

SAND8 POSTER PRESENTATIONS

- (1) Abbasi Asl, Reza [Deep Nets Meet Real Neurons: Pattern Selectivity in V4](#)
- (2) Aksenov, Alexandre [Group-wise feature selection using L1- and L0 Penalized N-way PLS in Brain Computer Interface study](#)
- (3) Ashmankamachali, Maboudi [Identifying memory-related temporal structures of neural data during rest periods using Hidden Markov Models](#)
- (4) Bar-gad, Izhar [Generalized linear model of basal ganglia neurons](#)
- (5) Bonner, Mick [Coding of navigational affordances in the human visual system](#)
- (6) Brodbeck, Christian [Neural source dynamics of brain responses to continuous stimuli with MEG: speech processing from acoustics to comprehension](#)
- (7) Bruett, Heather [Measure for measuring synchrony in the processing of visual information](#)
- (8) Carlos, Brandon [Testing a new method for assessing hemispheric lateralization using multi-voxel pattern analysis](#)
- (9) Fujita, Kazuki [Inferring the origin of nonstationary fluctuation in event occurrences](#)
- (10) Gerhard, Felipe [All-optical electrophysiology for neuroscience drug discovery](#)
- (11) Glaser, Joshua [Machine learning tools for neural decoding](#)
- (12) Greene, Patrick [A Bayesian approach to probabilistic spike detection and sorting](#)
- (13) Lisenbardt, David [Population encoding of the 'intent' to drink alcohol](#)
- (14) Morina, Rudina [Relationship between pairwise correlation and dimensionality reduction](#)
- (15) Mukherjee, Narendra [Taste-related cortical population dynamics are stochastic and behaviorally-relevant](#)
- (16) Narayan, Manjari [Statistical Methods to Investigate Interventional Changes to Interaction Networks](#)
- (17) Oswal, Urvashi [Inferring brain activity through pairwise similarities](#)
- (18) Sadagopan, Srivatsun [Optimal features for auditory recognition](#)
- (19) Silva Simões, Lucas [Mean Field studies of a society of Bayesian agents](#)
- (20) Timme, Nicholas [Improvements to information theory analysis techniques throughout neuroscience with MATLAB support](#)
- (21) Venkatesh, Praveen [Fundamental limits, algorithms, and instrumentation for novel non-invasive and minimally-invasive "ultra-resolution" EEG systems](#)
- (22) Vinci, Giuseppe [Adjusted Regularization of Cortical Covariance](#)
- (23) Williams-Garcia, Rashid [Unveiling the causal structure of living neural networks](#)
- (24) Zhang, Lijun [Dimensionality reduction schemes for understanding inter-trial neural response variations and their role in sensory coding](#)
- (25) Zhang, Qiong [Inter-Subject Alignment of MEG Datasets in a Common Representational Space](#)

Deep Nets Meet Real Neurons: Pattern Selectivity in V4

Reza Abbasi Asl
University of California, Berkeley
abbasi@berkeley.edu

Vision in humans and in non-human primates is mediated by a constellation of hierarchically organized visual areas. One important area is V4, a large retinotopically-organized area located intermediate between primary visual cortex and high-level areas in the inferior temporal lobe. V4 neurons have highly nonlinear response properties. Consequently, it has been difficult to construct quantitative models that accurately describe how visual information is represented in V4. To better understand the filtering properties of V4 neurons we recorded from 71 well isolated cells stimulated with 4000-12000 static grayscale natural images. We fit predictive models of neuron spike rates using transformations of natural images learned by a convolutional neural network (CNN). The CNN was trained for image classification on the ImageNet dataset. Our technique falls in the class of transfer learning methods. To derive a model for each neuron, we first propagate each of the stimulus images forward to an inner layer of the CNN. We use the activations of the inner layer as the feature (predictor) vector in a high dimensional regression, where the response rate of the V4 neuron is taken as the response vector. Thus, the final model for each neuron consists of a multilayer nonlinear transformation provided by the CNN, and one final linear layer of weights provided by regression. We find that models using the intermediate layers of three well-known CNNs provide better predictions of responses of V4 neurons than those obtained using a conventional Gabor-like wavelet model. We discover that the V4 neurons are tuned to a remarkable diversity of shapes such as curves, blobs, checkerboard patterns, and V1-like gratings. To arrive such results, we introduce new processes for interpreting our CNN based models. To characterize the spatial and pattern selectivity of each V4 neuron, we both explicitly optimize the input image to maximize the predicted spike rate, and visualize the selected filters of the CNN. Then we apply sparse PCA to visualize the diverse tuning properties of the whole population. Finally, we enhance the reliability of our results via stability analysis across different CNN structures.

This is a joint work with Yuansi Chen, Adam Bloniarz, Michael Oliver, Jack L. Gallant, and Bin Yu.

Group-wise feature selection using L1- and L0 Penalized N-way PLS in Brain Computer Interface study

Alexander Aksenov

alexander1aksenov@gmail.com

Multi-way (tensor based) analysis was reported recently as an effective tool for neural data processing. The advantage of this approach is the simultaneous treatment of data in several domains (ways of analysis) to improve information extraction. Spatial, frequency and temporal modalities are mostly considered. Extracted features are represented in a form of multi-way arrays (tensors). Among others, tensor methods have been applied for decoding limb movement trajectory from brain neural activity Electrocorticography (ECoG) recording. Decoding model was applied in real time in BCI experiments. Multiway analysis generally results in extremely high dimensional feature space. In the same time, real-time applications require efficient computations. Sparsification of decoding model is desirable.

Penalization is widely used for sparsification of the solution in regression analysis, classification, and other approaches of data analysis. L1 and L0 penalization terms are mainly considered for individual feature selection. Group-wise feature selection is less studied. At the same time, it is important in numerous applications including neural data processing. The particular example is real-time multi-frequency analysis of multi-electrodes recordings of neuronal activity. In this case, the features should be selected/excluded from the model by groups (e.g. all features related to a given electrode or/and to a given frequency) for the efficient computation. L1/L0 penalization can be applied for sparse tensor factorization. Sparse factors allow slice-wise feature section.

In the presentation L1- and L0-Penalized NPLS algorithms are considered for sparse tensor factorization and for group-wise informative feature selection. The cases of informative electrodes and informative frequency bands selections are studied and tested for the particular task of hand trajectory reconstruction from Electrocorticography (ECoG) recording.

Joint work with: Fabien Boux, Andrey Eliseyev, Guillaume Charvet, Alim-Louis Benabid, Tetiana Aksenova

Identifying memory-related temporal structures of neural data during rest periods using Hidden Markov Models

Kourosh Maboudi Ashmankamachali
University of Wisconsin-Milwaukee
Kourosh@uwm.edu

Finding about neuronal network activities and their dynamics have been an intriguing topic in the field of systems neuroscience, particularly study of brain memory mechanisms. When an animal is exposed to a novel environment, the naive network structure may change to encode the new encountered information. As an example, when a rat runs on a linear track, multiple assemblies are formed in the network of CA1/3 neurons, representing different locations (place fields) on the track. Activation order of these assemblies is concordant with the temporal order of locations that the animal passes during running. Moreover, the same sequence, but in more compressed time scales, occurs during offline rest or sleep periods (replays). The short time scale of these events makes them as candidate mechanisms for memory formation through STDP rules.

Identifying such activity patterns during offline periods is challenging regarding that no correlation with the simultaneous behavior exists. So far, methods were used for detection of replays based on measuring similarity with the activity during running. However, the total offline activity patterns are not limited to the replays. Patterns may exist due to some unexpected features of the environment, remote behaviors of the animal, and etc. Moreover, even the replay events may alter due to changes in the brain state or with the passage of time, which makes them harder to identify using the conventional methods. Therefore, an unsupervised method that identifies the patterns directly from the offline data itself is of special importance.

Recently, methods based on hidden Markov models (HMMs) have been used to model the activity during behavior. In these models, it is assumed that the network activity consistently transitions between a number of latent states. The states, transitions between them and their correspondence to neuronal activities are learned directly from data. We trained HMMs on the data during offline periods and tested the models with some criteria. First, we found that a model trained on well-structured data contains relatively low amount of randomness with transitions between the states. Second, we reasoned that as the offline period data usually contain lots of replays, as reported previously, at least some states should resemble the place fields on the track. We realized that this was the case, although there was not a one-to-one relationship between the states and the place fields. Finally, we measured the degree to which our model is able to explain individual replay events detected using some standard method. For majority of the replays this was the case. Moreover, we found that the model comparatively is more robust to false negatives. Joint work with Kamran Diba

Invariant statistical modeling of neuronal encoding

Izhar Bar-Gad

Bar-Ilan University, Ramat-Gan, Israel

izhar.bar-gad@biu.ac.il

The basal ganglia (BG) have an important role in the processing of motor, cognitive and limbic information through their reciprocal connections with the cerebral cortex. Abnormal BG activity has been related to a variety of severe psychiatric and movement disorders, such as Parkinsons disease. Characterizing the computational properties of individual BG neurons is crucial for understanding their contribution to normal brain function and its breakdown during different pathologies. We used a generalized linear model (GLM) to quantify the differences in the encoding of individual neurons by incorporating a linear stimulus filter, a spike history filter, and a bias term. The model was generated by fitting parameters to the in-vitro whole-cell responses of the neurons to repeated stimulation (frozen noise). These models accurately reproduced the responses of the experimentally recorded cells, however, the GLM parameters were found to be highly sensitive to the internal state of the neurons, such as the dependency on the baseline firing rates. Assuming that the underlying computation of the neuron is independent of the baseline firing rate, these different parameter sets deviate from the “real” encoding properties of the neuron. During intracellular recordings in the acute brain slices, the baseline firing rate of the neuron is manually tuned, by the level of a constant injected current, which simulates neuronal activity in the absence of external stimuli. Thus, the firing rate is arbitrary in in-vitro experiments; moreover, even during in-vivo experiments the firing rate typically fluctuates over multiple time scales, such as the variation of the rate between behavioral states. In order to characterize the inconsistency in the derived parameters, and to allow the extraction of an unbiased statistical model of the neuron, we use a combination of data derived from experimentally recorded neurons, and simulations of simple and compartmental neurons. Using this data we demonstrate the relation of the GLM parameters to the firing rate of the neuron and assess ways to deal with this variability and to create firing rate invariant models of the neurons. GLMs provide an exciting approach to modeling neurons; however, to utilize the full potential of this model, addressing potential caveats arising from the effect of the experimental properties should be carefully addressed.

Joint work with: Ayala Matzner, Lilach Gorodetzki, Alon Korngreen

Coding of navigational affordances in the human visual system

Michael F. Bonner

Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104, USA

michaфра@mail.med.upenn.edu

A central component of spatial navigation is determining where one can move within the immediate environment. For example, in indoor environments, walls limit ones potential routes, while passageways facilitate movement. In a set of fMRI experiments, we found evidence that the human visual system solves this problem by automatically identifying the navigational affordances of the local visual scene. Specifically, we found that the occipital place area (OPA), a scene-selective region of dorsal occipitoparietal cortex, automatically encodes the navigational layout of visual scenes, even when subjects are asked to perform perceptual tasks that are unrelated to navigation. These effects were found in two experiments: one using tightly controlled, artificially rendered scenes, the other using natural images of complex, real-world environments. A reconstruction analysis further demonstrated that the population codes of the OPA could be used to predict the affordances of novel scenes. Given the apparent automaticity of this process, we predicted that affordance identification could be rapidly achieved through a series of purely feedforward computations performed on retinal inputs. To test this prediction, we examined visual scene processing in a biologically inspired deep convolutional neural network (CNN) with a feedforward architecture. This CNN was trained for scene categorization, but previous work has suggested that its internal representations are general-purpose and transfer well to other scene-related tasks. We found that the CNN contained information relating to both the neural responses of the OPA and the navigational affordances of scenes.

This information arose most prominently in higher convolutional layers, following several nonlinear feature transformations. By probing the internal computations of the CNN, we found that the coding of navigational affordances relied heavily on visual features at high-spatial frequencies and cardinal orientations, both of which have previously been identified as low-level stimulus preferences of scene-selective visual cortex. These computations also exhibited a strong preference for information in the lower visual field, which is consistent with known retinotopic biases in the OPA. Visualizations of internal features from the CNN provide further insights into how affordance computation is achieved in the OPA. Together, these results reveal a previously unknown mechanism in the human visual system for perceiving the affordance structure of navigable space and they demonstrate the feasibility of encoding these affordances in a single forward pass through a hierarchical computational architecture.

Joint work with Russell A. Epstein

Neural source dynamics of brain responses to continuous stimuli with MEG: Speech processing from acoustics to comprehension

Christian Brodbeck

Institute for Systems Research, University of Maryland, College Park, Maryland

brodbeck@umd.edu

Reverse correlation of EEG and MEG data has been used to analyze neural processing of continuous stimuli such as speech, but the analysis is typically restricted to sensor space. We show that reverse correlation can be combined with source localization of MEG data to estimate the neural response to continuous signals in time as well as anatomical location. We first compute distributed minimum norm current source estimates for continuous MEG recordings, and then compute temporal response functions for the estimate at each neural source element, using the boosting algorithm with cross-validation. Permutation tests can assess significance of individual predictor variables as well as features of the corresponding spatio-temporal response functions. We demonstrate the viability of this new method by computing spatio-temporal response function for speech stimuli, using predictor variables reflecting different cognitive levels of speech processing. We show that processes related to comprehension of continuous speech can be differentiated anatomically: acoustic and lexical information are associated with responses in the posterior superior temporal gyrus, in the vicinity of auditory cortex, while semantic composition is associated with responses in classical higher level language areas, anterior temporal lobe and inferior frontal gyrus. This method can be used to study the neural processing of continuous stimuli in time and anatomical space and suggests new avenues for analyzing neural processing of continuous stimuli. This is especially relevant for language comprehension research, where event-related designs may heavily compromise the naturalness of the stimuli. To facilitate the use of this method we make the algorithms available in an open-source Python package.

Joint work with Alessandro Presacco, Jonathan Z. Simon

Informational Connectivity as a Measure for Measuring Synchrony in the Processing of Visual

Heather Bruett

University of Pittsburgh

heb52@pitt.edu

Multivoxel activity patterns in the brain can yield important information to researchers working with neuroimaging data, providing a way to decode patterns and reveal which areas are representing relevant information. On top of the increased sensitivity that multivoxel pattern analysis allows, an approach that examines the timeseries of pattern discriminability informational connectivity can help determine which regions are contributing significant decoding information to a particular comparison in the same trials - in other words, which

regions are acting in synchrony. I will present fMRI data that was analyzed via multivariate analysis tools and informational connectivity to determine how information synchrony plays a role in the processing of visual stimuli.

Joint work with Marc Coutanche

Testing a new method for assessing hemispheric lateralization using multi-voxel pattern

Brandon Carlos
University of Pittsburgh, Pittsburgh, PA
bjc89@pitt.edu

Hemispheric lateralization refers to the localization of a function in either the left or right hemisphere. In fMRI studies, univariate analyses have been employed to quantify functional lateralization, where mean activation is compared across equivalent regions in opposing hemispheres. Although reliable, this method has limitations, including in the kinds of questions that it can answer. The univariate method may not be able to answer research questions that require high sensitivity or dissociating between contrasts that provide similar activation. The present study tests a new method for assessing lateralization using multivariate analyses with machine learning techniques. In circumstances where the univariate method may fail to detect lateralization, this method can be used to examine patterns of activity, with a high degree of specificity.

Joint work with Marc Coutanche

Inferring the origin of nonstationary fluctuation in event occurrences

Kazuki Fujita
Department of Physics, Kyoto University, Kyoto 606-8502, Japan
fujita.kazuki.37n@st.kyoto-u.ac.jp

Existing time series data including neuronal firing in vivo are full of nonstationarity [1,2]. It has been revealed that the self-exciting process may exhibit nonstationary occurrence of events even if the system is not receiving any external stimulation [3,4]. By contrast, the system of weak self-excitation may also exhibit nonstationary occurrence if it receives external time-inhomogeneous stimulation. When receiving nonstationary time series, we wish to know whether the nonstationary fluctuation is internally generated by the system itself or externally stimulated. Here we developed a statistical model to make the inference for the cause of nonstationarity: We developed an Empirical Bayesian framework equipped with the self-exciting interaction term and applied it to spike trains generated by

the nonlinear Hawkes model of the GLM type to test if the model may infer the presence of intrinsic excitation. We will demonstrate the result in the poster.

Joint work with: Shinsuke Koyama, Shigeru Shinomoto

All-optical electrophysiology for neuroscience drug discovery

Felipe Gerhard

Q-State Biosciences Inc.

felipe.gerhard@qstatebio.com

Human stem cell-based models have become a powerful tool for modeling nervous system disorders for drug discovery applications. Human induced pluripotent stem (HiPS) cells derived from patient material can be differentiated into diverse neuronal cell types, which enable investigation of disease biology in the context of a human genetic background. HiPS cell systems will serve to complement existing rodent models in certain cases or supplant them where the human disease biology is not faithfully reproduced in the animal model.

Novel assays with high information content are needed to characterize the phenotypic response of disease-relevant cellular models and to detect pharmacological effects on the observed phenotypes. To this end, we have created an optogenetic platform called Optopatch that rapidly and robustly characterizes the electrophysiological response of electrically excitable cells. We elicit action potentials (APs) with a blue light-activated channel-rhodopsin (CheRiff). An Archaelhodopsin variant (QuasAr2) enables fluorescent readout of transmembrane potentials. A hybrid spatial-temporal PCA/ICA algorithm segments high-speed Optopatch recordings to identify active neurons. Optopatch assays yield more than 80 different functional parameters. Parameters are combined with information from other modalities such as morphological measures and average spike waveforms. Statistical significance testing is performed to identify the set of phenotype-specific properties. We use dimensionality reduction and regression models to maximize the sensitivity of the phenotype.

As a result, in neurons, the Optopatch platform can measure both intrinsic neuronal excitability and synaptic activity from hundreds of neurons with single-cell spatial resolution, milli-second temporal resolution and vastly higher throughput than manual patch-clamp. All-optical electrophysiology (Optopatch) provides a rapid and robust characterization of phenotypic response and provides an information-rich readout of pharmacological changes to the associated neurobiology. This approach will prove effective for profiling neurons from individual patients and opens the path towards precision medicine.

Joint work with: Ted Brookings, John Ferrante, Luis Williams, Kit Werley, Steven Nagle, Chris Hempel, David Gerber, Owen McManus, Graham Dempsey

Machine learning tools for neuraldecoding

Joshua I. Glaser

joshglaser88@gmail.com

While machine learning tools have been rapidly advancing, the bulk of neural decoding approaches still use last century methods. Improving the performance of neural decoding algorithms can help us better understand what information the brain represents, and can help for engineering applications such as brain machine interfaces. Here, we apply modern machine learning techniques, including recurrent neural networks (RNNs) and gradient boosting, to decode spiking activity in 1) motor cortex, 2) somatosensory cortex, and 3) hippocampus. We compare the predictive ability of these modern methods with traditional decoding methods such as Wiener and Kalman filters. Modern methods significantly outperform the historical ones. An LSTM decoder, a type of RNN, yields the best performance in all three brain areas. This approach typically explains over 30% of the unexplained variance from a Wiener filter (R2s of 0.86, 0.59, 0.50 vs. 0.75, 0.42, 0.28). Moreover, LSTMs are able to successfully decode from multiple tasks without a drop in performance, unlike Wiener Filters. These results suggest that modern machine learning techniques should be the default for neural decoding. We provide code so that everyone can utilize these methods.

Joint work with: Stephanie N. Naufel, Raaed H. Chowdhury, Matthew G. Perich, Lee E. Miller, Konrad P. Kording

A Bayesian approach to probabilistic spike detection and sorting

Patrick Greene

University of Arizona

pgreene@math.arizona.edu

Analysis of extracellular, multi-unit recordings typically requires that spiking events be distinguished from background noise (spike detection), and that individual spikes coming from different neurons be distinguished from each other (spike sorting). Spike detection is commonly done via simple thresholding, which may bias results by preferentially detecting neurons with large action potentials. Sorting is frequently done via clustering based on spike features, which often requires that the user visually inspect clusters and determine boundaries or number of clusters. Errors in detection and sorting can make replication difficult or potentially even lead to spurious findings [Navratilova 2016].

We address some of these outstanding issues in spike detection and sorting within a Bayesian framework, making use of a physical model of the spike detection process. Following Victor and Mechler [Mechler 2012], we model spiking units as current dipoles. Our probabilistic model can potentially account for overlapping spikes from multiple neurons,

and does not require careful tuning by the user.

Citations

Mechler F, Victor J. Dipole characterization of single neurons from their extracellular action potentials. *J. Computational Neuroscience*. 32, 73-100, 2012.

Navratilova Z, Godfrey KB, McNaughton BL. Grids from bands, or bands from grids? An examination of the effects of single unit contamination on grid cell firing fields. *J Neurophysiol*. 2016 Feb 1;115(2):992-1002.

Joint work with Kevin K. Lin

Population encoding of the 'intent' to drink alcohol

David Lisenbardt

dlinsenb@iupui.edu

Neural activity within the prefrontal cortex (PFC) is robustly altered by presentation of environmental stimuli associated with alcohol, and is correlated with alcohol craving, relapse, and mediated by genetic risk. However, we know very little about how genetic risk for excessive drinking influences the processing of alcohol-related cues across ensembles of neurons within the PFC, or how this population activity influences or represents drinking decisions. A behavioral model of cue-induced (Pavlovian) alcohol intake developed in our lab, dubbed 2-Way Cued Access Paradigm (2CAP) together with in vivo electrophysiology, was used to determine how genetic risk influences the neural dynamics that drive alcohol seeking and intake in the PFC. These studies used alcohol preferring P rats and their non-genetically predisposed (heterogeneous) founding Wistar population that were matched for ethanol intake history. During 2CAP, a light illuminated on either side of a rectangular operant box signaled the location and availability of alcohol. Extinction of responding for alcohol was evaluated by substituting water for alcohol. Animals were implanted with electrodes attached to moveable microdrives which were incrementally lowered through the PFC to maximize cell yield prior to electrophysiological recording and behavioral testing. Alcohol-associated cues elicited patterns of population-based neural activity that encoded the intent to drink (or abstain) in, but not P rats. In other words, only Wistar rats displayed unique patterns of neural activity in response to alcohol-associated cues that were predictive of future drinking/non-drinking. This suggests that cue-evoked encoding of information about the intent to drink (or not drink), is missing in the PFC of P rats. Additionally, only during extinction sessions, and only in P rats, population activity never differentiated (water) drinking trials from non-drinking trials. Thus, in P rats the PFC was biased toward encoding alcohol drinking, whereas in Wistar rats encoding of the intention to drink and drinking was present, regardless of fluid. These data provide novel

evidence that populations with a genetic vulnerability for excessive alcohol intake display altered processing of alcohol-associated cues and alcohol consumption, and suggest that the absence of cue-induced encoding of drinking intent in at risk populations may mediate continued excessive drinking and resistance to extinction.

Relationship between pairwise correlation and dimensionality reduction

Rudina Morina
Carnegie Mellon University
rmorina@andrew.cmu.edu

Spike count correlation, also known as noise correlation, has been used extensively to characterize the interaction between pairs of neurons. With the advent of multi-electrode array recordings, dimensionality reduction is being increasingly used to study the interaction of many neurons simultaneously. Here, we explore how the insights obtained from pairwise correlation metrics relate to those from dimensionality reduction. For a given population of neurons, we find that there is a systematic relationship between the spike count correlation distribution (created from all pairs of neurons) and the shared co-fluctuation patterns across the population, identified with factor analysis. This relationship depends on the number of shared co-fluctuation patterns identified as well as the strength of these patterns. We study this relationship both with simulated data and population recordings in macaque visual cortex. Our findings help to bridge results that utilize these different approaches for analyzing neural population activity.

Joint work with: Benjamin R. Cowley, Matthew A. Smith*, Byron M. Yu* (* indicates equal contribution)

Taste-related cortical population dynamics are stochastic and behaviorally-relevant

Narendra Mukherjee
Program in Neuroscience, Brandeis University, Waltham, MA
narendra@brandeis.edu

Dynamic processing of sensory stimuli by broadly distributed populations of neurons has been suggested to underlie the decision making processes that lead to the generation of appropriately timed behavioral responses. The mammalian gustatory (taste) system is particularly well suited for the study of the spatiotemporal neuronal dynamics that guides precisely timed behavior: taste ingestion-egestion decisions are implemented through a set of robust, rhythmic, brainstem-generated orofacial (mouth) movements, while single neurons in the gustatory cortex (GC) show, when analyzed through the lens of classical

trial-averaged peri-stimulus analyses, temporally rich responses that gradually transition from reflecting taste identity to taste palatability (the consumption decision variable). Ensemble analysis techniques (Hidden Markov Models), however, reveal the emergence of decision related firing to be much more sudden than can be detected in trial-averaged single neuron responses neural ensembles flip suddenly from an “identity coding” state to a “decision-related” state, at latencies that vary from trial-to-trial (Jones et al., 2007). Using EMG from the jaw muscles in combination with chronic multi-electrode recordings in adult rats, we recently demonstrated that that this variability in palatability-related dynamics of GC ensembles, hitherto dismissed as ‘noise’, correlates strongly with the onset of consumption-related behavior on a trial-by-trial basis (Sadacca, Mukherjee, et al., 2016). Our current work uses brief and precisely-timed optogenetic perturbations to test the functional significance of this correlation. Our preliminary results demonstrate that: 1) GC population firing states are temporally dissociated, with the disruption of an earlier state not affecting the onset of a later state; and 2) perturbations timed late in the taste trial is most potent in disrupting the onset of orofacial behavior. Together, these data imply that activity is the outcome of cross-talk between several different brain regions, an “output”? region that provides direct modulation of brainstem-generated taste-reactive orofacial movements as decision-related firing emerges. All in all, these results are allowing us to unravel the role of stochasticity in the processing of taste stimuli in GC ensembles and its behavioral significance in the context of brainstem controlled ingestion-egestion orofacial behavior.

Joint work with: Joseph Wachutka, Donald B Katz

Statistical Methods to Investigate Interventional Changes to Interaction Networks

Manjari Narayan
Stanford University
manjarin@stanford.edu

Across a variety of neuroimaging modalities, scientists observe brain activity at distinct units of brain function at mesoscopic or macroscopic levels, and seek to understand functional interactions between them. Typically we model such interactions from neural data, using a variety of probabilistic graphical models, where the nodes denote a unit of neural function and edges denote some form of statistical dependence between them. More recently, emerging techniques that simultaneously record and manipulate neural circuits in both humans and animals enable us to study such functional interaction networks under both passive and interventional regimes. In experiments that combine whole brain imaging with neurophysiological interventions, the macroscopic network effects of manipulating neural circuitry using interventions such as TMS or optogenetics often vary with site of

stimulation.

To explain site specific interventional effects, I will discuss recent work to adapt properties of network centrality from network science and matrix analysis to better understand the influence of a particular node in a network and explain the interventional effects of TMS. Unfortunately, popular measures of network centrality can be unstable to small changes to the adjacency. As a result, such properties cannot be used to summarize or compare multiple networks that vary across experimental subjects and conditions. To address these problems, I will discuss a general modification that yields stabilized alternatives to measures of node influence. Stabilized metrics can be applied to regularized estimates of high dimensional networks. Furthermore, stabilization enables the use of the nonparametric bootstrap to quantify the sampling variability of these metrics. Using empirically calibrated simulations and an interventional fMRI dataset, we illustrate both statistical benefits and scientific merits of this approach to studying interventional changes to functional networks.

Representational Similarity Learning with Application to Brain Networks

Urvashi Oswal
University of Wisconsin
uoswal@wisc.edu

Representational Similarity Learning (RSL) aims to discover features that are important in representing (human-judged) similarities among objects. RSL can be posed as a sparsity regularized multi-task regression problem. Standard methods, like group lasso, may not select important features if they are strongly correlated with others. To address this shortcoming we present a new regularizer for multitask regression called Group Ordered Weighted ℓ_1 (GrOWL). Another key contribution is a novel application to fMRI brain imaging. Representational Similarity Analysis (RSA) is a tool for testing whether localized brain regions encode perceptual similarities. Using GrOWL, we propose a new approach called Network RSA that can discover arbitrarily structured brain networks (possibly widely distributed and non-local) that encode similarity information. We show, in theory and fMRI experiments, how GrOWL deals with strongly correlated covariates.

Optimal features for auditory recognition

Srivatsun Sadagopan
University of Pittsburgh
vatsun@pitt.edu

A central challenge in the analysis of neural data is to understand how observed patterns of neural activity relate to or generate behavior. For example, neurons in primary (A1)

as well as higher auditory cortical areas exhibit highly nonlinear and surprisingly specific tuning properties. Our understanding of these responses is only at a descriptive level, and the critical question of how these responses might support behavior remains unresolved. Here, we show that these nonlinear responses encode essential mid-level features for the classification of ethologically-relevant sounds such as conspecific vocalizations. In vocal animals, increasing neural resources are committed for the processing of vocalizations (calls) as one ascends the auditory processing hierarchy. Therefore, the categorization of call types is a reasonable computational goal for the auditory cortex in these animals. We asked, using a theoretical information-maximization approach, how this goal can be best accomplished. Based on an earlier model for visual classification (Ullman et al., 2002), we first randomly generated a large number of mid-level features from marmoset calls, and used a greedy-search algorithm to choose the most informative and least redundant feature set for call categorization. We found that call categorization could be accomplished with high accuracy using just a handful of mid-level features. More interestingly, the responses of model feature-selective neurons predicted the observed nonlinear neural responses in marmoset A1 in astonishing detail. We further found that when using artificial call stimuli which were parametrically varied along multiple dimensions from a ‘mean’ vocalization, the performance of the model qualitatively mirrored marmoset behavior. These results demonstrate that the auditory cortex uses a mid-level feature based strategy for the recognition of complex sounds. These results further suggest that the tuning properties of neurons in higher auditory cortical stages are likely the result of goal-directed optimization. We argue that a goal-directed approach is essential for ascribing specific, behaviorally relevant roles for observed neural activity patterns.

Joint work with: Shi Tong Liu, Michael Osmanski, Xiaoqin Wang

Improvements to information theory analysis techniques throughout neuroscience with MATLAB support

Nicholas Timme
nmtimme@umail.iu.edu

Understanding how neural systems integrate and encode information is central to understanding brain function. An explosion in the availability of approaches that can be used to examine interactions across varying levels of brain function brings with it new challenges and opportunities. Information theory is well suited to the wide array of experiments and the challenging nature of data analysis typical to neuroscience. Frequently, data from neuroscience experiments are multivariate, the interactions between the variables are nonlinear, and the landscape of hypothesized or possible interactions between variables is extremely broad. Information theory is well suited to address these types of data as it possesses multivariate analysis tools, it can capture non-linear interactions, and it is model

independent (i.e., it does not require assumptions about the structure of interactions between variables). Methods currently exist to apply information theory analyses to many different types of data, including discrete data, continuous data, single trial data, trial based data, and aggregate data that result from dimensionality reduction techniques (e.g., principal component analysis). In total, information theory is a powerful tool for highlighting, detecting, and quantifying complex interactions in large systems with many types of variables. To reduce barriers to the use of information theory analyses, we have created a free MATLAB software package that can be applied to a wide variety of typical neuroscience data analysis scenarios. In addition to utilizing established analysis routines, this software package also includes several improvements to analyses of continuous data and trial based data. As demonstrations, we applied the software package to numerous model systems, including models of large Izhikevich networks, sensory habituation in Aplysia, location encoding in hippocampal place cells, movement direction encoding in primary motor cortex, and light stimulus encoding by center selective retinal ganglion cells. Among other things, analyses of these models showed time dependent information flow through networks, synergistic and redundant encoding by neurons, and encoding schemes modulated by inhibition, background activity, and stimulation correlation.

Mean Field studies of a society of Bayesian agents

Lucas Silva Simões

MSc in Physics, University of São Paulo - Brazil

lsimoes@if.usp.br

Humans, as social animals, learn their social constructs (for example what is morally right or wrong) interacting with their peers. How much does this 'learning from peers' behavior (constrained by experimental results from neuroscience and moral psychology) can describe the global characteristics of the whole society?

We approach this question modelling a society of agents who interact in pairs by exchanging for/against opinions about issues using an algorithm obtained using Maximum Entropy. The pair interaction can be described as a dynamics along the gradient of the logarithm of the evidence. This permits introducing an energy like quantity and an approximate global Hamiltonian. Knowledge of the expected value of the Hamiltonian is relevant information for the state of the society. We study the phase diagram of the society using a Mean Field approximation where a phase transitions separates ordered and disordered phases. These phases are interpreted in terms of Moral Foundation theory (MFT)

Joint work with: Nestor Caticha

**Fundamental limits, algorithms, and instrumentation for
novel non-invasive and minimally-invasive
“ultra-resolution” EEG systems**

Praveen Venkatesh
Department of Electrical and Computer Engineering, CMU
vpraveen@cmu.edu

What is the best possible spatial resolution attainable using EEG systems? We explore the updated spatial Nyquist rate results for EEG systems, and question whether the spatial Nyquist rate for reconstructing the scalp EEG is equal to the Nyquist rate for reconstructing the intracranial cortical potential. This requires us to undertake an analysis of the various sources of noise affecting the sensing of cortical potentials, and revisit the algorithms for source reconstruction. We provide provide fundamental limits on the best possible resolution achievable by EEG systems in the presence of noise. We then examine how existing source localization algorithms perform in comparison with the fundamental limits. We also discuss relevance in diagnosing various neural disorders. Finally, we also discuss experimental results that provide validation for use of these systems.

Joint work with: Pulkit Grover, Amanda Robinson, Marlene Behrmann, Ashwati Krishnan, Shawn Kelly, Jeff Weldon, Michael Tarr

Adjusted Regularization of Cortical Covariance

Giuseppe Vinci
Carnegie Mellon University
gvinci@andrew.cmu.edu

One of the most important challenges of computational neuroscience is estimating functional neural connectivity, that is inferring dependence structure among neural signals. Nowadays neuroscientists can record the activity of hundreds to thousands of neurons simultaneously, but only on limited numbers of trials. This high-dimensional setting requires regularized statistical methods to infer neural connectivity effectively. Sparse Gaussian Graphical Models (GGM), such as the Graphical Lasso, can provide sparse dependence structure estimates, but their performance in realistic scenarios of neural data can be unsatisfactory. We propose regularized GGM that incorporate neurophysiological information (e.g. inter-neuron distance) and provide better connectivity estimates. In realistic scenarios of neural data, we observed that our proposed approach can substantially outperform the Graphical Lasso and its variants. We apply the methods to estimate functional connectivity of neurons based of spike count data recorded with multielectrode arrays implanted in macaque visual cortex areas V1 and V4.

Unveiling the causal structure of living neural networks

Rashid V. Williams-Garcia

Departments of Neurobiology and Mathematics, University of Pittsburgh

rwgarcia@pitt.edu

Our nervous systems are composed of intricate webs of interconnected neurons interacting with each other in complex ways. These complex interactions result in a wide range of behaviors at the network level, which have implications for certain features associated with brain function, e.g., information processing and computational power. Under certain conditions, these interactions drive network dynamics towards critical phase transitions, where power-law scaling allows for optimal behavior. Recent experimental evidence is consistent with this idea and it seems plausible that healthy neural networks would tend towards optimality.

This hypothesis, however, is based on a number of problematic assumptions and potentially misleading statistical analyses of neural network data. Specifically, the observed power-law scaling supporting criticality is based on the analysis of neuronal avalanches, cascades of neuronal activations which conflate causal and unrelated activity, and thus confound important dynamical information. I will be presenting a novel method to unveil causal relations - which we call causal webs or c-webs for short - between neuronal activations, which contrasts and complements previous approaches based on neuronal avalanches. Using this method, we are able to separate cascades of causally-related events from unrelated events in multiunit recordings and unveil previously hidden features of the network dynamics.

When applied to mouse organotypic culture data, the c-webs method demonstrates that the observed neuronal avalanches are not merely composed of causally-related activations, and instead contain mixtures of concurrent but distinct cascades of activations, in addition to noisy spontaneous activations. Moreover, distributions of c-webs from these recordings do not feature power-law scaling - a result inconsistent with the criticality hypothesis.

Joint work with: John M. Beggs, Gerardo Ortiz

Dimensionality reduction schemes for understanding inter-trial neural response variations and their role in the sensory coding

Lijun Zhang

Washington University in St. Louis

lijunzhang@wustl.edu

Neural responses to sensory stimuli often change when the same cue encountered multiple times. Here we sought to understand what information encoded by a population of neurons

change in a trial-by-trial manner. Standard dimensionality reduction techniques such as linear principal component analysis (PCA) and non-linear local-linear embedding (LLE) have been used to visualize of such high-dimensional ensemble responses and how they vary over time. However, such techniques require smoothing data by averaging over different trials, thereby losing the information regarding inter-trial variability of population neural responses. To address this problem, we propose two approaches based on an assumption that there exists similarity in neural responses evoked across different trials response by the same stimulus. The first approach extends the minimum-error formulation of PCA with additional constraints, which can be solved using an iterative, alternating least squares scheme. To further empower the algorithm, in the second approach we reformulate the problem in the probabilistic PCA framework and present an Expectation-Maximization algorithm. We demonstrate the use of these techniques for variations in odor-evoked responses obtained from a relatively simple invertebrate model system. Our results reveal a simple scheme where adaptation does not confound intensity information but in fact optimizes the representation by encoding this information robustly but with fewer number of spikes.

Inter-Subject Alignment of MEG Datasets in a Common Representational Space

Qiong Zhang
Carnegie Mellon University
qiongz@andrew.cmu.edu

Pooling neural imaging data across subjects requires aligning recordings from different subjects. In magnetoencephalography (MEG) recordings, sensors across subjects are poorly correlated both because of differences in the exact location of the sensors, and structural and functional differences in the brains. It is possible to achieve alignment by assuming that the same regions of different brains correspond across subjects. However, this relies on both the assumption that brain anatomy and function are well correlated, and the strong assumptions that go into solving the inverse problem of source localization. In this paper, we investigated an alternative method that bypasses source-localization. Instead, it analyzes the sensor recordings themselves and aligns their temporal signatures across subjects. We used a multivariate approach, multi-set canonical correlation analysis (M-CCA), to transform individual subject data to a common representational space. We evaluated the robustness of this approach over a synthetic dataset, by examining the effect of different factors that add to the noise and individual differences in the data. On a MEG dataset, we demonstrated that M-CCA performs better than a method that assumes perfect sensor correspondence and a method that applies source localization. Lastly, we described how the standard M-CCA algorithm could be further improved with a regularization term that incorporates spatial sensor information.