal synapses are strengthened and proximal synapses are weakened, equalizing synaptic efficacies independent of synaptic location. This equalization is stable over time. We demonstrate this result for both unbranched cables and neuron models with realistic morphologies. We also show that it works with either passive or active dendrites. For active dendrites, synaptic efficacies are equalized under conditions that support back-propagating action potentials and in cases where dendritic action potentials are generated distally. Thus, whereas conventional STDP is useful as a form of competitive Hebbian learning, anti-STDP may be a useful homeostatic mechanism for eliminating the location-dependence of synapses.

Modeling Sensorimotor Adaptation as Linear Dynamical System

Philip N. Sabes and Samuel J. Sober
Department of Physiology and the W.M. Keck Foundation Center for Integrative Neuroscience,
University of California, San Francisco, CA 94143-0444

Many groups have suggested that sensory integration follows a minimum variance strategy in which less variable sensory signals are weighted more than noisier signals. This hypothesis states that the nervous system uses a single, optimal weighting of sensory signals. While most of the work supporting this hypothesis has focused on purely perceptual tasks, we have been studying sensory integration in the context of a sensorimotor task. We have previously shown that when planning a reaching movement, the nervous system relies on two separate estimates of arm position: one to plan the direction of the reach, and a second to convert that direction into a motor command (Sober and Sabes, 2003). These position estimates are based on the same peripheral sensory signals but rely on different combinations of visual and proprioceptive input, suggesting that the brain uses weighting criteria that depend on the computation performed at each planning stage. Here we show that altering the details of the sensorimotor task changes the relative weighting of vision and proprioception used in computing these estimates. Furthermore, these trial-by-trial changes affect each position estimate independently. The observed pattern of changes reflects a reduced reliance on signals that must be transformed between visual and proprioceptive coordinate frames, a previously undescribed strategy that minimizes the errors resulting from inherently noisy coordinate transformations. Our results therefore show that sensory integration is not determined solely by sensory variability but is significantly influenced by factors related to the demands of the task. These findings point to a view of sensory integration as a set of local, independently controlled processes optimized to improve sensorimotor performance.

NEURAL RESPONSE PARAMETERS AFFECTING ACCURACY AND PREDICTIVE POWER OF LINEAR STRF METHODS

Srivatsun Sadagopan and Xiaoqin Wang
Departments of Neuroscience and Biomedical Engineering,
Johns Hopkins University, Baltimore, MD 21205

Spectrotemporal response fields (STRFs) are general characterizations of the linear component of a neuron's response to complex auditory stimuli. In recent years, STRF methods have been widely applied to quantitatively characterize response properties of neurons at different processing stages of the auditory system with various degrees of success. However, the selection of the stimulus repertoire (i.e., its statistics) is crucial for STRF methods and must depend on the properties of the neurons under study. Because of the temporal dynamics of a neuron's response, stimulus statistics are often effectively altered when stimuli are presented continuously, and this may affect the accuracy of the STRF estimate. Therefore, a systematic study of the dependence of the STRF on neural response properties is essential to assess the applicability of these methods. In this study, we investigate the effects of parametric variation of response attributes such as duration, latency, spiking pattern (bursty vs. regular), jitter and adaptation on the accuracy and predictive power of a simple, linear STRF-based characterization of the neuron's response field. Neuronal selectivity is modeled as template-match filters that govern the probability of discharges (non-homogenous gamma or Poisson distribution). We then perturb the neuron's intrinsic firing properties to introduce systematic changes in response dynamics and we determine the sensitivity of the STRF estimate to these changes. We demonstrate that the application of simple STRF estimates is restricted to a subset of possible neural response types. We suggest a neuron-centric approach for designing STRF stimuli, where an initial estimate of a neuron's preferred region of stimulus space dictates choice of subsequent parametric stimuli. Though our response model is based on auditory cortical responses, we believe that the concepts we explore are applicable to other sensory neurons as well.

SINGLE CELL MEMBRANE POTENTIAL OSCILLATIONS OPTIMIZE STIMULUS REPRESENTATION IN NEURONAL SYSTEMS

Andreas T. Schaefer[1], Hartwig Spors[1] and Troy W. Margrie[2]

[1] WIN Group of Olfactory Dynamics
MPI for medical research; Jahnstr. 29, 69120 Heidelberg, Germany
[2] The Wolfson Institute for Biomedical Research
Dep. Physiol, UCL, Gower Street, London WC1E 6BT, UK

This study examines the question of whether a postsynaptic neuron can deduce specific patterns of processed input within a presynaptic cell. Within neurons, time-varying currents flow into the axon hillock and result in a temporally structured action potential (AP) train. In vivo, the temporal precision of the AP will be frustrated by various noise sources including channel noise and background synaptic activity that is not stimulus specific. Here we show that ongoing oscillatory rhythms in individual neurons permit postsynaptic cells to readout input-specific AP discharge patterns in their presynaptic counterparts.

Oscillatory activity ranging from theta to gamma frequencies has been observed across most brain regions and neuronal cell types both in vivo and in vitro. In the mammalian olfactory bulb in vivo, mitral cells display a subthreshold theta oscillation that is locked to the nasal inhalation cycle. Therefore, under experimental conditions when nasal airflow is decoupled from respiration by double tracheotomy, these subthreshold oscillations in mitral cells disappear. They can however be restored by rhythmic nasal inspirations. These can be coupled to the normal inspiration to the lung of the animal. We can therefore take advantage of this system to compare oscillation and non-oscillation dependent AP precision. Due to the clustering of AP discharge around the peak of the subthreshold oscillation, AP precision was substantially increased when subthreshold oscillations were restored. Additionally we find that within each oscillation cycle, AP jitter increased with AP number approaching values for APs evoked without the underlying oscillation.

To assess whether increased AP precision indeed enhances information transmission we investigated input representation in integrate-and-fire and Hodgkin-Huxley neuron models. Poisson spike trains were used as input stimuli while noise current was directly injected. We found that the presence of the subthreshold oscillation greatly enhances the discriminability of stimuli as assessed by a template matching scheme for the AP discharge. The same discrimination performance was achieved in the presence of 5 to 10 fold higher noise levels. This was insensitive to oscillation frequency (2-60 Hz), amplitude (2-20 mV), shape (skewed sine, sawtooth) or the details of the template matching scheme and increased with the size of the stimulus space.

Spike jitter in general is determined by the noisy current and the accumulated imprecision of earlier spikes. This accumulation is circumvented if a hyperpolarizing current injection results in a period of reduced spiking probability allowing the timing of the next spike to accurately reflect input stimuli and providing a mechanistic explanation for the increased discrimination performance. We therefore conclude that one immediate result of ongoing oscillatory activity is to improve the reliability of communication between pre- and postsynaptic neurons irrespective of assumed coding strategies.

FOLLOWING MOUSE TRAILS AND LISTENING TO BIRD SONGS: INSIGHTS INTO BEHAVIOR USING PREDICTIVE INFORMATION TECHNIQUES.

Katrin Schenk, S.J. Bonasera, L.H. Tecott, A.J. Doupe
Sloan-Swartz Center for Theoretical Neuroscience, Departments of Psychiatry and Internal Medicine, University of California, San Francisco

We present initial studies using predictive information (Bialek et al *arXiv-physics* 0007070, 2001) to examine mouse home cage exploration following treatment with the dopamine (DA) reuptake blocker GBR 12909, & individual differences in adult zebra & bengalese finch song syntactic complexity.

Predictive information correlates with the sub-extensive part of the Shannon entropy & is the mutual information between past & future observations of data. Predictive information also correlates with intuitive notions of system complexity; other complexity measures (e.g. Kolmogorov) relate to the Shannon entropy and thus correlate more with system randomness. We use a method similar to Strong et al (*Phys Rev Lett* **80**:197, 1998) to calculate measures of system randomness (entropy rate), & complexity (predictive information).

We initially applied this approach to analyze mouse home cage exploration (a behavior dependent on DA-ergic neurotransmission) following treatment with GBR 12909. Groups of mice received injections of GBR 12909 (vehicle, 3, 10, 30 mg/kg) in their home cage; position versus time data was sampled & binned. We then calculated entropy rate & predictive information. ANOVA found no effect on predictive information, but a statistically significant increase in entropy rate with increasing doses of GBR 12909. Thus, increasing doses of GBR 12909 does not alter the overall complexity of home cage exploratory behavior, but does increase the random noise that influences these paths. These findings may be interpreted to reflect the behavioral consequence of an inappropriately increased corticostriatal signal-to-noise ratio within striatal spiny neurons, leading to an inappropriate disinhibition of thalamic & cortical motor output systems.

We then applied this approach to the analysis of finch song syntactic complexity. Finch song is a complex, learned behavior produced by the avian forebrain. Bengalese & zebra finch songs were broken into syllables & analyzed as above. Preliminary results suggest that trends in the entropy rate & predictive information mirror the effects of known song features, such as syllable repitition & transition variability. Predictive information approaches may thus be useful in studying sensorimotor learning in birds by characterizing changes in song after manipulations such as deafening & CNS lesions. Supported by NIH F32MH67423-02 (KAS) & K08MH065983-01 (SJB).

PLANNING FOR UNCERTAINTY: BAYESIAN MODEL FOR HUMAN REACH AND GRASP

Erik J. Schlicht and Paul R. Schrater
Department of Psychology.
University of Minnesota, Minneapolis, MN 55455

If predictions about the expected quality of perceptual information are available, then the human motor control system can use that information to produce action plans that help maintain a desired level of task performance. In this study, we psychophysically investigate whether humans plan for sensory uncertainty when making reaching movements. To test this idea we had subjects repeatedly reach to the same object while gaze direction was varied, both when the target was in view and when it was occluded. We found, for visible targets, maximum grip aperture (MGA) was at a minimum near the target location and increased linearly away from the target. More surprisingly, MGA for reaches to occluded targets show dependence on eye position, even though there is no visual information specifying target location. This dependence takes the form of a U-shaped MGA function centered near the forward-view eye-position, irrespective of the target location. We developed a Bayesian model that can account for these changes in MGA. The model assumes that MGA is a measure of overall spatial uncertainty in the observers estimate of the target location (Wing et al., 1986). Spatial uncertainty is modeled as stemming from 4 major sources: noise in visual (1) and haptic (2) estimates of the target location, uncertainty resulting from sensorimotor transformation noise (3), and partial use of experience across trials (4). By making these assumptions, we are able to demonstrate that people appear to adjust their MGA to accommodate their estimate of the optimal overall uncertainty. In addition, we suggest that target locations are converted and stored in eye-centered coordinates, even when the information about target location is not visual.

HOW FEATURE DISCONTINUITIES ARE DETECTED IN CAT VISUAL CORTEX

Anita M Schmid and Kevan AC Martin
Institute of Neuroinformatics
Uni Zurich/ETH Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

Borders help define objects and thus give us important information about the structure of our visual environment. In cat area 17 (V1) neurons are especially responsive to luminance borders. How cortical neurons represent and process discontinuities defined by cues other than luminance, is less well understood.

We recorded the responses of single units in area 17 of anesthetized cats. The stimuli were gratings that included ’feature discontinuities’ defined by changes in phase or orientation or direction of motion. Nine of 32 neurons showed significantly enhanced firing rate when a feature discontinuity was close to their classical receptive field. The median of enhancement was about 180% for all stimuli. Contrary to previous results (1), the effect was as prominent with the phase discontinuity as with orientation and direction discontinuities.

To explore whether the enhanced firing rate was due either to facilitation or to release from suppression, we explored the role of the end-zone of the receptive field. In this second set of experiments 17 of 34 neurons showed enhanced firing to a feature discontinuity and 14 of these also showed end-zone inhibition (’end-stopped cells’). Four end-stopped cells showed the effect for all three stimuli shown, 3 neurons for combinations of two effects and 7 neurons showed the effect only for one of the three discontinuities. Four end-stopped cells did not show any effects. Thus the end-zone specificity for different features seems quite variable across cells.

Further analysis showed that the enhanced firing was principally due to release of a feature-specific suppression. Only one neuron showed facilitation. This result is different from the findings of Sillito *et al.* (2) who found that facilitatory mechanisms dominated the detection of orientation discontinuities.

The results show that end-stopped cells not only detect line endings and curvature, but can serve also to detect discontinuities defined by changes in phase or orientation or direction.

(1) De Angelis *et al.*, 1994, J. Neurophysiol., 71(1): 347-374
(2) Sillito *et al.*, 1995, Nature, 378: 492-496

Supported by EU and HFSP

PITCH REFLECTS THE STATISTICAL RELATIONSHIP BETWEEN AUDITORY STIMULI AND NATURALLY OCURRING PERIODIC SOUND

David A Schwartz and Dale Purves
Duke University

It has long been recognized that the pitch a listener hears in response to a sound stimulus is not a simple function of stimulus frequency structure or of any other physical parameter of the stimulus. We examined the hypothesis that the unexplained pitch phenomenon of a *spectral dominance region* reflects the auditory system's embodiment of the statistical characteristics of naturally occurring periodic sound in the human auditory environment. *Spectral dominance* refers to the empirical observation that when the frequencies of a subset of the harmonics in a stimulus are increased or decreased by a constant proportion -- such that the fundamental frequency of the altered harmonics differs from that of the unaltered set -- the pitch heard corresponds to the fundamental frequency of those harmonics that occupy a frequency band centered at ~600 Hz. Using a database of 68,000 speech sounds, the principal source of periodic sound energy for human listeners, we determined the relative likelihood that each speech sound could have contributed periodic sound energy to a given artificial complex tone stimulus. The conditional probability distribution of speech sound periodicities for a given test stimulus accurately predicted the pitch heard in response to that stimulus and the predominance of spectral components in the vicinity of 600 Hz in determining pitch; accurate pitch predictions could not be derived, however, from a similar analysis of periodic non-speech sounds (sawtooth waves). These findings support the hypothesis that pitch is based on the probabilistic relationship between an auditory stimulus and naturally occurring periodic sound.

STATISTICAL NORMALIZATION OF SENSORY REPRESENTATIONS

Odelia Schwartz[1] and Terrence J Sejnowski[1] and Peter Dayan[2]
[1]The Salk Institute and HHMI, La Jolla CA and [2]GCNU, UCL, London

A critical question in the neural representations of the sensory world, is how sensory systems can build multi-layer, hierarchical, representations whose semantics are as uniform as possible at every stage. Here, we demonstrate that a *statistically* normalizing transform on wavelet outputs can license effective hierarchical representations. The figure below outlines our technique in the case of images. (A) shows a classic bowtie dependence between two spatially adjacent wavelet coefficients; and (A: top and left) shows the sparse marginal distributions of the coefficients. We suggest performing the non-linear *marginal* transform in (B), giving rise to a standard Gaussian marginal distribution (hence statistically normalized; C: top and left), and a contrast enhanced bowtie (C), with a bimodal conditional distribution of one coefficient given the other. Applying ICA to the non-linearly transformed coefficients, yields higher level units capturing long range structure in the inputs (D; wavelet feedforward weights, black negative; white positive).

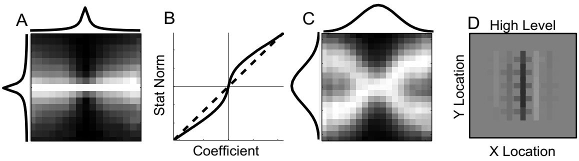

We consider statistical normalization in relation to three related frameworks: Gaussian scale mixture models (GSMs), divisive normalization and sparse coding. Crudely, GSMs model sensory input in terms of a hierarchical *scale* process multiplying a hierarchical standard Gaussian. Sparse coding models such as independent component analysis (ICA) gain purchase because of the leptokurtotic properties of GSMs, extracting wavelet-like filters. ICA typically cannot be iterated, since the untransformed coefficients are already strongly sparse. Divisive normalization can eliminate the bowtie dependence arising because of the common scale factors; but unsupervised learning based on the transformed outputs is difficult since they have strong joint Gaussian structure. In our technique, the contrast enhanced bowties arise because the normalizing transform ties the output to an estimate of the local scale factor, and hence to GSMs (we expect optimizing the target marginal in (C) to make this relation more precise). Our results suggest that this allows hierarchical ICA to succeed (D). We are exploring properties of the contrast enhanced bowties for visual and auditory inputs; and connections with adaptation and, via conditional rather than marginal standardization, divisive normalization.

DECODING ORIENTATION SELECTIVITY MODELS

Peggy Seriès[1], Peter Latham[2] and Alexandre Pouget[1]
1. Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY, USA, {pseries, alex}@bcs.rochester.edu.
2. Gatsby Computational Neuroscience Unit, Alexandra House, 17 Queen Square, London WC1N 3AR. England.

The origin of orientation selectivity in primary visual cortex (V1) is still a hotly debated issue, which often serves as a model problem for understanding cortical circuitry and computation. Two major classes of models have been proposed: (i) The feedforward or ‘no sharpening’ (NS) model, in which the narrow selectivity is the result of the convergence of LGN afferences onto cortical cells; (ii) The recurrent or ‘sharpening’ (S) model, in which the pooled input from the LGN is broadly tuned, and the activity is subsequently sharpened through cortical interactions. These models have been found to be difficult to distinguish experimentally. Moreover, because they generate similar tuning curves, they are commonly thought to produce similar codes for orientation. We show here on the contrary that these two models markedly differ in their network properties and statistical efficiency.

We have implemented a S model, which follows Somers et al (1995) and a NS model which is based on a simple Hubel & Wiesel model with feedforward inhibition (Ferster and Miller, 2000). We quantify information in these networks in terms of the smallest change in stimulus orientation that can be detected based on the networks population activities.

First, our results indicate that even when the tuning curves and spike train variability of individual neurons in the S and NS models are matched, the pairwise correlations of the fluctations of activity are markedly different. Second, the effect of the correlations is such that the population code in the NS model conveys far more information about orientation than the S model. Moreover, the majority of the information in the S model is conveyed by correlations, making the S model particularly inefficient for learning and computation. These results demonstrate that sharpening through lateral connections is not always as beneficial as generally believed. Our approach also provides new insights on the influence of correlated variability on the accuracy of population codes in realistic neural circuits.

References

Ferster, D and Miller, K. Neural mechanisms of orientation selectivity in the visual cortex. Annu Rev Neurosci. (2000), 23:441-71.

Somers DC, Nelson SB, Sur M, An emergent model of orientation selectivity in cat visual cortical simple cells. J Neurosci. (1995) Aug, 15(8):5448-65.

ESTIMATING ENTROPY RATES AND INFORMATION RATES IN SPIKE TRAINS WITH CONFIDENCE INTERVALS

J Shlens, M Kennel, HDI Abarbanel and EJ Chichilnisky
Salk Institute and UC San Diego, La Jolla CA 92037

Measurements of entropy and information rates in spike trains may be valuable tools for deciphering the neural code. However, direct estimates of these quantities lack reliable confidence intervals and exhibit bias, which can be particularly severe when analyzing real spike trains at high temporal resolution. Two opposing biases result from (1) temporal dependencies and (2) undersampled probability distributions. These biases significantly complicate the selection of word length – the free but crucial parameter in conventional entropy rate calculations.

We propose a new approach to estimating entropy and mutual information rates that substantially alleviates both problems. The approach derives from the perspective that choosing an estimate of entropy rate effectively involves selecting the appropriate model complexity for a finite data set and is best solved under the guiding principle of Minimum Description Length.

Following this principle we select an appropriate weighting of probabilistic models using a compression method known as Context Tree Weighting. We model a spike train as a weighting of finite-order, discrete Markov processes and use Bayesian estimators of entropy on this weighting of models to estimate the entropy rate directly. Using Monte Carlo techniques we sample the likelihood that a model could generate a range of entropy rates; this provides Bayesian confidence intervals on our estimate.

Using simulations with known entropy rates, we tested the performance of this estimator and compared its convergence and consistency to several other methods (e.g. direct method, string-matching, Lempel-Ziv compression). We also developed a series of diagnostic tests to aid in determining the convergence of information estimates. These diagnostics may be used to select stimulus parameters that minimize the overall variance of information estimates in an experiment of fixed duration.

Our first application of these tools is to investigate the temporal precision of mammalian retinal ganglion cell light responses by examining the dependence of information rates on the temporal resolution with which spike times are represented. We also assess the method on a simple Poisson model neuron with refractory dynamics. These results suggest the range of time scales over which retinal ganglion cells encode information as well as a conservative temporal limit beyond which ganglion cells encode no additional information.

NSF IGERT and La Jolla Interfaces in the Sciences (JS), McKnight Foundation, Sloan Foundation, and NEI grant 13150 (EJC).

OPTIMAL INFORMATION REPRESENTATION IN SENSORY RECURRENT NETWORKS THAT OPERATE NEAR BIFURCATION POINTS

Oren Shriki
Department of Neurobiology.
The Weizmann Institute of Science, Rehovot, Israel 76100

Recurrent connections are abundant in cortical circuitry but their functional role has been the subject of intense debates. In this work we shed light on this problem by developing an information theoretic approach to investigate the computational role of recurrent connections in sensory processing. The basic network model consists of two layers of neurons. There are feedforward connections from the input layer to the output layer and recurrent connections among the neurons in the output layer. The model neurons are characterized by positive continuous activities and a non-linear input-output relation. We use the principle of maximum information preservation to evaluate the optimal pattern of recurrent interactions and its dependence on the statistics of the external inputs.

As a case study, we focus on a hypercolumn in the visual cortex consisting of neurons that process contrast and orientation information in a local patch of the visual field. The output layer of the network serves as a simplified model of a hypercolumn. This is obtained by setting the feedforward connections to provide a weakly tuned input that mimics the input from the Lateral Geniculate Nucleus. The network is then presented with inputs that span all possible orientations and contrasts. In particular, we studied the effect of the contrast distribution. When the typical input contrast is low, as is the case in natural scenes, the optimal profile of interactions between orientation columns has a ”Mexican hat” shape, which helps to amplify the inputs to the network and span more efficiently the neuronal dynamic range. The maximal amplification is achieved when the amplitude of the interactions approaches a critical value, for which the network undergoes a bifurcation. Below this value the response to a homogeneous input is also homogeneous, while above it the response to the same input becomes inhomogeneous, peaking at a certain orientation.

We propose that operating near bifurcation points may be a general principle used by recurrent networks in order to increase their sensitivity to external inputs. Under certain conditions, such as high plasticity levels, networks may get ’too’ close to the critical point and even cross it to the ’pattern formation’ regime. This can be manifested as hypersensitivity and hallucinations. We hypothesize that such a mechanism plays a key role in a variety of neurological disorders.

A NEUROBIOLOGICAL SIMULATION OF DEDUCTIVE REASONING

Ray Singh and <u>Chris Eliasmith</u>
University of Waterloo

There remains a broad gap between accepted models in cognitive neuroscience and those in systems neuroscience, despite a similar interest in brain-based explanations of behavioral phenomena. In this paper we apply the neural engineering framework (NEF) described in Eliasmith and Anderson (2003) to help bridge this gap by constructing a biologically detailed model of a traditionally cognitive phenonema. Our focus is on deductive reasoning.

Deductive reasoning is an ideal exemplar because it has been widely suggested that neurally plausible architectures are unable to support the kinds of structure sensitive computations necessary for explaining deduction. Some have argued that this demonstrates that understanding neural computation is not relevant for understanding cognitive function (see Fodor and Pylyshyn, 1988; Jackendoff, 2002). However, since the early 1990s, there have been a number of suggestions as to how structure sensitive processing can be supportd by distributed representations (including Spatter Codes (Kanerva, 1992); HRRs (Plate, 1991); and Tensor Products (Smolenksy, 1989)).

Nevertheless, none of these suggestions have been implemented in a biologically plausible setting. The NEF provides principles by which this can be done. To model the most prevalent deductive reasoning task used in psychology (the Wason card selection task; Wason, 1966; Cheng and Holyoak, 1985, 1995), we present a model which integrates HRR representations and the relevant physiological and anatomical data from frontal cortices (Platt and Glimcher, 1998; and see below).

Specifically, we build a dynamic spiking network model that learns the relevant transformations of structured HRR representations in a context sensitive manner. The network consists of 3 main parts: bilateral VMPFC which provides familiarity/context information that is used to select the appropriate transformation (Adolphs et al., 1996); left language areas which provide HRR representations of the deductive rule to be tested (Parsons and Osherson, 1998, 1999); and right inferior frontal cortex which combines VMPFC and HRR information to select and apply the appropriate transformation (Parsons and Osherson, 2001), resulting in a set of actions on the cards. It is during the application of the transformation that learning also occurs. This biologically plausible network accounts for differences in the typical correct versus incorrect responses to content-dependent and content-independent versions of the Wason task respectively. In addition, it explains why those trained in logic do better on the content-independent tasks (Rinella et al., 2001).

A DYNAMIC MODEL OF WORKING MEMORY IN THE PFC DURING A SOMATOSENSORY DISCRIMINATION TASK

Ray Singh and Chris Eliasmith
University of Waterloo

Romo et al. have recently presented detailed measurements of neural activity in working memory during a somatosensory discrimination task (Romo et al. 1999; Brody et al. 2003). Their results show that the content of working memory is not only stimulus dependent but also strongly time-varying. We model this behavior using the Neural Engineering Framework (NEF) developed by Eliasmith and Anderson (2003).

Of particular importance to this model is the NEF's population-temporal encoding function: $\sum\limits_{n} \delta(t - t_{in}) = G_i[\alpha_i \left\langle \tilde{\phi}_i \mathbf{x}(t) \right\rangle_m + J_i^{bias}]$ This equation describes how the ith neuron in a population represents a parameter, **x**(t), as a spike train, $\sum_n \delta(t - t_{in})$. The maximal firing of a neuron is determined by its preferred direction vector $\tilde{\phi}_i$, the gain α_i, the background current J_i^{bias}, and the neural nonlinearity $G_i\left[\cdot\right]$ for which we use a LIF neuron. Here, we model the content of working memory, **x**(t), as a 2D vector which is the representation of stimulus frequency and expected delay time. We then define a neural integrator over this population representation. This simple model captures the dynamics of neural spiking observed by Romo et al.

Miller et al. (2003) have recently proposed a model using LIF neurons to simulate these experiments. However, our model more naturally captures the experimental data. Miller et al. dissect the problem into two independent parts, memory and expected delay time. They initially propose a neural integrator model of stable memory, avoiding the dynamics observed by Romo et al. To reproduce the dynamics, they subsequently propose a network of three neural integrator populations. Notably, this network only captures observations for positive monotonic responses.

In contrast, we use a single neural integrator to capture the dynamic phenomena for both positive and negative monotonic responses. By taking the neurons to be sensitive to both time and frequency, we show that all of the observed response curves are evident in the simulated population in virtue of their preferred direction vectors. Hence, our network naturally reproduces early, late, and persistent firing, as well as positive and negative monotonic responses. Moreover, this network explains the existence of non-monotonic curves which were explicitly excluded from Romo et al.'s data set.

In addition, the control theoretic nature of the NEF allows for easy separation of the stimulus and time components. Future work will exploit this feature in an analysis of the S2 where the discrimination of base and comparison stimulus is suspected to be performed (Romo 2003).

EFFICIENT CODING OF ACOUSTIC STRUCTURE WITH SPIKE TIMES

Evan C. Smith[1] and Michael S. Lewicki[2]
Departments of Psychology[1] & Computer Science[2]
Center for the Neural Basis of Cognition
Carnegie Mellon University

The representation of acoustic signals by the brain must serve a wide range of auditory tasks that require exquisite sensitivity in both time and frequency. Although properties of individual auditory nerve fibers, such as frequency tuning and phase-locking, contribute to this sensitivity, as do a complex cascade of non-linearities, the computational principles that underlie the encoding as a whole remain poorly understood. Here, we investigate a theoretical model of auditory coding in which the computational goal is to form an efficient, time-relative representation of the time-varying amplitude signal.

We apply a signal representation method introduced earlier by one of the authors (Lewicki,2002). In this model, the signal $x(t)$ is encoded with a set of kernel functions, $\phi_1 \ldots \phi_M$, that can be positioned arbitrarily and independently in time. The mathematical form of the representation with additive noise is

$$x(t) = \sum_{m=1}^{M} \sum_{i=1}^{n_m} s_i^m \, \phi_m(t - \tau_i^m) + \varepsilon(t), \tag{1}$$

where τ_i^m and s_i^m are the temporal position and coefficient of the i^{th} instance of kernel ϕ_m, respectively. The notation n_m indicates the number of instances of ϕ_m, which need not be the same across kernels. In addition, the kernels are not restricted in form or length.

We show that, compared to conventional signal representations such as Fourier and Wavelets, the theoretically motivated code allows greater precision in representing the non-stationary, time-relative acoustic features of natural sounds, particularly transient signals such as speech consonants. This framework also provides an objective means of assessing performance in terms of coding efficiency, which we use to compare both computationally and biologically motivated encoding algorithms against theoretical optima. The computational problem investigated here has direct relevance to the neural coding in the auditory system and the more general issue of encoding complex, time-varying signals with a population of spiking neurons.

PREPARATORY DELAY ACTIVITY IN THE PARIETAL CORTEX PREDICTS REACTION TIMES UNDER SOME CIRCUMSTANCES

Snyder LH, Dickinson AR, Lawrence BM. Washington Univ, St. Louis

Activity in the monkey parietal reach region (PRR) and the lateral intraparietal area (LIP) reflect both spatial and effector information for an upcoming reach or saccade task (Calton et al. 2002; Dickinson et al 2003). To further explore the role of these regions in sensory to motor transformations, and in particular, to determine whether they encode "motor set", we tested the hypothesis that increased neuronal activity during a preparatory delay period would lead to faster reach or saccade reaction times (RTs).

We tested preparatory activity in three delayed movement tasks. First, both the movement type (reach or saccade) and the spatial goal were cued prior to a delay. Movement was initiated when the fixation point dimmed at the end of the delay (memory reach/saccade task). Second, the movement type was cued before the delay and the spatial target appeared after the delay, triggering the movement (cue-delay-target task). Third, the spatial goal appeared before and the movement type cue appeared after the delay (target-delay-cue task). For each condition, we correlated the last 300 ms of delay activity with RT.

We found that PRR activity was correlated with reach but not saccade RT only in the cue-delay-target task. The average correlation coefficient was only 0.07, but the distribution of coefficients was significantly skewed towards positive values ($P<.01$, Wilcoxon test). No significant correlations were found for the target-delay cue task or for the memory task, and no significant correlations were found in LIP.

These data are at first surprising. They appear to suggest that PRR is critically concerned with movement type and not spatial location, yet posterior parietal cortex is well-known to code spatial locations. We believe the paradox is resolved by focusing not on the information that the animals are preparing, but on the information they await in order to initiate the movement. Significant correlations occur only when animals are waiting for spatial information regarding the goal of the reach (cue-delay-target task), and not when they wait for a trigger cue (memory task) or for movement type information (target-delay-cue task). These data demonstrate the important role of PRR in providing spatial information for an upcoming reach.

STATISTICS OF CONNECTIVITY AMONG LAYER 5 PYRAMIDAL NEURONS

Sen Song[1], P. Sjöstrom[2], Markus Reigl[1], S.B. Nelson[2], D.B. Chklovskii[1]
[1]Cold Spring Harbor Laboratory, [2]Brandeis University

In the quest to understand the basic rules of synaptic connectivity, we analyze statistics of connections among layer 5 pyramidal neurons from the rat visual cortex. Connectivity was established electrophysiologically in several hundred simultaneous quadruple whole-cell recordings. We report the following major findings:

i) We determined the probability for a pair of nearby (<100um) neurons to be unconnected - 83%, uni-directionally connected - 12%, bi-directionally connected - 5%. These results confirm previous reports that bi-directional connections are over-represented relative to the expectations based on the independence of synaptic connections.

ii) We found that synapses in bi-directional connections are stronger on average than those in uni-directional connections. Also, the strengths of reciprocal synapses in bi-directional pairs are weakly correlated. Yet, neither the dependence of synaptic strength on the pair connectivity, nor the decay of connection probability with distance between the neurons, can account for the observed excess of bi-directional connections.

iii) We determined the probability of encountering triplets belonging to each of the sixteen possible connectivity classes and compared it with the theoretical predictions based on the doublet connectivity statistics. We found several classes with significant deviations in the numbers of triplets. Most of the deviations fit the following rule: Triplets with all-to-all connections (either uni-directional or bi-directional) are over-represented by several-fold.

In conclusion, connectivity among layer 5 pyramids is highly inhomogeneous: All-to-all connected triplets and bi-directional pairs are highly over-represented. This result has implications for understanding of learning rules acting in the course of development and of neuronal dynamics in realistic cortical networks.

A MODEL OF A PORTION OF THE SOUND LOCALIZATION CIRCUIT SUGGESTS THAT COUNTERINTUITIVE MECHANISMS IMPROVE ACTION POTENTIAL PHASE-LOCKING.

W.J. Spain[1,2*]; P.C. Schwindt[1] 1. Physiol. & Biophys., U. Wash., Seattle, WA, USA 2. VAPSHCS, Seattle, WA, USA

Interaural time difference (ITD) in the arrival of tones is one cue used to localize sounds. CNS auditory neurons are specialized for accurate decoding of ITD. First order central auditory neurons (e.g. avian, nucleus magnocellularis, NM, neurons) are specialized to relay accurate information about the time of arrival of sound at the ipsilateral ear by producing action potentials (APs) that are phase-locked to the frequency of the tonal stimulus. These specializations include fast, secure synapses as well as post-synaptic potassium conductances which provide a fast membrane time constant that prevent repetitive firing to temporally summated EPSPs. We show here that a downside to a such a system is that contains the accuracy of phase-locking is extremely sensitive to the strength of the synapses. The sensitivity of phase-locking to variations in synaptic strength was tested in a model of NM neurons using Hodgkin-Huxley like conductances. The responses of the model neuron to pure tones was investigated. Phase-locking of model NM neuron spike out-put to virtual tones was quantified as vector strength. A dip in the vector strength vs. synaptic strength plot happens when the strength of one VIIIth nerve synapse is just large enough to trigger an AP. The steep dependence of accurate NM cell phase-locking on the strength of VIII N synaptic input results from the relation of NM neuron AP Latency to synaptic strength. This is because the period histogram for NM neuron action potential firing is narrow when VIIIth nerve synaptic strength is so small that it takes two simultaneous EPSPs to evoke an AP. However, when the synaptic strength is large enough so that one EPSP reaches action potential threshold, then NM neurons fire spikes at both long latency (i.e. evoked by one EPSP) and at short latency (i.e. evoked by two EPSPs), resulting in a bimodal period histogram. We hypothesized that variation in the size and/or timing of the synaptic inputs will lesson the dependence of precise phase-locking on the size of the synaptic inputs? Three known sources for such variation were tested: 1) the strength of each endbulb (VIIIth N. to NM) synapse was set different from the others. 2) jitter in the phase-locking of the VIIIth nerve firing to the tone frequency was added. 3) synaptic strength was made variable due to synaptic depression. The addition of synaptic depression had the greatest effect on decreasing the dependence of spike timing on synaptic strength. It works because the phase-locked firing of the VIII nerve is probabilistic which gives a distribution of VIIIth nerve interspike-intervals. With synaptic depression the distribution of interspike intervals results in a broad distribution of synaptic strengths. These simulations provide the counterintuitive result that factors which might be considered as "noise" that would degrade output timing (i.e. variation in timing and strength of synapses) actually make output timing more reliable. This results holds for the specific case where the summation of only a few synapses determines the output.

“POTENTIAL” CONNECTIVITY IN LOCAL CORTICAL CIRCUITS

Armen B. Stepanyants[1], Judith A. Hirsch[2], Luis M. Martinez[3], Zoltan F. Kisvarday[4], and Dmitri B. Chklovskii[1]

[1]Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724; [2]Bio.Sci., USC, Los Angeles, CA 90089; [3]Dept. Medicina U. A Coruña 15006 Spain; [4]Ruhr-Universitat Bochum, Bochum, Germany

Synaptic connectivity among cortical neurons may vary with time due to the growth and retraction of dendritic spines [Trachtenberg et al. Nature 420, 788-94 (2002)]. This suggests that the invariant description of cortical circuits should be formulated on the level of more stable features of connectivity, i.e. the layout of axonal and dendritic branches. To describe the relative branch layout for two neurons with overlapping axonal and dendritic arbors, we use the number of “potential synapses” [Stepanyants et al., Neuron 34, 275-288 (2002)]. A potential synapse is a location in the neuropil where an axon of one neuron is present within a certain distance from a dendrite of another. Particular values of this distance depend on the type of synapse. For synapses on spines, this distance is on the order of the average spine length (typically 2µm). For shaft synapses, it is roughly equal to the sum of dendritic and axonal radii (about 0.4µm). Defined this way, a potential synapse is a requirement for a physiological synapse.

We developed a method which allow us to determine the expected number of potential synapses between and within classes of excitatory and inhibitory neurons, depending on their laminar positions and relative separation in the direction parallel to the cortical surface. Based on the morphology of a number of 3D reconstructed spiny stellate cells, pyramidal cells, and basket cells from different cortical layers of the cat visual cortex, we are able to produce cortical connectivity maps for the expected number of potential synapses and the probability of potential connectivity.

The maps of potential connectivity have important implications for further research. First, the average number of actual synapses can be calculated by multiplying the number of potential synapses by the filling fraction (fraction of potential synapses that correspond to actual synapses) [Stepanyants et al., Neuron 34, 275-288 (2002)]. The resulting wiring diagram gives a foundation for realistic modeling of neuronal dynamics, beyond random networks. Second, by comparing the maps of potential connectivity with actual connectivity data, we will be able to determine potential for plasticity in cortical connectivity due to the formation and elimination of synapses.

FROM NEURON TO NETWORK: MAPPING THE STRUCTURE AND DYNAMICS OF CULTURED NEURAL TISSUE

Greg J. Stephens[1], Luís M. A. Bettencourt[1] and Guenter Gross[2]
[1] Los Alamos National Laboratory
Los Alamos, NM 87545
[2] University of North Texas
Denton, TX 76203

A central problem in neuroscience is how to make sense of the complex spatiotemporal spike patterns observed in neural systems. While much attention has focused on characterizing few-neuron activity, the enormous computational power implicit in the behavioral actions of every animal lies in the ability of multi-neuron networks to both code and process information. It is therefore important to develop and analyze *network* characteristics of neural activity. Here we report preliminary progress in the description of the spiking activity of cultured neural networks grown on a multi-electrode array. While cultured networks differ in important ways from intact neural systems, they provide unparalleled experimental access to network structure and network dynamics on time-scales ranging from milliseconds to months. Using spiketrain data from mature cultures of mouse spinal cord neurons, we derive functional network connectivity using a variety of criteria, including firing coincidence timings and cause-effect measures such as cross-correlation functions. We analyze the structure of the resulting graphs both from cultures during native activity and in the presence of specific pharmacological agents that control inhibition. Finally we compare these results with expectations derived from dynamical models of large networks of neurons.

QUANTITATIVE BAYESIAN MODEL OF HUMAN VISUAL MOTION PERCEPTION

Alan Stocker and Eero Simoncelli
Howard Hughes Medical Institute and Center for Neural Science
New York University

Humans do not perceive visual motion veridically: Psychophysical experiments have shown that the perceived motion of visual stimuli clearly depends on their contrast, with low contrast stimuli being perceived to move slower than high contrast ones. A probabilistic Bayesian model has been proposed that *qualitatively* accounts for a wide variety of the reported data in the literature [1]. Formulated as the *maximum a posterior* estimate $\hat{\boldsymbol{v}}$ of visual motion $\boldsymbol{v}$ given the image measurments $\boldsymbol{m}$,

$$\hat{\boldsymbol{v}} = \max[\underbrace{p(\boldsymbol{v}|\boldsymbol{m})}_{posterior} \propto \underbrace{p(\boldsymbol{m}|\boldsymbol{v})}_{likelihood} \underbrace{p(\boldsymbol{v})}_{prior}] , \tag{1}$$

the *prior* reflects the model's assumption that humans apply *a priori* information to estimate visual motion, namely that slow visual motions are more likely to be encountered than fast ones.

The model's quantitative behavior depends on both the likelihood and the prior, but it is very difficult to measure either of these directly. We show here that with an appropriate psychophysical experiment we are able to constrain the form of the likelihood and the prior individually.

We measured the contrast dependent *shift of perceived speed* and *speed discrimination thresholds* with a simple speed-matching experiment. While the shift of perceived speed is determined by the shift of the maximum of the posterior, the discrimination threshold is primarily dertermined by the width of the posterior. Width and shift are differently affected by the prior and the likelihood. Under the assumption of a Gaussian likelihood and given the collected data, we conclude that

- the standard deviation of the likelihood is monotonically increasing in speed and inversely proportional to stimulus contrast with saturation for low contrasts.
- the prior is to first approximation Gaussian (in particular for lower speeds) with zero mean and constant standard deviation.

---

[1] Weiss, Y. and Simoncelli, E. and Adelson, E.: Motion illusions as optimal percept; *Nature Neuroscience 5(6):598-604,* 2002

CORTICAL FEEDBACK IN THE VISUAL SYSTEM

Thomas J. Sullivan and Virginia R. de Sa
Department of Cognitive Science.
UC San Diego, San Diego, CA 92037

Cavanaugh and others[1] have performed electrophysiology experiments that explore the characteristics of surround suppression. In their experiments, the context stimuli that they use suppress the responses of active V1 neurons. They provide evidence (as does Angelucci[2]) that cortical feedback plays a role in this suppression. But most cortical feedback is made up of excitatory neurons connecting with excitatory neurons[3] (but also see [4]). It is not intuitive how feedback could achieve suppression with excitatory feedback.

Another experimental result that is difficult to explain comes from the fMRI work by Murray and others[5]. They show that the overall activity in V1 is reduced when a shape is perceived (compared to visual input when no shape perceived). Again, the excitatory nature of cortical feedback makes this result difficult to understand.

We create a model of cortical feedback that explains these findings. The model has the following characteristics: (1) the learned representations of features are encoded by localized groups of active neurons (neural assemblies), (2) the representations of features are learned by adapting to the statistics of the visual input, (3) if a neural assembly is active, excitatory drive is sent to lower level assemblies that have been associated with it in the past, (4) lateral inhibition results in competition among local neural assemblies, and (5) lateral excitation is involved in forming neural assemblies and facilitating co-occurring assembly activation.

The proposed model is shown to create suppression of active neurons with excitatory feedback and lateral connections. Also, through a sharpening of the neural population's response, the proposed model shows how a decrease in population activity might occur when a shape is perceived[6].

---

[1] James R. Cavanaugh, et.al., Journal of Neurophysiology, 88:2530, 2002.

[2] Alessandra Angelucci, et.al., The Journal of Neuroscience, 22(19):8633, 2003.

[3] R. Johnson, Journal of Comparative Neurology, 368:383, 1996.

[4] Kathleen S. Rockland, Ch.16 of *The Primate Visual System*, CRC Press, 2004.

[5] Scott O. Murray, et.al., PNAS, 99(23):15164, 2002.

[6] This work was supported by the National Science Foundation Career Grant No. 0133996. This work was also supported by an NIH Cognitive Neuroscience Predoctoral Training Grant.

UNSUPERVISED LEARNING OF COMPLEX CELL BEHAVIOR

Thomas J. Sullivan and Virginia R. de Sa
Department of Cognitive Science.
UC San Diego, San Diego, CA 92037

Understanding how complex cells develop their response properties is a crucial step toward understanding the human visual system. Recently, some progress has been made in this area by considering the temporal coherence of natural visual stimuli. A similar strategy is employed here. Additionally, it is shown that a temporal trace learning rule is compatible with an SOM model of lateral connectivity and that retinal waves are a sufficient input to learn complex cell response properties. It is shown that complex cell behavior emerges using unsupervised, local learning rules. The resulting learned neurons are phase invariant, tuned for orientation, tuned for spatial frequency, and form an orderly map.

This model employs two layers of neurons in a hypercolumn-sized region of V1. The first layer is pre-trained and made up of Gabor filters of various orientations, spatial frequencies, and locations. The second layer contains neurons with weights that are initially untrained and to be learned. The network is fully connected such that each neuron in the first layer is connected to every neuron in the second layer.

At each discrete time point, an input is presented to the model in the form of a video frame of an artificial retinal wave. The simple cells respond to the input by performing a dot-product of their weights with the input. The second-layer cells read the simple cell output as their input. A competition then takes place among the second-layer neurons. The learning rules employed are a combination of Hebbian, temporal trace, and self-organizing map. Through the learning process, the second-layer neurons develop complex cell properties: they become tuned for particular orientations and spatial frequencies, they become translation and phase invariant, and they arrange themselves into an organized map.

This work is significant because it shows that it is possible to combine a temporal trace rule with an SOM model of lateral connectivity. It also shows that retinal waves are a sufficient input for complex cells to learn their response properties. Unlike Földiák's work[1], the neurons here model both orientation selectivity as well as spatial frequency. Unlike Hyvärinen's work[2], the learning rules are local, and thus more biologically realistic[3].

---

[1] Peter Földiák, Neural Computation, 3:194, 1991.

[2] Aapo Hyvärinen and Patrik O. Hoyer, Vision Research, 0:1,2000.

[3] This work was supported by the National Science Foundation Career Grant No. 0133996. This work was also supported by an NIH Cognitive Neuroscience Predoctoral Training Grant.

Impaired and Enhanced Spatial Representations of the Post Synaptic Desnity-95 Knockout Mouse

Linus D. Sun†*, Albert K. Lee†, Jeenah Jung†, Seth G.N. Grant‡, Matthew A. Wilson†.

†Picower Center for Learning and Memory/Depts of Biology and Brain and Cognitive Science, MIT.
‡Center for Genome Research, and Center for Neuroscience, University of Edinburgh.
*Yale School of Medicine.

Postsynaptic density protein-95 (PSD-95) is the second most abundantly expressed synaptic protein in the postnatal forebrain. It is an integral part of the postsynaptic scaffolding complex and helps recruit receptors, channels and associated factors involved with synaptic transmission. A mouse whose wildtype gene was replaced with truncation mutant of PSD-95 preserving two PDZ binding domains causes a spatial learning and memory deficit and a dramatic enhancement of synaptic strengthening. Long Term Potentiation is enhanced at all frequencies of stimulation (1-100Hz), while Long Term Depression is absent in the mutants. This study explores CA1 pyramidal cell spatial representations in the PSD-95 mutant mice. Mutants are not significantly different than controls in running velocity. Nor are its pyramidal cells or interneurons different than controls in: place cell firing rates, sparsity of run active cells, bursting behavior, or theta modulated activity. However, mutants do exhibit significantly larger place fields and wider spike waveforms. Mutants also expressed enhanced directionality of place fields and increased post-run sleep correlation of firing for overlapping place fields. Mutants also exhibited disruption of asymmetrical place fields and phase precession, the first such observation reported in mice. In conclusion, LTP alone is not enough for the active process of encoding experience. Instead, bi-directional synaptic plasticity is necessary for proper place field formation, correlation, directionality, asymmetry, phase precession, and the formation of spatial memories.

SUPERVISED LEARNING RULES DERIVED FROM POPULATION STATISTICS

Christian D. Swinehart & L.F. Abbott
Volen Center for Complex Systems and Department of Biology.
Brandeis University, Waltham MA 02454-9110

Neural network learning of specific tasks is typically achieved through the use of a supervisor that evaluates performance and adjusts network properties appropriately. Modifications performed in this manner almost always consist of supervised changes of synaptic strength. We have recently developed models in which supervised learning takes place at the level of *neurons*, rather than synapses, through modulation of their response properties. Synaptic modification is still an important element of these models, but it occurs in an unsupervised, Hebbian manner. In previous work, we have shown that networks supervised in this way can learn to perform a variety of tasks.

Supervision of neurons rather than synapses has a number of advantages. The first is biological plausibility; it has never been clear how rapid supervision of individual synapses could occur in real neural circuits, but feedback to neurons is a universal feature of brain anatomy. Second, there are far fewer neurons than synapses so the dimension of the parameter space in which the supervisor must search is reduced. Despite this reduction, the dimension of the search space remains large due to the abundance of neurons in a typical network. Especially for reward-based supervision, it is essential to reduce the dimension of the parameter space as much as possible. This introduces a third advantage of supervising neurons rather than synapses that we exploit here. It is possible for the supervisory circuitry to monitor the activity of the neurons it is supervising and to use this information to develop optimal search strategies.

Consider, for example, two neurons that exhibit nearly identical patterns of activity. It is clearly not optimal for the supervisor to independently modulate the response properties of these neurons, they might as well be treated as a single unit. We extend this idea by using feedback pathways from the supervised network to the supervisor to convey information about correlations between network neurons. On the basis of this information, the supervisor performs a principal component analysis, extracting the eigenvectors of the correlation matrix with the largest eigenvalues. It then focuses its reward-based modification of neuronal response properties on the basis of this principal component analysis, first directing modification along the direction of the principal eigenvector, then the second principal eigenvector, and so on until an error criterion has been met. We explore how this strategy can optimize convergence speed for reward-based learning and speculate on how it can be realized in neuronal circuitry.

NOISE ENHANCES NEURONAL STABILITY: A STOCHASTIC MODEL OF ACTIVITY-DEPENDENT NEURONAL FEEDBACK SYSTEMS.

Michael R. Tadross and Kechen Zhang. Dept. of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21205.

Neuronal bursting is a nonlinear phenomenon that requires a delicate balance between the ionic concentration gradients across a cell membrane and the quantity and types of ion channels present on the membrane. Neurons from the stomatogastric (STG) ganglion of the crab Cancer borealis are capable of slowly modulating their "configuration" (i.e. number of each type of membrane ion channel) to maintain a stable bursting activity pattern in a changing environment. Interestingly, different biological STG neurons placed in identical environments can have vastly different configurations, yet very similar activity patterns. Both of these phenomena can be modeled by a nonlinear activity-dependent calcium-based feedback system that slowly modulates the neuron's configuration until the desired activity pattern is obtained (Liu et al. 1998). The system has numerous equilibrium states, allowing most randomly chosen initial configurations to converge to diverse final configurations that produce bursting. However, for 10% of the simulations, the original model does not stabilize, and instead either up-regulates all channels without bound or enters a periodic limit cycle. Inspired by the fact that noise can improve nonlinear optimization, such as in simulated annealing, and that noise exists ubiquitously in biological systems, a custom-written screensaver-based distributed super-computing cluster of approximately 50 PC computers was developed to investigate the effects of including various degrees of noise in the activity-dependent feedback system of the original model. Starting from hundreds of randomly chosen initial configurations, the original model, as well as a modified version of the model, was simulated using different levels of noise. For both models, it was found that an intermediate level of noise yields optimal stability. The existence of such optimal noise levels can be demonstrated analytically by simplified stochastic models. In our simulations, when compared to zero noise, the optimal level of noise reduces the incidence of system instability by a factor of four, and reduces the variability in final configurations to a range more consistent with biological neurons. The optimal noise level roughly corresponds to fluctuations consistent with the loss or gain of only a few ion channels, corresponding to errors in gene transcription, protein degradation, and other regulatory systems. These results imply that for biological neurons, the presence of these fluctuations may play a fundamental role in maintaining robust neuronal stability.

DISCRIMINATION OF SHORT TONES IN NOISE BY RATS

Lung-Hao Tai[1,2], Marta Moita[2], Anthony Zador[2]
[1]Dept. of Neurobiology and behavior, SUNY Stony Brook, NY 11790
[2]Cold Spring Harbor Laboratory, NY 11724

How do animals discriminate tones in background noise? As a first step to study neural correlates of auditory perception in rats, we trained Long-Evans rats to discriminate pure tones imbedded in a white noise background in a two alternative forced choice (2-AFC) paradigm, adapted from a similar olfactory task (Uchida and Mainen, 2003). 2-AFC tasks have been very important in understanding cortical correlates of perception in primates. By adapting this paradigm in rodents, we can readily utilize a rich collection of molecular and pharmacological tools available to explore the underlying cortical mechanism of perception and decision making. Subjects triggered the auditory stimulus by means of a nose poke to a center nosecone, and were rewarded (~50µl water) if they responded with a second nose poke into the appropriate nosecone (right, 1kHz or left, 15kHz). Training was fast (2-3 weeks), and performance on the easiest discriminations could exceed 95%. Typical reaction times were short (300~500 ms). Performance deteriorated in a systematic way with decreased signal to noise ratio and with decreased tone duration. This paradigm offers an ideal model system for studying the neural correlates of two competing perceptions whose representations are likely to be in distinct areas of any tonotopically organized region of auditory cortex.

Uchida, N. and Mainen, Z. (2003) Speed and accuracy of olfactory discrimination in the rat. Nature Neuroscience. 6(11):1224-9.

SPECTROTEMPORAL RECEPTIVE FIELDS AND DIRECTION SELECTIVITY

Andrew Y. Y. Tan and Christoph E. Schreiner
Coleman Laboratory and Keck Center for Integrative Neuroscience
University of California San Francisco

We provide simple examples of both separable and inseparable spectrotemporal receptive fields which can underlie direction selectivity. A nonlinearity is needed in conjunction with both separable and inseparable spectrotemporal receptive fields to produce direction selectivity.

Encoding and decoding representations of stimulus location in the leech

Eric E. Thomson and William B. Kristan
UCSD Neuroscience Graduate Program

In the local bend response, individual leeches can discriminate between touch locations that are a millimeter or more apart. Local bending is largely activated by four pressure sensitive mechanoreceptors (P cells). Previous work showed that P cell spike counts carry information about touch location (Lewis and Kristan 1998). We have recently discovered that the latency difference between spikes in P cells with overlapping receptive fields can be used to predict stimulus location with submillimeter resolution, much better than discriminability attained by spike counts alone (Thomson and Kristan 2003). In the present work we extend this analysis, applying metric space methods to compare stimulus estimation based on spike counts alone, latency difference alone, and entire spike trains.

Because of the simplicity of the system and the accessibility of the sensory neurons, we were able to experimentally determine how efficiently information contained in the latency difference and spike counts is decoded and used by the downstream neurons that guide behavior. Specifically, we used current clamp to parametrically change the spike counts and latency difference in two P cells while imaging the evoked behavioral response, calculating the minimum parameter change required to produce a different behavior.

References

Lewis, J.E. and Kristan, W.B. Jr. (1998). Representation of touch location by a population of leech sensory neurons. J. Neurophysiol. 80: 2584-2592.

Thomson, E.E. and Kristan, W.B. Jr. (2003) Mechanoreceptor latency encodes touch location in the leech. Society for Neuroscience Abstract 269.4.

CHAOS-INDUCED MODULATION OF RELIABILITY BOOSTS OUTPUT FIRING RATE IN DOWNSTREAM CORTICAL AREAS

Paul H.E. Tiesinga
Department of Physics & Astronomy, University of North Carolina,
Chapel Hill, NC 27599

The reproducibility of neural spike train responses to an identical stimulus across different presentations (trials) has been studied extensively. Reliability, the degree of reproducibility of spike trains, was found to depend in part on the amplitude and frequency content of the stimulus (Hunter and Milton, J. Neurophysiology 90:387 (2003)). The responses across different trials can sometimes be interpreted as the response of an ensemble of similar neurons to a single stimulus presentation. How does the reliability of the activity of neural ensembles affect information transmission between different cortical areas?

We studied a model neural system consisting of two ensembles of neurons with Hodgkin- Huxley-type channels. The first ensemble was driven by an injected sinusoidal current that oscillated in the gamma-frequency range ($40\ Hz$) and its output spike trains in turn drove the second ensemble by fast excitatory synaptic potentials with short term depression. We determined the relationship between the reliability of the first ensemble and the response of the second ensemble. In our paradigm the neurons in the first ensemble were initially in a chaotic state with unreliable and imprecise spike trains. The neurons became entrained to the oscillation and responded reliably when the stimulus power was increased by less than 10%. The firing rate of the first ensemble increased by 30%, whereas that of the second ensemble could increase by an order of magnitude. We also determined the response of the second ensemble when its input spike trains, which had non-Poisson statistics, were replaced by an equivalent ensemble of Poisson spike trains. The resulting output spike trains were significantly different from the original response, as assessed by the metric introduced by Victor & Purpura (J. Neurophysiology 76:1310 (1996)).

These results are a proof of principle that weak temporal modulations in the power of gamma-frequency oscillations in a given cortical area can strongly affect firing rate responses downstream by way of reliability in spite of rather modest changes in firing rate in the originating area.

A SYNTHESIS OF LOCAL AND GLOBAL METHODS FOR OPTIMAL FEEDBACK CONTROL OF NONLINEAR STOCHASTIC SYSTEMS

Emanuel Todorov; Department of Cognitive Science; UC San Diego

Optimal control theory has played a key role in theoretical investigations of biological movement. The majority of existing optimality models describe average behavior, by performing open-loop optimization in the space of control sequences or limb trajectories. A few recent models have gone further, and considered optimization over sensorimotor transformations or feedback control laws. The latter yield comparable or more accurate accounts of average behavior than average behavior models, and explain numerous additional phenomena such as goal-directed corrections, structured motor variability, sensorimotor synergies, task-dependent speed profiles, eye-hand coordination patterns.

Stochastic optimal feedback control models are presently implemented in the Linear-Quadratic-Gaussian (LQG) setting – which affords efficient model construction, but imposes severe limitations that prevent this very promising form of modeling from reaching its full potential. Another limitation is that the resulting control laws are time-varying, which is unsuitable for biologically-plausible implementation. The objective of the work presented here is to resolve both problems, and clear the way to optimality models of sensorimotor control on a qualitatively new level of biological realism. This is accomplished via a synthesis of two new methods that naturally complement each other.

Local method: iterative Linear-Quadratic-Gaussian (iLQG) design

Local trajectory-based approaches to optimal control are motivated by Pontryagin's Maximum Principle, which is an ordinary differential equation that allows us to find optimal trajectories without running into the curse of dimensionality. We extend these ideas to the stochastic case, by iteratively linearizing the nonlinear system around the average trajectory generated by the current control law, and using a generalized LQG method to improve that control law. The new method deals with control-dependent noise characteristic of the motor system, control constraints necessary to enforce non-negative muscle activations, and non-positive-definite quadratic costs which can arise when a general cost function is quadratized. We have applied iLQG to a realistic model of the human arm with 10-dimensional state space and 6-dimensional control space. For the task of reaching, the method converges in 5-10 iterations, takes ~7 sec using a Matlab implementation, and outperforms its nearest (deterministic) competitor by a factor of 10.

Global method: recurrent neural networks as optimal feedback controllers

Global methods in high-dimensional spaces require function approximation, which has had limited success for two reasons: (i) it is hard to find parameterizations that are compact enough to facilitate learning yet rich enough to represent complex time-varying control laws; (ii) the Reinforcement Learning (RL) methods used so far can be very slow (on non-toy problems) because of the difficult temporal credit assignment problem. Our method avoids both problems: (i) instead of a static mapping from states into actions, we use recurrent neural networks that can be trained to exhibit rich dynamics even though their parameters are few and static; (ii) instead of RL we use the efficient iLQG method with multiple initial states, and train the network to reproduce the outputs of the resulting locally-optimal control laws. We develop a new algorithm that makes the network learn not only a reference trajectory but also a local feedback control law – alleviating the generalization and dynamic stability problems associated with recurrent neural networks.

MODEL CELL ENCODES SPIKE-TIMING DEPENDENT SYNAPTIC PLASTICITY IN HIPPOCAMPAL CA1 NETWORKS.

Kazunobu Togashi, Norihiro Yoshida, Tatsuo Kitajima*, Kyonsoo Hong and Makoto Nishiyama
Department of Biochemistry, New York University School of Medicine
*Faculty of Engineering, Yamagata University

Activity-dependent synaptic modifications, long-term potentiation (LTP) and long-term depression (LTD), are essential in information processing and memory storage. γ-Aminobutyric acid (GABA)-mediated inhibition plays an important role in modulating hippocampal oscillatory activities at theta and gamma frequencies that are known to be linked to spatial memory. The potential importance of GABAergic activity in regulating synaptic modifications, however, is not well understood. We, therefore, investigated the role of GABAergic activity in regulating properties of spike-timing dependent synaptic plasticity (STDP) in CA1 neural networks employing computational analysis that incorporates temporal patterns of postsynaptic GABAA-mediated inhibition and paired-pulse responses of IPSPs/EPSPs. Using a 5-compartment model cell which consists of a soma, a spine and three dendrites, we first simulated membrane potential dynamics during each correlated pre- and postsynaptic activation. We then simulated postsynaptic Ca^{2+} dynamics at the spine head which were integrated as total postsynaptic Ca^{2+} elevation during the induction protocol. Postsynaptic Ca^{2+} elevations were finally converted to synaptic efficacy according to the Bienenstock-Cooper-Munro curve. Here, we report that spatiotemporally distinct GABAergic activity gates LTD at Schaffer collateral-CA1 glutamatergic synapses. At the theta frequency (5 Hz) of the spike-timing protocol, two distinct time windows for the induction of LTD, which flanked an LTP-time window, disappeared when postsynaptic GABAA-mediated inhibition was not applied. Furthermore, LTD appeared between the alpha and beta frequencies as the frequency of the LTP-timing protocol, which normally induces prominent LTP at 5 Hz, increased. This frequency-dependent appearance of LTD was due mainly to presynaptic GABAB-mediated desensitization of neurotransmitter release, where the LTD was flanked with twin peaks of robust LTP at theta and gamma frequencies. Taken together, our results support the critical importance of GABAergic inhibition in determining the timing- and frequency-dependent features of STDP.

INVESTIGATING THE PRESENCE AND INFORMATIONAL RELEVANCE OF PATTERNS IN NEURAL SPIKE TRAINS.

J.Vincent Toups and Paul H.E. Tiesinga
Department of Physics & Astronomy, University of North Carolina,
Chapel Hill, NC 27599

It is a canonical assumption in the information analysis of spike train data that spikes are produced according to a Poisson or renewal process. Recent work has revealed deviations from this assumption (Fellous, Tiesinga, Thomas & Sejnowski, in preparation) leading to the question of how to best identify these deviations, which represent themselves as repeated, statistically unlikely spike patterns, and to quantify their informational relevance.

We used the Victor and Purpura metric (VPM) to characterize short segments of spike trains, and then used fuzzy k-means clustering to identify different clustering schemes over the spike timing specificity parameter q provided to the metric. For small values of q, the VPM is essentially a spike count metric, which for short segments is a rate estimator. Larger values of q correspond to a greater and greater dependence on spike timing in the final clustering, until the metric becomes the number of non- simultaneous spikes in each train. We developed an algorithm that estimated the correct number of clusters for a given q value. By examining changes in clustering under parametric variation of q we often found only a small number (typically 1-5) of distinct clustering schemes. The local maximum in Region III corresponded to a ’good clustering,’ suggesting a simple method of selecting the proper q for analysis of information encoded in spike timing.

Once a clustering was found, we performed information analysis by breaking a spike train up into a series of segments each characterized by which spike pattern was expressed. This technique allows for analysis without choosing an arbitrary or poorly motivated bin size, and is a way to extend the standard information analysis techniques to non-Poisson spike trains. It could potentially improve the efficiency of entropy estimation compared with Strong’s method (Strong et al, Physical Review Letters 80:197 (1998)).

ESTIMATING COMPLEX CELL RECEPTIVE FIELDS USING NATURAL IMAGES

Jon Touryan, Gidon Felsen, and Yang Dan
Group in Vision Science; Helen Wills Neuroscience Institute; Division of Neurobiology, Department of Molecular and Cell Biology.
University of California, Berkeley, CA 94720-3200

Previous studies have shown that the receptive fields of complex cells in the primary visual cortex (V1) can be well represented by a small number of eigenvectors (typically two) of the spike-triggered stimulus ensemble (Touryan et al., 2002). However, this low-dimensional subspace representation has only been demonstrated in response to white noise stimuli. In the present study we use a modified spike-triggered correlation analysis to examine the cortical responses to natural images and find that the spatial structure of each complex cell receptive field is well represented by two eigenvectors. These eigenvectors contain separate on and off subregions, similar to the receptive fields of simple cells, and the two eigenvectors are in approximate spatial quadrature. In addition, we find that these eigenvectors give an accurate estimate of the cells preferred orientation as measured with drifting gratings. Thus, this analysis technique gives a fast and useful description of the complex cells receptive field using natural images.

ENCODING 3D STRUCTURE IN AREA CIPS

Doris Y.Tsao[1,2], Winrich Freiwald[1,2], Margaret S.Livingstone[1], Roger B.H.Tootell[2],
1. Dept. Neurobiology, Harvard Med. Sch., Boston, MA, USA, 2. MGH, Athinoula A. Martinos Ctr. for Biomed. Imaging, Charlestown, MA, USA

Although disparity-tuned cells have been reported in almost all areas of macaque visual cortex, the relative role of individual areas in generating the percept of 3D shape remains unclear. To investigate this, we performed fMRI (3T Siemens Allegra scanner) in three alert macaque monkeys while they viewed a sequence of cyclopean stimuli. We found that both dorsal (V3, V3A, and CIPS) and ventral (V4, TEO, and TE) areas were activated by changing disparity-defined 3D shapes compared to a flat plane. Surprisingly, however, we found that only dorsal areas (V3, V3A, and CIPS) were activated by a changing disparity-defined 3D landscape compared to a flat plane. This result indicates that the processing of 3D structure *per se* is not distributed throughout the entire cortex, but relies principally on dorsal cortical areas. In particular, 3D shape from binocular disparity is processed in areas V3, V3A and CIPS.

Anatomically, CIPS constitutes a key link in a chain of areas, V1->V2 ->V3->V3A->CIPS->AIP, along which an increasingly complex 3D representation is generated (Nakamura et al., J. Neurosci. 21: 8174-8187). To identify the 3D primitives used in CIPS, we reverse-correlated single-unit responses in CIPS to parametric sets of moderately-complex 3D stimuli (stimulus refresh = 33 ms/frame). The stimuli included: 1) disparity- and perspective-defined square grids oriented at 17 x 17 different slants and tilts, 2) disparity- and perspective-defined "book" stimuli in which the book angle varied from 0 to 180 degrees, and 3) triplets of rotating, disparity-defined spots (which generated an illusion of a rotating plane). Many cells in CIPS showed highly robust slant-tilt maps containing discrete regions of activation and suppression, which often occurred in oriented bands. Many cells also showed tuning to specific book angles. The tuning to book angle could not be predicted from the tuning to single planes, thus 3D angle may be an independent property coded by CIPS cells. *Support Contributed By: NIH, R01 MH67529 A01, R01 EB00790 A01, EY13135, and DFG, SFB 517*

CONCENTRATION-INVARIANT OLFACTORY CATEGORIZATION: BEHAVIOR AND RATION CODING

Naoshige Uchida and Zachary F. Mainen
Cold Spring Harbor Laboratory

Any odor, whether a single chemical or a complex mixture, activates a specific combination of odorant receptors. It has therefore been postulated that combinations of activated receptors can define odor quality ("combinatorial coding"). However, changes in concentration produced by changes in distance from an odor source will also lead to changes in the combination of activated receptors, but should be recognized as still belonging to the same source. [Furthermore, while intensity should not confound quality, intensity information must be independently available in order, for example, to localize odor sources.]

How are quality and intensity information extracted from the ensembles activity in the olfactory bulb in order to guide olfactory decisions? Here, we report behavioral results from rats performing an odor mixture discrimination task. In this task, rats are trained to categorize mixtures of two pure odorants in various ratios into two contiguous sets associated with different response ports. Do rats make concentration-invariant judgments about the stimuli? To answer this, after training, rats were tested with probe stimuli (randomly rewarded) with the same mixture ratios as in the training set but with different absolute concentrations [give %??]. The rats' performance matched the predictions of a ratio-based discrimination rule and was significantly less well-fit by alternatives such as an intensity-based rule. Thus, the ratios of the components of a mixture are the critical determinants of the discrimination strategy naturally adopted by rats in this task. In a more natural setting, this strategy would correspond to identifying a complex mixture by the molar ratios of its components, a [useful] characteristic that would remain unchanged with dilution.

What mechanisms could be used by the olfactory system to enable concentration-independent judgments about odors? Based on these behavioral results, we propose a simple coding mechanism: a set of ratios of the coverage of olfactory receptors can define a code for odor quality that is invariant to changes in concentration over the dynamic range of individual receptors ([we call this] "ratio coding"). We are currently testing the predictions of this model using intrinsic imaging and multi-electrode recording in anesthetized and behaving rats.

SYNAPTIC INTEGRATION IN ACCESSORY OLFACTORY BULB MITRAL CELLS

Nathan Urban
Carnegie Mellon University

One of the most striking anatomical differences between the accessory and main olfactory bulbs is the dendritic morphology of the mitral cells. In both accessory and main olfactory bulbs nearly all of their excitatory synapses received by mitral cells are made on sites on the distal tuft of the primary dendrites. The vast majority of mitral cells in the main olfactory bulb have primary dendrites that form a single dendritic tuft occupying a single glomerulus. By contrast, in the accessory olfactory bulb, most mitral cells are multi-tufted, with individual cells having between one and nine tufts associated with different glomeruli. This multi-tufted structure suggests that the individual dendritic tufts may function to integrate inputs independently, and thus may form discrete functional units.

We have begun exploring integration of inputs in the multi-tufted dendritic trees of AOB mitral cells using physiological and computational methods. Passive simulations of reconstructed accessory olfactory bulb mitral cells predict that individual tufts are electrotonically isolated from each other and from the soma. Steady state voltage changes were attenuated by more than 80% from tuft to soma and more than 95% from one tuft to another (N=4 reconstructed cells with 12 tufts). These passive simulations suggest that active properties of dendrites will be critical to integration in these cells. Calcium imaging in the primary dendrites of accessory olfactory bulb mitral cells indicates that the dendrites of these cells propagate signals robustly, without measurable attenuation throughout their entire length. Single action potential-evoked calcium transients in the most distal dendrites were not significantly smaller than those in proximal dendrites (distal/proximal ration = 1.25 ± 0.2 $p<0.01$, N=12 cells). We next tested whether individual tufts could support regenerative events by focally stimulating tufts of calcium dye-loaded mitral cells with an extracellular stimulation electrode. In 6 of 8 cells, isolated all-or-none events could be evoked in at least one mitral cell tuft under control conditions. In the remaining two cells, isolated calcium transients could be evoked following somatic hyperpolarization to prevent somatic action potential initiation. The amplitudes of these tuft calcium spikes were approximately the same as the amplitudes of the single action potential-evoked calcium transients recorded in these same cells, suggesting that the isolated calcium event may be generated intrinsically, rather than by calcium influx through calcium permeable synaptic channels. Further work underway is investigating the biophysical mechanisms of these isolated tuft calcium events, and determining the relationship between dendritic morphology and the degree to which individual tufts can fire calcium spikes independently of the soma and each other.

DETECTION SENSITIVITY AND TEMPORAL RESOLUTION OF PRIMATE RETINAL GANGLION CELLS NEAR ABSOLUTE THRESHOLD

Uzzell, V.J.,*[1,2] Field, G.D.,[3] Rieke, F.[3], and Chichilnisky, E.J.[1,2]

[1] Salk Institute, La Jolla, CA, USA
[2] University of California, San Diego, CA, USA
[3] University of Washington, Seattle, WA, USA
* All authors contributed equally to this work

Humans can detect flashes producing 5-20 photoisomerizations in a pool of hundreds of rod photoreceptors. How can the visual system detect whether and when a sparse signal occurs in an array of noisy detectors? We examined this question by comparing the detection sensitivity and temporal resolution of primate retinal ganglion cells (RGCs) to those of their pooled rod inputs.

Multielectrode extracellular recordings from RGCs (probably parasol and midget cells) were obtained from isolated retinas. Detection sensitivity and temporal resolution were characterized using a two-alternative forced choice discrimination of responses to dim flashes delivered at one of two times. The dependence of discrimination performance on flash strength defined the detection sensitivity. The dependence of discrimination performance on the separation between the two possible flash times defined the temporal resolution. The sensitivity and temporal resolution of RGCs were compared to the limits imposed by the pool of rods in the RGC receptive field by using the signal and noise properties of isolated rods measured with suction electrodes.

Some primate RGCs could detect the occurrence of 3-10 photoisomerizations in a pool of 2,000-9,000 rods. At flash strengths 50-100 times higher, ON RGCs could discriminate flashes offset by only 5-10 ms, a value much shorter than the 100-200 ms integration time of the rod light response. OFF cell temporal resolution was coarser. RGC sensitivity to the arrival and timing of dim flashes was nearly optimal given the noise in photoreceptor inputs, and retinal circuits appeared to utilize nonlinear processing to achieve this sensitivity.

HHMI Predoctoral Fellowship (VJU), McKnight Foundation (EJC & FR), Sloan Foundation (EJC), NIH grants EY13150 (EJC) and EY11850 (FR)

SELECTIVITY WITHOUT FUNCTIONAL ARCHITECTURE IN VISUAL CORTEX OF A HIGHLY VISUAL MAMMAL

Stephen D. Van Hooser[1], Jan-Alexander F. Heimel[1], Sooyoung Chung[1],
Sacha B. Nelson[1], and Louis J. Toth[2]
[1]Department of Biology, Brandeis University, Waltham, MA, 02454
[2]Department of Anatomy and Neurobiology,
Boston University School of Medicine, Boston, MA 02188

In primary sensory areas of mammalian neocortex, receptive field (RF) properties are thought to be organized along the cortical surface in smooth functional maps. In primary visual cortex (V1), neurons respond preferentially to bars or edges of a particular orientation, and in V1 of primates, carnivores, ungulates, and tree shrews, these orientation-selective cells are arranged in a semi-regular, smoothly varying map so that as one moves along the cortical surface, the orientation preference of neurons tends to change smoothly. Although orderly maps of RF properties are thought to be a fundamental organizing principle of neocortex, their functional role remains unclear. Curiously, while electrophysiological and imaging studies in V1 of many rodents (mice, rats, hamsters) and lagomorphs (rabbits) have identified orientation-selective neurons, no orderly map of orientation selectivity resembling those of primates, carnivores, ungulates, and tree shrews has been found in these species. Understanding why these animals do not have orientation maps may shed light on the functional roles and developmental mechanisms of orderly maps in mammalian sensory cortex.

Previously examined rodents lack the acuity and laminar definition of visual brain areas found in primates, carnivores, and tree shrews. In addition, V1 in these animals is relatively small. Perhaps all mammals possess the necessary mechanisms for developing orientation maps, but good acuity, well-defined visual areas, or large absolute V1 area are required for these maps to form. To examine these ideas, we used intrinsic imaging and single unit recording to assess organization of orientation selectivity in V1 of a highly visual rodent, the gray squirrel. Squirrels have large eyes, well-laminated LGN and V1, depend on V1 for orientation discrimination, and have a large V1 that is bigger than V1 in tree shrew and mink, which have orientation maps.

With both techniques, we found no evidence of orientation maps in squirrel, suggesting that formation of orientation maps depends on mechanisms not found in rodents. We did find robust orientation tuning in single unit recordings, and this tuning was invariant to stimulus contrast. Therefore it seems unlikely that orientation maps are important for orientation tuning or contrast invariance in V1. We suggest that an orderly arrangement of functional RF properties is not a universal characteristic of cortical architecture.

SIGNAL PROPAGATION IN LARGE NETWORKS OF INTEGRATE-AND-FIRE NEURONS

Tim Vogels & L.F. Abbott
Volen Center for Complex Systems and Department of Biology
Brandeis University, Waltham, MA 02454-9110

Signal propagation from one brain region to another is essential for cognitive function, but how it occurs reliably over an enormous dynamic range is not clear. Although anatomical studies help to identify the pathways of transmission, they cannot reveal how a neural network propagates information across multiple synapses without loss. In recent years, this problem has been studied computationally in feedforward networks. (Golumb, 1998; van Rossum et al, 2002; Reyes, 2003) Noise plays an important role in these models, but its source is external and not accounted for.

We study a large network of integrate-and-fire neurons with conductance-based synapses to address the issues of both noise and signal propagation. The model features sparse connectivity between a fixed ratio of excitatory and inhibitory neurons. This creates a highly irregular, self-sustained background activity that can be used to prevent synchronization without any external source of noise and allows the transmission of rate-encoded stimuli to directly neighboring neurons. The degree of accuracy in this transmission depends on the profile of the background activity. However, transfer of information beyond the first layer of neurons is not typically achieved for random patterns of connectivity.

To propagate a signal across multiple layers, it is necessary to change the overall connectivity of the network, going beyond the random model (Mehring et al, 2002, Avial et al, 2003). We explore how several rules that favor both local connectivity and clusters of long-range synaptic targets enhance polysynaptic signal propagation. We introduce continuous families of connectivity patterns that allow adjustment from totally random to highly structured as a function of model parameters. This allows us to determine where in this range of structures signal propagation can occur and to evaluate non-synchronous vs. synchronous modes of propagation through subgroups of network neurons. In general, we characterize the fidelity and range of input signal propagation in networks as a function of connectivity pattern, magnitude and variability of synaptic weights, ratio of excitatory and inhibitory neurons, network size, and other relevant parameters.

THALAMOCORTICAL MECHANISMS OF AUDITORY ADAPTATION

Michael Wehr and Anthony Zador
Cold Spring Harbor Laboratory

The auditory system is exquisitely sensitive to the timing of sensory information. For example, neurons in the auditory brainstem are sensitive to interaural time differences of only a few microseconds. Yet at higher levels of the auditory pathway, strong adaptation limits the ability of cortical neurons to faithfully follow rapid trains of stimuli. Auditory cortical neurons can rarely sustain responses to click repetition rates faster than 5-15 Hz. Perceptually, repetition rates lower than this produce a sensation of rhythm, while higher rates produce a sensation of pitch. In contrast, thalamic neurons can often follow trains of up to 100-200 Hz. This suggests that some aspects of sound perception share properties with auditory cortical neurons. Furthermore, the rapid neuronal adaptation that might underlie this perceptual transition appears to originate in the auditory cortex. What are the cellular and synaptic mechanisms of this adaptation? Inhibitory circuitry within the cortex is likely to play a role in this phenomenon, as is synaptic depression at excitatory intracortical synapses. The major contribution, however, is likely to be synaptic depression at the thalamocortical synapse, as seen in the somatosensory system. To distinguish between these potential mechanisms, we are characterizing the degree to which thalamocortical synaptic depression accounts for adaption in cortical responses to repeated sounds. In order to isolate the monosynaptic thalamocortical component of sound-evoked responses, our approach is to combine intracellular recordings with cortical inactivation. We used extracellular application of muscimol (a $\mathrm{GABA_A}$-ergic agonist) to completely silence the auditory cortex. In the presence of muscimol, all cortical neurons are prevented from spiking, but thalamic afferents still deliver synaptic inputs. We then used *in vivo* whole-cell methods to record subthreshold PSPs from cortical neurons, which under these conditions should be solely thalamocortical in origin. These subthreshold responses showed strong adaptation to trains of clicks, with a frequency dependence that roughly matched that seen in the intact cortex. This suggests that a major component of adaptation is present even without the effect of cortical circuitry. However, while this implicates thalamocortical synaptic depression as the key mechanism for rapid adaptation, the time course of adaptation and recovery in the subthreshold PSPs was not well described by standard (exponential) models of synaptic depression. This suggests that the adaptation we observed in thalamocortical inputs is due not only to thalamocortical synaptic depression, but also to the accumulated adaptation from all subcortical levels of the auditory hierarchy. To isolate the contribution of thalamocortical synaptic depression, we placed stimulating electrodes in the thalamus, and recorded the thalamic inputs to cortical neurons evoked by trains of electrical stimuli. By comparing the click-evoked responses to the responses evoked by electrical stimulation of the thalamus, we should be able to quantify the contribution of synaptic depression at the thalamocortical synapse to rapid adaptation in the auditory cortex. This may represent a link between the dynamics of a specific synapse and an overt feature of sound perception.

SYNAPTIC DYNAMIC MEDIATES FORM-CUE INVARIANT SENSITIVITY TO MOTION

Ralf Wessel[1], Reza Khanbabaie[1], and Harald Luksch[2]
[1]Dept of Physics, Washington University, St. Louis, MO 63130; [2]Inst. Biology II, RWTH, 52074 Aachen, Germany

We, and other animals, generally perceive motion independent of the cues that define the moving object. In an effort to understand the underlying mechanisms of this form-cue invariant generalization of stimulus attributes, we have examined the cellular properties of wide-field tectal neurons that are sensitive to a variety of moving but not to static stationary stimuli. This in vitro study revealed phasic signal transfer at the retino-tectal synapse and binary dendritic responses to synaptic inputs that interact in a mutually exclusive manner in the postsynaptic tectal neuron. A model of the tectal circuitry predicts that the two observed cellular properties mediate sensitivity to a wide range of dynamic spatiotemporal stimuli, including moving stimuli, but not to static stationary stimuli in a tectal neuron. The form-cue invariant computation for dynamic spatiotemporal stimuli is initiated by tectal neurons and apparently completed by rotundal neurons that integrate outputs from multiple tectal neurons in a directionally selective manner.

DELAY PERIOD MICROSTIMULATION PERTURBS MEMORY GUIDED SACCADE BEHAVIOR

Robert L. White and Lawrence H. Snyder
Department of Anatomy & Neurobiology.
Washington University, St. Louis

Previous experiments have shown that microstimulation above a threshold at sites in the frontal eye field (FEF) or the lateral intraparietal area (LIP) evokes saccades to specific locations in visual space. We were interested in whether subthreshold stimulation during a delay period could be used to probe spatial memory storage and processing in these areas. To do so, we performed stimulation below threshold while an animal planned a saccade to a remembered target and asked whether this influenced the subsequent saccade endpoint. We hypothesized that stimulation would systematically bias memory-guided saccades towards the RF of neurons recorded at the stimulation site.

Statistically significant deviations between stimulated and control saccades were observed at 27 of 36 FEF sites, but only 4 of 33 LIP sites (one animal for each area). Surprisingly, most of these effects were in the direction away from the recorded RF. In FEF, effects did not depend on whether significant ($p<0.05$) memory activity was recorded at the site ($n=23$). However, in LIP, when pooled across the population of sites with significant memory activity ($n=11$), a significant deviation in a direction away from the RF was observed. No effect was observed in the pooled population of LIP sites with only visual responses.

Why are stimulated saccades directed away from the RFs recorded at the site of stimulation? A repulsive effect could be explained if stimulation silenced a subpopulation of neurons that were involved in a distributed code of the spatial information. However, stimulation is known to cause excitatory effects in almost all other contexts. An alternative explanation is that stimulation may counterfeit an efference copy signal. The circuitry may respond to subthreshold stimulation as if a saccade has been executed in the suprathreshold evoked direction. The result is a deviation of the remembered location that is antiparallel to the evoked vector. Analysis of the direction of stimulation effects across all target locations revealed that the systematic effect of stimulation for the population was better explained by a vector antiparallel to the preferred direction observed at the site (or the suprathreshold evoked direction) than a vector antiparallel to the RF center. These results demonstrate that delay period microstimulation is a useful method for probing spatial information processing in the brain.

REVISITING HUBEL AND WIESEL: CLASSIFICATION OF SIMPLE AND COMPLEX CELLS IN A SPIKING NEURON MODEL OF MACAQUE STRIATE CORTEX

Jim Wielaard and Paul Sajda
Department of Biomedical Engineering, Columbia University, NY, NY 10027

Hubel and Wiesel (*J. Physiol (Lond)*, 1962)) postulated the existence of simple and complex cells as reflecting the existence of different classes of cortical circuitry and/or LGN axon connectivity within primary visual cortex (V1). Since their original work there has been significant debate on whether the simple and complex cell classification does in fact reflect a difference in neuronal circuitry or may be an artifact of non-linear measures on an otherwise single cell class (Mechler and Ringach, *Vision Research*, 2002). We use a large-scale spiking neuron conductance-based model of V1, layer 4, to investigate whether the simple/complex cell classification does in fact provide insight into the underlying neuronal circuitry. Using the model we compute both extracellular (spike) and intracellular modulation indices ($m(s)$ and $m(v)$) in response to drifting gratings. The extracellular modulation index has been a traditional metric for classifying cells, with complex cells having $m(s) < 1$ and simple cells $m(s) > 1$. This classification has been proposed because of the profound bimodality of the distribution of $m(s)$ over a large population of cells. The extracellular response of a cell is, however, a rather poor measure for its synaptic input because of the drastic nonlinearity introduced by the spike threshold. Intracellular responses are much better, but they are more difficult to obtain experimentally.

We consider two architecturally distinct cell classes in our model: those that receive direct LGN input and those that do not. Under the Hubel and Wiesel hierarchical model one would expect those cells receiving LGN input to be classified as simple and those that do not as complex. For the extracellular modulation index, our simulation results show that about 30% of cells that do not receive LGN input have $m(s) > 1$ (misclassified "simple") while about 30% of the cells that do receive LGN input have $m(s) < 1$ (misclassified as "complex"). The joint distribution of $(m(s), m(v))$ computed for the cellular responses in the model clearly displays a profound bimodality as well as exhibiting long "tails".

The simulation results make a clear prediction for the distribution of the intracellular modulation index. The marginal distribution of the intracellular modulation index is not, or is only weakly, bimodal though we find intracellular criteria for the simple/complex classification that are reasonably consistent with extracellular classification (e.g. 10-20% error). Concerning any hypotheses about whether a difference in the underlying circuitry of cells in V1 can be inferred based on the distribution of the modulation indices, our work seems to suggest that this metric does not necessarily provide a very precise picture of the circuitry, for example in our simulations it cannot accurately classify which cells do or do not receive direct LGN input.

SPATIAL ATTENTION SHIFTS RECEPTIVE FIELDS IN MACAQUE AREA MT TOWARD THE ATTENTIONAL FOCUS

Thilo Womelsdorf and Stefan Treue
Cognitive Neuroscience Laboratory
German Primate Center, Kellnerweg 4, 37077 Goettingen, Germany

Attention modulates the magnitude of sensory responses of single neurons in visual cortex. This has frequently been reported as a gain change, i.e. a multiplicative influence. Modulation is particularly strong when more than one stimulus is present at different positions within a neuronal receptive field (RF), but only one is attended. One hypothesis has suggested a space-based mechanism involving shrinkage and shift of the RF toward the attended location. Alternatively, the response could reflect gain changes in the neuron under study, leaving shape and position of the RF unchanged. We examined these hypothesis directly by obtaining quantitative RF profiles of macaque area MT neurons while attention was directed at different positions within the RF.

Our task required spatial attention to one of three moving random dot patterns (RDPs) in order to detect a direction change of the target RDP. Two RDPs were always placed within the RF of direction-selective neurons in macaque area MT. While one of the three stimuli was attended, quantitative RF maps were obtained by recording the responses to irrelevant, moving probe RDPs flashed at the intersections of a dense virtual grid covering the RF.

The resulting maps reveal a systematic shift of the sensitivity profile (the RF) toward the location of attention. We did not observe a corresponding shrinkage when attention was directed inside versus outside the RF.

Our results provide the first direct evidence for a dynamic change of spatial RF position with selective attention to different subregions of the RF - on a trial-by-trial basis and under identical visual stimulation and task difficulty: Attention enhances sensitivity close to the attended location and reduces it at distant position. This spatial weighting of responses corresponds to a non-multiplicative modulation - albeit without a shrinkage. It is, however, likewise accomodated with an attentional gain mechanism acting multiplicatively on pre-synaptic neurons with smaller RFs. Integration of these gain modulated responses would reproduce the observed non-multiplicative spatial sensitivity change. On a population level the reported spatial shift toward attended locations reflects the recruitment of additional neuronal resources devoted to feature analysis. It is consistent with behavioral evidence of spatial gradients of perceptual facilitation centred on attended locations.

ENCODING IN RAT PRIMARY SOMATOSENSORY CORTEX OF ELECTRICAL STIMULATION IN VPL THALAMUS

Shaohua Xu , John K. Chapin

State University of New York Downstate Medical Center,
Department of Physiology and Pharmacology, 450 Clarkson Ave,
Brooklyn, NY 11203

Neuronal ensemble responses to discrete natural touch of the forepaw and electrical stimulation in somatosensory thalamus (ventroposterior lateral nucleus, VPL) were recorded in the primary somatosensory cortex (S1) of anesthetized and awake rats. Multivariate statistical analyses were used to determine neural code for stimulation parameters such as stimulus location and intensity. An ideal space map was constructed to represent the pattern of cortical response to each different stimulus using factor analysis. We then tested a range of monopolar, bipolar and multi-electrode stimulation parameters in VPL and statistically analyzed the S1 cortical responses to each. Both natural and thalamic stimulation tended to produce excitatory response lasting from 10-25 msec followed by inhibition. Statistical comparison showed certain stimulation parameters yielding cortical neural responses that were remarkably similar to those evoked by natural touch, with the correlation coefficient (r) around 0.8. The cortical responses to these thalamic stimuli suited well in the space map mentioned above, while other thalamus stimuli poorly correlated (r generally below 0.4) to natural stimulation were far away from the space map. Our results indicated an optimal stimulation in VPL thalamus might generate a sensory perception close enough to the natural one for somatosensory prosthesis purpose. They might also have provided an electrophysiological explanation for previous finding that electrical stimulation in human thalamus induced mostly paresthesia but occasional natural tactile sensations.

NEURONAL ACTIVITY OF MIDBRAIN VENTRAL TEGMENTAL AREA (VTA) NEURONS IN SINGING BIRDS IS MODULATED BY SOCIAL CONTEXT.

Shin Yanagihara and Neal A. Hessler
RIKEN Brain Science Institute, Saitama, Japan

Birdsong, like human speech, is a learned vocalization used to communicate with other individuals. Male zebra finches sing a similar song in two distinct social contexts. “Directed song” is used during courtship of females, while “undirected song” is not targeted to any particular recipient. Previous studies have shown that singing-related neural activity (Hessler & Doupe, 1999) and immediate early gene expression (Jarvis, et al. 1998) in the anterior forebrain pathway (AFP) nuclei LMAN and Area X depend strongly on the social context in which a bird sings. In contrast, singing-related neural activity in the motor nuclei HVC and RA is similar in any social context (Hessler, et al. unpublished). Therefore, the AFP could receive a selective modulatory input dependent on social factors or motivational level. Given the role of dopamine in other systems in reward processes, we have begun to test whether this transmitter may mediate the social modulation of forebrain activity. In this study, we recorded activity of neurons in the midbrain dopaminergic nucleus VTA while birds sang both directed and undirected songs.

Consistent with a role of dopamine as a source of forebrain modulation in this system, we have found two distinct types of VTA neurons whose activity is consistently dependent on social context. The firing of “type-1” neurons phasically increased during directed singing, but did not change during undirected singing. The spontaneous firing of these neurons was also clearly dependent on the behavioral state of the bird - when a finch was quietly awake or asleep, such neurons ceased firing. Another type of VTA neuron fired in a pattern opposite to that of type-1 neurons. Such “type-2” neurons ceased firing during directed singing, but fired tonically during undirected singing. We also observed that responses of type-2 neurons were often habituated after repetitive trials of directed singing. In a previous study, a similar habituation of activity modulation was seen in LMAN. These results support the idea that VTA could directly modulate forebrain activity, dependent on the social context in which a bird sings. This system should be very useful for investigation of the role of dopamine in modulation of behavior and neural activity during rewarding behavior.

Context-adaptive maximum entropy coding of light intensity

Zhiyong Yang and Dale Purves
Center for Cognitive Neuroscience and Department of Neurobiology, Box 90999, LSRC Building, Duke University, Durham, NC 27708, USA
yang@neuro.duke.edu, purves@neuro.duke.edu

How the visual system encodes light patterns on the retina and transforms this information into the full range of percepts of the physical world is the central problem of vision. A simple starting point in understanding this process is to examine how the intensity of light reaching the retina is encoded and in relation to the corresponding perception (brightness), which is arguably the most basic aspect of vision. Validating any proposal about this scheme depends primarily on its ability to account for a wide range of perceptual phenomena in which the same amount of light from a visual target elicits different, sometimes radically different, brightness percepts as a function of context.

These phenomena, many of which date back a century or more, show that a code that simply generates a linear neuronal response to light intensity cannot be the code used by the brain. The debate about an appropriate alternative, initiated by Helmholtz, Hering, Mach and others, has shaped the development of modern vision science, and continues today. One common assumption at present is that the neural code for light intensity that ultimately elicits brightness percepts arises from a hierarchical scheme that involves low-level (retinal adaptation, filtering), intermediate-level (edge detection and grouping) and high level (inferred material properties and 3D spatial relations) neural processing. This approach, however, has not led to a satisfactory account of how light intensity is coded and brightness percepts are generated.

Here, we present an analysis of light intensity coding based on information theory. There is now much evidence that visual systems are optimized with respect to the natural environments that various visual animals inhabit, and that the neural code in early visual processing generates an efficient representation of visual stimulus. It follows that the neural coding of light intensity should be based on the probability distribution of light in the natural stimuli as a means of dealing in an efficient way with all possible occurrences of luminance. Since as a result of the statistical regularity of the physical world the light at one location in the retinal image is always statistically related to the light at another location, the coding of light intensity must be context-dependent. We have thus proposed that the code for light intensity is based on the conditional probability distribution of light at one location given the contextual pattern of light, always achieving the maximum entropy in the coding space and thus making use of the full coding capacity of the system. The brightness percept experienced by the observer simply should thus correspond to the value of the code of the underlying light intensity.

To test this hypothesis, we analyzed a large database of natural images to obtain the conditional probability distributions of particular target luminance values within specific contextual luminance patterns. The neural code of light intensity based on the probability distributions obtained this way accounts for a wide range of long standing brightness phenomena that have been previously assumed to arise from low-level, intermediate-level, and/or high-level of visual processing. This evidence supports the conclusion that the visual system codes light intensity in an optimal way on the basis of the probability distribution of the co-occurring luminance experienced in natural environments, and that the perception of light intensity is a direct consequence of this coding strategy.

IMPROVING NEURAL PROSTHETIC SYSTEM PERFORMANCE FOR A FIXED NUMBER OF NEURONS

Byron M. Yu[1], Stephen I. Ryu[1,2], Mark M. Churchland[1], Krishna V. Shenoy[1,3]
[1]Department of Electrical Engineering, [2]Department of Neurosurgery,
[3]Neurosciences Program
Stanford University, Stanford, CA 94305, USA

The performance of neural prosthetic systems can be improved by increasing the number of neurons used, developing better decode algorithms, or increasing the amount of model-fitting data. Although increasing the neuron count has been the principal approach, there is a practical limit to the number of neurons that can be obtained. To further increase decode performance, we assessed the benefits of employing the other two strategies.

We investigated point-to-point reaches in a delayed center-out task. We recorded from the dorsal pre-motor cortex (PMd) of a macaque monkey during reaches to visual targets at 10 different distances along a line (30 to 120 mm in each direction). Most PMd neurons emitted action potentials during both the delay period (*plan activity*) and the movement period (*peri-movement activity*). We recently proposed a decode algorithm that incorporates both plan and peri-movement activity. Here, we apply this idea to neural data and show how to improve performance for a fixed number of neurons.

We considered three decode algorithms, all of which employ maximum-likelihood classification with Gaussian data fitting. The first fits a Gaussian to the distribution of firing rates averaged over a single window that includes both the plan and movement periods (*Undifferentiated Rate*). The second estimates firing rates in the plan and movement periods separately (*Plan Rate / Move Rate*). The third is based on plan firing rate and the first principal component of the peri-movement spike train (*Plan Rate / Move PC*).

We found the average distance error between the true and estimated reach targets to be 15.8, 14.6, and 13.3 mm, respectively, for the three algorithms above using 20 neurons and 19 model-fitting trials per reach distance. By differentiating between the plan and peri-movement periods and by taking into account the time-varying structure of peri-movement activity (*Plan Rate / Move PC*), we obtained a 16.3% improvement in performance compared with *Undifferentiated Rate*. Furthermore, we found a 31.2% performance difference for *Plan Rate / Move PC* between training set sizes of four and 19.

Our results show that, even with a fixed neuron count, decode performance can be significantly improved by combining information from the plan and movement periods. The amount of training data is also an important, and possibly limiting, factor. Taken together, these ideas should enable higher-performance neural prosthetic systems for any fixed number of neurons.

DISCOVERING THE RECEPTIVE FIELD STRUCTURE OF COMPLEX CELLS WITH A MAX MODEL OF SYNAPTIC SUMMATION

Hsin-Hao Yu*, Zeph Landau† and Virginia de Sa*
*Cognitive Science, UC San Diego, La Jolla, CA 92093
†Mathematics, City College of New York, New York, NY 10031

Reverse-correlation is the most widely used automated method for mapping the receptive field structure of early visual neurons. A stimulus of white noise is used, and a linear model of the neuron can be estimated by calculating the average of the frames that triggered spikes (ie. *spike-triggered average*). For neurons with significant nonlinear summation of input signals, such as complex cells in the primary visual cortex, higher-order kernels must be considered. The *Spike-triggered covariance* technique estimates a second-order model of the neuron by taking the biggest eigenvectors of the covariance matrix of the stimuli triggered by spikes. This technique is closely related to the energy model of complex cells, which assumes that complex cells sum the squares of inputs from simple cells.

In this poster, we consider an alternative nonlinearity: the output of a complex cell is assumed to be the max of several linear filters (Riesenhuber & Poggio, 1988). There are three main motivations for using this model: 1. Spike threshold, an important nonlinear neural behavior, is naturally captured by the MAX operator. 2. Layered networks of neurons performing the MAX operation can be reduced to one MAX neuron receiving inputs from many linear filters. This allows us to model (to first order) the computation of several layers of cortex in a single MAX neuron. 3. The MAX model has the potential to capture more general distributions of the spike-triggering stimuli, such as disjointed manifolds, instead of the elliptical manifold assumed by the PCA-based spike-triggered covariance. Also the linear filters are not constrained to be orthogonal.

An iterative learning rule for the linear filters is derived by performing gradient descent on the Kullback-Leibler divergence between the observed probabilities of spikes and the estimation of the model. The non-continuity of the derivatives through the max function is not a problem for the same reason it is not a problem in Competitive Learning (Rumelhart & Zipser, 1986): at the discontinuities the value of the error function is the same on both sides of the discontinuity. We have successfully fitted the MAX model to simulated complex cells made by the energy model. We are currently applying our algorithms to physiological data recorded from the complex cells in cats.

Acknowledgement This work was supported by the National Science Foundation Career Grant No. 0133996

Saddle-point property of nonlinear sensory response

Kechen Zhang, Michael J. Anderson, and Eric D. Young
Department of Biomedical Engineering
Johns Hopkins University School of Medicine
Baltimore, Maryland 21205
{kechen, anderson, eyoung}@bme.jhu.edu

Randomly generated stimuli have been used to probe the response properties of neurons in various sensory modalities. The simplest nonlinear model beyond a linear kernel is a quadratic or second-order model in power series expansion. Using a quadratic model, we analyze how the firing rate of an auditory neuron depends smoothly on the parameters of random stimuli and, in particular, where the optimal stimuli are located in the stimulus space. An optimal stimulus, as one would expect intuitively, should correspond to a local maximum of response so that any small deviation from the optimum in the stimulus space always decreases the response. A quadratic model allows a maximum only when all the eigenvalues of the second-order kernel are non-positive. We have found both positive and negative eigenvalues, often of comparable magnitudes, in quadratic models fitted to neurons recorded from the cochlear nucleus of unanesthetized decerebrate cats in response to random spectra stimuli across different sound pressure levels. This result implies a surprising saddle-point property in the stimulus-response relationship; that is, the response to the same stimulus can be both a maximal and a minimal response, depending on which direction the stimulus is perturbed. In other words, the optimal stimulus as a true maximum cannot exist here; among all stimuli allowable in an environment or all stimuli that are producible by the experimental setup, the stimulus that elicits the largest response must always lie on the boundary of the allowable stimuli. To further verify the quadratic theory, we examined how the average responses depended on the overlaps of random stimuli with the eigen-stimuli defined by the second-order kernel, and confirmed both convex and concave curves as predicted by the signs of the eigenvalues. We have also tested the theoretical predictions directly by using eigen-stimuli generated online and applied to the same neurons after their quadratic model were obtained. Finally, because saddle point behavior was very common in the neurons tested, we sought to understand its mechanism by studying standard rate models based on known anatomy and physiology, and found analytically that the eigenvectors of the second partial derivatives or the Hessian matrix, which should correspond to the second order kernel, are linear combinations of the synaptic weight vectors involved, and the associated eigenvalues can indeed have opposite signs.

CORTICAL NETWORKS THAT LEARN TO SYNCHRONIZE

Valentin Zhigulin
Physics Department, California Institute of Technology, Pasadena, CA 91125
Institute for Nonlinear Science, UCSD, La Jolla, CA 92093-0402

Synchronous neural activity plays an important role in the functioning of the brain. It is a robust phenomenon, frequently observed across populations of neurons with diverse membrane properties and intrinsic frequencies. In the light of such diversity it remains unclear how can precise synchronization be achieved in heterogeneous networks. Several mechanism were suggested and many of them require unreasonably high degree of network homogeneity or very strong connectivity to achieve coherent neural activity. It was recently demonstrated in simulations of two synaptically coupled excitatory neurons that spike timing dependent synaptic plasticity (STDP) enhances phase locking of neural oscillations and makes it robust to noise and variability of the membrane properties [1, 2]. These effects are caused by the dynamic self-adaptation of the synaptic conductance to the value that is optimal for the entrainment of postsynaptic neuron under given mismatch of pre- and post-synaptic frequencies. This entrainment is with non-zero phase lag between neurons because model neurons with type I excitability were used. Such neurons have positive phase response curve and their in-phase synchronization is not possible. However, most observations of synchronization in the brain reveal very small or zero phase lag between oscillations of neurons from different, sometimes very distant, cortical areas. In this work we show that, due to a mechanism similar to the described above, STDP leads to the robust in-phase synchronization in simulations of heterogeneous networks of excitatory neurons with type II excitability. Taking into account the fact that most neurons in the brain are believed to be of the type I, we also show that the same phenomenon of robust in-phase synchronization is observed in heterogeneous excitatory-inhibitory networks of type I neurons. The ability of inhibitory interneurons to effectively change excitability of their postsynaptic targets from type I to type II explains these observations. Finally, we show that such networks are able to represent different external stimuli by synchronous spiking activity in different parts of the network.

[1] V. P. Zhigulin, M. I. Rabinovich, R. Huerta, and H. D. I. Abarbanel. Robustness and enhancement of neural synchronization by activity-dependent coupling. *Physical Review E*, **67**, 021901 (2003).

[2] T. Nowotny, V. P. Zhigulin, A. I. Selverston, H. D. I. Abarbanel, and M. I. Rabinovich. Enhancement of synchronization in a hybrid neural circuit by spike timing dependent plasticity. *J. Neuroscience* **23**, 9776-9785 (2003).

UNCERTAINTY IN SENSORIMOTOR CONTROL

Daniel Wolpert

Sensory and motor uncertainties form fundamental constraints on human sensorimotor control. I will first describe how signal-dependent noise on the motor output places constraints on performance. Given these constraints features of goal-directed movement arise from a model in which the statistics of our actions are optimized. I will then describe how prediction of the consequences of our actions can be used to reduce uncertainty and present experiments on tickling and force escalation which elucidate the predictive mechanisms. Finally, I will show that the CNS reduces the uncertainty in estimates about the state of the world by using a Bayesian combination of prior knowledge with an estimate of the uncertainty of its own sensors. Together these studies provide a probabilistic framework for sensorimotor control in which prediction plays a key role.

A MODEL FOR CONTEXT-DEPENDENT SELECTION OF VISUO-MOTOR MAPS

Melanie Wyder and Emilio Salinas
Dept. of Neurobiology and Anatomy, Wake Forest University School of Medicine, Winston Salem, NC 27157, USA

The relationship between sensory information and behavior is complex and highly labile. Thus, in different contexts, the same stimulus might elicit very different behaviors. For instance, the response to the ringing of a doorbell depends heavily on whether one is at home, or at someone else's house. The neural basis of this flexibility is not known. One possibility is based on a nonlinear representation in which contextual information regulates the gain of sensory responses. Here we present three neural network models for three hypothetical oculomotor tasks, each requiring a flexible association between stimulus and response. While the precise tasks are not critical, we present the results using behavioral paradigms similar to those that have been used in experiments with awake-behaving monkeys. Within each experiment, the same physical stimulus is presented in several conditions, with a different motor output (saccade) required under each condition. Crucially, the cue that signals context contains no explicit spatial information. Nevertheless, in each case, the contextual cue can quickly reconfigure the sensory-motor maps, effectively changing the functional connectivity of the network and thus the association between sensory stimuli and behavior. The model networks consist of a first layer of gain-modulated sensory units that drives a second layer of output units, with each sensory unit connected to all output units. The activity of the output units represents a population code of the vector of an impending saccade. The firing rate of each output unit is determined by a simple weighted sum of the gain-modulated sensory-unit firing rates, with weights corresponding to synaptic connections. The models predict sensory responses that are nonlinearly modulated by contextual cues in a wide variety of behavioral situations. Because any relevant circumstantial information can be part of the context, this mechanism should allow for flexible associations between sensory stimuli and behaviors, and may partly explain the complex and rich behavioral repertoire of higher organisms.

PREDICTING ACTION DURING FREE CHOICE: REACH PLANS IN POSTERIOR PARIETAL CORTEX

Bijan Pesaran and Richard A. Andersen
Division of Biology
California Institute of Technology, Pasadena, CA 91125

In natural behavior, reaching to visual targets is the result of a selection process and occurs as a coordinated hand-eye movement. However, there are few studies of hand-eye coordination in monkeys, and most studies employ instructed behaviors without choice. To study the development of reach plans in PPC under more natural conditions, we trained two monkeys to perform a task in which they were allowed to choose where to look and where to reach. Monkeys were presented with three identical reach targets one of which resulted in a reward. The reward target was assigned randomly each trial and they made reaches to search for the reward until they found it. We chronically implanted multielectrode arrays in the medial bank of the intraparietal sulcus (mIPS; 64 electrodes) and area 5 (A5; 32 electrodes) in each of two monkeys and we recorded spiking, local field potential (LFP) activity and eye position as they performed the task. To characterize the selection and planning process before the first reach we used a predictive framework in which the monkeys choice was predicted from neural activity and we compared performance of spiking and LFP activity in each cortical area.

After many trials, each monkey learned a strategy for performing the task summarized as probabilities of selecting a given target. Analysis of hand-eye coordination showed that on average the monkeys scanned the targets with 4 saccades before making the first reach and looked at their hand before reaching. Spiking and LFP activity in both mIPS and A5 showed strong spatial tuning for the reach direction but had different temporal dynamics. Spike rate in both areas and LFP activity in A5 became significantly tuned after the saccade to the hand while LFP activity in the mIPS was spatially tuned before the saccade to the hand. Decoding spike activity, we found reach direction (1 of 8) could be most accurately decoded at reach onset (mIPS:97%; A5:84%). This declined to chance 400 ms before the reach. Decoding LFP activity, we found reach direction (1 of 8) could also be most accurately decoded at reach onset (mIPS:90%; A5:95%). Surprisingly, LFP activity in mIPS predicted the reach choice up to 300 ms earlier than LFP activity in A5 or spiking activity in either area. In summary, under these conditions we find monkeys prefer to look to the hand before reaching which could have advantages for coordinate transformations for reaching. We also find LFP activity in mIPS predicts the reach choice earlier than spiking suggesting this area may receive top-down inputs before computing the visual-motor transformation.

THE CONTRIBUTION OF INVERSE PLASTICITY MECHANISMS TO CEREBELLUM-DEPENDENT LEARNING

Edward S. Boyden[1], Talal A. Chatila[2], and Jennifer L. Raymond[1]
1. Department of Neurobiology, Stanford University
2. Department of Pediatrics, Washington Univ. School of Medicine

Many plasticity mechanisms have been characterized *in vitro* in circuits important for learning. We are beginning to parse out the dependence of motor learning upon different plasticity mechanisms in the cerebellum and related structures. The most influential model of cerebellum-dependent learning has focused on a single plasticity mechanism, long-term depression (LTD) at parallel fiber to Purkinje cell (pf-Pk) synapses (Albus 1971, Marr 1969), but our studies of the vestibulo-ocular reflex (VOR) suggest that multiple plasticity mechanisms are involved. Applied to motor learning in the VOR, the Marr-Albus model attributes both increases and decreases in VOR gain to pf-Pk LTD (Ito 1982). Here we present two lines of evidence that increases and decreases in VOR gain depend upon different plasticity mechanisms.

First, increases in VOR gain are reversed more readily than decreases in VOR gain. This difference in reversal properties suggests that increases and decreases in VOR gain are mediated by different plasticity mechanisms. Furthermore, the behavioral asymmetry suggests that these plasticity mechanisms reverse each other with unequal efficacy (Boyden & Raymond 2003). The plasticity mechanisms at the pf-Pk synapse (LTD and two forms of LTP) seem to possess such an asymmetric reversal property. Thus we propose a new model in which pf-Pk LTD contributes primarily to increases in VOR gain, whereas pf-Pk LTP contributes primarily to decreases in VOR gain. We are testing this model with mutant mice deficient in either pf-Pk LTD or one of the forms of pf-Pk LTP. Mice lacking Ca^{+2}/CaM-kinase IV (CaMKIV), a molecule required for the late phase of pf-Pk LTD are selectively impaired in retention of an increase in VOR gain. Acquisition of increases and decreases in VOR gain is normal, as is retention of a decrease in VOR gain (Boyden et al 2003). Therefore, long-term memory for an increase in VOR gain relies upon a CaMKIV-dependent process (such as pf-Pk LTD), whereas long-term memory for a decrease in gain does not. Our results suggest that models of cerebellum-dependent motor learning should be revised to consider the role of inverse plasticity mechanisms at pf-Pk synapses.

EFFICIENT REPRESENTATION OF PREDICTIVE INFORMATION

William Bialek (Princeton University)

The use of information theory as a framework for thinking about biological systems has always been problematic because Shannon's formulation does not seem to have any notions of relevance, value or meaning for the information being measured. While it might make sense to say that our sensory systems are constructed (perhaps even optimized) to gather information, surely at some point we have to face the fact that some bits are more interesting than others, and that our brains almost certainly discard the uninteresting bits. I've been interested in the possibility that what is relevant, valuable and perhaps even meaningful for the organism is the ability to make predictions, and hence that the goal should be not the efficient representation of all information but rather the efficient representation of **predictive information**.

The idea that it is useful to make predictions hardly seems radical (although it is radical to think that concepts such as relevance or value can be identified precisely with any single mathematical construct). What is surprising is how far we can go by formalizing this idea in information theoretic terms. One can show that the amount of predictive information in a stream of data provides a compelling quantitative measure of what we intuitively call the complexity or richness of those data. For data streams with appropriate structure, efficient representation of the predictive information corresponds to well known strategies in signal processing, such as matched filtering to separate signal from noise or the identification of hidden states for an underlying Markov model of the data. At a higher level, generalization in learning is exactly the problem of prediction, and efficient representation of predictive information corresponds to extracting the parameters of models that describe the incoming data stream. Generalizing to the case where we don't just collect data but can act back upon the world to shape or select the data stream, the same efficiency principle generates a form of exploration/exploitation tradeoff and may contain as a special case a theory of active learning.

To connect these theoretical ideas back to experiment, I'll discuss several possibilities: Could the timing-dependent plasticity rules at synapses be derivable as a strategy for optimizing the extraction of predictive information from the stream of synaptic inputs? Could we measure the efficiency with which animals (including humans) extract predictive information when they learn, especially in a probabilistic context? Finally, if all goes well, I'll describe the design and analysis of a preliminary experiment which measures directly the information that a neuron carries about the future of its sensory inputs.

This presentation draws on joint work with Ilya Nemenman, Rob de Ruyter van Steveninck, Susanne Still, and Naftali Tishby.

RECEPTIVE FIELDS IN CAT PRIMARY VISUAL CORTEX OBTAINED AS MOST INFORMATIVE DIMENSIONS FOR NATURAL STIMULUS ENSEMBLE

T. Sharpee[1,2], S. Rebrik[1], H. Sugihara[1], M.P. Stryker[1,2], and K.D. Miller[1,2]
Department of Physiology[1] and Sloan-Swartz Center for Theoretical Neurobiology[2],
University of California, San Francisco

We have obtained receptive fields of cells in cat V1 by stimulation with natural time-varying images and compared them to receptive fields obtained from a white noise stimulus. In addition, the optimal orientation and spatial frequency were measured with moving gratings for each cell. Neural responses were analyzed in the framework of a linear-nonlinear model, in which spike probability is a nonlinear function of the stimulus convolved with the cell’s spatiotemporal receptive field (STRF). In the case of white noise stimulation, the cell’s STRF is given by the spike triggered average (STA). In the case of the natural stimulus ensemble, the STRF is found as the maximally informative dimension (MID), i.e. the dimension in the stimulus space whose convolution with the stimulus ensemble accounts for the maximal possible part of the overall information carried by the arrival of one spike [1]. We find that, on average, the STRF obtained in response to natural scenes predicts the cell’s preferred orientation to moving gratings somewhat better than the RF obtained as the STA in response to white noise. Two possible reasons for this better performance include (1) The STRFs obtained as the MID to natural stimuli show, on average, higher signal-to-noise ratios than those obtained as STA's to white noise; (2) Although the main features of STRFs, such as spatial size, preferred orientation and preferred spatial frequency, are similar between the white noise STA and the MID for natural stimuli, the deviations between these two STRF estimates exceed the noise level. Despite these differences, the white noise STA more closely resembles the MID for natural stimuli than another estimate of the STRF obtained from responses to natural stimuli, namely the STA of the natural stimuli responses corrected for the second-order correlations of the natural ensemble. The latter vector would be a correct estimate of the cell's STRF if it had a linear input-output function or if the statistics of natural scenes were Gaussian. These results confirm that computation of the MID is both practically feasible and yields equal or superior results compared to other existing methods. Furthermore, this method is applicable to cells that do not respond well to noise stimuli or do not show a structured STA in response to noise stimuli. For cells that do show a structured STA in response to white noise, comparison between the MID to natural scenes and the white noise STA allows us to study possible changes that occur in the receptive field parameters as a result of adaptation to the probability distribution of inputs.

[1] T. Sharpee, N.C. Rust, and W. Bialek, “Analyzing neural responses to natural signals: maximally informative dimensions” to appear in Neural Computation, 16 (2) 2004.

DISENTANGLING SENSORY TRANSDUCTION CHAINS *IN VIVO*

Tim Gollisch and Andreas V.M. Herz
Institute for Theoretical Biology, Humboldt University Berlin, Germany

Signal-processing chains underly the proper function of all sensory systems. We here show that it is possible to extract fine temporal details of individual processes within such chains from observing the output activity alone. This progress results from a new experimental approach that is based on identifying specific iso-response sets, i.e., sets of stimuli that cause the same response. The key to the method is to compare different stimuli within one iso-response set in such a way that single processing steps can be dissociated.

We apply this strategy to the auditory transduction chain of locusts, and obtain a precise characterization of the mechanical resonance of the ear drum and the electrical integration of the attached receptor neuron *in-vivo*. New details of these processes are revealed with a resolution far below one millisecond. This was possible due to four attractive properties of the new framework:

First, the method effectively decouples temporal sensitivity on the input side from temporal precision on the output side by focusing on spike probabilities. In our measurements, e.g., spike latencies varied by about one millisecond within a single recording set due to cell-intrinsic noise. Still we were able to probe the system with a resolution down to a few microseconds. This would have been impossible using classical correlation-based techniques.

Second, the method is robust against moderate levels of spontaneous output activity, as this affects all stimuli within one iso-response set in the same way. Methods that require measurements at different response levels, on the other hand, are likely to be systematically affected because the same internal noise level may have a different influence at different levels of output activity.

Third, the method can be used to investigate processing sequences that include multiple nonlinear steps; our system corresponds to a LNLN cascade. Finally, in many input-output systems, the last stage of processing can be described by a monotonic nonlinearity. Here, this is the relation between the effective stimulus strength S and the spike probability p. By always comparing stimuli that yield the same output activity, our analysis is independent of the actual shape of $p(S)$. Preceding integration steps may thus be analyzed without any need to model $p(S)$. This feature is independent of the specific output measure and applies to spike probabilities, firing rates, etc.

The technique also allows one to dissociate feedforward contributions from output-dependent feedback, which is necessarily constant within one iso-response set. With increasingly available high-speed computer power for online analysis and stimulus generation, our method therefore seems well suited to identify even complicated signal-processing chains in many systems.

REDUNDANT, INDEPENDENT AND SYNERGISTIC SYMBOLS IN A POPULATION NEURAL CODE

Elad Schneidman, Jason Puchalla, William Bialek & Michael J. Berry II
Department of Molecular Biology, Department of Physics and
the Lewis-Sigler Institute for Integrative Genomics
Princeton University, Princeton, NJ 08544 USA
{elads,puchalla,wbialek,berry}@princeton.edu

Understanding the neural code of an ensemble of neurons relies on the characterization of the activity patterns of groups of neurons, and the information that these activity patterns encode. Ideas of redundant, independent and synergistic encoding have played an important role in theoretical and experimental studies of population coding. On one hand, redundancy reduction, decorrelation and independent coding channels have been suggested as design principles for efficient information coding. On the other hand, it has been argued that redundancy is fundamental for robustness of information coding by neurons and for learning how to decode neural activity.

Our intuitive notions of redundancy, synergy and independence of neurons, can be made precise using information theory: redundancy (or synergy) among cells is defined as the difference between the information that the cells carry together and the sum of information values that each of the cells carry alone. Combining the redundancy (synergy) with information measures of the joint activity of cells and the information among cells given the stimulus, results in a consistent framework for the description and interpretation of population encoding, which overcomes known difficulties with the interpretation of cross-correlation functions. These measures have been used for the study of population coding in sensory and cortical neurons. Interestingly, in some cases, neurons have been found to be close to independent, whereas in other cases, significant values of redundancy or synergy have been reported.

However, mutual information, as well as the synergy or redundancy measure, are average quantities of the full ‘coding dictionary’. Rather than asking about the average information that is carried by a single neuron or a group of neurons, we can ask about the information about the stimulus that specific code words convey, thus giving a more complete characterization of the neural code.

We have applied this approach to study the information content of specific code words of pairs of ganglion cells in the tiger salamander in response to different stimuli. Using natural movies as well as simplified artificial stimuli (randomly flickering checkerboards and spatially uniform flicker), we explore the distribution of information contents of compound symbols and their relation to the nature of the stimulus.

YET ANOTHER ANALYSIS OF THE ROLE OF CORRELATIONS IN THE NEURAL CODE

Peter E. Latham and Sheila Nirenberg
Department of Neurobiology.
University of California at Los Angeles, Los Angeles, CA 90095

One of the main goals of neural coding is to to understand how spike trains encode information. For sensory processing, which we will concentrate on here, “understand” is synonymous with “able to translate spike trains into the stimuli that produced them”. One of the biggest potential obstacles to this is the existence of correlations – the fact that, for a given stimulus, spikes both within and across neurons are not independent. This is because correlations make the transformation from stimulus to response high-dimensional, and thus essentially impossible to estimate from data.

There is, however, mounting evidence in mammalian retina [1], somatosensory cortex [2], supplementary motor area [3], and visual cortex [4], that correlations among pairs of neurons can be ignored without much loss of information. If this result were to extend beyond pairs, to populations, it would greatly simplify the problem of decoding spike trains. Here we discuss two general approaches for assessing the role of correlations, one based on decoding and the other on information theory. We then show that a recently proposed information-theoretic cost function [1, 2, 5] provides an upper bound on the information lost when correlations are ignored. This last result provides us with a code-independent measure that can be used to evaluate the importance of correlations for transmitting information.

[1] S. Nirenberg, S.M. Carcieri, A.L. Jacobs, and P.E. Latham (2001). Retinal ganglion cells act largely as independent encoders. *Nature* **411**:698-701.

[2] S. Panzeri, G. Pola, F. Petroni, M.P. Young, and R.S. Petersen (2002). A critical assessment of different measures of the information carried by correlated neuronal firing *Biosystems* **67**:177-185.

[3] B.B. Averbeck and D. Lee (2003). Neural noise and movement-related codes in the macaque supplementary motor area. *J. Neurosci.* **23**:7630-7641.

[4] H.D. Golledge, S. Panzeri, F. Zheng, G. Pola, J.W. Scannell, D.V. Giannikopoulos, R.J. Mason, M.J. Tovee, and M.P. Young (2003). Correlations, feature-binding and population coding in primary visual cortex. *Neuroreport* **14**:1045-1050.

[5] S. Nirenberg and P.E. Latham (2003). Decoding neuronal spike trains: how important are correlations? *Proc. Natl. Acad. Sci.* **100**:7348-7353.

NEW PLACES AND NEW REPRESENTATIONS: PLACE FIELD FORMATION AND STABILIZATION IN THE HIPPOCAMPUS

Loren M. Frank[1,2,3], Garrett B. Stanley[2], and Emery N. Brown[1]
[1]Department of Anaesthesia and Critical Care, Massachusetts General Hosptial, Boston, MA. [2]Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA..[3]Current Address: Keck Center for Integrative Physiology, Department of Physiology, University of California, San Francisco, San Francisco CA.

Numerous studies have demonstrated that neurons in the hippocampus are active in a place specific manner and that hippocampal activity and hippocampal plasticity are essential for the learning of spatial tasks in rodents. At the same time, little is known about the precise time course of place field formation or the types of changes that are seen over multiple days of exposure to a new place. To understand better hippocampal neural dynamics and the formation of spatial representations we trained four rats on a continuous alternation task in three arms of an eight-arm maze. Each day we recorded from multiple neurons in the CA1 region of the hippocampus and in adjacent cortical regions while the animals performed the task in both a familiar configuration and in a novel configuration involving a previously unexplored arm of the maze. Here we focus on the results from CA1. We found that substantial plasticity was present during the first day in a new place, including numerous instances where a previously inactive neuron developed a new place field over the course of a single traversal of the environment. The activity of simultaneously recorded putative interneurons increased over the course of the novel exposure, suggesting the rapid emergence of place fields was not due to decreased inhibition. Surprisingly, the largest changes in place cell activity occurred during the second day when the animal had had relatively little experience in the novel arm on the first day (< 3 minutes). This difference suggests that in situations where the hippocampus has not had sufficient time to form a stable representation on the first exposure to a new place, the circuit enters a highly plastic state on the second exposure. This pattern of results was not visible when the data were organized by the number of traversals through the environment, indicating that the total time spend in a new place, rather than the number of visits, was related to the observed plasticity. First day experiences longer than 3 minutes were not associated with amplified plasticity on day 2, and our data suggest that the hippocampal representation stabilized after between 3 and 6 minutes of experience. These results suggest that the hippocampus requires a minimum amount of experience in a place to form a stable representation, and that an experience shorter than that minimum may produce memories that are vulnerable to interference.

LOCAL COMPETITIVE INTERACTIONS IN RAT ORBITOFRONTAL CORTICAL MICROCIRCUITS DURING A DECISION TASK

M.C. Quirk, C.E. Feierstein, Z.F.Mainen.
Cold Spring Harbor Laboratory, Cold Spring Harbor, New York, USA.

While much is known about the cellular and synaptic components of cortical circuits *in vitro*, considerably less is known about how these circuits behave in the intact functioning brain. Using tetrode recordings it is possible to isolate action potentials arising from multiple cells within a small local volume of tissue (diameter <100 um) and thus to simultaneously monitor multiple single units in a region comparable to the size of the cortical micro-column. We used tetrodes to record ensembles of neurons within the orbital frontal cortex (OFC) of rats performing a two-alternative olfactory discrimination task (Uchida and Mainen, 2003) in order to study the representation of information within cortical microcircuits during the formation of a binary perceptual decision.

In sharp contrast to the typical columnar organization of primary sensory areas, nearby cells in OFC displayed diverse response properties. The heterogeneity of responses across cells recorded in the same tetrode was comparable to that of cells across tetrodes, indicating that the salient task variables represented in OFC (including sensory, motor and reward information) are available locally within individual microcolumns.

Because cells within a cortical column receive largely overlapping afferent input, the diversity we observed is likely to arise from properties local to the microcircuit. To investigate the contribution of intrinsic cell properties to these response characteristics, we used criteria including spike waveform and interspike interval distributions to define classes of neurons including putative fast spiking interneurons, bursting cells, regular spiking cells, and long refractory period cells. Differences in these biophysical characteristics appear to account for some differences in OFC response properties in the discrimination task.

We next used cross-correlation analysis to investigate the contribution of synaptic interactions to functional diversity in OFC. The interactions between neurons were influenced not only by the distance between cells but also by cell type and the behavioral state of an animal. In particular, nearby cells were often anti-correlated with each other during task performance. Furthermore, strong trial-by-trial correlations were sometimes observed in the firing rates of nearby neurons, suggesting competition between neighboring neurons whereby functionally distinct subsets of cells actively suppress the firing of other local neurons.

These competitive interactions within OFC cortical microcircuits suggest a mechanism for the creation of diverse local responses and provide a potential substrate for winner-take-all computations that may be essential components of neural mechanisms of decision making.

RAPID LEARNING BY REINFORCEMENT OF VARIATION DUE TO INJECTED NOISE IN A SPIKING NEURAL NETWORK

Ila R. Fiete[1] and H. Sebastian Seung[2]
[1] Department of Physics, Harvard University, Cambridge, MA 02138
[1,2]Brain and Cog. Sci. and [2]H.H.M.I., M.I.T., Cambridge, MA 02139

The adage that "practice makes perfect" refers to the gradual learning of complex skills through an iterative process of trial, error, and improvement. Here we hypothesize that such behavioral learning is based on blind variation, in the form of noisy neural inputs, and selective reward-based modification, guided by the correlation of a postsynaptic noise input with a global reinforcement signal. It can be shown that the learning rule performs stochastic gradient ascent on the reward. However, stochastic gradient learning is often criticized as too slow to be biologically plausible. We demonstrate with a combination of numerical simulations and theoretical arguments that, contrary to conventional wisdom, the learning time in a test case of birdsong acquisition in a spiking network model with delayed reward is fast enough to be compatible with that of real zebra finches. This is because learning time scales up with the complexity of the task to be learned, and not with the size of the network; since the information content of song is reasonably small, learning proceeds rapidly. Our results suggest that stochastic gradient learning can, in at least some contexts, be fast enough to be taken seriously as a biological mechanism.

Participant List

Dr. Sandra Aamodt
Nature Neuroscience
s.aamodt@natureny.com

Dr. Gloster Aaron
Columbia University
ga264@columbia.edu

Dr. Larry Abbott
Brandeis University
abbott@brandeis.edu

Mr. Corey Acker
Boston University
cacker@bu.edu

Mr. Emre Aksay
Princeton University
eaksay@princeton.edu

Mr. Mark Andermann
Program in Biophysics, Harvard University
anderman@fas.harvard.edu

Dr. Michael Anderson
Johns Hopkins
anderson@bme.jhu.edu

Dr. Jorn Anemuller
University of California San Diego
jorn@salk.edu

Dr. Kevin Archie
University of Southern California
karchie@lnc.usc.edu

Mr. Dmitriy Aronov
Columbia University
da2006@columbia.edu

Dr. Gurinder Atwal
Princeton University
gatwal@princeton.edu

Dr. Bruno Averbeck
University of Rochester
baverbeck@cvs.rochester.edu

Dr. Stephen Baccus
Harvard University
baccus@fas.harvard.edu

Mr. Steven Badelt
UCLA
sbadelt@ucla.edu

Mr. Jonathan Baker
Montana State University, Bozeman - CCB
jbaker@nervana.montana.edu

Dr. Matthew Banks
University of Wisconsin
mibanks@wisc.edu

Mr. Boris Barbour
Ecole Normale Superieure
barbour@ens.fr

Dr. Edward Bartlett
Johns Hopkins University
bartlett@bme.jhu.edu

Dr. Sue Becker
McMaster University
becker@mcmaster.ca

Dr. Anthony Bell
Redwood Neuroscience Institute
tbell@rni.org

Ms. Marina Belova
Columbia University
mab2058@columbia.edu

Mr. Jan Benda
University of Ottawa
j.benda@biologie.hu-berlin.de

Mr. Tom Bennett
University College London
t.bennett@ucl.ac.uk

Prof. Michael Berry II
Princeton University
berry@princeton.edu

Dr. Matthias Bethge
Redwood Neuroscience Institute
mbethge@rni.org

Dr. William Bialek
NEC Research Institute
wbialek@princeton.edu

Dr. Elie Bienenstock
Brown University
elie@dam.brown.edu

Prof. Kwabena Boahen
University of Pennsylvania
boahen@seas.upenn.edu

Dr. Sander Bohte
University of Colorado, Boulder
bohtes@cs.colorado.edu

Mr. Alireza Boloori
Harvard University
boloori@deas.harvard.edu

Dr. Stephen Bonasera
University of California, San Francisco
bonaser@itsa.ucsf.edu

Dr. Tobias Bonhoeffer
Max-Planck-Institute of Neurobiology
tobias.bonhoeffer@neuro.mpg.de

Mr. Vincent Bonin
Smith-Kettlewell Eye Research Institute
vincent@ski.org

Dr. Leanne Boucher
Vanderbilt University
leanne.boucher@vanderbilt.edu

Mr. Edward Boyden
Stanford University
boyden@stanford.edu

Dr. David Brainard
University of Pennsylvania
brainard@psych.upenn.edu

Dr. Naama Brenner
Technion - Israel Institute of Technology
nbrenner@tx.technion.ac.il

Mr. Steven Bressler
Florida Atlantic University
bressler@fau.edu

Mr. Kevin Briggman
University of California - San Diego
kbriggma@ucsd.edu

Dr. Ken Britten
UC Davis
khbritten@ucdavis.edu

Dr. Carlos Brody
Cold Spring Harbor Laboratory
brody@cshl.edu

Mr. Bede Broome
Caltech
bede@caltech.edu

Mr. Eric Brown
Princeton University
ebrown@math.princeton.edu

Dr. Marina Brozovich
Catech
brozovic@vis.caltech.edu

Dr. Nicolas Brunel
CNRS
nicolas.brunel@univ-paris5.fr

Mr. Douglas Buchanan
Columbia University
drb2004@columbia.edu

Dr. Daniel Butts
Harvard Medical School
daniel_butts@hms.harvard.edu

Dr. Matteo Carandini
The Smith-Kettlewell Eye Reserach Institute
matteo@ski.org

Mr. Stijn Cassenaer
California Institute of Technology
stijn@caltech.edu

Mr. Paul Chadderton
University College London
p.chadderton@ucl.ac.uk

Ms. Beth Chen
Cold Spring Harbor Laboratory
chenb@cshl.edu

Dr. Sen Cheng
UC San Francisco
chengs@phy.ucsf.edu

Mr. Denis Chigirev
Princeton University
chigirev@princeton.edu

Dr. Dmitri Chklovskii
Cold Spring Harbor Laboratory
mitya@cshl.edu

Dr. Jinhoon Choi
Sogang University
choi@sogang.ac.kr

Dr. Hirsh Cohen
Alfred Pl Sloan Foundation
cohen@sloan.org

Dr. Erik Cook
Harvard University
epcook@fas.harvard.edu

Mr. Greg Corrado
Stanford University
gcorrado@stanford.edu

Ms. Gabriela Czanner
University of Pittsburgh
gamst24@pitt.edu

Dr. Rava da Silveira
Harvard University
rava@physics.harvard.edu

Mr. Kofi DARKO
obengkorankye@yahoo.com

Dr. Akira Date
Communications Research Laboratory
date@crl.go.jp

Mr. Stephen David
University of California, Berkeley
stephen@socrates.berkeley.edu

Dr. Nathaniel Daw
University College London
daw

Dr. Peter Dayan
Gatsby Computational Neuroscience Unit
dayan@gatsby.ucl.ac.uk

Ms. Saskia de Vries
Harvard University
sdevries@fas.harvard.edu

Dr. Valeria Del Prete
King's College London
vale@mth.kcl.ac.uk

Dr. Sophie Deneve
University College London
sdeneve@gatsby.ucl.ac.uk

Dr. Mike DeWeese
Cold Spring Harbor Laboratory
deweese@CSHL.edu

Dr. Kamran Diba
California Institute of Technology
diba@klab.caltech.edu

Dr. Jochen Ditterich
Univ. of California
jd@shadlen.org

Dr. Melissa Dominguez
McMaster University
domingm@mcmaster.ca

Dr. Rodney Douglas
Institute of Neuroinformatics
simone@ini.phys.ethz.ch

Mr. Patrick Drew
Brandeis University
drew@brandeis.edu

Prof. Martin Egelhaaf
Bielefeld University
martin.egelhaaf@uni-bielefeld.de

Dr. Veronica Egger
Cold Spring Harbor Laboratory
mail@veronicaegger.de

Dr. Michael Eisele
Keck Center, UC San Francisco
eisele@phy.ucsf.edu

Dr. Chris Eliasmith
University of Waterloo
eliasmith@uwaterloo.ca

Ms. Yasmine El-Shamayleh
New York University
yasmine@cns.nyu.edu

Dr. Monty Escabi
University of Connecticut
escabi@engr.uconn.edu

Dr. Adrienne Fairhall
University of Washington
fairhall@u.washington.edu

Dr. Michael Farries
Univ. of Washington
farries@u.washington.edu

Ms. Claudia Feierstein
Watson School of Biological Sciences, CSHL
feierste@cshl.edu

Mr. Gidon Felsen
University of California, Berkeley
gidonf@uclink.berkeley.edu

Ms. Ila Fiete
Harvard University/MIT
prasad@physics.harvard.edu

Dr. Alon Fishbach
Northwestern University
fishbach@northwestern.edu

Dr. Leonardo Franco
University of OXford
Leonardo.Franco@psy.ox.ac.uk

Dr. Loren Frank
University of California, San Francisco
loren@phy.ucsf.edu

Dr. Winrich Freiwald
Harvard Medical School
freiwald@brain.uni-bremen.de

Mr. Gene Fridman
UCLA
gfridman@ucla.edu

Dr. Rainer Friedrich
Max Planck Institute for Biochemistry
rainer.friedrich@mpimf-heidelberg.mp

Ms. Juan Gao
Princeton University
JGAO@PRINCETON.EDU

Prof. Daniel Gardner
Cornell University
dan@med.cornell.edu

Ms. Caroline Geisler
Brandeis University
geisler@brandeis.edu

Prof. Stuart Geman
Brown University
geman@dam.brown.edu

Mr. Dileep George
Redwood Neuroscience Institute, Stanford Univ
dgeorge@rni.org

Ms. Gail Gilboa
Technion - Israel Institute of Technology
gail@tx.technion.ac.il

Mr. Amir Globerson
Hebrew University, Jerusalem
gamir@cs.huji.ac.il

Dr. Joshua Gold
University of Pennsylvania
jigold@mail.med.upenn.edu

Dr. Mark Goldman
Wellesley College
mgoldma1@wellesley.edu

Mr. Arnulf Graf
Max Planck Institute for Biological Cybernetics
arnulf.graf@tuebingen.mpg.de

Dr. Bradley Greger
Caltech
greger@vis.caltech.edu

Mr. Jan Grewe
Bielefeld University
jan.grewe@uni-bielefeld.de

Dr. Kalanit Grill-Spector
Stanford University
kalanit@psych.stanford.edu

Dr. Jack Grinband
Columbia University
jg2269@columbia.edu

Mr. Aaron Gruber
Northwestern University
a-gruber1@northwestern.edu

Mr. Ming Gu
Caltech
gu@its.caltech.edu

Mr. Rudy Guyonneau
CerCo
rudy.guyonneau@cerco.ups-tlse.fr

Mr. Bilal Haider
Yale University School of Medicine
bilal.haider@yale.edu

Dr. Vincent Hakim
cnrs
hakim@lps.ens.fr

Mr. Farid Hamzei-Sichani
State University of New York Health Sciences Centr
fh81@alumni-mail.gs.columbia.edu

Dr. Kenneth Harris
Rutgers University
kdharris@andromeda.rutgers.edu

Ms. Andrea Hasenstaub
Yale University Medical School
snark@ugcs.caltech.edu

Dr. Jeff Hawkins
Redwood Neuroscience Institute
jhawkins@rni.org

Dr. Robert Hecht-Nielsen
Institute for Neural Computation, UCSD
r@fairisaac.com

Dr. Andreas Herz
Humboldt University Berlin
a.herz@biologie.hu-berlin.de

Dr. Neal Hessler
RIKEN Brain Science Institute
hessler@brain.riken.jp

Dr. Judith Hirsch
University of Southern California, Los Angeles
jhirsch@usc.edu

Dr. John Hopfield
Princeton University
jhopfield@princeton.edu

Dr. Greg Horwitz
The Salk Institute for Biological Studies
horwitz@salk.edu

Dr. Toshihiko Hosoya
RIKEN Brain Seience Institute
hosoya@brain.riken.jp

Dr. Patrik Hoyer
Helsinki University
patrik.hoyer@helsinki.fi

Mr. Tomas Hromadka
Watson School of Biological Sciences, CSHL
hromadka@cshl.edu

Ms. Anne Hsu
UC Berkeley
showen@uclink4.berkeley.edu

Dr. Yue Hu
Genetics Dept. Univ. of Penn
yueh@mail.med.upenn.edu

Mr. Kai Hynna
University of Pennsylvania
kmhynna@seas.upenn.edu

Mr. Elias Issa
Graduate Student/Johns Hopkins University
eissa@jhu.edu

Mr. Shimshon Jacobi
Weizmann Institute
shimshon.jacobi@weizmann.ac.il

Dr. Adam Jacobs
UCLA
ajk@ucla.edu

Mr. Santiago Jaramillo
National University of Ireland, Maynooth
sjara@ieee.org

Mr. Vivek Jayaraman
California Institute of Technology
vivek@caltech.edu

Dr. Dezhe Jin
Massachusetts Institute of Technology
djin@mit.edu

Dr. Daniel Johnston
Baylor College of Medicine
dan@mossy.bcm.tmc.edu

Prof. Jorge Jose
Northeastern University
jjv@neu.edu

Mr. Igor Kagan
Caltech
igor@vis.caltech.edu

Dr. Jagmeet Kanwal
Georgetown University
Kanwalj@georgetown.edu

Ms. Tara Keck
Boston University
tk@bu.edu

Dr. Richard Kempter
Humboldt University Berlin
r.kempter@biologie.hu-berlin.de

Dr. Adam Kepecs
Cold Spring Harbor Laboratory
kepecs@cshl.edu

Mr. Peyman Khorsand
Northeastern University
pkhorsan@lynx.neu.edu

Mr. Eyal Kimchi
Yale University
e_kimchi@yahoo.com

Dr. Francois Klam
University of Rochester
fklam@bcs.rochester.edu

Mr. Philipp Knuesel
Institute of Neuroinformatics
pknuesel@ini.phys.ethz.ch

Ms. Jennifer Ko
McMaster University
jenko@ieee.org

Dr. Kilian Koepsell
Redwood Neuroscience Institute
kkoepsell@rni.org

Dr. Adam Kohn
New York University
adamk@cns.nyu.edu

Mr. Hiroshi Kohsaka
University of Tokyo
kosaka@bio.phys.s.u-tokyo.ac.jp

Dr. Michael Kositsky
Northwestern University
kositsky@northwestern.edu

Dr. Alexei Koulakov
CSHL
akula@cshl.edu

Dr. Alexay Kozhevnikov
MIT
alexay@mit.edu

Dr. Bart Krekelberg
The Salk Institute
bart@salk.edu

Dr. Jutta Kretzberg
The Salk Institute
jutta@salk.edu

Dr. William Kristan
University of California, San Diego
wkristan@ucsd.edu

Mr. Mauktik Kulkarni
Johns Hopkins University
mkulkar1@jhmi.edu

Dr. Eduard Kuriscak
assistant professor
ekuri@lf1.cuni.cz

Mr. Stephen Larson
MIT Computer Science Artificial Intelligence Lab
stelar@alum.mit.edu

Dr. Peter Latham
UCLA
pel@ucla.edu

Mr. Brian Lau
New York University
blau@cns.nyu.edu

Dr. Mark Laubach
Yale University
mark.laubach@yale.edu

Mr. Jean Laurens
CNRS
jean.laurens@college-de-france.fr

Dr. Gilles Laurent
California Institute of Technology

Mr. Thomas Lauritzen
Redwood Neuroscience Institute
tzl@phy.ucsf.edu

Dr. Aurel Lazar
Columbia University
aurel@ee.columbia.edu

Dr. Daeyeol Lee
University of Rochester
dlee@cvs.rochester.edu

Mr. Nicholas Lesica
Harvard University
lesica@fas.harvard.edu

Dr. Alex Lewis
University College London
alex.lewis@ucl.ac.uk

Ms. Lei Li
SUNY Downstate Medical Center
lei.li@downstate.edu

Dr. Konstantin Likharev
SUNY, Stony Brook
klikharev@notes.cc.sunysb.edu

Dr. Susana Lima
Memorial Sloan Kettering Cancer Center
limas@mskcc.org

Mr. Shih-Chieh Lin
Duke University
sclin@neuro.duke.edu

Dr. Christiane Linster
Cornell University
CL243@cornell.edu

Dr. Robert Liu
UCSF
liu@phy.ucsf.edu

Dr. Jing Liu
HHMI and Stanford University
jingliu@stanford.edu

Dr. Yonatan Loewenstein
MIT
yonatanl@mit.edu

Dr. Michael Long
Massachusetts Institute of Technology
miclong@mit.edu

Dr. Wei Ji Ma
Caltech
ma@klab.caltech.edu

Dr. Jason MacLean
Columbia University
jm2107@columbia.edu

Dr. Jaime Mancilla
Brown University
Jaime_Mancilla@brown.edu

Mr. Valerio Mante
Smith-Kettlewell Eye Research Institute
valerio@ski.org

Mr. Sashi Marella
Purdue University
smarella@purdue.edu

Mr. Julien Mayor
Swiss Federal Institute of Technology
julien.mayor@epfl.ch

Dr. Andrei Medvedev
Georgetown University
am236@georgetown.edu

Dr. Markus Meister
Harvard University
meister@biosun.harvard.edu

Dr. Paul Miller
Brandeis University
pmiller@brandeis.edu

Dr. Samat Moldakarimov
University of Pittsburgh
sam47@pitt.edu

Dr. James Muller
Stanford University
jim@monkeybiz.stanford.edu

Mr. Shreesh Mysore
California Institute of Technology
shreesh@cds.caltech.edu

Mr. Nandakumar Narayanan
Yale University MD / PhD program
kumar.narayanan@yale.eud

Mr. Reza Nassiri
University of Connecticut
RNassiri@engr.uconn.edu

Ms. Maria Neimark
Harvard University
neimark@fas.harvard.edu

Dr. Ilya Nemenman
KITP, UCSB
nemenman@kitp.ucsb.edu

Mr. Ehren Newman
Princeton University
enewman@princeton.edu

Dr. William Newsome
Stanford University
bill@monkeybiz.stanford.edu

Dr. Quoc Thang Nguyen
UCSD
qnguyen@physics.ucsd.edu

Mr. Cristopher Niell
Stanford University
cris@stanford.edu

Dr. Sheila Nirenberg
University of California at Los Angeles
sheilan@ucla.edu

Dr. Makoto Nishiyama
New York University School of Medicine
nishim01@med.nyu.edu

Ms. Yael Niv
Hebrew University, Jerusalem
yaeln@freud.tau.ac.il

Dr. Kenneth Norman
Princeton University
knorman@princeton.edu

Dr. Akinao Nose
University of Tokyo
nose@phys.s.u-tokyo.ac.jp

Mr. Daniel O'Connor
Princeton University
doconnor@princeton.edu

Dr. Frank Ohl
Leibniz Institute for Neurobiology
frank.ohl@ifn-magdeburg.de

Dr. Itsaso Olasagasti
Wellesley College
iolasaga@wellesley.edu

Dr. Bence Olveczky
Harvard University
olveczky@fas.harvard.edu

Dr. Donald O'Malley
Northeastern University
d.omalley@neu.edu

Dr. Gonzalo Otazu

Dr. Thomas Palmeri
Vanderbilt University
thomas.j.palmeri@vanderbilt.edu

Prof. Lucas Parra
City College of NY
parra@ccny.cuny.edu

Dr. Klaus Pawelzik
University of Bremen
pawelzik@neuro.uni-bremen.de

Dr. Barak Pearlmutter
National University of Ireland Maynooth
barak@cs.may.ie

Dr. Amandine Penel
Cold Spring Harbor Laboratory
penel@cshl.edu

Dr. Mark Penney
Novartis Pharmaceuticals AG
mark.penney@pharma.novartis.com

Dr. Bijan Pesaran
California Institute of Technology
bijan@caltech.edu

Mr. Jonathan Pillow
New York University
pillow@cns.nyu.edu

Mr. Xaq Pitkow
Harvard University
pitkow@fas.harvard.edu

Dr. Alexandre Pouget
University of Rochester
alex@bcs.rochester.edu

Dr. Nicholas Priebe
Northwestern University
nico@northwestern.edu

Dr. Astrid Prinz
Brandeis University
prinz@brandeis.edu

Dr. Jason Puchalla
Princeton University
puchalla@princeton.edu

Mr. Michael Quirk
Cold Spring Harbor Laboratory
quirk@cshl.org

Mr. Sridhar Raghavachari
Brandeis University
sraghava@brandeis.edu

Ms. Kanaka Rajan
Volen Center, Brandeis University
kanaka@brandeis.edu

Dr. Roger Ratcliff
Ohio State University
ratcliff.22@osu.edu

Dr. Heather Read
University of Connecticut
heather.read@phy.ucsf.edu

Dr. Jenny Read
National Institutes of Health
jcr@lsr.nei.nih.gov

Prof. Hermann Riecke
Northwestern University
h-riecke@northwestern.edu

Dr. Dima Rinberg
Monell Chemical Senses Center
drinberg@monell.org

Dr. Dario Ringach
University of California, Los Angeles
dario@ucla.edu

Dr. Jason Ritt
MIT
jritt@mit.edu

Dr. David Robbe
robbe@axon.rutgers.edu

Dr. Lizabeth Romanski
University of Rochester Medical Center
liz_romanski@urmc.rochester.edu

Mr. Alan Rorie
Stanford
rorie@stanford.edu

Dr. Alexander Roxin
University of Paris V
alexander.roxin@univ-paris5.fr

Dr. Simon Rumpel
Cold Spring Harbor Laboratory
rumpel@cshl.org

Mr. Clifton Rumsey
Brandeis University
rumsey@brandeis.edu

Ms. Nicole Rust
New York University
rust@cns.nyu.edu

Dr. Phillip Sabes
UCSF
sabes@phy.ucsf.edu

Mr. Srivatsun Sadagopan
Johns Hopkins University
vatsun@jhu.edu

Prof. Paul Sajda
Columbia University
ps629@columbia.edu

Dr. Emilio Salinas
Wake Forest University School of Medicine
esalinas@wfubmc.edu

Dr. Massimo Scanziani
University of California San Diego
massimo@biomail.ucsd.edu

Mr. Andreas Schaefer
MPI f. med. Forschung
schaefer@mpimf-heidelberg.mpg.de

Dr. Katrin Schenk
University of California, San Francisco
schenk@phy.ucsf.edu

Mr. Erik Schlicht
University of Minnesota
schl0360@umn.edu

Ms. Anita Schmid
ETH Zuerich
anita@ini.phys.ethz.ch

Dr. Elad Schneidman
Princeton University
elads@princeton.edu

Dr. Odelia Schwartz
Salk Institute
odelia@salk.edu

Dr. David Schwartz
Duke University
schwartz@neuro.duke.edu

Dr. Terrence Sejnowski
The Salk Institute for Biological Studies
terry@salk.edu

Dr. Peggy Series
University of Rochester
pseries@bcs.rochester.edu

Dr. Michael Shadlen
University of Washington
shadlen@u.washington.edu

Dr. Tatyana Sharpee
University of California
sharpee@phy.ucsf.edu

Mr. Jonathan Shaw
University of Rochester
jshaw@cs.rochester.edu

Mr. Jonathon Shlens
Salk Institute / UC San Diego
shlens@salk.edu

Dr. Oren Shriki
The Weizmann Institute of Science
oren.shriki@weizmann.ac.il

Dr. Angus Silver
University College London
a.silver@ucl.ac.uk

Mr. Ray Singh
University of Waterloo
ray440@email.com

Mr. Lucas Sjulson
Sloan-Kettering Institute
sjl2001@med.cornell.edu

Mr. Evan Smith
Carnegie Mellon University
evan@cnbc.cmu.edu

Dr. Lawrence Snyder
Washington University, School of Medicine
larry@eye-hand.wustl.edu

Dr. Sara Solla
Northwestern University
solla@nwu.edu

Dr. Friedrich Sommer
Redwood Neuroscience Institute
fsommer@rni.org

Dr. Sen Song
cold spring harbor laboratory
songs@cshl.edu

Dr. William Spain
University of Washington
spain@u.washington.edu

Dr. John Spiro
Nature Publishing Group
j.spiro@natureny.com

Dr. Garrett Stanley
Harvard University
gstanley@deas.harvard.edu

Dr. Armen Stepanyants
Cold Spring Harbor Lab.
stepanya@cshl.edu

Dr. Greg Stephens
Los Alamos National Laboratory
gstephen@physics.umd.edu

Dr. Alan Stocker
Howard Hughes Medical Institute/NYU
alan.stocker@nyu.edu

Mr. Tom Sullivan
UC San Diego
tom@sullivan.to

Dr. Linus Sun
Yale SOM/MIT Picower Ctr. For Learning and Memory
linus@cortical.mit.edu

Mr. Christian Swinehart
Brandeis University (Volen Center)
cds@brandeis.edu

Mr. Michael Tadross
Johns Hopkins University
mtadross@jhmi.edu

Mr. Lung-Hao Tai
SUNY, Stony Brook

Mr. Andrew Tan
University of California San Francisco
atyy@alum.mit.edu

Dr. David Tank
Princeton University
dwtank@molbiol.princeton.edu

Dr. Ching-Ling Teng
UCSF
ling@phy.ucsf.edu

Dr. Frederic Theunissen
UC Berkeley
fet@socrates.berkeley.edu

Mr. Eric Thomson
UCSD
ethomson@ucsd.edu

Dr. Paul Tiesinga
Univesity of North Carolina
tiesinga@physics.unc.edu

Dr. Emanuel Todorov
UCSD
todorov@cogsci.ucsd.edu

Mr. Kazunobu Togashi
New York University
togask01@med.nyu.edu

Mr. Jonathan Toups
University of North Carolina at Chapel Hill
toups@physics.unc.edu

Mr. Jonathan Touryan
University of California, Berkeley
touryan@uclink4.berkeley.edu

Dr. Stefan Treue
German Primate Center
treue@gwdg.de

Dr. Todd Troyer
University of Maryland, College Park
ttroyer@glue.umd.edu

Ms. Doris Tsao
Harvard Medical School
doris@nmr.mgh.harvard.edu

Dr. Naoshige Uchida
Cold Spring Harbor Laboratory

Dr. Nathan Urban
Carnegie Mellon University
nurban@cmu.edu

Ms. Valerie Uzzell
Salk Institute/ UC San Diego
uzzell@salk.edu

Mr. Stephen Van Hooser
Brandeis University
vanhoosr@brandeis.edu

Mr. Tim Vogels
Brandeis University
vogels@brandeis.edu

Dr. Robert Waltereit
University of Tuebingen, Center for Neurology
robert.waltereit@uni-tuebingen.de

Dr. Xiaoqin Wang
Johns Hopkins University School of Medicine
xwang@bme.jhu.edu

Mr. Brendon Watson
Columbia University
bow4@columbia.edu

Dr. Michael Wehr
Cold Spring Harbor Laboratory
wehr@cshl.edu

Dr. Quan Wen
Cold Spring Harbour
wen@cshl.edu

Dr. Ralf Wessel
Washington University
rw@wuphys.wustl.edu

Mr. Robert White
Washington University School of Medicine
rob@eye-hand.wustl.edu

Mr. Sebastian Wills
University of Cambridge
saw27@mrao.cam.ac.uk

Dr. Matt Wilson
Massachusetts Institute of Technology
wilson@ai.mit.edu

Prof. Daniel Wolpert
University College London
wolpert@ion.ucl.ac.uk

Dr. Thilo Womelsdorf
German Primate Center
twomelsdorf@dpz.gwdg.de

Dr. Sarah Woolley
University of California, Berkeley
swoolley@socrates.berkeley.edu

Ms. Melanie Wyder
Wake Forest University School of Medicine
mwyder@wfubmc.edu

Mr. Shaohua Xu
SUNY Downstate Medical Center
xus@downstate.edu

Dr. Shin Yanagihara
RIKEN
syanagihara@brain.riken.jp

Dr. Zhiyong Yang
Duke University Medical Center
zhyyang@duke.edu

Mr. Byron Yu
Stanford University
byronyu@stanford.edu

Mr. Hsin-Hao Yu
University of California San Diego
hhyu@cogsci.ucsd.edu

Dr. Tony Zador
Cold Spring Harbor Laboratory
zador@cshl.edu

Dr. Richard Zemel
University of Toronto
zemel@cs.toronto.edu

Dr. Ketchen Zhang
Johns Hopkins University
kechen@bme.jhu.edu

Mr. Valentin Zhigulin
California Institute of Technology
zhigulin@caltech.edu

VISITOR INFORMATION

EMERGENCY	CSHL	BANBURY
Fire	**(9) 742-3300**	**(9) 692-4747**
Ambulance	**(9) 742-3300**	**(9) 692-4747**
Poison	**(9) 542-2323**	**(9) 542-2323**
Police	**(9) 911**	**(9) 549-8800**
Safety-Security	**Extension 8870**	

Emergency Room **Huntington Hospital** 270 Park Avenue, Huntington	**631-351-2300** **(1037)**
Dentist Dental Referral Service	**800-511-8663**
Doctor mediCenter 365 W. Jericho Tpke., Huntington	**631-423-5400** (**1034**)
Drugs - 24 hours, 7 days Eckerd 391 W. Main Street, Huntington	**631-549-9400** (**1039**)

Free Speed Dial
Dial the four numbers (****) from any **tan house phone** to place a free call.

GENERAL INFORMATION

Books, Gifts, Snacks, Clothing, Newspapers
BOOKSTORE 367-8837
Located in Grace Auditorium, lower level.

Photocopiers, Journals, Periodicals, Books, Newspapers
MAIN LIBRARY 8:30 a.m.–10:00 p.m. daily
10:00 a.m. – 6:00 p.m. Saturday
Helpful tips - Obtain PIN from Meetings & Courses Office to enter Library after hours. See Library staff for photocopier code.

Computers, E-mail, World Wide Web
Grace Auditorium
Upper level: E-mail only; Lower level: Word processing and printing.
To access your E-mail, you must know the name of your home server.

Dining, Bar

Blackford Hall

Breakfast 7:30–9:00, Lunch 11:30–1:30, Dinner 5:30–7:00
Bar 5:00 p.m. until late

Helpful tip - If there is a line at the upper dining area, try the lower dining room

Messages, Mail, Faxes

Message Board, Grace, lower level

Swimming, Tennis, Jogging, Hiking

June–Sept. Lifeguard on duty at the beach. 12:00 noon–6:00 p.m. Two tennis courts open daily.

Pool Table, Ping Pong, Television

Blackford Hall Rec Room - downstairs

Russell Fitness Center

Dolan Hall, west wing, lower level

Concierge

On duty daily at Meetings & Courses Office.

After hours – From tan house phones, dial x2554 (Call Center) for immediate assistance.

Pay Phones, House Phones

Grace, lower level; Cabin Complex; Blackford Hall; Dolan Hall, foyer

Helpful tip - Numbers preceded by a (9) may be dialed free from the tan house phones.

	1-800 access #'s:
AT&T	**800-321-0288**
MCI	**800-674-7000**

Local Interest

Fish Hatchery	631-692-6768
Sagamore Hill	516-922-4447
Whaling Museum	631-367-3418
Heckscher Museum	631-351-3250
CSHL DNA Learning Center	x 5170

New York City

Helpful tip -

Take Syosset Taxi to Syosset Train Station
($7.00 for 1, $8.00 for 2. 15 minute ride),
then catch Long Island Railroad to Penn Station (33rd Street & 7th Avenue). Train ride about one hour.

TRANSPORTATION

Limo, Taxi

Classic Limousine	631-567-5100 (**1033**)
Syosset Limousine	516-364-9681 (**1031**)
To head west of CSHL - Syosset train station	
Syosset Taxi	516-921-2141 (**1030**)
To head east of CSHL - Huntington Village	
AAA Taxi	631-423-2800 (**1048**)
Orange & White Taxi	631-271-3600 (**1032**)
Executive Limo	631-696-8000 (**1047**)

Trains

Long Island Rail Road	822-LIRR

Schedules available from the Meetings & Courses Office.

Amtrak	800-872-7245
MetroNorth	800-638-7646
New Jersey Transit	201-762-5100

Ferries

Bridgeport / Port Jefferson	631-473-0286 **(1036)**
Orient Point/ New London	631-323-2525 **(1038)**

Car Rentals

Avis	631-271-9300
Enterprise	631-424-8300
Hertz	631-427-6106
Premier	631-385-7777

Airlines

American	800-433-7300
America West	800-247-5692
British Airways	800-247-9297
Continental	800-525-0280
Delta	800-221-1212
Japan Airlines	800-525-3663
KLM	800-374-7747
Lufthansa	800-645-3880
Northwest	800-225-2525
TWA	800-221-2000
United	800-241-6522
US Airways	800-428-4322