Sloan-Swartz Centers for Theoretical Neurobiology
Annual Summer Meeting 2006

Book of Abstracts

Self-motion perception: Multisensory integration in extrastriate visual cortex

Dora E. Angelaki (Washington University School of Medicine)

Optic flow patterns generated during self-motion provide a strong cue for the perception of our own movement through space (heading). However, accurate judgments of heading often require integration of visual and nonvisual cues, including vestibular, kinesthetic, and eye movement signals. This sensory integration is complicated by the fact that signals from different modalities may originate in different coordinate frames (e.g., eye-centered or head-centered). To investigate the neural basis of self-motion perception, we record from neurons in area MSTd of macaque monkeys during optic flow (*Visual* condition), real motion (*Vestibular* condition) and congruent combinations of the two (*Combined* condition) using a novel virtual reality system that can move animals along arbitrary paths through a 3D virtual environment.

To examine how visual and vestibular signals in MSTd contribute to heading perception, we train animals to perform a fine self-motion discrimination task. Heading directions are varied in small steps around straight forward, and monkeys report whether their heading is to the right or left of straight ahead. Psychophysical thresholds averaged 1-3º and, although the most sensitive MSTd neurons had thresholds close to behavior, the average neuron was much less sensitive than the monkey under both single-cue conditions. In the Combined condition, psychophysical thresholds were significantly lower than in the single-cue conditions and very close to predictions of optimal cue integration theories. According to whether the visual and vestibular heading preferences were well matched or nearly opposite, MSTd neurons could be divided into two distinct groups: ‘congruent’ and ‘opposite’ cells. We found that neuronal thresholds in the combined condition were strongly dependent on congruency of heading preferences, such that 'congruent' neurons showed substantially lower thresholds under cue combination, whereas ‘opposite’ cells showed elevated thresholds. The effect of cue combination on ‘congruent’ cells was very similar to that seen behaviorally, suggesting that this subset of MSTd neurons may contribute to sensory integration for heading perception.

To assess functional coupling between MSTd responses and perceptual decisions, we compute 'choice probabilities' (CPs), which characterize the trial-to-trial covariation between neural responses and the animal's choices. CPs were strongest under the Vestibular condition, in which ~30% of MSTd neurons showed significant positive effects and the mean CP (0.56) was significantly greater than chance (0.5). In contrast, CPs in the Visual condition depended strongly on the congruency of visual and vestibular heading preferences. ‘Congruent' neurons were positively correlated with choices in the Visual condition (mean CP=0.59), whereas ‘opposite’ neurons tended to be negatively correlated (mean CP=0.45) with heading percepts. These results provide the first evidence that links cortical vestibular signals with heading perception. In addition, our Visual CP data suggest strongly that the visual responses of MSTd neurons are read out differently depending on the congruency of visual and vestibular tuning preferences. This result supports the notion that ‘congruent’ cells may have a privileged role in cue combination, and highlights the importance of studying multi-sensory integration for understanding spatial perception.

Gustatory processing changes within sessions, dependent on attention

Donald B. Katz (Brandeis University)

Gustatory cortical networks reliably transition through taste-specific states

Lauren M. Jones (Brandeis University)

Distinct states of firing patterns in the primary visual cortex of awake ferrets

József Fiser (Brandeis University)

In a series of three talks, we will present a coherent view of brains functioning as a dynamic network processing complex multi-dimensional input. The specifics of this approach include a) large scale multi-electrode recordings as the appropriate tool for data collection, b) performance of behavioral tasks as the appropriate preparation, c) data analysis methods based on the principle that processing in such networks depends on the prior state of the networks and involves complex coherent patterns of activity in neural ensembles, and d) the use of multiple modalities and species for validation of the generality of results. In the first talk, we show that the processing of tastes in gustatory cortex varies across a recording session dependent on mass action in cortical networks, such that as the rat switches from being attentive (desynchronized EEG) to being inattentive (oscillatory EEG) the nature of taste coding changes; in inattentive rats, neural and behavioral responses to a battery of experimenter-administered tastes are coded more simply in relation to palatability—pleasing tastes are more distinct from noxious tastes, and noxious tastes are more similar to each other. In the second talk, we show that taste responses vary from trial to trial, dependent on coherent ensemble processes. Ensembles of taste cortex neurons go through a series of coherent states—defined as a period of time in which each neuron has a particular firing rate—with minimal switching time between states; a particular series of 2-3 states is specific for a particular taste stimulus, but the time spent in each state varies widely across trials. A given state series is a more effective predictor of taste identity than PSTH-based methods. Finally, in the third talk we switch to a new modality (vision), a more unconstrained task (passive viewing), and a different species (ferret). While the visual system that is thought to work very differently from taste, and the passive viewing task does not require stimulus processing and response (ferrets sat in complete darkness or watched either a natural scene or a random noise movie), cortical functioning still exhibits similar hallmarks of dynamic state-dependent functioning. We suggest that rather than attempting to reproduce simple output or trial-averaged responses of single neurons, analyses that capture the underlying structure of the dynamic behavior may provide useful information for theoretical/computational attempts to unravel the function of cortical networks.

Analysis of spontaneous and sensory-driven activity in ferret V1

J. Zhao[1,], G. Szirtes[1], M. Eisele[1], J. Fiser[2,3], C. Chiu[3], M. Weliky[3], and K.D. Miller[1].
[1]Columbia, [2]Brandeis, [3]Univ. of Rochester

We analyze multiunit recordings from linear arrays of 16 electrodes spanning 3 or 9 mm in awake ferret V1, as in Fiser et al. Nature 431:573 (2004). Recordings were made at ages ranging from 29 to 168 days postnatal. Fiser et al. 2004 found that activity from P30 to P90 was dominated by similar activity patterns whether in dark or when stimulated by white noise or a natural movie. They showed that temporal correlations on a single electrode were long at early ages but became progressively shorter, while spatial correlations at a single time were short-ranged at early ages but became long-range at later ages. Correspondingly, activity patterns became dominated by bursts spanning all electrodes.

We find the principal components of simultaneous activity across the electrodes. At later ages, most of the variance is in the first component, which is uniform across electrodes (each electrode deviates by the same number of standard deviations from its mean activity). This component's autocorrelation shows some tendency to oscillate, with a bump of power in the range 10-17Hz. This temporal structure is quite similar for dark and movie stimuli. However, for noise stimuli, particularly at ages >= P120, very long-lasting oscillatory autocorrelation at 11-12 Hz is seen. This may represent alpha activity, which has been argued to represent an "idle" or "disengaged" state, suggesting the awake animal may disengage from the noise stimulus. More generally, this dominant first component seems likely to represent a global state rather than specific visual input. Subtracting off the principal component, the remaining activity shows correlations that are much more localized in space and time. Power in the remaining activity seems to fall off as a power of spatial frequency, suggesting that it might have no characteristic spatial scale. Supported by RO1-EY13595 (KDM) and RO1-EY 012494 (MW) from the NEI.

Models of Background Activity in V1

Brendan Murphy (UCSF)

Responses in primary visual cortex are determined both by the visual stimulus and by ongoing cortical spontaneous activity. Experiments using voltage sensitive dyes have found that spontaneous activity patterns show significant spatial correlations with visually-evoked orientation maps that decay over times ~80 msec. Another important aspect of V1 spontaneous activity is that neurons show large ongoing voltage fluctuations and high spike time variability. Model networks with strong but balanced excitation and inhibition and sparse connectivity have been shown to have these two characteristics. These networks have most commonly been studied with random connectivity. We show that a balanced network of integrate-and-fire neurons with V1-like synaptic connectivity can exhibit spontaneous activity patterns similar to those observed experimentally. We also study a simple linear rate model to better understand the underlying mechanisms by which these patterns can be generated.

Fine Structure of Human Sleep Unraveled

Philip Low (Salk)

I will show that, contrary to current scientific belief, human waking and sleep stages can unambiguously be distinguished using a single channel of EEG. Moreover, I will show that Slow Wave Sleep is not Slow, that Rapid Eye Movement (REM) sleep is not "awake-like", "paradoxical " or "desynchronized." In addition, I will present preliminary analysis revealing the presence of a novel human sleep state.

Discriminating temporal patterns: spiking neurons and 'ideal observers'

Haim Sompolinsky (The Hebrew University)

Animals need to detect or discriminate stimuli which are rich in temporal structure. Our recently developed tempotron model demonstrates that an integrate and fire neuron can learn to discriminate between stimuli which are encoded in spatio-temporal spike patterns. This raises the interesting question of how well simple neuronal circuits perform in temporal processing relative to an 'ideal observer'.

A standard ideal observer model which is unlimited in its memory capacity may provide unrealistic bounds of performance for temporally restricted neural systems. I will describe a Dynamic Ideal Observer model which incorporates temporal locality and other dynamical-system features into the ideal observer decision process. This model provides useful bounds for the performance of simple neural systems in discriminating temporally extended stimuli.

Creating Maps in the Cortex

Michael P. Stryker (UCSF)

What are the effector mechanisms that organize maps in the brain? In the virtual absence of evidence, it has been a truism that gradients of molecular cues must set up a coarse map, which might then be refined under the influence of neural activity. We now have evidence for something like this view.

We studied the Ephrin-As and their receptors, EphAs, as potential molecular cues because they are expressed in gradients in the developing cortex (1). We mapped the visual cortex (V1) in triple knockout mice deficient for ephrin-A2, -A3, and -A5 functionally, using intrinsic signal optical imaging and microelectrode recording, and structurally, by anatomical tracing of thalamocortical projections. V1 was shifted medially, rotated, and compressed, and its internal organization was degraded. Ectopic expression of any one of these ephrins lateral to the visual cortex shifted the map of V1 medially, and expression within V1 disrupted its internal organization. These findings indicate that interactions between gradients of EphA/ephrin-A in the cortex guide map formation, but that factors other than redundant ephrin-As are responsible for the remnant map. Together with earlier work on the retinogeniculate map, the current findings show that the same molecular interactions may operate at successive stages of the visual pathway to organize maps.

We found that structured neural activity in the eyes was necessary for the formation of precise maps in V1 using a mouse in which the waves of spontaneous activity in the retina during the first postnatal week were disrupted genetically (2). This anatomical mapping defect was fully expressed by postnatal day 8 and had functional consequences, as revealed by optical imaging and microelectrode recording in adults. Pharmacological disruption of the retinal waves during the first week phenocopied the mapping defect, confirming both the site and the timing of the disruption in neural activity that was responsible for the mapping defect. Analysis showed that the geniculocortical miswiring was not a trivial or necessary consequence of the retinogeniculate defect. These findings demonstrate that disrupting early spontaneous activity in the eye alters thalamic connections to the cortex.

Until very recently, studies of activity-dependent plasticity in the developing visual cortex have focused on the critical period of susceptibility to the effect of monocular occlusion beginning near the end of the third week of life. This time corresponds to a phase in cortical development very different from the critical period of the primary somatosensory ‘‘barrel’’ cortex in the first postnatal week. Our new experiments provide evidence that retinal activity during the first week of life is essential for organizing precise connections to the visual cortex. This form of activity-dependent plasticity, which aids in the establishment of a precise topographic order, appears likely to correspond to the critical period for barrel cortex. It takes place at the same point in the life history of the cells and serves a similar function. In both systems, a defect in the sensory periphery does not abolish thalamocortical topography, but instead prevents it from refining to its normal precision. The role now revealed for patterns of spontaneous neural activity in the establishment of maps during the initial ingrowth and elaboration of connections to the cortex may be general.

What is the interaction between neural activity and the molecular cues that guide map formation? Are these two signals the whole story? And what aspects of the structure of normal activity are important in map formation, and how are they decoded? Experiments in progress address these issues.

With thanks to my co-authors listed below, to the NIH for experimental support, and to the UCSF Sloan Swartz Center for the intellectual environment and for supporting the creation of tools essential for this work.

(1) Cang*, J.C., Kaneko*, M., Yamada, J., Woods, G., Stryker, M.P., and Feldheim, D.A. (2005) Ephrin-As Guide the Formation of Functional Maps in the Visual Cortex (* co-first authors). Neuron 48: 577-589.

(2) Cang, J.C., Renteria, R.C., Kaneko, M.., Liu, X., Copenhagen, D.R. and Stryker, M.P. (2005) Development of precise maps in visual cortex requires patterned spontaneous activity in the retina. Neuron 48: 797-809

Response Variability in LGN

M.Eisele
Co-authors: C.Weng, J. Z. Jin, C. I. Yeh, J.-M. Alonso, K.D. Miller

Ken & Michael: Columbia University
Other authors: State University of New York

Neural response properties vary not only between cell classes, but also, more gradually, within each class. Some of these variations presumably help in processing the diversity of sensory stimuli, while others are just biological noise. We argue that response variations need to reach a certain size before they become useful for signal processing and that this can be used to distinguish them from smaller variations that are just noise. We demonstrate how to make this distinction in the lateral geniculate nucleus (LGN) of the cat. Receptive fields of X- and Y-cells were mapped in space and time using reverse correlation with dense white noise stimuli of high or low contrast or with sparse noise. These receptive fields were then analyzed with a modified version of principal component analysis (PCA). Apart from the obvious diversity of preferred spatial positions, the diversity response time-courses also seems large enough to play a functional role, while other forms of diversity are much weaker and probably just noise. This result agrees well with what we know about how LGN outputs are processed in primary visual cortex. By comparing the diversity in LGN to the responses of simple cells, which were recorded under similar conditions, we find that the actual diversity in LGN is near the optimal diversity for cortical signal-to-noise. These results show how gradual, apparently random variations can play a role in optimal encoding.

How MT cells compute pattern motion

Tony Movshon (NYU)

Neurons in MT (V5) are selective for the direction of visual motion. In addition, many MT neurons are selective for the motion of complex patterns independent of their component orientations, a behavior not seen in earlier visual areas. Using a novel mutlicomponent "hyperplaid" stimulus, we show that the responses of MT cells can be captured by a linear model that operates not on the visual stimulus, but on the afferent responses of a population of nonlinear V1 complex cells. We fit this cascade model to responses of individual MT neurons and show that it robustly predicts the separately-measured responses of MT cells to gratings and plaids. The model captures the full range of pattern motion selectivity found in MT. Cells that signal pattern motion are distinguished by having broadly tuned excitatory input from V1, strong motion opponent suppression, and a tuned normalization at the V1 stage that may correspond to surround suppression.

Sparse Bayesian learning of brain activity patterns

Rey R. Ramirez (UCSD)

Although the use of anatomical location priors to spatially constrain the current density estimates to grey matter is highly beneficial, use of other priors including fixed source pdfs, and exact model order assumptions may lead to inaccurate solutions. The Bayesian evidence framework represents a clear departure from the classical objective of simply introducing a subjective prior to the inference process (i.e., MAP estimation). Alternative hypothesis corresponding to different priors are invoked for model comparison, and are objectively ranked by evaluating their Bayesian evidence, the probability of the data given the model (MacKay, 1992). Bayes rule automatically embodies Occam's razor to prefer simpler models that predict the data more strongly than overcomplex models. For regression and classification, this sparse Bayesian learning (SBL) methodology takes the form of the relevance vector machine (RVM) (Tipping, 2001), a Bayesian alternative to the support vector machine (SVM). Recently, SBL has been expanded to solve underdetermined inverse problems under sparsity constraints (Wipf and Rao, 2004). The

SBL algorithm assumes a Gaussian likelihood model and a source distribution parameterized by a vector of hyperparameters that are learned from the data by integrating out the unknown source parameters (e.g., the current density) and performing evidence maximization (i.e., type-II maximum likelihood) using an Expectation Maximization (EM) algorithm or a fixed-point gradient update rule that minimizes the same cost, but more rapidly (Ramirez, 2005). Importantly, SBL finds the maximally sparse solution in cases where other algorithms, such as FOCUSS (Rao et al., 2003), converge to local minima. However, maximal sparsity is not always desired since sources have variable spatial extents.

Here, a new framework for neuroelectromagnetic source imaging is developed in which the fundamental data representation atoms are concatenated MEG/EEG gain vectors generated by locally distributed multi-resolution current density functions on cortical patches, instead of single dipoles. An overcomplete lead-field dictionary is constructed using these distributed neural bases at multiple scales. Cortical geodesic distances and orientation constraints are exploited to discriminate between gyri or sulci that are near in Euclidian but not geodesic space. A sparse solution to the multiscale-transformed inverse problem is computed using SBL. Although it is maximally sparse in terms of the transformed system, it represents a distributed current density estimate once it is back-transformed with the a priori neural basis matrix. Computer simulations demonstrate that we can reconstruct the spatial current density distribution of many simultaneously active sources in the presence of noise and modeling error. We analyze MEG data with SBL and infomax ICA to characterize the large-scale neuronal networks involved in cued visuospatial covert attention and target detection.

From disgust to compassion: distributed EEG dynamics of emotion

Julie Onton (UCSD)

Emotions are a fundamental part of being human, yet recording brain activity during genuine emotional states is difficult to achieve in a laboratory setting. Many studies have attempted to recreate natural emotional reactions by showing subjects pictures of intensely emotional or disturbing facial expressions or scenes. While such pictures can trigger a cascade of subjective and neural events leading to sustained emotional experiences, the responses of viewers depend in large part on the attitudes they bring to the viewing. As well, such methods ignore the broad natural range of emotion and feeling states that pervade our everyday experience. In particular, a wide variety of positive feelings--contentment, happiness, love, compassion, awe, etc.--experienced in milder forms during daily life are not normally addressed in such studies. This study attempted to discover the brain dynamic correlates of imagined and embodied emotional feeling. Subjects were asked, via a voice recording, to recall and/or imagine a series of scenarios in which they had felt or would feel a series of suggested emotions, in each case allowing the imagery and somatic feeling sensations to become as vivid as possible. A series of fifteen suggested positive and negative emotions were separated by brief relaxation periods. During these sessions, we recorded 256 channels of EEG data from the scalp, neck and face. After decomposing the EEG data using Independent Component Analysis (ICA) into maximally independent time courses and associated spatial maps, we used a novel approach to identify independent power spectral modes active during the experiment. From this analysis, 8 major spectral patterns were identified: 3 classes of alpha (8-12 Hz), 2 classes of beta (15-30 Hz), and 3 classes of gamma (>30 Hz) modulations. Different emotional experiences were differentially associated with these major spectral modulations and, while results differed considerably across subjects, some consistent associations between brain patterns and emotion were identified. The results suggest that mood/emotion states bias EEG dynamics towards identifiable patterns of activity that could possibly be used in the future for mood, response and/or therapeutic monitoring.

Fluctuations in neuron pairs and populations - modelling and mechanisms

Eric Shea-Brown (NYU)

It has been proposed by Zohary and collaborators that a primary source of fluctuations in population-averaged firing rates is spike-to-spike correlation amongst the constituent neurons, and these and future authors have highlighted a variety of implications for the neural code. However, a general mechanistic understanding is still lacking of how these correlations develop in neural populations, and, critically, how their strength depends on the population operating range (that is, the rate and regularity of its inputs or its spiking). Our new results toward such an understanding include (i) the surprising finding that these correlations increase dramatically with firing rate over ranges of up to approximately 20-30 Hz. in spiking neuron models (independently of of the accompanying regularity -- or lack thereof -- in single-cell spiking), (ii) an intuitive explanation of this fact based on the linearity of input-output relationships of individual cells, (iii) and an analysis of implications for encoding and transmission of static stimuli via population firing rates. This is joint work with Jaime de la Rocha, Brent Doiron, and Kresimir Josic.

Analysis of intrinsic properties in a multidimensional parameter space

Adam L. Taylor and Eve Marder (Brandeis)

What is the mapping between the maximal conductances of a neuron and its functional properties? We examined this question using a multicompartmental model of a neuron in the stomatogastric ganglion of the crab Cancer borealis. Neurons in this system can exhibit a variety of different behaviors, including intrinsic bursting. Which behavior a particular neuron exhibits is determined by the particular mix of maximal conductances in that neuron and by the presence or absence of various neuromodulators. We examined how varying the different maximal conductances changed such emergent quantities as the burst period, the burst duty cycle, and the number of spikes per burst. Understanding how the maximal conductances relate to behaviorally important quantities provides a basis for understanding what conductances may be coregulated by homeostatic mechanisms and by neuromodulation.

Dynamic clamp half-center oscillator construction as an assay of variability in intrinsic properties

Rachel Grashow (Brandeis)

I use the dynamic clamp to construct two-cell reciprocally inhibitory circuits out of gastric mill (GM) neurons from the crab *Cancer borealis.* Using the dynamic clamp I control both the synaptic connectivity strength (g_{syn}) between the two neurons, as well as the amount of hyperpolarization-activated inward conductance (g_h) in each cell. I have systematically varied the g_h and g_{syn} for each biological circuit and searched the parameter space for half-center oscillator behavior. I have found that within the parameter space there exist regimes where cells are silent, bursting, tonically spiking or show an irregular network behavior. In parameter regimes where the cells form half-center oscillators, our findings confirm those of Sharp et al. (1996) in terms of the effects of individually varying g_h or g_{syn} on the burst frequency of the network. Finally, I have found that disparate amounts of g_h and g_{syn} can produce similar network output within one experiment. Across experiments the half-center activity regimes within the parameter vary. My hope is to use this method as a functional assay of the variability of the intrinsic properties of neurons across animals.

Optimal Information Storage in Noisy Synapses

Dmitri 'Mitya' Chklovskii (CSHL)

Experimental investigations have revealed that synapses possess interesting and, in some cases, unexpected properties. We propose a theoretical framework that accounts for three of these properties: typical central synapses are noisy; the distribution of synaptic weights among central synapses is wide; and synaptic connectivity between neurons is sparse. We also comment on the possibility that synaptic weights may vary in discrete steps. Our approach is based on maximizing information storage capacity of neural tissue under resource constraints. Based on previous experimental and theoretical work, we use volume as a limited resource and utilize the empirical relationship between volume and synaptic weight. Solutions of our constrained optimization problems are not only consistent with existing experimental measurements but also make non-trivial predictions.

Optimal Neural Decision Boundaries for Maximal Information Transmission

Tatyana Sharpee[1] and William Bialek[2]

[1]UCSF, [2] Princeton

We consider here how to optimally encode multidimensional signals using a spike or its absence as two possible outcomes in order to maximize the mutual information transmitted about those signals. Our goal is to understand the optimal shape of contours in the input space that would separate stimuli leading to a spike from stimuli that do not elicit spikes, and how this optimal shape changes with the average rate of firing a spike. Our assumption is that most of the variability in neural response will be caused by stimuli near the spiking decision boundary, and that the width of the uncertainty region is much smaller than the characteristic length scales over which the probability distribution of inputs changes. Stimuli that are far away from the decision boundary will elicit an almost certain response.

For a single neuron, the problem of maximizing information transmission given an average spike probability is equivalent to minimizing the noise entropy. The noise entropy is proportional to an integral over the probability distribution of inputs along the decision contour. Solving the variation problem with respect to the contour's shape, we derive a general equation for the decision boundary valid for arbitrary probability distributions that relates its curvature, scaled by the noise level, to the probability distribution of inputs. Solving this equation for Gaussian inputs, we show that the optimal contours in this case are straight lines. In other words, neurons that are optimally designed to process Gaussian inputs should be sensitive to only one stimulus dimension, which is in the direction perpendicular to the decision boundary.

However, neurons in several sensory modalities, as well as Hodgkin-Huxley model neurons, have been shown to be sensitive to several stimulus dimensions. Is this a sign of neural non-optimally? Signals derived from natural environment can sometimes be more closely approximated by an exponential, as pposed to a Gaussian, distribution. As an example, we considered two-dimensional exponentially distributed inputs $\sim\exp(-|x|-|y|)$. The optimal decision contours in this case can be solved for exactly. They have different shapes depending on the average spike probability, but are almost always curved. At extreme spike probabilities (either close to 0 or to 1), the optimal contours are similar to squares with rounded corners. For spike probabilities near ½, the optimal contours extend to infinity, resembling a wedge with a rounded corner. The ubiquity of non-Gaussian signals in nature, particularly of the exponential distributions considered here, makes these results relevant for neurons across different sensory modalities.

This work was supported by the Swartz Foundation and NIMH.

Prefrontal pathways for executive control

Helen Barbas (Boston University)

The prefrontal cortex in primates guides behavior by selecting relevant information, disregarding irrelevant information, and accessing motor control systems for action. In rhesus monkeys, distinct prefrontal cortices are linked with structures underlying sensory perception, cognition, and emotions. These topographically organized cortical pathways show remarkable regularity in their laminar organization, in patterns that can be predicted by a few structural features in the linked areas. Prefrontal axons terminating in the superficial versus the middle layers of superior temporal cortex are distinct at the synaptic level, in a pattern suggesting differences in efficacy of synaptic transmission across cortical layers. Moreover, prefrontal pathways project onto cortical excitatory neurons and onto distinct neurochemical classes of inhibitory neurons that also show regularity in their laminar distribution and differ in efficacy of inhibitory control. In addition, prefrontal pathways target subcortical structures, terminating onto excitatory as well as inhibitory systems in the amygdala and the thalamus, in patterns suggesting gating and recruitment of structures that may underlie the selection of relevant signals and suppression of irrelevant signals within a behavioral context. Quantitatively based maps of connections and their patterns may form the basis for investigating the dynamic recruitment of specific pathways underlying flexible behavior.

A Self-Organized Vortex Array of Hydrodynamically Entrained Sperm Cells

Ingmar H. Riedel, Karsten Kruse, Jonathon Howard (Caltech)

Many patterns in biological systems depend on the exchange of chemical signals between cells. We report a spatio-temporal pattern mediated by hydrodynamic interactions. At planar surfaces spermatozoa self-organized into dynamic vortices resembling quantized rotating waves. These vortices formed an array with local hexagonal order. Introducing an order parameter that quantifies cooperativity, we found that the array appeared only above a critical sperm density. A model allowed us to estimate the hydrodynamic interaction force between spermatozoa to be ~0.03 pN. Thus, large-scale coordination of cells can be regulated hydrodynamically, without chemical signals being required. (Science, 2005)

Spike timing dependent plasticity

Ruadhan O'Flanagan (UCSD)

The weight-dependence of a spike timing-dependent plasticity learning rule determines the computational properties of the resulting model neurons. Neurons sum their synaptic inputs and compare the sum to a threshold, and emit spikes if the sum exceeds the threshold. If one interprets the summation as a summation of evidence in favor of a hypothesis, and the spike as a decision that the total evidence provided by the synaptic inputs is sufficient to deduce the truth of the hypothesis, then a specific form of the STDP learning rule can be derived, which fixes the weight-dependence. The resulting learning rule shows stable learning and synaptic competition, in addition to admitting an information-theoretic interpretation. Properties of the learning rule, and the validity of the interpretation which leads to it, will be discussed.

On the agnosticism of spikes: attention, intention, and salience in the monkey lateral intraparietal area

Michael E. Goldberg (Columbia)

Attention is the process whereby the brain filters out sensory information unimportant for behavior. Clinical studies show that the parietal lobe is important for the attentional processes. Neurons in the lateral intraparietal area (LIP) filter out visual stimuli that are behaviorally unimportant , for example stable objects in the environment, although they do respond to those same stimuli when they appear abruptly in the environment.

Although LIP filters out behaviorally irrelevant visual stimuli, it does not filter out salient objects that are not the targets of a planned saccade. When a monkey plans a memory-guided saccade away from the receptive field of a neuron, the abrupt onset of a distractor in the receptive field evokes an enhanced response relative to the case when the monkey plans a saccade to the receptive field and the distractor subsequently appears in the receptive field. In these cases the distractor had no effect on the performance of the saccade.

Attention, as measured by an improvement in contrast sensitivity at the attentionally advantaged site, lies at the goal of a memory-guided saccade during the delay period, but it can be briefly captured by the abrupt onset of a distractor. The activity of neurons in LIP correlates with the monkey's attention both when it lies at the saccade goal and when it lies at the distractor site, and the time at which attention returns from the distractor to the saccade goal is predicted by the activity of neurons in LIP.

Most studies of eye movements in awake, behaving monkeys demand that the animal make specific eye movements. We have developed a new paradigm in which the monkey performs a visual search for an upright or inverted T among 7, 11, or 15 cross distractors, and reports the orientation of the distractor with a hand movement. The search array is radially symmetric around a fixation point, but once the array appears the monkey is free to move its eyes. The monkey's performance in this task resembles that of humans in similar tasks (Treisman and Gelade, 1980) : manual reaction time shows a set size effect for difficult searches (the crosses resemble the T's) but not for easy searches (the T pops out). Saccades are made almost exclusively to objects in the array, and not to intermediate positions, but fewer than half of the initial saccades are made to the T. We recorded from neurons in the lateral intraparietal area (LIP) while the monkey performed the search. LIP neurons distinguish the saccade goal at an average of 86 ms after the appearance of the array. The time at which neurons distinguish saccade direction correlates with the monkey's saccadic reaction time, suggesting that most of the jitter in reaction time for free eye movements comes from the discrimination process reflected in LIP. However, they also distinguish the T from a distractor on an average of 111 ms after the appearance of the array even when the monkey makes a saccade away from the target, suggesting that LIP has access to cognitive information about the target independent of the saccade choice.

We suggest that LIP provides a salience map which can be used by multiple systems. The salience map is constructed from independent signals (visual, cognitive, saccadic) which are summed in a linear fashion. When a saccade is appropriate, the oculomotor system can use the peak of the salience map to drive a saccade. The visual system uses the same spikes to determine the locus of attention. The source of the spikes, whether from a saccade plan or the visual system reporting the abrupt onset of a visual stimulus, is irrelevant to the use to which the recipient area puts the signal.

<u>Olfactory Coding in Flies</u>

Sean Luo (Columbia)

How does the nervous system categorize thousands of distinct odors? In the adult Drosophila, odor molecules bind to the olfactory receptors on the first order sensory neurons. The sensory neurons expressing the same olfactory receptor protein converge onto a single glomerulus in the antennal lobe and target a few excitatory projection neurons. Here we first present the latest calcium imaging and electrophysiological data and summarize the general features of the olfactory code. We then describe a simple feedforward model that examines olfactory learning and coding. In particular, we investigate the role of temporally synchronous spikes on olfactory learning and odor discrimination.

<u>The brain considered as a one-layer network between neurons and synapses.</u>

Tony Bell (Salk)

Feedforward information theoretic unsupervised learning algorithms have been developed for both rate-coding models and spike-coding models. However, it has been very hard to get them past learning a single layer of weights. (This is because of the partition function gradient). But taking a closer look at the physiology of plasticity (STDP in particular) suggests that the information integrator is the post-synaptic density, not the axon hillock. This means that the relevant information mapping is from neurons to synapses, not neurons to neurons. This introduces the new concept of an *inter-level* information flow. I will argue that if it could be optimised, it would answer most of the complaints against unsupervised learning principles (that they can't have goals, and so on).

<u>Olfactory processing in the *Drosophila* Brain</u>

Rachel I. Wilson (Harvard)

The chemical world has a higher dimensionality than any other class of sensory stimuli, and the olfactory system receives input from an unusually large number of unique information channels. This suggests that aspects of olfactory processing may differ fundamentally from processing in other sensory modalities. The large number of processing channels (~1000 glomeruli) in the rodent olfactory bulb has so far precluded a systematic analysis of olfactory processing in these systems. The fruit fly, by contrast, has only ~40 glomeruli, but the basic circuit architecture of the *Drosophila* and vertebrate olfactory systems is remarkably similar. *Drosophila* olfactory neurons are also electrophysiologically accessible *in vivo*, and can be mapped to identified glomeruli corresponding to specific olfactory receptor genes. We have exploited these advantages in performing a detailed analysis of the input/output function of six representative glomeruli. Using extracellular recording from six types of identified primary receptor neurons in the antennae, we have mapped their receptive fields using a diverse odor panel. Using whole-cell patch-clamp recording, we have also characterized the receptive fields of the six types of second-order neurons directly postsynaptic to each of these receptor neuron types. These experiments reveal a substantial transformation of olfactory representations in the *Drosophila* brain. Four features are prominent: the odor responses of second-order neurons are (1) amplified, (2) more tuned to stimulus change, (3) more broadly tuned, and (4) less noisy, as compared to primary receptor cells. These results are discussed in terms what we know about the connectivity of second-order neurons to their downstream targets.

Dynamics of context modulations in the "reach" network

Marina Brozovic (Caltech)

The pro-reach/anti-reach task is well suited for studying the way in which context configures cortical operations. The sensory inputs are the same, only the context/state of the animal is different. Yet these context differences lead to completely different neural processes and responses. Our neural network study is based on the experimental evidence which shows that the neurons in the parietal reach region (PRR) predominantly encode the position of motor goal as opposed to the position of the visual cue. Furthermore, the motor goal representation in PRR is delayed by ~60 ms in anti- compared to pro-reaches, implying the need for an additional step of cortical processing. Based on our recurrent network models, we propose that the basic reach network (visual areas ->PRR->PMd) has mainly feed-forward nature for the instructed pro-reaches, because the context information is not necessary. On the contrary, the reach network strongly depends on the feed-back projections from the PMd to the PRR for the case of anti-reaches, because these pathways bring in the context information for the anti-reaches.

An integrated microcircuit model of attentional processing in the neocortex

Xiao-Jing Wang (Brandeis)

I will discuss a joint work with Salvador Ardid and Albert Compte on a spiking neuron model of feature-based selective attention. We found that a wide range of physiological phenomena induced by selective attention arise naturally in a reciprocally connected loop of two (sensory and working memory) networks. Our model instantiates the `feature-similarity gain principle', and provides a synthetic account for biased competition and multiplicative gain modulation. The underlying circuit mechanism critically depends on an interplay between selective top-down excitation from the working memory network, feedback inhibition in the sensory network, and power-law input-output relationship of neurons. This work supports the notion that a common circuit subtrate subserves both working memory and selective attention.

Attentional state modulation of neural responses in rat auditory cortex

Gonzalo H. Otazu, Anthony M. Zador (CSHL)

Primary auditory cortex is involved in the representation of auditory information. However, an animal's attentional state can also affect the neural representation of a stimulus. Here we have studied the neural correlates of modulation during a simple attentional task.

We compared neural responses recorded while animals were performing a two-alternative choice auditory task ("attending condition") to those recorded while the animal was awake, but not engaged in any task ("non-attending condition"). We implanted tetrodes to allow chronic recordings in rat auditory cortex. We delivered auditory stimuli through earphones in the unrestrained animal, allowing us to measure cortical responses to identical acoustic stimuli while the animal was free to move.

We observed that the multiunit activity evoked responses were consistently larger in the non-attending condition than in the attending condition. There was no change in the spontaneous activity, and the effect could also be observed at the level of the local field potential (LFP).

We also recorded in the medial geniculate body of the thalamus of animals performing the same task. We did not observe a difference in the evoked responses between the attending condition compared with the non-attending condition. However, we did observe an increase in the spontaneous activity while the animal was performing the task. This increase is consistent with a model in which the elevated spontaneous firing rate in thalamus depresses the thalamic-cortical synapses, reducing the effects of acoustical stimulation on the engaged auditory cortex.

Noise-induced alternations in an attractor network model of perceptual bi-stability.

Rubén Moreno-Bote , John Rinzel and Nava Rubin (NYU)

When a stimulus supports more than one possible interpretation, perception alternates in a haphazard manner between them. What causes the bi-stable perceptual switches remains an open question. We develop a new, attractor-based framework in which alternations are induced by noise, and are absent without it. Our framework starts with an energy minimization model and then realized with rate-based and spiking network models, providing different levels of description of perceptual bi-stability. The model behavior is compared with experimental results from binocular rivalry, the most extensively studied bi-stable phenomenon. The model reproduces two long-standing experimental results, Levelt (1968) propositions II and IV. It also predicts a reduction of activity during rivalry compared to non rivaling vision, consistent with recent neuroimaging findings. The model can be implemented in an architecture that includes inhibitory populations locally connected to the competing excitatory populations and driven globally by an excitatory pool. This architecture readily generalizes to several competing populations, providing a natural extension for multi-stability phenomena.