Experimentation in Neuroscience
Abstracts

Talks – Friday and Saturday

Alison Barth (Biological Sciences, Carnegie Mellon University)
Dynamic connectivity: roadmaps for information flow in the neocortex
Abnormal temperature and pain sensation is a common symptom of patients after stroke, and in some cases can lead to central pain syndrome, a debilitating disorder. Although rodents have been a good model for some neurological disorders, it has been disputed whether they are appropriate for pain studies because of anatomical differences in central somatosensory areas. Here we will use sophisticated molecular tools to identify the primary neocortical area representing pain and temperature stimuli in mice, focusing on cold sensation mediated by the TrpM8 receptor. TrpM8 is the sole receptor required for thermal sensation from ~10-24^{o} C, and animals genetically lacking TrpM8 have a lack of temperature sensation in this range. In addition to cold temperature, TrpM8-expressing neurons are also excited by the chemical ligand menthol, offering complementary routes for specific receptor activation. We find that TrpM8 stimulation specifically activates neurons in the posterior insula, the same region that has been identified as cold-responsive in humans. In addition, TrpM8 stimulation also induces neural activity in the hypothalamus, the amygdala, and the posterior piriform cortex. This activation is reduced or absent in TrpM8 receptor knock-out mice. These experiments lay a critical foundation for understanding how pain and temperature circuitry in the neocortex can be modified by experience and disease.

Aaron Batista (Bioengineering, University of Pittsburgh)
New Neural States, New Mental States
New mental states necessitate new neural states. The new thoughts we are capable of are somehow made possible by the spectrum of thoughts we have had in the past. How are knowledge and experience manifest in neural activity states? Why are some new thoughts accessible to us while others are beyond our reach? My laboratory studies how neural activity patterns reorganize as animals learn to control a brain-computer interface (BCI). A BCI enables us to request that our nonhuman primate research subjects generate specific patterns of neural activity, and observe whether they are able to achieve those patterns. We accomplish this by specifying the pattern of neural activity that is required to move the BCI cursor, and observing whether the animals can control the cursor. In this talk I will present our recent experimental results showing computational techniques that allow us to predict which new neural states are within the grasp of our subjects. These observations may have relevance to other cognitive domains such as learning, visual imagery, and creativity.

Marlene Behrmann (Psychology, Carnegie Mellon University)
Face perception: where neuropsychology and fMRI do not see eye-to-eye
A topic of much interest in cognitive neuroscience concerns the brain-behavior correspondences that support face recognition. The ability to discriminate friend from foe rapidly and accurately is of critical importance for survival and of great evolutionary significance. There continues to be considerable debate concerning the psychological and neural mechanisms that mediate the recognition of individual faces and some of this may be a consequence of the fact that different methodologies tap into different aspects of the underlying system. Neuropsychological data, on the one hand, have favored a modular perspective, in which damage to a specific cortical region is taken as evidence for the primacy of a circumscribed region, notable the 'fusiform face area'. fMRI data, on the other hand, invoke a distributed system with multiple nodes, all of whose action contributes to face recognition. I will argue that the rapid and accurate discrimination of faces implicates a distributed neural network and that converging evidence from multiple methods, while laudable, may be muddying the waters.

Mazviita Chirimuuta (History & Philosophy of Science, University of Pittsburgh)
"Cool Experiments" and "False Models": Intervening and Representing in Neural Systems
In this talk I focus on a particular field of experimental neuroscience – brain computer interface (BCI) research – and discuss the complicated relationship between theory and practice which is characteristic of this field. At the heart of any BCI system one finds neural decoders. These are algorithms that allow one to interpret patterns of brain activity in terms of associated perceptual or motor states. Decoders are the basis of some of the most impressive feats of neurotechnology, such as the reconstruction of moving images perceived by subjects from fMRI scans (Nishimoto et al. 2011) and the use of neural signals to control external robotic hardware (e.g. Hochberg et al. 2012; Collinger et al. 2013). While such technologies apparently offer scientists a valuable means to test whether they have cracked the neural code, they also raise interesting questions about how best to glean understanding of brain function from the computational models used in successful technological applications.

Carl Craver (Philosophy-Neuroscience-Psychology, Washington University, St. Louis)
***Graphing the Brain's Dark Energy: How Network Analysis Contributes to Our Mechanistic Understanding of Complex Systems*.**
Neuroscientists at all levels of organization are increasingly borrowing formal methods from network analysis to analyze brain structure and function. After reviewing a few such applications, I focus specifically on recent, promising uses of network analysis to study resting state fMRI data (Power et al. 2012, *Neuron*, 79:798-813; Warren et al. forthcoming, *PNAS*). This example provides no evidence for the claim that network analysis brings a new, topological, style of explanation to the neurosciences (Huneman 2010, *Synthese*). It suggests, more interestingly, that network analysis is a tool for describing and discovering novel aspects of the causal structure of complex systems. The work I discuss emphasizes data-driven methods for discovering the parts of complex

systems and descriptive tools for representing novel (yet causal) kinds of difference-makers. The methods identify aspects of cortical organization that explain, among other things, why lesions to certain highly circumscribed regions of cortex result in symptoms spanning many different cognitive domains. My take-home point is that the interesting philosophical questions about the application of network analysis to neuroscience (and other sciences) are not questions about explanation; they are questions about discovery and the ontology of complexity. I close by considering possible directions for future philosophical work on this topic.

David Kaplan (Philosophy, Monash University)
Noise and signal in neuroscience
In a well-known series of papers, Bogen and Woodward (1988; Woodward 1989, 2000) describe the problem of detecting a phenomenon (i.e., a stable, repeatable pattern or effect) as one of detecting a signal embedded in noise. More recently, Bogen (2010) has suggested that noise can itself be treated as a signal or phenomenon to be explained. However, this produces a tension since noise is standardly defined as random or unpredictable disturbances that are not part of a signal. In this paper, I explore the tension, and Bogen's schematic answer to it, in the context of neuroscience where cellular, synaptic, sensory, and motor noise are routinely treated as objects of prediction and explanation. I argue that in order to make sense of noise as a genuine phenomenon variability must be distinguished from noise. This requires accounting for the causal sources of variability as well as appreciating its effects on brain structure and function.

Arnon Levy (History & Philosophy of Science, Hebrew University, Jerusalem)
Against the Layered Picture of Neuroscientific Explanation
I wish to argue against the idea explanation in neuroscience is organized into levels. First, I'll attempt to specify what levels talk means: roughly, I take it that levels are (vertically) orderable and (horizontally) uniform, in a way that is explanatorily informative and fairly general. I'll then look at the recent idea, put forward by several mechanist writers, that levels are best understood in mechanistic-constitutive terms. While the mechanistic view nicely captures nesting relations in neuroscientific explanation, it leads to an overly partial and local notion of levels. I'll then discuss some ways of accounting for levels that start not from the hierarchical structure of neuroscientific explanations, but from the intrinsic properties the entities and processes that enter into such explanations, such as size or structural features. Relying on some recent work by (among others) Potochnik and McGill, McGivern and Rueger as well as Craver, I will argue that this type of approach too, at least in the context of neuroscience, won't support a bona fide notion of levels. I shall close on a more positive note: abandoning levels outright is rather costly, since it threatens to leave us with a picture of neuroscience as a "dappled", theoretically unstructured domain. Briefly surveying two recent developments in the field – canonical computations and network motif models – I will raise the possibility of a general yet non-level-based form of theoretical organization.

Edouard Machery (History & Philosophy of Science, University of Pittsburgh)
A plea for reverse inference
In this talk, I will discuss the criticism of reverse inference put forward by Russ Poldrack, and I will argue that properly understood reverse inference is valid. Other defenses of reverse inference will also be discussed.

Jackie Sullivan (Philosophy, University of Western Ontario)
Experimentation and Construct Stabilization in Neuroscience
Scientific explanation requires stable explanatory targets. However, experimental practice in neuroscience presents obstacles to establishing stable explanatory targets. One obstacle is the freedom afforded investigators to produce, measure and detect phenomena they desire to understand and explain. A second obstacle stems from differences among investigators working within the same as well as across different areas of neuroscience on the emphasis placed on stabilizing explanatory targets. The aim of this talk is to further clarify the nature of these obstacles and to tease out their implications. In the positive part of the talk, I propose a strategy for overcoming these obstacles, consider its merits and indicate areas for further work.

Posters – Saturday afternoon

Nina Atanasova (Department of Philosophy, University of Cincinnati)
Neurobiological Knowledge Validated
Experimental neurobiology abounds with multiple experimental protocols and procedures aiming to study identical phenomena. This state of affairs puts the validity of knowledge claims produced on the basis of neurobiological experimentation in jeopardy, especially when it comes to animal experimentation (Sullivan 2007, 2009). However, scientists seem to favor the diversity of experimental protocols and procedures for neurobiological experiments involving animal models of human conditions (Wahlsten 2001).
I argue that the multiplicity of experimental protocols and procedures in fact fosters rather than hinders the validation of neurobiological knowledge claims. In support of my claim, (1) I show that neurobiological experiments employing animal models typically involve multiple different experimental tests integrated in standardized "test batteries". (2) I analyze the validation procedures used to establish animal models of neurophysiological disorders and conditions. I take these procedures to exemplify what Campbell and Fiske (1959) call convergent validity. What matters is producing compatible and converging results on the basis of multiple experimental arrangements. (3) Inherent in the procedure of establishing convergent validity is the requirement for obtaining converging results from at least two different experimental arrangements (Campbell and Fiske 1959). Therefore, a multiplicity of experimental protocols and procedures is actually beneficial for establishing convergent validity.

Frédéric-Ismaël Banville (Philosophy, University of Western Ontario)
Making Problems out of Problems: Representational Vocabulary in Neuroscience as a Methodological Heuristic
This project attempts to provide the foundations for a systematic account of heuristics in science, which is lacking in current views of explanation in neuroscience. This lacuna is unfortunate, because current accounts of explanation make reference to the fact that explanation in neuroscience requires cognitive labor, but never spell out precisely what such labor entails (Bechtel, 2008). One suggestion is that this labor involves heuristics, viz. procedures that enable cognitive agents with limited computational and temporal resources to solve problems in economical, fast, and relatively reliable ways. Heuristics make problems more tractable by providing constraints on the search for a solution, thus reducing the number of possible "paths" to a solution (Newell & Simon, 1972; Gigerenzer, 1999; see also Nickles, 1981 and Wimsatt, 2007 for an application of these ideas to science). Because of their lack of flexibility, however, heuristics are liable to biases. In the case of methodological heuristics this means that a given heuristic may be inadequate in certain contexts, such that it will need to be revised following some form of evaluation. The current project has three specific goals: 1) to provide a clear understanding of how heuristics actually enhance tractability, 2) to account for how heuristics shape inquiry and experimental design and 3) how they are evaluated and revised. Recent work (Bechtel, in preparation) on the status of representational vocabulary in neuroscience provides both a starting point and an informative case study.

Trey Boone (History & Philosophy of Science, University of Pittsburgh)
Conway's Game of Life and the Limitations of Multilevel Integration
Causal-mechanical accounts of explanation have dominated recent thinking in philosophy of neuroscience. But confusion persists regarding the role of abstraction within the mechanist framework. This confusion is rooted in the fact that most mechanists take decomposition between levels of analysis to be required for adequate explanation. My aim in this project is both to show that this is a mistake—i.e. that multilevel integration often entails distinct sacrifices in epistemic value—and to point toward alternative ways of thinking about the epistemology of levels of analysis in neuroscience. I use Conway's Game of Life as a simplified example to demonstrate the tradeoffs in epistemic virtues incurred as one moves across, and integrates between, different levels of analysis. On the basis of this example, I argue that forcing models into a dichotomy of either "genuinely explanatory" or "merely descriptive" is counterproductive; instead a multidimensional approach characterizing tradeoffs in epistemic virtues at different levels of analysis is called for. I extrapolate from this example to model construction in computational neuroscience and suggest directions for future study.

Daniel C. Burnston (Philosophy and Interdisciplinary Cognitive Sciences, UC San Diego)
Perceptual Context and the Nature of Neural Function

Traditional accounts of neural function are *absolutist*—they seek to find a univocal functional description for any given neural unit (a cell, population, or system), such that that description holds in any context in which the unit functions. Historically, these theories have sought to demonstrate function for a particular unit by uncovering a privileged relationship between the unit and a specific type of information in the world, or a specific cognitive/perceptual task, or (usually) both. Criticisms of this general view of function have argued against its *empirical adequacy*—they claim that the informational influences on and task involvement of any given unit are too multifaceted to sustain the traditional absolutist approach. Generally, these theorists (e.g., Michael Anderson, Vincent Bergeron, Charles Rathkopf) have argued that we should embrace absolutism of a different form. Specifically, they suggest that an account of function should *abstract* from particular informational or task context, and claim that each neural unit performs a particular computation on any sort of information it receives, and contributes that operation to any task in which it is involved. I undertake a detailed case study of perceptual area MT to argue that neither absolutist position is adequate. Recent developments have shown MT to be much less specific in its functioning than traditionally supposed. Far from being simply the "motion area," MT in fact shows sophisticated responses to motion, color, and depth in different contexts. However, I argue that abstract absolutism neither describes these developments adequately nor provides an epistemically satisfying account of the functions being uncovered. I thus claim that either form of absolutism should be abandoned.

Bryan Chambliss (Philosophy, University of Arizona)
Experiments in Second-Person Neuroscience

The experimental designs of *Standard Social Neuroscience* have utilized only a restricted range of social exchanges, e.g. exclusively spectatorial designs which involve observing, not interacting with, other people. By doing so, social neuroscience can be incorporated into cognitive neuroscience without changing how we understand experimentation, data or explanation in the latter. Closer inspection reveals this to be untenable. *Interactionist Social Cognition* claims that social cognition operates differently in second-person interaction than it does in first or third-person observation. Because Standard Social Neuroscience cannot test this thesis, Interactionists call for a *Second- Person Neuroscience* that can investigate both spectatorial and interactionist proposals concerning social cognition. Thus, Interactionists both expand the domain of social cognition, and propose modifications to neuroscientific practice to adequately investigate it.

Importantly, pursuing a Second-Person Neuroscience changes how we understand the data, experiments, and explanations offered by social neuroscience. Second-Person neuroscience involves rejecting the individualistic heuristics operative when social neuroscientists understand group psychology in terms of the individuals involved, and understand those individuals without appealing to the social context in which they are situated. Such heuristics lead Standard Social Neuroscience to both construct their experiments and interpret their data in ways that understand the neural processes of the individual without recourse to anything interestingly "social." Once such heuristics are rejected, old data and experiments must be reinterpreted. Moreover, if Second-Person Neuroscience explains the social cognitive processes involved in social interaction via a coupling of mental states between different subjects, its data and experiments issue in a different kind of explanation than we see in standard neuroscientific practices.

Phoebe Friesen (Philosophy, CUNY Graduate Center)

For Accuracy and For Honesty: Embracing the Epistemic Side of Neuroscientific Explanations

With regards to the question of what constitutes a neuroscientific explanation, Carl Craver has recently suggested that we should think of explanations in an ontic sense, as existing in the world as opposed to our minds (2007, 2013). He argues that such a conception will allow us to overcome many of the issues that have surfaced along with epistemic accounts that see explanations as representations developed by cognitive agents. The poster I hope to present at 'Experimentation in Neuroscience' will offer an alternative view that attempts to redeem the epistemic stance on neuroscientific explanations. I will demonstrate how acknowledging the critical epistemic inputs that contribute to the development of an explanation is necessary for a theory of neuroscientific explanations that is both accurate and honest. Accuracy results from recognizing the crucial role of choice that is involved in three stages in the development of an explanation. These three stages include 1) selecting the phenomenon to be explained 2) constructing an experimental design and 3) developing one's explanation into a particular form. The role of choice in each of these three stages will be demonstrated by way of an example from neuroscientific investigations in memory. Honesty arises in an explanation as a result of reflecting on what these choices mean with regards to the larger scientific goals of reductionism, unification, and justification. By further examining neuroscientific research in memory, this poster will illustrate the way in which recognizing epistemic contributors leads to a scientific explanation characterized by honesty.

Philipp Haueis (Independent Researcher; Member of the Research Group "Neuroanatomy and Connectivity" Max-Planck-Institute for Human Cognitive and Brain Sciences, Leipzig)

Neuroscientific Experimentation at the Mesoscopic Scale: The Case of the Cortical Column

Despite the increased philosophical attention neuroscience has received in the past two decades, scholars rarely ask whether our current cognitive and psychological concepts are adequate to describe and explain neural phenomena, or to find generalizations that express the principles governing cerebral organization. The reason is that most writers share with human neuroscientists the assumption that "the mind is what the brain does" (e.g., Kosslyn 1999), whether philosophically understood as a reductionist statement (Bickle 2003), in terms of a constitutive relationship between neural mechanisms and psychological phenomena explained by them (Craver 2007), or as a heuristic theory of identity (Bechtel 2008). In the project “meeting the brain on its own terms”, I contend that the assumption that “mind is what brain does” is dispensable in neuroscientific practice. Dropping the assumption allows researchers to search for new, more adequate concepts to describe the human brain, by conducting exploratory experiments without the use of cognitive vocabulary

I propose a case study that applies the Heideggerian metaphysics of science and the demand for investigating neural entities at the mesoscopic scale to the history of neuroscientific experimentation. The topic of the case study is the cortical column, an alleged organizational unit of the brain that guided 50 years of electrophysiological research practice. The cortical column was discovered by Mountcastle et al. (1955), who defined it as a vertical organizational structure that spans all cortical layers and whose cells all respond to a single perceptual dimension (e.g., skin stimulation). Further empirical research based on such an understanding, however, challenged whether the concept of the column can be generalized to the whole cerebral cortex. Most columnar structures respond to multiple stimuli, vary in their presence within or across species, and do not decompose into or form together functionally significant sub- or superstructures (i.e., mini- and hypercolumns). Such discrepancies have led skeptical commentators to the conclusion that the layer-transgressing cellular bands found in Nissl-stained sections (so-called ontogenetic columns) represent a “structure without function” (cf. Horton and Adams 2005). The supposed functional significance of ontogenetic columns would then be in part an artifact of the demand to find stable single-unit recordings at an ‘appropriate’ scale.

Tobias Huber (Independent Researcher)
Experimentation and Explanation
The paper discusses mechanistic and model-based philosophical approaches to neuroscientific practice. On mechanistic accounts, mathematical models are taken as explanatory if they describe mechanisms (Kaplan and Craver 2011; Bechtel 2011). In contrast, dynamic-mathematical accounts promote the view that not all neuroscientific phenomena are mechanistically explicable (Stepp et al. 2010; Silberstein and Chemero 2012). I will investigate the experimental practice of neuroscience and show that an adequate philosophical theory of the nature of neuroscientific explanation must integrate both mechanistic and computational models.

Several cases are discussed to illustrate that research in cognitive and systems neuroscience combines experimental methods, which provide lower-level mechanistic information about brain structure and function, with mathematical models, which target the higher-level complex dynamical behavior of multiscale cognitive systems. Whereas abstract and idealized computational models explain complex neurobiological phenomena in quantitative terms, mechanistic explanations render explicit how a neural mechanism’s organization of qualitatively specified structural components and operations generates the cognitive or behavioral phenomena under investigation. I argue that an analysis of the experimental practice in cognitive and systems

neuroscience shows that neuroscientific research integrates both qualitative mechanistic and quantitative computational explanations.
I further argue that this integration of causal-mechanisms and dynamical-mathematical models demonstrates the need for a philosophical account of neuroscientific research, which can accommodate the experimental, methodological, and explanatory plurality of neuroscientific research.

Maurice Lamb (Philosophy, University of Cincinnati)
Understanding Dynamical Models
Neo-mechanists argue that in order for a claim to be an explanation in cognitive science it must reveal something about the mechanisms of a cognitive system. Recently, they claimed that JAS Kelso and colleagues have begun to favor mechanistic explanations of neuroscientific phenomena. We argue that this view results from a failure to understand dynamic systems explanations and the general structure of dynamic systems research. Further, we argue that the explanations by Kelso and colleagues cited are not mechanistic explanations and that neo-mechanists have misunderstood Kelso and colleagues' work, which blunts the force of one of the neo-mechanists' arguments.

Joseph McCaffrey (History & Philosophy of Science, University of Pittsburgh)
Varieties of Neural Multi-Functionality
Many philosophers and cognitive scientists worry that widespread evidence of multi-functionality in the brain undermines traditional forms of functional attribution in cognitive neuroscience. Cathy Price and Karl Friston, on one hand, and Colin Klein, on the other, have recently offered competing accounts of the nature of neural multi-functionality and its significance for the design and interpretation of neuroimaging experiments. Price and Friston claim that neural multi-functionality is "apparent" (i.e. it is a contingent product of existing cognitive ontologies) while Klein claims multi-functionality is a "genuine" feature of neural architecture. In this poster, I claim that neither account is likely to succeed as a general treatment of multi-functionality in the human brain. I argue it is likely that the brain contains different kinds of multi-functional parts, some of which are more amenable to Klein's analysis and some of which are more amenable to a particular interpretation of Price and Friston's analysis. This suggests that neuroscientists will need to adopt different strategies for characterizing different functional systems in the brain.

Lauren N. Ross (History & Philosophy of Science, University of Pittsburgh)
Dynamical Models and Explanation in Neuroscience: the Canonical Model Approach
Kaplan and Craver argue that all explanations in neuroscience are descriptions of mechanisms and that models in this field are explanatory to the extent that they "reveal the causal structure of a mechanism" (Kaplan and Craver 2011, 602). In recent work they have extended their mechanist position to the use of mathematical models in neuroscience. They make two main claims regarding the explanatory status of mathematical models in neuroscience: that they must meet a model-to-mechanism-mapping (3M) constraint to be explanatory, and that their explanatory status increases

as they contain more relevant mechanistic details. Kaplan and Craver rely on their 3M constraint to oppose the claim that dynamical models, a type of mathematical model, can be explanatory even though they do not reveal the underlying causal structure of system-level dynamics. They state that dynamical models are at best descriptions of complex mechanisms and that those who consider them explanatory for any other reason "fundamentally misidentify the source of explanatory power in their models" (Kaplan and Craver 2011, 602). In this paper I argue against Kaplan and Craver's claims that mathematical models in neuroscience must meet their 3M constraint in order to be explanatory and that their explanatory status increases as they include more details. I support this argument by describing a dynamical model in neuroscience, referred to as a canonical model, that provides certain explanations despite failing to meet these mechanist requirements. I argue that this explanation cannot be ac- commodated by the mechanist framework and that it is well characterized by Batterman's account of minimal model explanation.

Jessey Wright (Philosophy, University of Western Ontario)
The Epistemic Status of Multi-Voxel Pattern Analysis: Towards a Deeper Philosophical Understanding of Neuroscientific Methodology
Philosophers of science have expressed concern about whether functional magnetic resonance imaging (fMRI) is a knowledge-generating technology (e.g., Bogen 2002, 2004, Roskies 2010, Uttal 2011). Some of these philosophers have indicated that answering this question requires a closer examination of the statistical analysis procedures used in conjunction with fMRI. However, to date no one has undertaken the task of providing a fine-grained analy- sis of these procedures. This is my aim here. Specifically, I will examine a new method of statistically analyzing fMRI data, multi-voxel pattern analysis (MVPA) and compare it with well entrenched voxel-based methods of analysis (VBM). Through this contrast I demonstrate that the use of both MVPA and VBM techniques requires investigators to make specific assumptions about the function and structure of the brain–assumptions that have specific implications for the epistemic status of these analysis techniques. In particular, I show that data analysis is not a purely objective process, and that the assumptions each method requires investigators to make about the system of interest (ie. the brain) shapes what the results of the analysis can be said to indicate about the localization of specific functions in the brain. This suggests that the concerns raised by some philosophers of science about fMRI are correct and that we should be suspicious of whether fMRI methodology is knowledge-generating. I conclude the poster by suggesting that, since each method relies on di↵er- ent assumptions, they can be treated as partially-independent detection tech- niques and that the convergence of di↵erent analyses on similar results does improve the epistemic status of the methodology. While this does not com- pletely dissolve worries about the epistemic status of fMRI, it suggests that a more complete analysis of how the use of the methodology fits within the larger context of neuroscience is required to understand how to conceive of fMRI as contributing to the generation of knowledge.