

THIRTY-SEVENTH ANNUAL INTERDISCIPLINARY CONFERENCE

The Village at Breckenridge, Colorado

January 29 - February 3, 2012

Organizer: George Sperling, University of California, Irvine

ABSTRACTS

=====

Benjamin Backus
SUNY College of Optometry

Motion From Structure
Authors: B. Backus & B. Caziot

Surfaces at different disparities in a static stereogram appear to move relative to one another when the observer moves relative to the stereogram. How does a stimulus with no moving parts give rise to apparent motion? We call this a “motion from structure” (MFS) inference because in a real scene the absence of relative motion (e.g. absence of dynamic occlusion) in the proximal stimulus requires that the near surface move with the observer, relative to the far surface. But what mechanism is responsible for this inference? MFS looks like smooth motion, not a sequence of discrete standstill scenes. It is visible for minute head movements. These facts hint that MFS could be supported by a dedicated mechanism that combines 2D motion with represented depth structure. Paradoxically, in the real world dynamic occlusion is almost always present when parts move: 3D motion without 2D motion is highly accidental (it is no accident that stereopsis and motion parallax combine at threshold). Yet 2D relative motion is detected extremely well. Thus, MFS occurs when a trusted inferred scene structure—including information about its overall movement—combines with a reliable measurement of 2D motion. We will show, accordingly, that discrimination for MFS is most sensitive when proximal relative speed is zero. We characterize MFS along dimensions that have previously been used to distinguish between first order (luminance), second order (pattern contrast), and third order (saliency-tracking) motion (Lu and Sperling, 1995). MFS makes use of first- and second-order motion mechanisms to detect the absence of 2D motion and, in addition, depends on something that is different from (in addition to) third-order motion.

=====

Farran Briggs
Dartmouth Medical School

Synaptic and Circuit Mechanisms for Spatial Attentional Modulation
Authors: F. Briggs & W. M. Usrey

Visual spatial attention modulates activity in early visual pathways by changing neuronal firing rates. The mechanisms underlying these effects, however, are not known. We examined the effects of spatial attention on synaptic communication and correlated activity patterns in geniculocortical circuits. In two awake-behaving monkeys, we semi-chronically implanted stimulating electrodes in the lateral geniculate nucleus such that we could stimulate the presynaptic inputs to individual identified postsynaptic neurons recorded in layer 4C of primary visual cortex (V1). Animals were instructed to attend to optimal drifting sinusoidal gratings placed within or outside the receptive fields of recorded geniculocortical-recipient neurons in alternating trial blocks. Within 150 msec prior to a contrast change to the attended stimulus, the time window in which attention is thought to reach its peak, we stimulated geniculocortical inputs to recorded postsynaptic neurons in V1. Importantly, stimulation levels were set such that half of the stimulation trials resulted in post-synaptic spikes. We found that postsynaptic neurons demonstrated an increase in the probability of postsynaptic spiking in response to stimulation during trials in which the monkey attended to the stimulus within the recorded neuronal receptive field. Attention-mediated increases in geniculocortical synaptic efficacy correlated with attentional modulation of overall firing rate during the peak attentional window. In some cases, we were able to record simultaneously from multiple layer 4C neurons postsynaptic to common geniculocortical inputs. In these cases, we examined the structure of correlated activity across postsynaptic neurons with and without attention. We found that correlated activity patterns depended on functional

network organization and attention differentially modulated these networks. Our results provide the first evidence that spatial attention modulates neuronal activity at the synaptic level and suggest possible mechanisms by which attention influences communication across neuronal networks.

Patricia Cheng
University of California, Los Angeles

A Causal Explanation of Functional Basic-Level Categories
Authors: N. James, J. Saiki, & P. Cheng

Lawrence Cormack
The University of Texas at Austin

The Binocular Processing of Three-Dimensional Motion

The binocular sensing of motion through depth (motion towards and away from the observer) can be accomplished using one of two cues – specifically, by either tracking the changes in retinal disparities over time, or by computing a motion signal separately in each eye’s image, and then subtracting the two motion signals. The former cue is widely appreciated, due in no small part to the existence of elegant stimuli that preserve changing disparity information while eliminating the motion signal altogether (Julesz, 1971). A body of recent evidence, however, indicates that the cue based on interocular velocity differences actually dominates over a wide range of stimulus conditions. Moreover, adaptation experiments indicate that motion through depth is processed by a specialized mechanism that is both distinct from frontoparallel motion circuits and uniquely sensitive to the interocular velocity difference cue. In this talk, we will review this evidence in order to update the current zeitgeist concerning dynamic depth processing.

Barbara Doshier
University of California, Irvine

Visual Working Memory and Attention

Collaborators: W. Chu, R. Najima, Z.-L. Lu

Visual working memory (VWM) has generally been tested under conditions of high visibility. However, VWM is a perceptual and memory function that presumably must operate on low contrast and noisy stimuli as well. We have evaluated VWM under external noise and contrast variation to measure change-detection and cued continuous feature report, the two major VWM test paradigms, under many conditions of visibility. Memory report was systematically affected by display size, Gabor contrast, and external noise, varying from poor to excellent. Often, results from change-detection are interpreted as discrete high-threshold ‘slot’ models, while more often continuous reports that measure the variability of feature memory seem consistent with continuous limitations in capacity. Here we show that an integrated perceptual template model of the observer, including task-relevant decision modules, provides a strong account of Gabor orientation VWM performance in both tasks.

James Elder
York University

Ecological Relevance of Exogenous Attention

In the standard exogenous cueing paradigm, a peripheral visual pre-cue affects the time to detect a subsequent peripheral target. This exogenous attention effect is thought to be a reflexive process based on simple properties of the cueing stimulus. This is complicated somewhat by recent experiments in which human faces are used as pre-cues, and their effects depend upon facial expression. These results have been interpreted in terms of a reflexive process selective for threat-related signals. Here we examine whether exogenous cueing may be based on more general ecological principles, by comparing the efficacy of human face pre-cues with random-phase controls. The target was a bright 0.25 deg disk. The face cues were 2 deg natural images of faces with neutral expressions. The control cues had identical amplitude spectra but randomized phase. The intensity, hue and saturation of the face images and their controls were matched in both mean and variance. Subjects were asked to maintain fixation on a central cross. After 500 msec, a cue was flashed for 20 msec, 8 deg randomly to the left or right of fixation. Following a variable SOA, the target was displayed 6 deg randomly to the left or right of fixation until response. The location of the cue was not predictive of the location of the target. Forty-five observers each completed 960 randomly-interleaved trials. Faces were found to be significantly more effective as exogenous attentional cues than random-phase controls. Interestingly, this effect was lateralized in the invalid-cue condition. Specifically, the effect of an invalid face cue was significantly less pronounced when the cue was presented in the left hemifield and the target was presented in the right hemifield, than vice-versa. This finding can be interpreted in terms of a specialization for faces in the right hemisphere of the human brain.

=====
Casper Erkelens
Utrecht University

Perceptual Rivalry in Monocular and Binocular Vision

Human vision follows from two systems that operate according to different principles: the monocular and binocular visual systems. Monocular vision makes use of assumptions. Assumptions prescribe how 2D images must be segmented into cohesive and meaningful patterns and objects. Additional assumptions about viewpoint and layout of the scene lead to 3D perception that is surprisingly often veridical. Numerous visual illusions illustrate the many assumptions. Binocular vision is based on epipolar geometry, which is the projective geometry between the two eyes' views. Applying epipolar geometry gives rise to unique disparity fields. Stereovision from random-dot patterns shows that the binocular visual system can exploit the disparity fields for the unambiguous reconstruction of 3D scenes. In conclusion, the monocular and binocular visual systems are based on different cognitive and mathematical principles. In monocular vision, there is room for interpretation and ambiguity whereas in binocular vision there is not. By this fact, monocular visions should be superior to monocular vision. However, this is not always the case. For instance, Necker cubes keep alternating when we view them with both eyes. To investigate the interaction between monocular and binocular vision, I measured the distribution of perceptual alternations during the viewing of various monocularly visible and hidden Necker cubes. The results give rise to a new model of monocular and binocular interaction.

=====
Wilson Geisler
The University of Texas at Austin

Toward a Neurophysiology-Based Model of Detectability in Natural Images Authors: W. Geisler, C. Bradley, & M. Michel

The human visual system implements an elegant compromise between the competing goals of maximizing spatial resolution, maximizing field of view, and minimizing neural resources. It combines a small high-resolution fovea and a large low-resolution periphery with a saccadic eye movement system that can rapidly direct the fovea toward points of interest in the visual scene, based on information encoded in the low-resolution periphery. Although this is a beautiful system, it is complex. In particular, rigorous analysis of behavior and physiology in most natural tasks requires first characterizing the variation in the spatial resolution of early visual processing across the visual field. Thus there is a critical need for a model that adequately predicts target detectability across the visual field for naturalistic and natural backgrounds. I will describe our recent attempts to develop a practical working model of target detectability that is based on the optics, anatomy, and neurophysiology of the early stages of processing in the primate visual system.

Justin Halberda
Johns Hopkins University

An Ensemble Group is Selected and Processed as a Single Item

Visual systems, both human and computer, are taxed to provide an interpretation of the large amounts of noisy evidence present in a visual scene. Summary statistics (or ensemble representations), which collapse across multiple objects or pixels to provide a statistical description of the appearance of a part of a scene, have proven computationally efficient for providing a general interpretation of visual evidence and much excitement has surrounded recent demonstrations that such representations may play an important role in how the human visual system interprets visual evidence in the periphery, under conditions of crowding, outside the focus of attention, and when attention selects groups of items. In this talk I highlight an area for further theorizing in this domain. Every summary statistic generated for an image first requires a selection of the relevant data to be described by the statistic. Instances of such selection range from physiological filters, through segmentation and top-down parsing. This selection requirement constrains which aspects of a scene are amenable to being described by summary statistics and provides a scaffold for beginning to describe the stages of information processing that must support the generation of summary statistics. I present this argument and evidence that, when this selection is performed by effortful top-down attention, an ensemble group of items functions as a single object for visual attention, visual working memory, and cognitive processing.

Xiaoping Hu
Emory University

Time-Frequency Dynamics of Brain Resting State Effective Connectivity

Authors: X. Hu & G. Deshpande

Resting state (functional MRI) connectivity of the brain holds an enormous potential for studying the brain in both its normal and diseased states. A number of resting state networks have been consistently identified and form the basic signatures of the brain at resting state; these networks include the default mode network, dorsal attention network, fronto-parietal control network, and hippocampal cortical memory network. To date, most of the studies of the resting state networks have relied on the implicit assumption of stationarity, i.e. the functional MRI signal fluctuations that are the basis of the identification of these networks and which have time-invariant statistical characteristics. Recently, evidence is emerging that demonstrates the nonstationary and dynamic nature of these fluctuations. We hypothesized that the effective connectivity between resting state fMRI time series may dynamically evolve across time and frequency. Using a dynamic, multiscale, vector, autoregressive model based on Kalman filtering, we obtained effective connectivities between 33 regions in the aforementioned four resting state networks as functions of time and frequency. We observed that in the 0.1 to 0.01 Hz frequency range, the resting state networks exhibit a finite number of distinct states and the network toggles between them. While further understanding of this observation is needed, it indicates both stability and variability of the resting brain.

Alex Huk
The University of Texas at Austin

Encoding and Decoding of Signals for Perceptual Decisions in Posterior Parietal Cortex

Authors: A. Huk, M. Meister, J. Pillow, & I. Park

When primates perform a moving-dot direction-discrimination task, neural activity in the posterior parietal cortex (area LIP) ramps upwards or downwards. This ramping activity, and its dependence on the direction and strength of the visual motion stimulus, suggests it is a direct and explicit neural correlate of the accumulation of evidence over time. Here we test this hypothesis by making subtle changes to decision-irrelevant components of the task. These manipulations reveal that decision-related activity in LIP co-exists with decision-irrelevant signals, and that the two types of information are often

multiplexed in idiosyncratic ways that do not intuitively map on to the formation of decisions. We then show how a generalized linear model can extract and isolate these signals, offering a more thorough framework for understanding the enigmatic encoding and decoding of LIP activity during perceptual decision-making.

Holly Jimison
Oregon Health & Science University

Big Data: The Modeling and Inference Challenges of Monitoring Health Behaviors in the Home
Authors: H. B. Jimison & M. Pavel

There is a growing scientific interest in addressing the challenges of modeling and summarizing the large influx of new types of medical and behavioral data. Several governmental agencies have initiatives around the topic of what is being called “Big Data.” Although much of the focus thus far has been on the large datasets associated with genetic data, there are also new types of data being collected as part of home and mobile health (mHealth) monitoring that also require innovation in analysis techniques for modeling, inference and summarization. These new data from home and mHealth monitoring are heterogeneous, of varying quality, and sampled at widely varying rates. For example, the data may come from video sensors used to monitor gait, microphones used to measure voice stress, computer keyboards used to measure motor speed, motion sensors to measure physical activity, etc. These data types are frequently generated by inexpensive sensors to provide unobtrusive but indirect measures of important latent health variables of interest (behavioral phenotyping). In this presentation I will describe a framework for inferring health indicators from models based on noisy and indirect, but inexpensive and frequent, sensor data from the home. This approach will be contrasted with standard clinical measures.

Shaw Ketels
University of Colorado at Boulder

Attentional Focus and Complex Motor Skill Acquisition: The Case of Snowboarding
Authors: S. L. Ketels & K. R. Lohse

Experimental research on instructional design has shown that performance (in the short-term) and learning (in the long-term) depend critically on the nature of the instructions being given. An excellent example of this is how verbal instructions direct a learner's attention to different aspects of the task (see Wulf, 2007, for a review). Instructions encouraging learners to focus on the effects of their actions improve performance and learning in a number of ways: increased accuracy, more efficient muscular recruitment, and decreased preparation time (Lohse, Sherwood, & Healy, 2010). Conversely, verbal instructions encouraging learners to focus on the movements of their actions have the reverse effects and generally worsen performance and learning. Although these effects have been demonstrated in a number of laboratory tasks (Maddox et al., 1999; Shea & Wulf, 1999; Wulf, Shea, & Park, 2001; Wulf, HöB, & Prinz, 1998; Wulf, Lauterbach, & Toole, 1999), there has been little to no research on the practical significance of these effects outside of the laboratory in naturalistic settings. Here we present results of a naturalistic investigation focusing on the instruction of novice snowboarders.

David Knill
University of Rochester

Learning Scene Statistics for Perception and Action

Humans often perceive, act and make decisions based on noisy and ambiguous sensory information. In order to optimize behavior, the CNS should learn and use knowledge of the statistics of objects and events in the environment to reduce uncertainty. Consider the problem of estimating object velocity for purposes of guiding action - avoiding an approaching car, timing the swing of a bat, etc. The sensory information available for such actions is typically quite sparse. In order to

accurately estimate velocity, an observer / actor should use an accurate statistical model of object motion. A popular modern theory of motion processing holds that early stages of motion estimation incorporate a prior bias toward slow velocities - a form of statistical knowledge. We have been studying the learning and application of statistical models of object speed for guiding a simple motor behavior - timing movements to hit a moving object. We use hierarchical Bayesian analysis techniques to fit the implicit statistical models that subjects use to plan their movements. Our results show that subjects learn accurate models of the first-order statistics (means and std. deviations) of object speeds, but overestimate the temporal correlations across trials. Subjects adapt their implicit models to stimuli sets with different temporal correlations, but maintain a strong bias toward positive correlations. When subjects generate the object motions, by "throwing" the objects themselves, the bias toward positive correlations disappears, suggesting they have prior knowledge of the statistics of their own movements. I will also discuss results on subjects' abilities to learn more complex (e.g. bimodal) distributions of object speeds. Our results show that the CNS is remarkably adaptive to the statistics of its inputs and that it incorporates different models for different sources of variability (self-generated vs. external), though it retains strong biases even when the local statistics (in time) deviate from its prior assumptions.

Adam Kohn
Albert Einstein College of Medicine

Adaptation in the Primate Motion Processing Stream
Authors: A. Kohn, C. Patterson, & S. Wissig

Adaptation has been shown to affect neurons at different stages of the visual hierarchy, from the retina to temporal cortex. It remains unclear, however, how a particular adapter alters encoding and processing across the cortical hierarchy. Addressing this requires understanding which neurons adapt for a particular input and, when effects occur early in the system, how this changes responses and computations downstream. We studied these questions using extracellular recordings in the primate motion processing stream —primary visual cortex (V1) and area MT—using identical stimulus conditions. We found that adaptation with sinusoidal gratings can have a range of effects in both areas, depending on the size and duration of the adapter: neuronal responses to gratings can suppress or facilitate, and tuning can shift toward or away from the adapter. These effects can be explained by a model in which adaptation reduces drive both to the receptive field and the suppressive surround. Across adaptation conditions, effects in MT mirrored those in V1. To test how altered V1 responses influence computations unique to MT, we measured the effect of prolonged adaptation with gratings on plaid tuning. Adaptation lowered pattern selectivity in MT, in a manner consistent with recently proposed models of V1-MT circuitry. Our findings suggest that the effects of grating adaptation occur in V1, but that these altered feedforward signals can disrupt computations performed in MT.

Tom Landauer
University of Colorado, Boulder

The New Word Maturity Metric: How Words Develop Their Meanings by Exposure to Language

For human learners, the meanings of words and passages grow with cumulative exposure to language. Importantly, what can be learned at any point depends on what has already been learned. As a result, every word develops over a different trajectory, from unknown to a nearly constant level at adulthood. We simulated this process by a computational model that successively added groups of 40,000 paragraphs taken randomly from a general corpus of ~500,000 words. Each word and paragraph was represented by a vector of 300 independent real numbers. At each step we recomputed the cosines between each vector and every other in the corpus. The following algorithm was then used to determine the resulting "maturities" of the words and paragraphs. "A word's meaning is assumed to equal to the sum of the meanings of all the paragraphs it has occurred in and a paragraph's meaning the sum of all the words that contain it." Quite accurate emulations of human learning have been obtained. For example, the average maturity of its words correlated ~ 0.8 with the school grade in which a given reading is usually assigned.

Michael Lee
University of California, Irvine

Two Misuses of Probability Theory in Cognitive Modeling

Zhong-Lin Lu
Ohio State University

Functions and Mechanisms of Perceptual Learning
Authors: Z.-L. Lu & B. Doshier

Perceptual learning -- the improvement of performance through practice or training -- has been observed over a wide range of perceptual tasks in adult humans. The high degree of plasticity of the adult perceptual systems suggests that perception and perceptual learning cannot be studied separately. In this talk, we will review some major functions and mechanisms of perceptual learning, including specificity of perceptual learning, the law of practice in perceptual learning, mechanisms of perceptual learning, the level and mode of perceptual learning, optimal training procedures, and computational models of perceptual learning. Studies of these various aspects of perceptual learning have greatly enhanced our understanding of the information processing limitations of the human observer, and how the state of the observer changes with training, with strong implications for the development of potential noninvasive training methods for perceptual expertise in normal populations and for the amelioration of deficits in challenged populations.

References:
Doshier, B., & Lu, Z.-L. (1998) PNAS.
Lu, Z.-L., Liu, J., & Doshier, B. A. (2010) VisRes
Liu, J., Lu, Z.-L., & Doshier, B. A. (2010) JOV

Cory Miller
University of California, San Diego

Responses of Marmoset Prefrontal Cortex Neurons During Natural Vocal Communication

There is little doubt that selection for neural mechanisms that permit individuals to effectively navigate the complexities of their respective social landscape was central in the evolution of primate cortex. There remains, however, little known about the neural mechanisms that underlie behaviors that mitigate social interactions in nonhuman primates. Here we describe work aimed at addressing this topic. We examined the neural basis of a vocal behavior known as antiphonal calling in common marmosets, a natural behavior involving the reciprocal exchange of vocal signals. At the behavioral level, we employ software-based interactive playback experiments to dissect the sensory and motor processes underlying this behavior. Neurophysiology experiments combine this software with a technique for recording the activity of single neurons in prefrontal cortex while marmosets are freely-moving and engaged in this behavior as well as in other related contexts. Analyses indicated that neural responses were strongly context-dependent. Individual units exhibit different responses depending on whether the animal is restrained, freely-moving or engaged in active communication. Moreover, we found that neural responses were correlated with idiosyncrasies of the behavior, such as the duration of antiphonal calling bouts. The ongoing research on antiphonal calling has the potential to provide unique insights into the neural processes underlying natural behavior in primate cortex.

Jeff Mulligan
NASA Ames Research Center

Cognitive Engineering for Next Generation Air Traffic Control Towers

Today a controller in a typical air traffic control tower accomplishes his/her tasks with the aid of a large number of "tools," each of which provides information or facilitates communication with aircraft and other controllers. Each tool usually consists of a separate microcomputer, with the collection of tools representing the full spectrum of operating systems. As part of its effort to modernize the National Airspace System (NAS), the Federal Aviation Administration (FAA) is attempting to combine the tools into an integrated workstation, and recognizes the need to incorporate knowledge from cognitive psychology into the design process. Our group at NASA Ames is assisting in the development of functional human factors requirements for integrated workstations in air traffic control towers. In this talk, we present some of the open issues and describe the process we are using to generate requirements.

Erica Okada
University of Hawaii

The Private and Public Benefits of Green Products and the Price/Quality Tradeoff
Authors: E. M. Okada & E. L. Mais

Consumers generally make choices to make themselves better off. But the main reason why many people choose ecologically sustainable or "green" products is because they are better for all of us. Theoretically, we contrast between the private and public benefits. Drawing on construal level theory (Liberman and Trope 1998; Trope and Liberman 2010), we present that when consumers make green purchases for their own benefit, they put more weight on more concrete product features and decision criteria, such as price and how much money they can save. And when they make green purchases for the benefit of all, they put more weight on more abstract product features and decision criteria, such as quality and how much time they can save.

Practically, purchases become more attractive to consumers when they can get more and/or pay less. When they buy green for their own benefit, paying less becomes more important, and when they buy green for the good of all, getting more becomes more important.

Consumers pay a combination of time and money in product acquisition, and the two currencies are typically traded off. One pays in time for low price, and in money for convenience. When consumers buy green for their own benefit they are willing to pay relatively more in time, and when they buy green for the good of all they are willing to pay relatively more in money.

We explore this topic through two field surveys and one behavioral experiment.

Wolfgang Pauli
University of Colorado at Boulder

Changepoint Detection versus Reinforcement Learning: Separable Neural Substrates for Different Forms of Bayesian Inference

Foraging animals in nonstationary environments must adapt to changes in payoffs expected from each location. Competing normative approaches to learning in nonstationary environments have been founded on two alternative assumptions regarding temporal dynamics. One approach relies on the assumption that reward probabilities drift continuously over time. Under this assumption, optimal inference is described by the Kalman filter and is closely approximated by reinforcement learning mechanisms that incrementally adjust value estimates following each action. An alternative assumption is that the environment undergoes periods of stationarity punctuated by change points. Optimal behavior under this assumption entails abrupt resetting of knowledge at likely change points and little learning otherwise. We will report results of a reversal-learning experiment with rats, using concurrent variable interval schedules on two levers. Before the contingency reversal we pharmacologically impaired either the dorsomedial or dorsolateral striatum of the basal ganglia. It is predicted that the dorsomedial striatum is better characterized by change-point inference, whereas the dorsolateral system is better characterized by reinforcement learning.

Misha Pavel
National Science Foundation

Fundamental Scientific Challenges in Transforming Healthcare: Model-Based Approaches

Transforming healthcare to be evidence-based, patient-centered and proactive will require substantial technical advances. Recognizing these challenges, NSF has developed a program in Smart Health and Wellbeing that is focused on stimulating relevant research. In this presentation I will discuss a subset of the relevant technical issues that underlie the healthcare transformation and advocate approaches based on computational (mathematical) modeling. I will then discuss in some detail a number of specific approaches in the domains such as sparse sampling and information fusion.

Zygmunt Pizlo
Purdue University

Definition of Shape

Shape has been studied for centuries despite an absence of a commonly accepted definition. The absence of a definition was a major obstacle in the progress of understanding the mechanisms underlying shape perception. In particular, there has been no consensus about whether shapes are perceived veridically. I will review the main definitions of shape and discuss their shortcomings. My new definition is based on the following two assumptions: (i) not all objects and patterns have shape, and (ii) shape is related to geometrical self-similarities (symmetries) of an object, rather than to similarities among objects. Adopting this new definition leads to several important insights: (i) it becomes clear how veridical shape perception can be achieved; (ii) shape can be "assigned" equally well to rigid and non-rigid objects; (iii) symmetries are universal shape priors in the sense that they do not have to be abstracted from objects or updated based on experience; and (iv) symmetries can be used for efficient memory indexing (content addressable memory).

Nicholas Port
Indiana University

The Development of Microsaccades from Toddling to Retiring
Authors: N. Port, J. Daniels, T. Hoffa, S. Hitzeman, & S. Beckerman

The evolution of the primate visual system has solved multiple conundrums over the millennium. One such conundrum is the competing need to have (A) rapidly adapting photoreceptors (in order to compensate for extreme changes in absolute light levels) and (B) very high acuity in our central fovea (where single cone photoreceptors are mapped onto single retinal ganglion cells). Without a solution to this problem, the visual image would instantly fade whenever an image remained still on the photoreceptor mosaic. The ocular motor system accomplishes this feat through simultaneously moving the fovea from location to location while wiggling the overall image with micro-eye movements in order to keep the photoreceptors from adapting. There are three forms of micro-eye movements: drift, tremor, and microsaccades.

Drift and tremor can only be studied in the laboratory with the invasive scleral eye coil technique. Fortunately, recent advances in high-speed and high-resolution video-based eye trackers now allow the study of microsaccades noninvasively. However, despite the renewed interest in microsaccades and how they interact with visual attention and perception, there are, to the best of our knowledge, no studies of the development of microsaccades. Using a visual search task that has been shown to maximize microsaccade frequency (Otero-Millan et al 2008), we studied the microsaccades of people of all ages while solving the classic "Where's Waldo" puzzles. Specifically, we studied the eye movements of 435 children and 47 adults utilizing our portable eye tracker laboratory while people played "Where's Waldo" for about 7 minutes. We found children as young as 3 years old make microsaccades, but at a significantly reduced rate. In general, microsaccade rate appears to plateau around the age of 12 and possibly diminish around age 60; however, puzzle solving accuracy and response time in young children (5-10) were generally better than in adolescents or adults.

Roger Ratcliff & Gail McKoon
Ohio State University

Diffusion Decision Model Assumptions and Using the Model to Map Stimulus Properties
Authors: R. Ratcliff & G. McKoon

Reviewers often ask about the distributional assumptions that are made about across-trial variability in model parameters in the diffusion model. In this talk, we present examples in which the typically-used normal distribution for drift rates across trials is replaced by beta and uniform distributions, examples in which the typically-used uniform distribution for starting point is replaced by a beta distribution, examples in which starting point variability is replaced by independent variability in the two decision criteria, and examples in which the typically-used uniform distribution for nondecision time is replaced by a normal distribution. The conclusion (as of November 2011): The main parameters of the model – drift rate, boundary separation, and nondecision time – are well-recovered in all cases (with poorer recovery only for extreme parameter values). In a second part of the talk, we examine the form of the drift rate function across levels of difficulty in numerosity discrimination, brightness discrimination, motion discrimination, line-length discrimination, and area discrimination. For some of these tasks, the form of the drift rate function is linear with the variable manipulated; in others, it looks like an s-shaped psychometric function. Importantly, the diffusion model provides estimates of drift rate when accuracy is at ceiling or floor, and so produces psychometric functions that could not be obtained with only accuracy measures.

=====
John Reynolds
Salk Institute

Neural Mechanisms of Attention

=====
Eyal Seidemann
The University of Texas at Austin

Attentional Modulations Related to Spatial Gating But Not to Allocation of Limited Resources in Primate V1

Attention can modulate neural responses in sensory cortical areas and can improve behavioral performance in perceptual tasks. The nature and purpose of these modulations, however, remain under debate. We used voltage-sensitive dye imaging (VSDI) to measure V1 population responses while monkeys performed a difficult detection task under focal or distributed attention. We found that despite improved behavioral performance under focal attention, V1 responses at attended locations were indistinguishable between focal and distributed attention, inconsistent with the hypothesis that an important goal of attention is to allocate limited representational resources in V1. However, V1 responses at all attended locations were significantly elevated relative to ignored locations, consistent with the hypothesis that an important goal of attention in V1 is to gate task-irrelevant information. This elevation, which is widespread and starts shortly before stimulus onset, could bias competition in subsequent processing stages in favor of attended stimuli.

=====
Steve Shevell
University of Chicago

Beware Your Inner Chromatic Compass: Binding Errors of Color and Motion-Direction

In natural viewing, the retinal stimulus typically is a mosaic of many distinct regions, each of which has a shape, size, location, direction of motion and color. We see integrated objects with all features bound together (imagine a soaring green inbound frisbee). Two hypotheses of feature binding were tested by measuring binding *errors* of motion-direction and color, in which red dots moving upward and green dots moving downward were seen going in the opposite directions

(red downward, green upward; Wu, Kanai & Shimojo, 2004). The first hypothesis, from Barlow (1981), posits that the commonality of features other than color and motion (for example, shape) regulates the frequency of color/motion binding errors. The second posits that color/motion binding errors reflect resolution of an ambiguous neural representation of the direction of motion, not erroneous pairing of each color with a well represented direction of motion (so-called illusory conjunction). Experiments supported both hypotheses.

Richard Shiffrin
Indiana University

A Dynamic Model for Recognition Decisions
Authors: G. Cox & R. Shiffrin

We present a new dynamic model of decision making in recognition memory and apply it to accuracy and response time data from a study using widely varying kinds of stimuli (e.g. word, dot pattern, toaster, face, scene, blob, name, snowflake). Supposing that absolute levels of activation or familiarity vary widely for different materials, how can a criterion for old-new recognition be set properly when there is no opportunity to learn because each stimulus type is studied and tested just once? We depart from standard signal detection models that predict recognition accuracy and propose a model in which the decision is made by accumulating changes in the activation profile as features are successively extracted from the test stimulus. Increases in activation accumulate in one counter and decreases in another. The counters each have a threshold. The first to be reached determines the choice and the response time. Our study shows that participants are remarkably well tuned to differences in stimulus types, having little trouble in producing reasonable judgments for all (hits and false alarms more or less centered) despite large differences in levels of performance. The model accounts for both accuracy and response times.

George Sperling
University of California, Irvine

The Computational Architecture of Visual Selective Attention: A Review

Bosco Tjan
University of Southern California

Object Identification in the Periphery -- A Case of Mistaken Priors
Authors: B. S. Tjan, A. S. Nandy, M. Kwon, & S. T. L. Chung

Object recognition in clutter is severely impaired in normal peripheral vision beyond what can be accounted for by the lower spatial resolution. We hypothesize that this phenomenon of “visual crowding” is caused by a consistent and systematic error in acquiring the statistical priors about the visual world in the peripheral field. To form a visual percept, outputs from neurons at the early stages of visual processing must be selectively grouped into progressively complex elements in a manner congruent with the statistical regularities of the visual world. Visual experience is needed for the visual system to acquire the statistical priors, implemented as connections and interactions among neurons. We assume that the acquisition of image statistics occurs predominantly at attended spatial locations. Since peripherally deployed attention often elicits a saccade towards the attended location, if the attention spotlight is not fully extinguished before the eyes move, the acquisition of image statistics at the attended location will be affected by the saccade. We show with simulations that this theory provides a quantitative explanation of the shape of the spatial extent of crowding. The theory makes specific predictions, among which we have empirically demonstrated that (1) radially oriented Gabor flankers yield smaller crowding zones than tangentially oriented flankers, (2) crowding zones are more elongated in the lower visual field, and (3) when saccades are re-referenced to a preferred peripheral location (PRL) due to a central scotoma, the crowding zone at the PRL becomes isotropic; elsewhere they are the superposition of two elongated ellipses, with the major axis of one pointing towards the fovea and the other towards the PRL.

Funding acknowledgement: NIH/NEI Grants R01-EY017707

John K. Tsotsos
York University

Revisiting Ullman's Visual Routines: Towards an Executive Controller for Visual Attention

This presentation will present the first steps we have taken towards re -visiting, -defining and -formulating what was a wonderful concept, but one that is now significantly outdated: The Visual Routines framework of Shimon Ullman. Ullman based his conceptualization on knowledge of human vision and attention of the early 1980s, and most followers of VRs do the same. Among them, Ballard & Hayhoe make strong arguments about the need for non-saliency methods (the ones Ullman used) for attention but do not propose an alternative. Ullman's work is based completely on how Marr viewed visual processing in the brain. Attentive operations are critical and based on Koch & Ullman (1985) who provided a computational view of Feature Integration Theory (Treisman & Gelade 1980). A great deal has changed, however, in our understanding of human vision since 1984. We briefly present those changes, show how a more appropriate attention model is our own Selective Tuning model. We describe how the original Visual Routines framework must change in order to be updated to the present and end with our current formulation and examples of its operation.

Andrew B. Watson
NASA Ames Research Center

Video Acuity: A Metric to Quantify the Effective Performance of Video Systems

There is a widely acknowledged need for metrics to quantify the performance of video systems. Existing metrics are either difficult to measure, or are largely theoretical, or do not reflect the full range of effects to which video may be subject, or do not relate easily to performance in real world tasks. We propose a new empirical metric –Video Acuity – that is simple to measure and is likely to relate directly to task performance. Video acuity is determined by the smallest letters that can be automatically identified through the video system. It is expressed most conveniently in letters per degree of visual angle. We provide examples of application of the metric to several video systems.

Deyue Yu
Ohio State University

Crowding in Peripheral Reading

In the United States, macular degeneration is one of the major eye diseases causing low vision due to the loss of central vision. Unfortunately, there are no preventive measures, and no cure for this eye disease. Since people with central vision loss must use their peripheral vision for daily tasks such as reading newspapers, identifying faces, and recognizing objects, understanding sensory determinants of peripheral object recognition is vital for the rehabilitation of these patients. In this talk, I will use reading as an example and review how crowding, a major sensory factor, limits the performance of letter and word recognition in peripheral vision. Crowding refers to the impairment of target recognition due to the close proximity of neighboring objects. A full understanding of this visual phenomenon is crucial for both practical and clinical purposes.