AIC 49, Abstracts

Collin Allen, UCSB

Attention is not all you need (two kinds of structures for transformers)

Attempts to say what transformer-based systems don't do well are rapidly rendered obsolete as the models and architectures are fine-tuned and otherwise tweaked to improve performance in specific areas. A genuine question arises as to whether these modifications achieve task-specific success by papering over fundamental shortcomings in the transformer approach, perhaps exemplifying Goodheart's Law that any measure that becomes a target ceases to be a good measure. I argue that despite the obvious advances made through the iterations of LLMs and other transformer-based systems, there remain stubborn indicators of an underlying problem. Because of multi-headed attention, transformers are excellent at detecting structure in the data but frequently fail in often unpredictable ways to represent structures underlying the data, leading to outputs that are often locally coherent but globally incoherent. While LLMs can be prompted in ways that elicit descriptions of underlying structures, they routinely fail to apply such "knowledge" effectively.

Ariana Andrei, Houston Methodist

Cortical organization of visual information destined for perception

Only a small fraction of the information processed by the brain ever reaches sensory awareness. Understanding the selection process will allow for precisely targeted neuroprosthetic interventions to restore sensation lost by disease or accident. In my endless quest to discover the essential circuits elements and computational transformations required for sensory information to be rendered perceivable, I present a compendium of my research results and the clues they have yielded about this mysterious process. First, using in vivo electrophysiological recordings in macaque monkeys, combined with optogenetic stimulation as a causal activation tool, I present evidence that perceivable visual information from primary visual cortex (V1) is integrated across similarly tuned populations. Secondly, using optogenetic inhibition of excitatory neurons in V1 in superficial layers, perception is only altered when responses across cortical columns are homogeneous rather than heterogeneous. Thirdly, exploring the temporal domain of visual perception, I present results from a study of visual statistics of cinema showing that virtually all movies made in the past 100 years rely on an average shot duration of 2 seconds, suggestive that human perceptual and attention mechanisms are optimally engaged by this duration. Lastly, exploring the temporal aspects of learning and memory, I present evidence an experimentally novel form of homeostatic plasticity that stabilizes visual cortical networks, preventing Hebbian plasticity during wakefulness, but not during sleep.

Kristin Andrews, York U.

Animal Rules

Learning theory demonstrates that across species, animals are able to learn associations. But do they follow rules? In this talk I argue that many species also follow rules with deontic content, and illustrate by considering the dimensions of social normativity found in invertebrates and vertebrates.

Jake Beck & Kevin Lande, York U.

Hidden Movements

Most objects are partially obscured by other objects, sometimes disappearing entirely from sight and then reappearing again—a deer passing behind a boulder; a car through a tunnel; a baseball behind a foul pole. In such cases, we perceive the object as moving continuously even as it has disappeared from view, a phenomenon sometimes called "the tunnel effect." The tunnel effect raises fundamental questions about what it is to see or perceive. Intuitively, you cannot see what is fully occluded; so, you do not see the baseball when it is behind the foul pole. This suggests that your perception of the flying ball is discontinuous: you see it, then you momentarily don't see it, and then you see it again. Yet there are also strong reasons to think that you can continuously visually track the baseball behind the foul pole. Successful tracking plausibly depends on having an accurate perceptual representation of the ball behind the pole and not merely a cognitive representation such as a belief. We have something of a puzzle here: How do you continuously see the ball passing behind the pole if you do not continuously see the ball? We argue that a good answer can be recovered by making two distinctions: first between the perception of an object and the perception of an event, and second between seeing something and merely being visually aware of it.

Joshua Brown, IU

Scalable neural network models of human vmPFC activity, goal-directed planning, and problem solving

Cognitive control has largely focused on constructs including inhibition and rule-guided behavior, but we argue that goal-directed behavior is a more powerful and encompassing construct. We demonstrate a series of two computational neural models of goal-directed behavior. The first conceptualizes goal-seeking as gradient ascent on a flexible value landscape and has capabilities that RL variants lack, but suffers from the scalability limitations of one-hot representations. The second model overcomes this by incorporating a goal-reducer deep network, which learns by random exploration and can achieve goals efficiently by removing loops in state trajectories experienced during learning and then recursively finding efficient intermediate subgoals. The model outperforms deep RL. Representational similarity analysis between the model and humans doing a goal-seeking task with fMRI shows that the goal and goal reduction mechanisms have representations that match those of human ventromedial prefrontal cortex.

Matt Chafee, U. Minn.

Cortical and subcortical computations for cognitive control

Prefrontal cortex (PFC) supports cognitive control, the ability to use neural representations of abstract information (such as goals or rules) to guide behavior. To carry out the computations required, prefrontal cortex communicates with the caudate nucleus (CN) and mediodorsal thalamus (MD). We have only a partial understanding of the nature of this communication or its computational significance. To address those questions, we trained monkeys to perform tasks requiring cognitive control, and recorded neural activity from the PFC, CN, and MD during task performance. In one task, sequences of stimuli dictated the correct response direction. We identified neural signals in PFC that appeared to encode the occurrence

of rare or unanticipated sensory events occurring at different times within the stimulus sequence. We trained an artificial neural network to perform the task by minimizing sequence prediction error and found that the hidden units exhibited neural activity patterns closely resembling those we found in PFC. This suggested that PFC generated an internal model of the temporal structure of the task, such that breaches of expectation were registered by the activation of PFC neurons. We found that similar temporal expectation signals were present in MD thalamus, indicating that thalamus contributes to the cognitive functions of PFC, and that these contributions extend to the neural representation of temporal expectation. Neural signals encoding the behavioral response were markedly stronger and earlier in MD in comparison to PFC, suggesting that thalamus plays a preeminent role in response selection and execution. In a separate cognitive control task, we found that neural signals reflecting cognitive control were distributed between PFC and the caudate nucleus, including neural signals that reflected cognitive strategy, suggesting that PFC may engage reinforcement learning mechanisms in the basal ganglia to learn cognitive control.

Sam Clarke, USC

Can we see value? Spatiotopic "visual" adaptation to an imperceptible dimension

Spatially indexed adaptation effects are often considered uniquely perceptual. As Block (2022) puts it, no cognitive aftereffect “has ever been shown to be retinotopic or spatiotopic”, hence why spatiotopic adaptation to non-obviously perceptible properties, like number, has been taken to establish that these non-obviously perceptible properties are “primary visual attributes”: represented in vision, alongside color and shape, not just post-perceptual thought. In this talk, I'll challenge these assertions, describing the results of three experiments which found spatiotopic adaptation to a seemingly imperceptible dimension: the cumulative value of arbitrarily valued “coins” in a display. After describing our results, I'll ask whether they should be construed as a case of adaptation to a surprisingly high-level content, or – as I'll proceed to argue – a refutation of the widely held view that spatiotopic adaptation is somehow uniquely perceptual.

This is joint work with Sami Yousif at UNC Chapel Hill

Zac Davis, Utah.

Traveling waves reflect fundamental cortical computations

Fluctuations in waking cortical activity are not independent or synchronous across the cortex but often exhibit structured spatiotemporal patterns that can be described as traveling waves. These waves have been observed in numerous species and cortical systems in various experimental conditions and spatiotemporal scales, however the mechanisms that give rise to them, and what they can tell us about cortical function are subject to debate. Through a combination of modeling and experimental observation, we find traveling wave dynamics can be explained as an emergent property of large-scale synaptic activity coordinated by action potentials traversing the highly recurrent topology of anatomical connectivity that defines cortical networks. The flow of activity through these networks can subserve fundamental cortical computations as they relate to the state of cortical activity in space and time. Consistent with this premise, we find that the states of waves reflect relative levels of excitability across cortical populations and are predictive of sensory-evoked gain and perceptual sensitivity over the representation of visual space and across tuning domains.

Valentin Dragoi, Rice U.

Information processing and coding in cortical circuits during natural behavior.

The long-range goal of my lab is to understand the mechanisms underlying state and experience-dependent changes in the function of cortical populations and how the coordination of distributed networks of neurons influences behavior. To accomplish these goals, we combine electrophysiological (multi-electrode recording in restrained and freely moving non-human primates), optogenetic and electrical stimulation, behavioral approaches, and computational methods. Our basic strategy is to help develop new tools for modulating and recording population activity across cortical circuits in restrained and unrestrained animals and then apply these techniques to examine the neural computations and coding principles across cortical circuits. My seminar will focus on recent work from my lab involving multi-electrode recordings of population activity in visual and prefrontal cortex in restrained and freely moving macaques to examine the neural network underpinnings of natural behavior. This includes information processing and coding principles in cortical networks underlying changes in behavioral performance after sleep and during complex behavior, such as foraging and social interactions.

James Elder, York U.

Fourier analysis of shape perception in humans and machines

Fourier analysis has been a useful tool for vision science: What can it tell us about the perception of shape? Any closed contour in the plane can be linearly recoded by taking the Fourier transform of its coordinate function. In this shape frequency representation, lower harmonics code the coarse features of the shape while higher harmonics code the finer features. Measurements of visual response to stimuli that are low-pass filtered in shape frequency space can thus provide insights into the relative sensitivity to coarse and fine shape features. We measured 9-way classification performance of lowpass-filtered animal shapes by humans and AI object recognition models. While AI models are competitive with humans for unfiltered shapes, analysis of performance for lowpass-filtered shapes reveals striking differences in the information used to achieve this performance. For human observers, performance rises quickly as the lowest harmonics (coarse-scale features) are added to the stimuli, while AI models rely much more on higher frequency shape components (fine-scale features). Both the phase and amplitude of these shape components contribute to recognition, for both humans and AI models. These results suggest that Fourier analysis is a useful tool for understanding limitations of current AI models in accounting for human shape perception.

Michael J Frank, Brown U.

Curriculum effects and compositionality emerge with in-context learning in neural networks

Human learning is sensitive to rule-like structure and the curriculum of examples used for training. In tasks governed by succinct rules, learning is more robust when related examples are blocked across trials, but in the absence of such rules, interleaving is more effective. To date, no neural model has simultaneously captured these seemingly contradictory effects. Here we show that this same tradeoff spontaneously emerges with ``in-context learning" (ICL) both in neural networks trained with metalearning and in large language models (LLMs). ICL is the ability to learn new tasks ``in context" -- without weight changes -- via an inner-loop algorithm implemented in activation dynamics. Experiments with pretrained LLMs and metalearning transformers show that ICL exhibits the blocking advantage

demonstrated in humans on a task involving rule-like structure, and conversely, that concurrent in-weight learning reproduces the interleaving advantage observed in humans on tasks lacking such structure.

Brian Garibaldi, Northwestern U.

Reinvigorating the Bedside Clinical Encounter

As physicians spend less time in direct contact with patients, fundamental bedside skills such as the physical exam are in decline. This decline leads to diagnostic error, increased healthcare costs, an erosion in the doctor-patient relationship and a rise in physician stress and burnout. Despite an increasing reliance on technology, bedside skills remain fundamentally important to clinical diagnosis, and will remain so even in the face of artificial intelligence. We have used 5 strategies to reinvigorate the bedside clinical encounter in partnership with the Society of Bedside Medicine.

1. Practice and teach an evidence-based approach to the clinical encounter.
We need to approach the physical exam just as any other diagnostic test – by using Bayesian probability in a threshold decision making process. By doing so, we can restore clinicians' confidence in high-yield physical exam maneuvers and encourage them to spend time learning accurate and reproducible bedside skills.

2. Create Opportunities for Intentional Practice.
We need to make the most of the limited time we have in direct to contact with patients to model and teach appropriate bedside clinical skills. A number of teaching tools have been developed to aid in this endeavor including the 5-Minute Moment.

3. Use technology to teach/reinforce clinical examination skills.
Point of care technologies such as digital stethoscopes and hand-carried ultrasound (POCUS) enhance diagnostic capabilities at the bedside and allow calibration of traditional physical exam techniques. Perhaps most importantly, they draw learners to the bedside where they engage with patients and learn how other findings on history and physical provide context for an ultrasound exam. Even AI tools will require a human physician to gather and interpret information, and to engage in shared decision making with patients.

4. Seek and provide feedback on clinical skills.
Direct observation and feedback are rare in US graduate medical education which instead relies on summative multiple choice exams to assess clinical competence. If we care about clinical skills we must develop scalable assessments that rely on direct observations of trainees in contact with real patients.

5. Acknowledge the power of the bedside beyond diagnosis.
There is more to the bedside encounter than determining diagnosis and prognosis. The bedside builds human connection and is a source of comfort and fulfilment for patients and physicians alike. We need to remember this important aspect of care as we continue to deploy newer technologies in healthcare delivery.

Indie C. Garwood[1], Alex J. Major[2], John Choi[3], Jonathan Viventi[4], Adam Charles[5], Earl K. Miller[2], Polina Anikeeva[2], Emery N. Brown[2,6], Bijan Pesaran[1]
[1] Penn, [2]NYU, [3]MIT, [4]Duke U. [5]Johns Hopkins, [6]Massachusetts General Hospital.

Characterizing multiregional dynamic responses to neuromodulation with multifunctional Neurotechnology

Alterations of multiregional neuronal communication underlie nearly every neuromodulation technique, from microstimulation to pharmacology. However, studying the impact of neuromodulation on multiregional neural processing presents a major experimental challenge. Standard electrophysiology approaches cannot probe neurochemical alterations and are inherently limited by sampling only a small portion of the brain. In my doctoral research, I developed multifunctional fiber-based neurotechnology capable of probing cortical and deep brain neurochemical signaling in non-human primates. In conjunction with multiregional electrophysiology, I applied this technology to study the mechanisms of ketamine anesthesia. Recent research indicates that disruption of multiregional communication with anesthetics results in loss of consciousness but compared to other anesthetics, ketamine has divergent effects that have eluded mechanistic description. High dose ketamine results in a state of dissociative anesthesia, as well as a dramatic oscillatory cortical signature characterized by bursts of gamma power alternating with high amplitude slow waves. We hypothesized that "gamma bursts" reflect a systemic restructuring of neuronal communication, which disrupts sensory processing and ultimately leads to loss of consciousness. We recorded bilateral, multi-area, laminar intracranial electrophysiology in 6 rhesus macaques following 10-20 mg/kg IM ketamine boluses. We found that ketamine results in rigid, coordinated structure across all recorded brain areas, and that ketamine disinhibition in part results from elevated cholinergic tone. Despite ketamine's disruption of cortical communication, we surprisingly found that both sensory and cognitive processes were maintained in some, but not all animals. Consistent with clinical reports of awareness during ketamine anesthesia, our research provides insight into cortical biomarkers that can be readily observed with intraoperative EEG.

A key goal of characterizing the mechanisms of anesthesia is to inform more targeted neuromodulation approaches. Cortical microstimulation has shown increasing promise for modulating multiregional neural circuits to achieve significant clinical benefit. Understanding the mechanistic relationship between cortical microstimulation and large-scale neural dynamics is critical to optimizing stimulation therapies. In my postdoctoral work, I am studying the relationship between cortical microstimulation and concurrent cellular calcium imaging, a high resolution recording modality unaffected by stimulation artifacts. Simultaneous recordings provided precise measures of response latency, revealing that the recorded response to stimulation was predominantly mediated by synaptic mechanisms. Our results delineate the circuit level responses to varying stimulation protocols, including monopolar, bipolar, and patterned stimulation. Our future work will focus on optimizing electrical stimulation protocols for targeted modulation of cortical circuits.

Justin Halberda, Johns Hopkins.

If there is no such thing as a JND, then why do all people believe that there is such a thing? AND Most thoughts are magnitude thoughts.

The construct of a Just Noticeable Difference (JND) was introduced by Weber and Fechner over 160 years ago. It captures the idea that when two stimuli become too similar (e.g., too close in weight, brightness, loudness, etc.) we can no longer distinguish them (i.e., we will not be able to report which is heavier, brighter, louder, etc. better than chance). An alternative approach, from Signal Detection Theory,

suggests that discriminability is limitless, with no JND, and that performance will never drop to chance, even for the most difficult comparisons. I review the results of a series of experiments showing that both naïve adults and undergraduate Psychology students hold beliefs consistent with the existence of a JND. And, I review evidence across the five senses from human adults, children, individuals blind from birth, rats, mice, seals, bats, and butterflies that consistently shows that there is no such thing as a JND. In each case, subjects' discrimination performance never drops to chance, even for the most difficult comparisons. In short, while intuitively appealing, the JND does not exist, for any mental dimension. I suggest that people are sympathetic to the notion of a JND because it is consistent with our naïve intuitive theory about how perception works.

Monika Jadi, Yale.

Brain state and visual stimulation differentially modulate inter-layer communication subspace in V1

Visual perception emerges from intricate interactions among neuronal populations across the brain's sensory processing hierarchy. Neuronal correlations in the visual cortex can be modulated by both external visual stimuli and internal states. Recent studies emphasize the role of communication subspaces as a population-level routing mechanism that serve as "channels" for transmitting information, facilitating coordinated activity across brain areas (e.g., V1-V2) or between cortical layers (e.g., input-superficial). It is however not known if this co-ordinated activity has directed or shared origins. In this study, we investigated the inter-laminar communication in macaque V1 neural populations. Our findings reveal that layers interact through a low-dimensional subspace, whose structure is modulated by both visual stimuli and wakefulness states (eyes open/closed). Additionally, efficacy of the subspace is enhanced by visual stimulation and by eye-closing in the spontaneous period. Notably, the efficacy modulation due to visual stimulation is more prominent in the cross-layer direction, whereas efficacy modulation in the spontaneous period shows little direction preference. Furthermore, we found that evoked communication between input and superficial layers exhibits a time lag of approximately 10 milliseconds, which is absent without visual input. We conclude that co-ordinated activity during visual stimulation is due to directed communication, while that during wakefulness states of the spontaneous period is likely due to shared inputs.

Michael Landy, NYU

The processing of 3D motion in the human visual cortex

I will discuss three studies that strive to delineate the processing of three-dimensional motion in the human brain. (1) The first study uses a random-dot display of stimuli moving in various directions in the horizontal plane (toward, rightward, away, leftward and oblique directions). The binocular stimuli had changing disparity and inter-ocular velocity differences. Control stimuli moved vertically, containing similar motion energy but no 3D motion percept. Areas hMT and IPS0 yielded responses to 3-D motion over and above the raw motion signals in the control stimuli. (2) In monkey, a key area in the analysis of 3D motion is area FST, but no clear method for isolating a human homolog has been found. We used several approaches to isolate putative human FST (pFST) including 2D and 3D motion, population receptive fields, motion opponency (paired dots) and measures of myelination. We observed consistent anterior and inferior activation relative to hMT and MST in response to stimuli that contained coherent 3D, but not 2D, motion. Motion opponency and myelination measures further validated the functional and structural distinction between pFST and hMT/MST. This provides a robust framerwork for localizing pFST in human and demonstrating its importance for the analysis of 3D motion. (3) Using 3D-motion

stimuli containing monocular and/or binocular cues, we investigated the hierarchy within the human motion complex to understand the neural mechanisms underlying motion perception. On each trial we decoded 3D motion direction (towards/away) based on the BOLD response in primary visual cortex (V1), and areas MT, MST, and FST within hMT+. We found that 3D-motion direction could be reliably decoded from all four areas, but with distinct patterns of cue preference. MT showed greatest accuracy with monocular cues, whereas FST showed greatest accuracy with binocular cues. While motion direction could also be decoded in V1 and MST, these results could be explained based on retinotopic variation in the BOLD response that depended on motion direction. In contrast, MT and FST were less impacted by retinotopic biases in the BOLD response. A subset of participants had near-chance performance on the 3D task with binocular cues alone. Good behavioral performance with binocular cues was accompanied by better decoding performance in FST but not in MT. A control experiment that eliminated 3D-motion percepts for binocular stimuli, but not monocular stimuli, also revealed that, unlike MT, decoding accuracy in FST was influenced by perceptual components of 3D motion.

Julio Martinez-Trujillo, Western University, London, Canada

The marmoset Hippocampus is a GPS, but G is for gaze

The role of the hippocampus in spatial navigation has been primarily studied in nocturnal mammals, such as rats, that lack many adaptations for daylight vision. Here we demonstrate that during 3D navigation, the common marmoset, a new world primate adapted to daylight, predominantly uses rapid head-gaze shifts for visual exploration while remaining stationary. During active locomotion marmosets stabilize the head, in contrast to rats that use low-velocity head movements to scan the environment as they locomote. Pyramidal neurons in the marmoset hippocampus CA3/CA1 regions predominantly show mixed selectivity for 3D spatial view, head direction, and place. Exclusive place selectivity is scarce. Inhibitory interneurons are predominantly mixed selective for angular head velocity and translation speed. Finally, we found theta phase resetting of local field potential oscillations triggered by head-gaze shifts. Our findings indicate that marmosets adapted to their daylight ecological niche by modifying exploration/navigation strategies and their corresponding hippocampal specializations.

Cory Miller, UCSD.

Neural mechanisms for solving Cocktail Party Problem

Resolving the Cocktail Party Problem (CPP) is amongst the most significant challenges for the auditory system. While most research efforts to explicate the underlying neural mechanisms have focused on cortical substrates, notably prefrontal and auditory cortex, the pervasiveness of this fundamental challenge for audition since this sensory system emerged early in vertebrate evolution suggests that subcortical substructures may play a key role. Here we present evidence that mechanisms for resolving the CPP are evident in both the Anterior Cingulate Cortex and Hippocampus in marmoset monkeys. Notably, we examined this issue while freely-moving animals engaged in their natural behavior suggesting that naturalistic experimental paradigms may be integral to resolving the CPP.

Anirvan Nandy, Yale

Behavioral & neural dynamics of cooperation in marmoset dyads

Advanced social cognition enables individuals to apply complex and flexible strategies when interacting with others. Cooperative strategies may be at the apex of the complex social strategies used by highly social species. Common marmosets display strong prosociality and high social tolerance, making them a great model for studying cooperative interaction mechanisms. Here, we used the Marmoset Apparatus for Automated Pulling (MarmoAAP) to examine the behavioral dynamics and strategies that marmoset dyads used to facilitate mutual cooperation. Freely moving marmoset dyads spontaneously cooperated for rewards at a high rate after showing converging learning patterns over time. Cooperative performances depended on social relationships, including dominance, kinship, and sex. Crucially, control sessions without social visual information or with automatically moving partner's lever demonstrated the necessity of social vision in successful cooperation. Dynamic causal dependencies between social gaze and pulling behaviors at multiple timescales, captured from markerless head-frame tracking and dynamic Bayesian network modeling, generated quantitative representations of employed strategies. Marmoset dyads dynamically used both gaze-independent and gaze-dependent strategies to facilitate mutual cooperation. Remarkably, cooperative strategies were highly flexible, where the same marmosets implemented distinct adaptive strategies depending on social partners. These findings demonstrate diverse and flexible behavioral patterns used in mutually cooperative exchanges in common marmosets, providing compelling evidence of sophisticated social interaction by a species evolutionarily distant from humans. Preliminary neural data from the medial prefrontal cortex (mPFC) suggests that distinct neural trajectories in mPFC population state space represent distinct task-related strategies.

Ximena J Nelson, University of Canterbury

Cognition in jumping spiders

I will present an overview of my work on cognition in jumping spiders (Salticidae), noting that these animals are excellent models for studies on cognition due to their visual, active, and predatory lifestyle. Due to their need to leave a nest to explore their complex surrounds for food and mates, which are detected and classified visually at a distance, these animals exhibit not only navigational prowess, but also spatial ability that seems to include forward-planning in their decision-making. The sophistication of their behaviour is like that seen in some mammals, with a fraction of the 'brain power'. It is for this reason that salticids allow us to explore, in depth, hypotheses pertaining to the evolution of cognition, and the trade-offs between speed and accuracy that animals make when making decisions. I argue that the ability to accurately detect objects from a distance in real time is a critical selection factor leading to complex problem-solving ability, and that vision is unique in this respect.

Jean-Paul Noel, U. Minnesota

Aberrant updating of internal models in Autism Spectrum Disorder.

Both the phenotypes and biological bases of autism spectrum disorders (ASD) are heterogenous, limiting our ability to understand and ultimately treat the condition. To face this diversity, recent years have seen the advent of computational psychiatry, wherein researchers attempt to broadly ascribe varied phenotypes to a set of underlying computations that may have gone awry in disease. Inspired by this approach, I will first suggest that ASD is characterized by an impairment of causal inference – the domain general ability to ascribe causes (e.g., an interlocutor) to observed data (e.g., auditory speech and visual mouthing). Next,

I will demonstrate that three different monogenetic mouse models of ASD (Fmr1, Cntnap2, Shankb3) show a blunted update of their expectations (i.e., “Bayesian priors”), mimicking the finding of inflexible internal models in humans with ASD. To understand the biological basis of this anomaly, I performed a brain-wide survey of extra-cellular single neuron activity with silicon probes (150+ brain areas, 53k+ neurons) in each of these mouse models of ASD. The differentiating factor between animals flexibly (wildtype) and inflexibly (ASD) updating their priors was a shift in the weighting of prior encoding, form sensory to frontal cortices. Further, in mouse models of ASD frontal areas showed a preponderance of units coding for deviations from the animals’ long-run prior, and sensory responses did not differentiate between expected and unexpected observations. These findings establish a common computational and neural deficit across different genetic mouse models of ASD. More broadly, they demonstrate a degree of biological convergence: different genetic perturbations may result in a common neurophysiological and computational account of ASD.

Franco Pestilli, U. Texas.

Putting big data to good use in neuroscience

Neuroscience is experiencing unprecedented growth in dataset size both within individual brains and across populations. Large-scale, multimodal datasets are transforming our understanding of brain structure and function, creating opportunities to address previously unexplored questions. However, managing this increasing data volume requires new training and technology approaches. Modern data technologies are reshaping neuroscience by enabling researchers to tackle complex questions within a Ph.D. or postdoctoral timeframe. I will discuss cloud-based platforms such as brainlife.io, that provide scalable, reproducible, and accessible computational infrastructure. Modern data technology can democratize neuroscience, accelerate discovery and foster scientific transparency and collaboration. Concrete examples will illustrate how these technologies can be applied to mapping brain connectivity, studying human learning and development, and developing predictive models for traumatic brain injury (TBI). By integrating cloud computing and scalable data-sharing frameworks, neuroscience can become more impactful, inclusive, and data-driven.

Jonathan Phillips, Dartmouth.

Domain-general modal cognition.

Much of high-level cognition relies on the ability to determine what the relevant possibilities are in a particular situation. To judge that someone is morally responsible for a given action requires assessing what other actions were available to that person. To decide what caused something to happen requires determining what else could have happened instead. An important but unanswered question is whether each different high-level judgement relies on a domain-specific representation of the possibilities relevant for that judgement, or whether humans have a domain-general way of representing possibilities that is recruited across various forms of high-level cognition. Here, we provide evidence for the latter hypothesis. In this talk, I'll introduce a general method for empirically measuring the set of possibilities that people consider to be relevant in a particular situation (through sequential sampling), quantitatively characterize what this set of possibilities looks like, and show that the same representation of possibility is recruited across distinct forms of high-level reasoning, namely force judgments, causal judgments and blame attribution. I'll then illustrate that this domain-general modal representation of possibilities has a particularly interesting set of features: it is both highly flexible (reflecting even relatively small changes in context), and it also seems to be represented in a compact, heuristic manner, rather than being computed on the fly.

Zyg Pizlo, UC Irvine

A guide through symmetry: math, physics, perception

There were two fathers of psychophysics in the 19th Century: Fechner and Mach. Fechner viewed perception as measurement, whereas Mach viewed perception as an inference. Here, I elaborate Mach's approach. He pointed out the role of invariance and symmetry in perception. He also pointed out the role of simplicity principle, known in physics as a least-action principle. Symmetry has been established by Klein as the foundation of Geometry. Symmetry was also recognized as fundamental in physics by Einstein and Noether. In visual perception, redundancy aspect of symmetry becomes essential. The talk makes two new observations: First, 3D shape perception is a form of a conservation law analogous to conservations in physics. Second, all forms of depth cues represent the operation of symmetry.

Nicholas Port, Indiana U.

Diffusion Tensor Analysis of White Matter Tracts is Prognostic of Persisting Post-Concussion Symptoms in Collegiate Athletes

Background and objectives: After a concussion diagnosis, the most important issue for patients and loved ones is how long it will take them to recover. The main objective of this study is to develop a prognostic model of concussion recovery. This model would benefit many patients worldwide, allowing for early treatment intervention.

Methods: The Concussion Assessment, Research and Education (CARE) consortium study enrolled collegiate athletes from 30 sites (NCAA athletic departments and US Department of Defense service academies), 4 of which participated in the Advanced Research Core, which included diffusion-weighted MRI (dMRI) data collection. We analyzed the dMRI data of 51 injuries of concussed athletes scanned within 48 h of injury. All athletes were cleared to return-to-play by the local medical staff following a standardized, graduated protocol. The primary outcome measure is days to clearance of unrestricted return-to-play. Injuries were divided into early (return-to-play <28 days) and late (return-to-play >=28 days) recovery based on the return-to-play clinical records. The late recovery group meets the standard definition of Persisting Post-Concussion Symptoms (PPCS). Data were processed using automated, state-of-the-art, rigorous methods for reproducible data processing using brainlife.io. All processed data derivatives are made available at https://brainlife.io/project/63b2ecb0daffe2c2407ee3c5/datasset. The microstructural properties of 47 major white matter tracts, 5 callosal, 15 subcortical, and 148 cortical structures were mapped. Fractional Anisotropy (FA) and Mean Diffusivity (MD) were estimated for each tract and structure. Correlation analysis and Receiver Operator Characteristic (ROC) analysis were then performed to assess the association between the microstructural properties and return-to-play. Finally, a Logistic Regression binary classifier (LR-BC) was used to classify the injuries between the two recovery groups.

Results: The mean FA across all white matter volume was negatively correlated with return-to-play ($r=-0.38$, $p=0.00001$). No significant association between mean MD and return-to-play was found, neither for FA nor MD for any other structure. The mean FA of 47 white matter tracts was negatively correlated with return-to-play ($r_\mu = -0.27$; $r_\sigma = 0.08$; $r_{min} = -0.1$; $r_{max} = -0.43$). Across all tracts, a large mean ROC Area Under the Curve (AUCFA) of 0.71 ±0.09 SD was found. The top classification performance of the LR-BC was AUC=0.90 obtained using the 16 statistically significant white matter tracts.

Discussion: Utilizing a free, open-source, and automated cloud-based neuroimaging pipeline and app (https://brainlife.io/docs/tutorial/using-clairvoy/), a prognostic model has been developed, which predicts athletes at risk for slow recovery (PPCS) with an AUC=0.90, balanced accuracy=0.89, sensitivity=1.0, and specificity=0.79. The small number of participants in this study (51 injuries) is a significant limitation and supports the need for future large concussion dMRI studies and focused on recovery.

Merri Rosen, NEOMED (Northeast Ohio Medical University)

Characterizing behavioral effects of early-life stress in an animal model of auditory processing

Animal models provide significant insight into the development of typical and disordered sensory processing. The Mongolian gerbil (Meriones unguiculatus) is a well-established model for auditory processing, with a hearing range similar in frequency to that of humans and an easily accessible cochlea. Development is a time period of heightened plasticity and susceptibility to environmental insults. Thus developmental hearing loss during postnatal maturation causes auditory processing deficits, in both gerbils and humans.

An additional early insult that may affect auditory processing is postnatal early-life stress (ELS), which acts on critical period elements known to be the source of auditory deficits following early hearing loss. ELS is well-known to impact cognition and higher neural regions (e.g., amygdala, hippocampus, prefrontal cortex), but has barely been investigated for effects on sensory processing. We have thus developed the Mongolian gerbil as a sensory model for early-life stress, evaluating auditory processing deficits.

This presented a challenge. Measuring deficits in auditory perception – e.g., using operant conditioning – involves behavioral elements such as attention, memory, learning, and emotion. All of these elements are well-known to be impacted by ELS, and may affect behavioral measurements in ways that could be misconstrued as sensory deficits. Thus we characterized which behavioral elements are affected by ELS in our gerbil model.

We induced ELS during maturation of the auditory cortex (postnatal (P) days 11-20), and measured behavior in juveniles (~P30). ELS had no effect on overall locomotor activity but reduced anxiety-related behavior, impaired recognition memory, and improved spatial memory, with some sex-specific effects. In a passive avoidance task involving aversive conditioning, ELS improved learning but reduced memory retention. These effects are likely to influence the ability of gerbils to learn and retain operant training used to evaluate auditory perception, particularly if anxiety-provoking reinforcement is used. We thus used gap-inhibition of the acoustic startle response (GPIAS) as an automatic measure. GPIAS does not involve training and minimizes the influence of attention. We assessed behavioral sensitivity to brief gaps in noise, which is critical for speech perception. ELS gerbils had higher behavioral thresholds for detecting short gaps. In parallel, auditory cortical neurons had higher gap detection thresholds, and a population ideal observer model of auditory cortical responses misclassified neural responses to short gaps. We are following up by examining developmental windows of susceptibility to ELS, and neural mechanisms by which auditory processing is disrupted.

Adina Roskies, UCSB

Free will and the Readiness Potential

The question of whether or not we have free will has long puzzled philosophers. In the last 50 years, some neuroscientists have claimed to be able to weigh in on the problem, and they typically argue that their evidence undermines free will. I begin by reviewing the landscape of free will, and by discussing one of the ways in which neuroscientists have tried to provide evidence that we lack it. I sketch the evidence from a category of EEG studies originally performed by Libet, which purport to show that the readiness potential begins before awareness of decision. I argue that new ways of interpreting that evidence undermines the conclusions that the classic view reaches, and explain why the new interpretation is not a threat to conscious free will.

Thomas Serre, Brown.

Visual perspective taking in machines vs. humans

Visual perspective taking (VPT) is the ability to perceive and reason about the perspectives of others. It is an essential feature of human intelligence, which develops over the first decade of life and requires an ability to process the 3D structure of visual scenes. A growing number of reports have indicated that deep neural networks (DNNs) become capable of analyzing 3D scenes after training on large image datasets. We investigated if this emergent ability for 3D analysis in DNNs is sufficient for VPT with the 3D perception challenge (3D-PC): a novel benchmark for 3D perception in humans and DNNs. We tested human participants (N=33) and linearly probed or text-prompted over 300 DNNs on the challenge. We found that nearly all DNNs approached or exceeded human accuracy in analyzing object depth order. Surprisingly, DNN accuracy on this task correlated with their object recognition performance. In contrast, there was an extraordinary gap between DNNs and humans on VPT-basic. Our challenge demonstrates that the training routines and architectures of today's DNNs are well-suited for learning basic 3D properties of scenes and objects but are ill-suited for reasoning about these properties like humans do.

Laura Soter, York U.

Investigating Children’s Theory of Belief Change: The Role of Prior Belief Strength

Much prior work shows that young children have theories of other agents as rational actors, and also that children’s own belief updating practices are highly rational. But do children also have theories of other agents as rational believers? We investigate how children ages 4-8 understand the role of prior belief strength and novel evidence in belief updating. In Study 1, we find that by age 5, children judge it to be more difficult to change a strongly held belief than a weakly held one, regardless of the domain of that belief. In Study 2, we find that children use information about prior belief strength to inform their predictions of how a character will update their belief in response to counterevidence. These studies offer evidence that children understand that belief updating is crucially dependent on the strength of one’s initial belief.

Kari Theurer, Trinity College

On the role of development and environment in explanations of psychiatric disorder

Psychiatry, as a branch of medicine, has two intertwined goals: to explain and to treat behaviors and mental states that the profession regards as disordered, and which often cause considerable distress to patients. Those behaviors and mental states cluster together into characteristic patterns that constitute the familiar taxonomy of psychiatric disorders outlined by psychiatry's major diagnostic manuals, the DSM and the ICD. The psychiatric profession's approach to explaining disorder is at present predominantly causal/mechanical: explanation primarily involves elucidating the neurobiological mechanisms that generate psychiatric symptoms. Those explanations, in turn, presumably yield effective interventions into the cognitive and affective systems that produce those symptoms. And yet, mounting evidence suggests that this research program has stalled. There is no shortage of studies attempting to link disorders to dysfunctional neurobiological mechanisms, yet many of the most common and debilitating psychiatric disorders remain stubbornly resistant to the pharmacological interventions that the profession has championed in recent decades. The result is a profession in crisis, which is reflected by its recent ambitions to rebuild its taxonomy by starting with neurobiological mechanisms rather than symptom clusters. I argue that this focus on neurobiological mechanisms necessarily excludes information about development and environment that is essential to the goal of understanding and treating psychiatric symptoms. Real understanding of psychiatric symptoms – the kind that leads to meaningful and sustainable improvement in the life of the person – requires more than an explanation of the proximate brain mechanisms that produce those symptoms. It requires information about the history and environment of the person that those mechanisms do not and cannot capture. Were psychiatry to refocus its approach accordingly, the result would be a shift much more radical than a mere reconfiguration of taxonomy; indeed, it would require interventions that extend far beyond the brain.

Jessica Witt, Colorado State.

Ensemble Perception to Improve Decisions under Uncertainty

The visual system can easily extract statistics from a visual scene. We leveraged this visual ability to create data visualizations that improve decision making over current graphics that are used to convey hurricane forecasts. Our new visualization is called ARTs (Animated Risk Trajectories). The current product (the Cone of Uncertainty) leads to a bias called the Containment Effect for which people see all areas located within the cone as being at risk and all areas beyond the cone's boundary as being safe. Furthermore, the ARTs can convey additional information by using visual annotations such as color to communicate hurricane risks like storm surge.