Robustness in Neurological Systems

13 - 15 November, 2015
Center for Philosophy of Science

ABSTRACTS

Alison Barth, Professor, Department of Biology, and Head, Barth Lab, Carnegie Mellon University

Functional Connectivity is Regulated by SOM Interneurons Spontaneous Activity

Understanding the dynamic range for synaptic transmission is a critical component of building a functional circuit diagram for the mammalian brain. We find that excitatory synaptic strength between neocortical neurons is markedly suppressed during network activity in mouse somatosensory cortex, with smaller EPSP amplitudes and high failure rates than previously reported. This phenomenon is regulated by tonic activation of presynaptic GABAb receptors via the activity of somatostatin-expressing interneurons. Optogenetic suppression of somatostatin neural firing was sufficient to enhance EPSP amplitude and reduce failure rates, effects that were fully reversible and also occluded by GABAb antagonists. These data indicate that somatostatin-expressing interneurons can rapidly and reversibly rewire neocortical networks through synaptic silencing, and suggest a critical role for these neurons in gating perception and plasticity.

Emilio Bizzi, Department of Brain and Cognitive Sciences and McGovern Institute for Brain Research, MIT

Muscle Synergies, Concept, Principles, and Potential use in Neurorehabilitation

When the central nervous system (CNS) generates voluntary movement, many muscles, each comprising thousands of motor units, are simultaneously activated and coordinated. Computationally, this is a daunting task, and investigators have strived to understand whether and how the CNS's burden is reduced to a much smaller set of variables. In the last few years, we and our collaborators have searched for physiological evidence of simplifying strategies by exploring whether the motor system makes use of motor modules, to construct a large set of movement.

The core argument for the neural origin of motor modules rests on studies of the spinal cord in several vertebral species, conducted using a variety of techniques. With these approaches, we and others were able to provide the experimental basis for a modular organization of the spinal cord circuitry in vertebrates. A spinal module is a functional unit of spinal interneurons that generates a specific motor output by imposing a specific pattern of muscle activated with a muscle synergy.

Muscle synergies are neural coordinative structures that function to alleviate the computational burden associated with the control of movement and posture. In my presentation, I will address two critical questions: 1) are muscle synergies explicitly encoded in the nervous system? And, 2) how do muscle synergies simplify movement production? I will argue that shared and task-specific muscle synergies are neurophysiological entities whose combination, orchestrated by the motor cortical areas and the afferent systems, facilitates motor control and motor learning.

Trey Boone, Department of History and Philosophy of Science, University of Pittsburgh

Temporal Organization and Robustness in Neural Systems

Neural circuits exhibit a delicate balance between sensitivity to and robustness over neuromodulatory inputs. Sensitivity to neuromodulatory inputs allows a single circuit to switch between a range of different functional states, while robustness enables a circuit to maintain stable performance of a particular function over variations in circuit parameters. Understanding the balance between sensitivity and robustness is fundamental to proper understanding neural circuit function. However, developing such understanding requires taking into account different sets of mechanisms that operate over a wide range of timescales. In particular, the biochemical mechanisms that modulate synaptic strength and intrinsic properties of cells often occur over much longer timescales than the relevant timecourse over which a circuit performs a particular function. As a result, proper understanding of both robustness and sensitivity of circuit function requires careful consideration of the temporal organization of those circuits and the processes that modulate them. The concept of temporal organization—i.e. which processes occur at what times and how processes occurring at different timescales interact—has received some attention in philosophy of neuroscience, but extant accounts—e.g. of ‘active’ or ‘dynamic’ organization in mechanistic explanation—do not seem up to the task when multiple timescales are involved. My talk will be largely exploratory. After discussing some of the basic issues of temporal organization implicated in the interplay between sensitivity and robustness, I will discuss some limitations of extant accounts of temporal organization in philosophy of neuroscience, and conclude by offering some of positive remarks about what sort of framework would be needed to address these issues.

Raffaella Campaner, Department of Philosophy, University of Bologna

Robustness Notions and Physiological Adaptability: Philosophical and Biomedical Reflections on the Neurological Basis of Disorders

Notions of robustness and invariance have recently been the objects of a pretty wide philosophical debate. In this joint paper we question whether and to which extent some philosophical reflections on the topic can profitably intertwine with reflections on the robustness of neuropsychiatric disorders as conceived of from within neurophysiology itself. As a matter of fact, the complex arrangement of causal factors bringing about, and reinforcing, neuropsychiatric disorders is in many cases still poorly understood. Such disorders lie at the crossroad of a number of disciplinary interests (e.g. physiology, psychology, neuroscience, psychiatry, genetics, epidemiology, pharmacology, etc.), which appeal to a variety of conceptual and methodological tools and whose models are supported by different kinds of evidence, which ought to be integrated, especially for explanatory purposes. With respect to this scenario, we will question whether the notion of robustness – and which notion of robustness – can be taken as a "picklock" to deal with the very idea of *disease*, its relations with the organism's overall behaviour and with its coping with changing conditions and environmental variations. Our reflections will be developed along two main lines:

- Does some notion of robustness help identifying a given neurobiological disorder over and above disciplinary-specific approaches, distinct models and ways to gather and integrate evidential data? How does the notion of "interactome" contribute to grasp the status of a disorder as a cluster of interplaying factors? Which idea of robustness does it assume and vehicle? To what extent can it – by showing overlappings and disease-disease relationships – help unravel invariant, "trans-diseases", relations?

- How does invariance relate to the intrinsically dynamic equilibrium (homeostasis) that characterizes every living systems? Which relations hold between stability of the organism and stability of the disorder, and between stability of the organism and arousal of the disorder? How does reversibility or irreversibility of a disorder impact on the robustness of a system and its capacity to adapt to changes and variations?

These issues will be tackled from both the vantage points of philosophy and physiology, looking for fruitful interactions between them and some joint rethinking of the very idea of neuropsychiatric disorder.

Mazviita Chirimuuta and Sandra Mitchell, Department of History and Philosophy of Science, University of Pittsburgh

Robustness in Contemporary Science and Philosophy

Robustness is a system-level property that allows the maintenance of function in response to external and internal perturbations. Kitano (2004) argues that, "Robustness is an inherent property of evolving, complex dynamic systems." What challenges do the study of robustness raise for our understanding of causal models, experimental protocols, and clinical practices? We will explore some of the ways robustness is characterized in biological systems generally and in neuroscience more specifically. We will present some examples to illustrate robust system behavior and the trade-offs between robustness and fragility. In particular, we will discuss the features of brain systems which make discussions of robustness particularly interesting in this area of biology. First, one of the most functionally significant features of the brain is its plasticity. The capacity of neural circuits and connections to reorganize themselves in response to new experience, and in the aftermath of damage to the brain, is what makes possible both lifelong learning and recovery from strokes and other brain injuries. Yet of course in many cases recovery is minimal or partial. A key question for advancing clinical neuroscience is to understand what the limits of the brain's robust mechanisms are, and how they can be altered through clinical manipulations. Second, theoretical neuroscientists following Barlow (1961) have often argued that evolution has worked to minimize the cost of biological information processing by reducing the number of redundant computations. This would, it seems, result in less robust systems. So an important question is how robustness in neural systems is achieved without sacrificing efficiency, and vice versa. Answering such questions would lead to insights into the general "design principles" (Sterling and Laughlin, 2015) of neural systems.

Flavio Keller M.D., Nicola Di Stefano Ph.D.
Laboratory of Developmental Neuroscience and *Istituto dell'Agire Scientifico e Tecnologico* , Università Campus Bio-Medico

Robustness of Musical Language: The Role of Motor Systems

To achieve perceptive stability and avoid perceptual ambiguities, perception must solve the problem of instability that arises both from changes in the external world, and from the intrinsic properties of the sensory organs. To bring just an example: if the eyes could assume different orientations for each gaze direction, it would be impossible to know whether a change in visual perception arises from a change in the world, or from a change in eye orientation. On the contrary, the eye takes always the same orientation for each gaze direction. This phenomenon, known as Donders' law, is but one of the many strategies by which sensorimotor systems solve the problem of perceptual ambiguities. In other situations, perceptive systems rely on perception-action loops to achieve stability.

One field that is very interesting in the context of robustness in perception is music: music is almost the quintessence of integration between perception and action. The human hear can tolerate small changes in the frequency of two simultaneous sounds without perceiving a change of harmony. Within a specific chord, one note is usually more critical than others for the character of the chord, and changing this note changes the harmonic character of the whole chord. In this paper we will present experimental data on the development of the discrimination between consonant and dissonant chords in children, exploiting the perception-action principle to establish when and how such capacity arises.

Arnon Levy, Department of Philosophy, Program in History and Philosophy of Science, The Hebrew University of Jerusalem

Causal Order and Types of Robustness

This talk is part of a project dealing with the notion of *causal order*. I use this term to signify two kinds of parts-whole dependence. Orderly systems have rich, decomposable, internal structure: specifically, parts play differential roles, and interactions are primarily local. Disorderly systems are aggregates of the activities of their parts, such that internal causal distinctions are of minor significance.

My focus here will be the connection between order and robustness, i.e. resilience in the face of internal or environmental perturbations. I distinguish three varieties of robustness. Ordered robustness is grounded on the system's specific organizational pattern, such as integral feedback or proofreading. In contrast, Disorderly robustness stems from the aggregate outcome of many similar parts, performing a form of trial and order on a massive scale. In between, we find semi-ordered robustness, wherein a messy ensemble of elements is subjected to a selection or stabilization mechanism. I'll discuss examples in each category and look at the prospects of this taxonomy to illuminate robustness across biological contexts and the relations between robustness and plasticity.

Timothy O'Leary, Marder Lab, Volen Center for Complex Systems, Brandeis University

Reconciling Variability with Robust Behavior at the Single-Neuron Level

Nervous system function depends on the electrical properties of neurons and these properties are largely determined by the types and distribution of signaling proteins, such as receptors and ion channels found in neuronal membranes and the connectivity between neurons. Perhaps surprisingly, the combination of important signaling components is found to be extremely variable even in neurons that have very stereotyped properties. How can this variability be reconciled with stable function? And what does variability tell us about the underlying regulation mechanisms in neurons? I will discuss recent theoretical work that offers simple conceptual answers to these questions and which implicates regulation, or 'homeostasis' in nervous system dysfunction as well as its robustness. Furthermore, I will argue that robust function is entirely compatible with significant variability in nervous system composition.

Acknowledgements
This work is supported by the Charles A King Trust and the NIH.

POSTER SESSION ABSTRACTS

Philipp Haueis, Independent Researcher, Max-Planck-Institute for Human Cognitive and Brain Sciences

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. The summary section of my former research will show how such experiments are possible and philosophically significant. The notion of "meeting the brain on its own terms" receives support by Martin Heidegger's Aristotelian metaphysics of "ways of being" (MacDaniel 2009). In Heidegger's view, every kind of entity (in this case brains and minds) demands its own vocabulary to adequately describe it. Demanding adequacy is crucial because the (scientific) description of entities articulates the criteria for what it is to be an entity of a certain kind, rather than not to be an entity at all. But because further empirical findings can drastically change our understanding of the investigated domain itself, it may be that certain entities disappear from scientific practice although they fruitfully guided former research (cf. Haugeland 2007). In the case of neuroscientific experiments, finding new concepts and generalizations that describe the structure and explain the function of neural entities (e.g., neurons, cortical areas can be regarded as autonomous, if they are counterfactually stable independently of psychological generalizations (cf. Lange 2000, 2008).

The second part of my research investigated how neuroscientific experiments can be used to meet the brain "on its own terms". I therefore applied criteria for exploratory experiments from the HPS literature (e.g., Steinle 1997, Burian 1997, Franklin 2005) to neuroscientific methods (intervention, functional localization) and instruments (electrode recordings, fMRI). One major result of my application was that more such hypothesis-free exploratory experiments have to be directed towards the mesoscopic scale of cerebral organization. Because the manipulation of mesoscopic entities such as cortico-cortical microcircuits is causally ambiguous

(Goense et al. 2012), micro- and macroscopic evidence about cerebral organization cannot be unified so far. Although attempts to understand the human brain through the concept of "connectivity" and mathematical network analysis are important steps in finding a unified descriptive vocabulary for cerebral organization, they rarely address the mesoscopic scale either (an exception is Bohland et al. 2009). In the section of the poster, 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. The cortical column thus is a showcase for the "historic life" of scientific entities. Its rise inspired the groundbreaking, though oversimplified 'ice-cube' model of cortical organization by Nobel Prize winners Hubel and Wiesel (1972), while its "death" (cf. Costa and Martin 2010) indicates that the mesoscopic scale was and remains.

Sanjeev Khanna, Dept. of Neuroscience, University of Pittsburgh

Structure of Neuronal Correlation During Eye Movement Planning in FEF

Pairs of nearby neurons in the visual cortex exhibit correlated spiking activity. There is a strong link between the correlation among neurons and the amount of information that can be represented in a neuronal population. This sensory information can be used to guide motor output, such as eye movements, but very little is known about the population activity in areas that bridge the sensory and motor divide. We investigated correlated variability in the frontal eye fields (FEF), a part of the prefrontal cortex that has both visual and motor functions. We predicted that the structure of neuronal correlation in FEF would be consistent with that found in the visual cortex, reflecting a conserved architecture across the neocortex. We used a laminar probe to record from groups of FEF neurons in alert rhesus macaque monkeys performing a conventional memory guided saccade task, and measured neuronal correlations on long time scales (spike count correlation). We found that the structure of neuronal correlation in FEF mirrored that seen previously in the visual cortex. Additionally, correlated spiking activity was strongest in pairs of neurons with similar visuosaccadic properties, and depended on the planned saccade direction. Future work will be aimed at understanding how population activity in FEF is modulated during active visual perception.

Emily Oby , Systems Neuroscience Institute, University of Pittsburgh(Presenter), Alan Degenhart, Elizabeth Tyler-Kabara, Byron Yu, Aaron Batista

"Coaching" Facilitates Brain-Computer Interface Learning

Some new skills are more difficult to learn to perform than others. We hypothesize that difficult- to-learn skills may necessitate the formation of new neural activity patterns, and that may be a gradual process. We used a Brain-Computer Interface (BCI) learning paradigm to test this idea. In a BCI, we can specify the mapping from neural activity to cursor kinematics, and then allow our Rhesus monkey subjects to learn that mapping. We recently reported that BCI mappings which are consistent with pre-existing patterns of population neural activity can usually be learned within a few hours, while BCI mappings that would require new patterns of neural activity are usually not learned to proficiency within a few hours. Here, we presented a monkey with the same difficult-to-learn BCI mapping for several consecutive days. Nine new BCI mappings were presented, for an average of four days each. For four of these mappings, learning occurred, but for the other five, performance remained fairly poor. We then examined whether we could facilitate learning with "coaching": Over several days we presented the animal with a series of successive approximations to the full BCI mapping. This coaching strategy does appear to facilitate learning. Finally, we examined whether such coaching can facilitate learning within a single experiment session. We found that it could not. In sum, BCI mappings that require subjects to exhibit new patterns of neural activity resemble skill learning in that both take extensive practice, and coaching can help.

Evan Pence, History & Philosophy of Science, University of Pittsburgh

Sensory Substitution

Sensory substitution is among the most surprising results in the history of perceptual psychology. One may, it seems, teach the blind to "see" simply by allowing them to explore visual environments with a video camera and transforming its inputs into tactile or auditory stimuli. Sensory substitution provides a striking example of perceptual plasticity and for years has served as one of the principle supports for enactive theories of perception. In particular, it has been used to show that perception and cortical function depend on sensorimotor knowledge. In defending these points, however, enactivists have tended to rely on studies conducted in the early days of sensory substitution research. When one examines more recent studies, I argue, the enactivist position appears far less plausible. Sensory substitution depends on seemingly innate neural mechanisms common with vision and appears possible even when subjects lack any credible source of sensorimotor feedback. The marvel of sensory substitution is not the malleability of the brain but its ability to maintain its functions across such wide variation.

Mark Povich, Philosophy-Neuroscience-Psychology, University of Washington in St. Louis

Implications of Robustness for the Theory of Scientific Explanation

The ontic-epistemic conception debate in the philosophy of explanation has shifted twice since Salmon (1984, 1989). Salmon framed the debate in terms of what explanations *do*. After Salmon, the debate was framed metaphysically, as a debate about what explanations *are*. The ontic conception was associated with the claim that scientific explanations are causal dependencies in the world; the epistemic conception was associated with the claim that scientific explanations are epistemic or representational states (Illari 2013). Craver's (2014) most recent formulation of the ontic conception backs away from the claim that explanations are ontic structures in the world and focuses on demarcatory and normative constraints on explanation. Under this most recent framing, the ontic conception holds that attention to ontic structures, rather than representational form, is required in order to demarcate explanation from other scientific achievements, like prediction, and to distinguish good from bad explanations, how-possibly from how-actually explanations, and explanatorily relevant from irrelevant features.

I argue that while explanations of robustness (Huneman, 2010; Sporns 2011) pose problems for causal-mechanical accounts of explanation, they can be accommodated by a generalized ontic conception. According to a generalized ontic conception, attention to more of the ontic than just the causal-mechanical is required to achieve the philosophical objectives of explanatory demarcation and normativity. I follow an emerging trend that locates explanatory power in the ability to answer what-if-things-had-been-different questions about the explanandum (Bokulich 2011; Craver 2007; Rice 2013; Saatsi and Pexton 2013; Woodward 2003). I argue that this requires accurate representation of anything on which the explanandum depends, causally or otherwise. There are kinds of ontic dependence other than causal dependence, knowledge of which allows us to answer what-if-things-had-been-different questions about the explanandum. For example, the dependence of an explanandum on the mechanism that underlies it is not plausibly causal, but constitutive (Craver and Bechtel 2007). As another example, Saatsi and Pexton (2013) present an explanation of Kleiber's law, an allometric scaling law that relates an organism's body mass to a biological observable (West, Brown, and Enquist 1999). The precise details of the explanation are irrelevant for our purposes. What matters here is that there is a feature, the scaling exponent, that counterfactually depends on the dimensionality of the organism. It is plausible that this counterfactual dependence relation contributes explanatory power, yet it is implausible that the dimensionality of organisms is a causal variable that can, in practice or in theory, be intervened upon (Saatsi and Pexton 2013, 620).

Similarly, explanations of robustness in ecological and neurological systems exploit a kind of (arguably) noncausal dependence on broadly topological properties. A generalized ontic conception can therefore account for explanations of robustness.

Zachary Tosi, Cognitive Science, University of Indiana, Bloomington

Self-Organizing Mechanisms for Growing Micro-circuits

Through the use of electrophysiological, optical, and information theoretic techniques we have begun the process of understanding the structural and functional connectivity of neural circuits at the network level: the so-called *micro-connectome*. Though there is still much we do not know, certain features have revealed themselves. For instance, we know that synaptic connectivity is highly nonrandom, containing over represented motifs, hubs, and heavy-tailed synaptic efficacy and degree distributions. Other features like heavy-tailed firing rate distributions further indicate that neurons are not selecting incoming synaptic connections randomly. How these features manifest themselves is largely a mystery in spite of the fact that many mechanisms of neural plasticity have been known of and understood for decades. The following work presents a model tying together known mechanisms of neuronal and synaptic plasticity and the more recently discovered structural and phenomenological features found in neural circuits. The model demonstrates that a combination of Spike-timing dependent plasticity (STDP), homeostatic plasticity, synaptic regulation, and synaptic pruning results in a highly nonrandom network possessing many known features of biological circuits. This includes heavy tailed firing rate, synaptic efficacy, and synaptic degree distributions. Additionally, the resulting network self-organizes many putative or expected (but not outright proven) micro-connectome features including modular/hierarchical and rich-club organization, and a high level of computational power. Interestingly, by allowing neurons to *self-organize preferred firing rates* other, more well understood plasticity mechanisms produce the confluence of aforementioned features. The model addresses one of the major goals of computational neuroscience and in particular neuromorphic artificial intelligence, which is to determine the minimal number of abstract mechanisms or behaviors which account for the observed features in neural circuits, and furthermore has been shown to organize them when given input from *real living neural circuits*. Ultimately it is hoped that the model will help contribute to future research in neuromorphic artificial intelligence and brain machine interface software.