Carnegie Mellon University

Annual Progress Report: 2009 Formula Grant

Reporting Period

July 1, 2013 – December 31, 2013

Formula Grant Overview

The Carnegie Mellon University received $910,547 in formula funds for the grant award period January 1, 2010 through December 31, 2013. Accomplishments for the reporting period are described below.

Research Project 1: Project Title and Purpose

Mid-Level Feature Representation in Human Visual Cortex - Despite five decades of research, remarkably little is known about the representation of objects in "high-level" human visual cortex. Although both neurophysiology and neuroimaging have informed us about the spatio-temporal properties of early visual cortex (V1->V4), the subsequent encoding of objects within the ventral pathway is largely an unknown. Advances in neuroimaging technologies, including functional Magnetic Resonance Imaging (fMRI), Diffusion Tensor Imaging (DTI), Magnetoencephalography (MEG), and Electroencephalography (EEG) have altered the way vision scientists are able to study object representation. To that end, this project will develop a variety of novel, real-time neuroimaging paradigms and analysis methods designed to study and decode the representation of objects in terms of visual features.

Duration of Project

1/1/2010 – 12/31/2013

Project Overview

Human visual object recognition abilities are profoundly better than those of even the most powerful artificial vision systems, yet our understanding of high-level vision in humans remains embarrassingly vague. One of the reasons is a failure to articulate some notion of representation beyond early visual processing. That is, although it is well established that early vision codes for oriented local edges, there is almost no notion of how said edges (and other image features) are combined to represent parts, objects, and scenes. One reason for this lack of theory is that there are, as yet, inadequate computational tools to learn compositional (i.e., hierarchies comprised of reusable parts) structures from real-world images. A second reason is that neurophysiology, including both single-unit and multi-unit recording, is far too narrow in scope to provide a clear picture of how the primate visual system represents visual information. To address these problems in visual coding, we propose a research program that leverages two neuroimaging methods that allow broad coverage of the human brain in action – fMRI and MEG. We will

combine the particular strengths of each of these research tools with state-of-the-art machine learning methods that offer a much more powerful means for exploring and understanding high-dimensional, complex datasets. We will introduce new experimental designs that enable fMRI to function in a real-time system for adaptively assessing neural responses in the visual system and that will enable much better source localization in MEG. For both techniques, coverage of the entire visual system combined with exquisite source localization in fMRI and exquisite temporal resolution in MEG, will allow us to gain a better understanding of the distributed codes used in object and scene representation. There are three facets of this project. First, fMRI and MEG will be used to examine the nature of mid-level features that form the core of object representations. Second, both technologies will be improved with respect to their spatial (MEG) and temporal (fMRI) resolution. Third, a variety of machine learning methods will be applied to more effectively and efficiently identify the optimal stimulus/stimuli for specific brain regions to better identify and understand how the resultant features are combined to support high-level vision.

Principal Investigator

Michael J. Tarr, PhD
Professor and Co-Director, Center for Neural Basis of Cognition
Carnegie Mellon University
115c Mellon Institute, CNBC
4400 Fifth Avenue
Pittsburgh, PA 15213

Other Participating Researchers

None

Expected Research Outcomes and Benefits

Addressing the root causes and concomitant impairments arising from both endogenous (e.g., autism, dyslexia, Alzheimer's, etc.) and exogenous (e.g., TBI, stroke, tumors, etc.) brain disorders requires a detailed account of basic perceptual and cognitive processes from a neural perspective. With respect to the study of high-level vision, the process of object recognition is at the nexus of how we see, interpret, and think about the world around us. Yet there is no comprehensive account of how the human visual system accomplishes this routine and critical action. Thus, a better understanding of how the human brain is able to robustly recognize and interpret objects and scenes in an invariant manner is critical to allowing clinical scientists to more accurately diagnose and more effectively treat a wide range of neural diseases and disorders. For example, several different developmental disorders, including autism, dyslexia, and congenital prosopagnosia manifest, in part, as impairments in visual functioning – advances in understanding the underlying visual mechanisms that are disrupted by such diseases will enable the next generation of behavioral interventions. Likewise, many stroke and TBI patients have impaired visual recognition abilities, the most salient of which is prosopagnosia in which face recognition is rendered non-functional. Elucidating the neural processes whereby face recognition is realized will lead to clearer models as the basis for potential treatments. More generally, the development of improved neuroimaging methods will enhance the study of the

brain across many different cognitive domains: neuroimaging has already revolutionized how we study and ultimately understand the neural bases of cognition, and as such, improvements in these varied techniques have the potential to further enable better discernment of mental processes. It is only through more precise and more coherent models that we will be able to develop better diagnosis and analysis tools for many different brain disorders.

Summary of Research Completed

Real-time fMRI. Object perception recruits a cortical network that encodes a hierarchy of increasingly complex visual features. While the earliest stages of human vision have been reasonably well-modeled using local oriented edges, the encoding of visual properties in higher-level cortical regions is less well understood. Prior work has been able to characterize neural selectivity for mid-level features in IT and V4 using simple parts-based, parametric approaches. Although these studies identify individual shapes maximizing neural firing for selected neurons, they provide little insight as to the principles used to distinguish preferred stimuli from anti-preferred stimuli. Here, we explore the grouping of properties salient to cortical perception by defining and using Euclidean feature spaces based on complex visual properties of natural and synthetic object classes. Natural objects are drawn from real-world photographs and are projected into a space reflecting visual grouping based on a SIFT bag-of-words distance metric. Synthetic objects are Fribbles - animal-like objects composed of geometric shapes - and projected into a space capturing the morphs between Fribble appearances. These visual feature spaces present a new tool for analysis of cortical activity underlying object perception. Using fMRI and pre-selected regions of the ventral visual cortex, we find that object stimulus pairs near one another in feature space can produce responses at the opposite extremes of the measured activity range. Suppression of response to properties slightly altered from preferred properties is observed for different object classes and centered on different locations in a given feature space. These observations may demonstrate extension of surround suppression observed in lower levels of vision - such as the Gabor-like models of V1 receptive fields and studies of location-based suppression effects for visual texture perception.

Approach. Object stimuli were projected to coordinates in Euclidean feature spaces derived from complex visual properties of "natural" real-world photographs (Figure 1c) and synthetic Fribbles (Figure 1d). For natural objects, a pairwise distance matrix was computed for all available stimulus pairs, based on a SIFT bag-of-words distance metric. Multi-dimensional scaling was used to find a four-dimensional space maximally preserving the computed pairwise distances. Recent work supports the use of SIFT as a model of cortical representation of visual objects, further bolstered by its representational parallels with the biologically-popular HMAX model. For each of four classes of synthetic Fribble objects, a space was defined in which movement along an axis corresponded to morphing the shape of an associated appendage between two fixed geometric shapes. These classes of synthetic objects and their corresponding spaces provide a more controlled environment for probing cortical object representations, while employing visual forms similar to those of real-world objects. The novel visual feature spaces present a new tool for analysis of cortical activity underlying object perception.

Stimuli displayed to subjects were drawn from a pool of ~300 images, selected in real-time for each viewing trial based on measured cortical responses to prior stimuli shown. Similar to past

studies of single-neuron activity in primates, selection of stimuli was intended to identify images maximizing the response of a pre-determined voxel region. This brain-guided "image search" was pursued through a modified form of simplex simulated annealing. Image searches found to be stable across scanning sessions, for a fixed subject and predetermined brain region, were evaluated for the grouping in feature space of properties salient to perception. Real-time fMRI stimulus selection itself is a novel method we develop and discuss more deeply in other venues.

Human cortical regions under study were chosen in ventral visual cortex prior to the performance of image searches. These regions of interest were chosen based on their selectivity for classes of visual objects used in our study. This selection was intended to maximize measured cortical response signal over background noise in fMRI recordings.

Results. Observing cortical activities over the defined visual feature spaces, we find ventral cortical voxel populations producing high responses for multiple sets of visual properties, i.e., for two or three locations in space (Figure 1a-b). We also find voxel regions suppress their responses for stimuli adjacent in space - and slightly varied in visual appearance - from those stimuli evoking markedly high cortical activity, exhibiting high-level representational parallels to "surround suppression".

Learning fine-grained visual categories. Human expertise in learning and individuating faces is facilitated by neural mechanisms that flexibly encode novel facial identities. Previous research has suggested that regions in the ventral stream are central to face processing, but little is understood about how these regions encode faces during the rapid discrimination time course, and critically, how their properties evolve when new faces are learned. Using magnetoencephalography (MEG), we recorded with millisecond resolution from the cortex when participants learned trial-by-trial to discriminate between two face categories that differ in facial parts. To accurately reconstruct cortical activities from MEG sensors – a procedure made difficult by its under-constrained nature, we developed a trial-partitioning method that uses trials in midst of learning for estimating a spatially-constrained source model, which makes source localization more precise during early- and late-learning. Using this methodology, we probed spatiotemporal properties of the ventral stream and found that initially, information about faces is increasingly more prominent in time course from posterior to anterior regions of the ventral stream, in line with a hierarchical organization. Interestingly, we found this pattern to shift as learning progressed, whereby face discriminability improves significantly in the inferior occipital gyrus (IOG), possibly due to strategic reliance on informative facial parts. Furthermore, we found this enhanced discriminability to be concomitant with better synchrony between the IOG and the middle fusiform gyrus (mFus), supported by increasing gamma-band (30-50Hz) synchrony between the two regions. Overall, our methodological framework helped elucidate spatiotemporal plasticity of the ventral stream key to acquired proficiency in face individuation.

Approach. Two face categories were created in a parameterized space. Each category included a unique set of 364 face images that are slight variations of a prototype. The two prototypes were identical except for the eye size and mouth width. These two dimensions were systematically varied in a grid-based design. All face images were rendered in 3D using the Face-Gen software (http://www.facegen.com/index.htm).

The main experiment involved an online learning task where participants were trained to discriminate between the two face categories with trial-and-error. The experiment consisted of 728 trials and was divided into four blocks (182 trials in each) with self-paced breaks in between to reduce fatigue. One additional 30-second break was introduced in the middle of a block to allow for eye blinking. During each trial, an A or B face (subtending a visual angle of 6 degrees vertically and horizontally) was presented for a 900msec-duration. The sequence of A and B faces was randomized, and the occurrence of each category was balanced so that there were equal numbers of A and B faces in every 20 trials. The presentation of a face was preceded by a fixation cross in the screen center. Meanwhile, a machine-generated sound (630msec) that reads either "A" or "B" was played, followed by a jitter (120-150msec). This audio label was randomly assigned to prompt the categorical decision, but it did not indicate the category membership of a stimulus face. Such a scheme helped decouple the sound categories with the categories of faces. The numbers of "A" and "B" sounds were maintained equally within every 20 trials. The participant was instructed to respond with "yes" or "no" to indicate whether the sound label reflects the true category of a face during the period while the face was shown. The "yes" and "no" signs appeared in the left or right bottoms of the screen with their positions counterbalanced for each session. After the 900msec period, a fixation cross was shown for a random duration (100-120msec), which was followed by a feedback sign of "correct", "wrong" or "too slow" in the screen center for 750msec to inform the participant whether the response is correct, incorrect or missing within the deadline. The inter-trial-interval was 400msec. Cortical activities were recorded using MEG while participants performed the face category learning task.

Results. We presented a framework for exploring spatiotemporal plasticity of the human ventral stream. Methodologically, we proposed a source localization method for MEG tailored to a continuous, online learning paradigm. By partitioning trials into different stages of learning, we utilized data in the middle of learning – otherwise often under-used in evaluating hypothesis about learning – to estimate a source model constrained to cortical regions relevant to face processing. This procedure increased precision in the reconstruction of cortical activities during early- and late-learning. Theoretically, our spatiotemporal approach helped confirm previous theories about a hierarchical architecture in the ventral stream. We find that facial identity information is made more available in discrimination time course from posterior to anterior ventral stream in the IOG, mFus and aIT (Figure 2). Importantly, we also find that face discriminability improves significantly in the IOG towards the end of learning, suggesting increased strategic reliance on informative facial parts during face individuation. This finding also suggests that the cortical basis for configural (e.g. whole-face) and component-based processing may be implemented in separate circuitries, whereby anterior and posterior parts of the ventral stream are oriented toward whole-face and part-based processing respectively.

Complementary to cortical time course, we found activities in the IOG and the mFus are strongly coupled in the alpha and beta bands (≤25Hz) during 100-200msec post-stimulus throughout learning (Figure 3). This time window overlaps with waveforms (e.g. M100 and M 170) reported in face perception, suggestive of a low-frequency interactive mechanism possibly due to long-range cortical feedback loops and attention. Critically, we discovered that learning increases IOG-mFus synchrony in the gamma-band (30- 50Hz) during the same post-stimulus period. This high-frequency phase-locking – concomitant with the enhanced face discriminability in the IOG – is likely to be associated with increased attention to facial parts. Additionally, it suggests top-

down and feed-forward mechanisms between the IOG and the mFus may be key to part-based encoding for faces. We attribute this to the fact that the mFus should provide top-down supervision in directing the IOG to process informative facial parts, whereas the IOG feed-forwards information about facial parts processed at the entry level in return. Our result thus provides empirical evidence for neural computational models that propose the mFus to serve to magnify differences across categories.

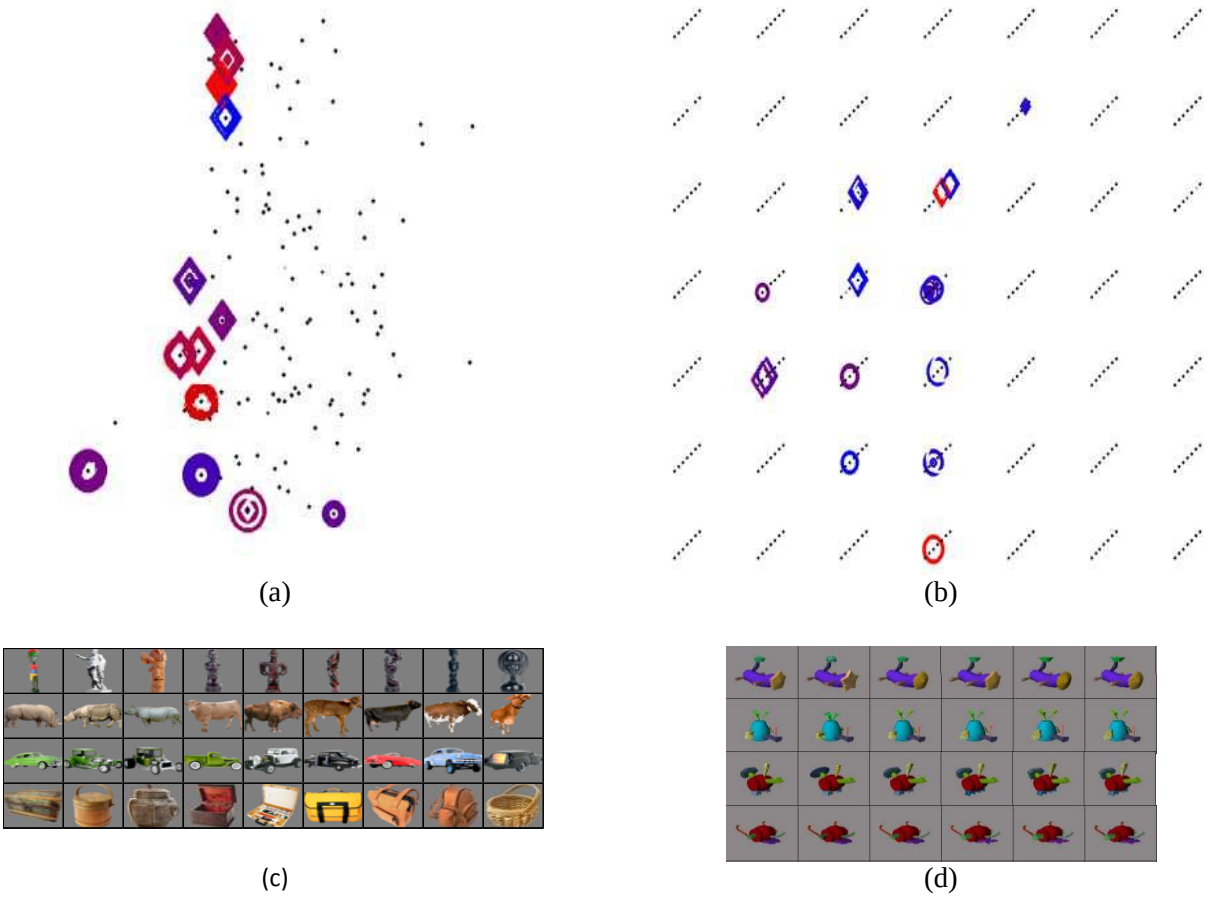

Figure 1. (a, b) Stimuli displayed in an image search for pre-selected brain region using **(a)** natural and **(b)** synthetic objects. Location of all potential stimuli in feature space shown as black dots. Stimuli displayed three of more times during image search shown as circles and diamonds. Colors span blue–dark blue–dark red–red for low through high responses. Size of shape corresponds to time each point was displayed during image search, with larger shapes corresponding to later points in search. **(c, d)** Example images from classes of natural and synthetic object stimuli shown to subjects.

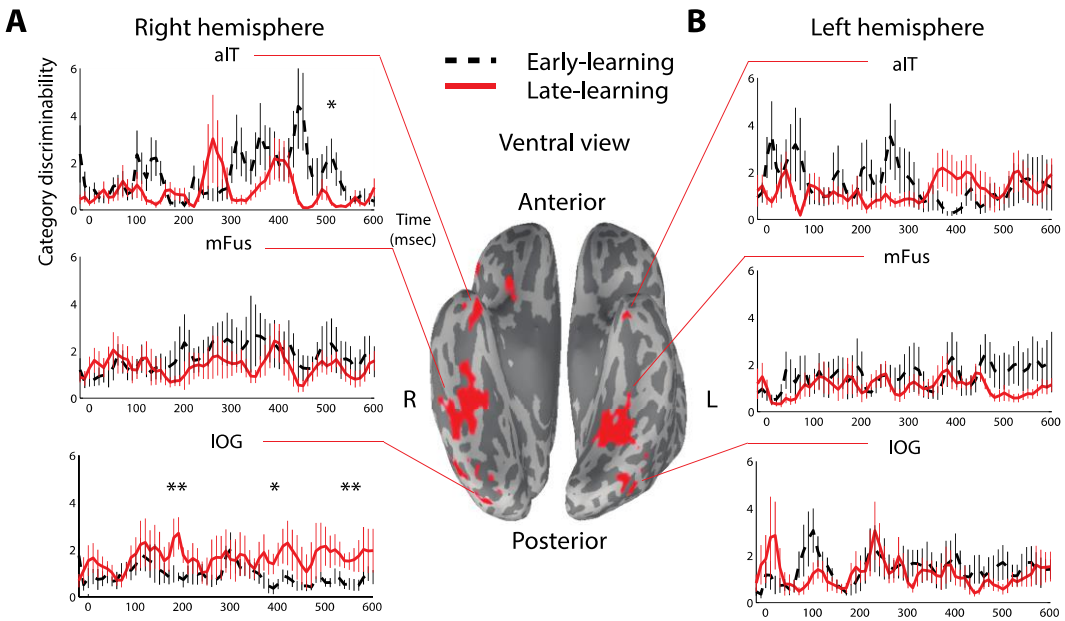

Figure 2. Category-discriminative time course in the ventral stream during face learning. **(A)** Group-level face-category-discriminative time course in the inferior occipital gyrus (IOG), middle fusiform gyrus (mFus) and anterior inferior temporal gyrus (aIT) along the ventral stream in the right hemisphere. **(B)** Discriminative time course in similar positions from the left hemisphere. Inflated cortical surface is taken from a representative subject. "0" on the time axis marks the visual stimulus onset. Vertical bars indicate standard errors of the mean. Asterisks correspond to significant difference in discriminability between early- and late-learning at $p < 0.05$.

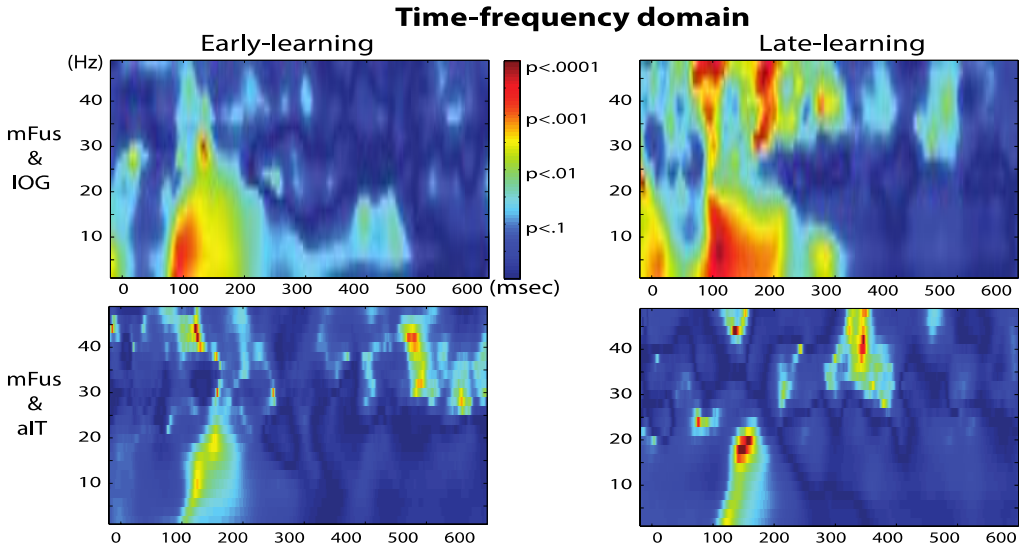

Figure 3. Group-level statistical maps of phase-locking (synchrony) between the mFus and the IOG (upper row) or the aIT (lower row) across time and frequency during early- and late-learning. "0" on the time axis marks the visual stimulus onset.

Research Project 2: Project Title and Purpose

Understanding Schizophrenia through Logical and Empirical Analysis - Schizophrenics have persistent experiences of hearing voices (auditory hallucinations) and thoughts that are paradoxically taken to belong to another (inserted thoughts). To explain these symptoms, a common model postulates a mechanism whereby a person's thoughts or actions are "marked" as self-generated. Under this theory, trouble arises when this mechanism fails: a person's thoughts appear "alien" because they are not marked as self-generated. This project suggests an alternative hypothesis: a central problem in patients who hear voices or have alien (inserted) thoughts is not the misattribution of central thoughts, but rather that thinking and imagining are abnormal because they *automatically* generate thoughts and images. This project develops and tests this new hypothesis regarding the origins of these specific symptoms of schizophrenia.

Duration of Project

1/1/2010 – 12/31/2013

Project Overview

This project focuses on two symptoms of schizophrenia: auditory verbal hallucinations (AVH) and inserted thoughts. Two of the most common symptoms of schizophrenia are persistent experiences of hearing voices and thoughts that schizophrenics paradoxically take to belong to another individual (i.e., not themselves). Within the psychiatric literature, these symptoms are commonly explained as a dysfunctional "monitoring" mechanism whereby a person's thoughts or actions are usually "marked" as self-generated. Trouble arises for schizophrenic individuals (i.e., they exhibit symptoms associated with schizophrenia) when this putative mechanism fails: a person's thoughts or internal voices appear "alien" because they are not marked as belonging to that person. That is, prevailing theories of AVH invoke notions of controlled processes that involve active monitoring of inner speech. The problem with such models (which dominate the field) is that they fail to answer the right questions about these symptoms. This is due to an inadequate conceptualization of the symptoms themselves; that is, they are taken at face value rather than considered within the larger framework of cognitive processing. Careful analysis of these symptoms leads to an alternative hypothesis regarding their origin: that a central problem in schizophrenic patients who hear voices or have alien (inserted) thoughts is that *normal* processes of thinking and imagining are abnormal because they *automatically* (and unconsciously) generate thoughts and images (i.e., not that they are mislabeled). This precise and testable hypothesis is reached only by understanding the widely used, but poorly defined, notion of automaticity drawn from cognitive psychology. To investigate the role of automaticity in these two symptoms of schizophrenia this project will involve collaboration with psychologists and psychiatrists at two levels. Conceptually, we shall provide a rigorous, well-specified theoretical account of the notion of automaticity as used in cognitive psychology. Empirically, we will formulate and test more precise alternative models and explanations for some aspects of schizophrenia, specifically investigating mechanisms of AVH, with the aim of exploring both spontaneous cortical oscillatory activity as well as that elicited by external auditory stimuli, with the general hypothesis that increased cortical excitability may mediate the automatic nature of AVH in schizophrenia. Elucidating such mechanisms will further our understanding of one of the

core symptoms in schizophrenia and provide the basis for developing novel treatment approaches for the illness.

Principal Investigator

Wayne Wu, PhD
Assistant Professor and Associate Director, CNBC
Carnegie Mellon University
115 Mellon Institute, CNBC
4400 Fifth Avenue
Pittsburgh, PA 15213

Other Participating Researchers

Raymond Cho, MD - employed by University of Pittsburgh Medical Center
Mark Wheeler, PhD - employed by University of Pittsburgh
Lori Holt, PhD - employed by Carnegie Mellon University

Expected Research Outcomes and Benefits

Schizophrenia is a mental disorder affecting approximately 1% of the US population – over 2 million individuals. Although a variety of antipsychotic medications have been developed to treat the disease, these have many side effects and are only partially successful in ameliorating schizophrenic symptoms. Psychosocial treatments are likewise limited in their efficacy. Successful treatment and, ultimately, developing more effective medications and behavioral interventions must rely on more accurate characterizations of the disease and its concomitant symptoms. Thus, a better model of the root causes of two of the most common symptoms that is grounded in cognitive mechanisms will provide a clearer basis for advanced drug and behavioral therapies.

The broader analysis of automaticity and a better model of its role in cognition will also have implications for a wide range of disorders, for example, dyslexia and ADHD, as well as improving our understanding of the symptoms of schizophrenia.

Summary of Research Completed

A central goal has been to map more clearly different causal models of Auditory Verbal Hallucination in schizophrenia (AVH) and, in doing so, suggest more definitely and clearly experiments that can distinguish between the models. Two models are at issue: (1) the standard self-monitoring model, which emphasizes a defect in a control mechanism, and (2) a less discussed sensory model, which emphasizes automatic activation of perceptual mechanisms.

In the current six-month research period and in collaboration with Raymond Cho (University of Pittsburgh, Psychiatry), we have published a paper in Frontiers in Psychiatry: Schizophrenia, *Mechanisms of Auditory Verbal Hallucination.* Using their "Altmetric" impact assessment, *Frontiers* reports that the "article is amongst the highest ever scored in this journal (ranked #43

of 327)" in impact based on their measures. The paper contrasts two competing models of AVH, something that has not been made clear in the literature and which we believe will lead to more effective experimental work on understanding the mechanisms of AVH. Science works best when there is competition.

In the recent period, we have formed two hypotheses to test the models. If self-monitoring involves inner speech, then certain sorts of experience should be less plausible or indeed should be ruled out:

(a) rapid back and forth dialog where patients answer their AVH with their own inner speech;
(b) simultaneous inner speech and AVH.

(a) is implausible because of the implications of inner speech models of self-monitoring. AVH arises because self-monitoring is broken with respect to inner speech. The patient generates inner speech, but because self-monitoring is faulty, experiences inner speech as AVH. It seems to us implausible then that a patient could have a dialog with its AVH using inner speech. The reason is that the dynamics of the inner conversation would require a precise oscillation of the self-monitoring mechanism between being broken and working normally, the former with AVH, the latter with normal inner speech. We think that patient experience provides a wealth of information about their phenomenology that is not only clinically relevant but *theoretically* relevant.

We have begun to extract this information as part of clinical assessment of their symptoms. In one patient, OC, we extracted the following dialog:

> **AVH**: Oh you're going to let him get away with that (this voice is apparently speaking to another entity about OC)
> **OC**: God loves them too
> **AVH**: I got you a 100 times (taunting OC that it affected OC somehow)
> **OC**: yeah, you did.
> **OC**: Judge not and you wouldn't be judged (paraphrased from the Bible)

This shows that patients experience a back and forth between their own inner speech and AVH, demonstrating that normal dialog structure can occur between AVH and inner speech. Again, while the self-monitoring explanation of such episodes is not impossible, it seems to us implausible. These observations have occasionally been made in passing, but no systematic examination of patient reports have been published, to our knowledge. Our current work aims to pay more attention to patient experience in a way to assess the underlying mechanisms of their symptoms.

(b) strikes us as *impossible* to adequately explain, at least on a single mechanism version of self-monitoring. The reason is that if failure of self-monitoring is needed to explain how inner speech becomes AVH and yet adequate self-monitoring is needed to explain how inner speech is experienced as one's own, then cases where AVH and inner speech occur simultaneously, as when one tries to talk over another, should be impossible. After all, this implies that the

mechanism is both working and not working simultaneously. We are investigating whether such simultaneous experiences occur, and there is some suggestion that they might. The challenge here is getting patients to provide information about the time-course of their experiences.

The upshot is that the clinical data that psychiatrists use everyday to assess their patients' mental states provide rich data that is important in testing models.

We have provided an analysis of the sorts of experiments that should be done to test which model is correct and made predictions of the likely outcomes. This puts us, and others, in a good position to conduct experiments that can decide which mechanism holds (and it is logically possible that both mechanisms might be true, explaining different cases and thus types of AVH). The resulting framework for experiments is given in the following three tables covering (1) proof of concept experiments; (2) tests of self-monitoring models and (3) tests of spontaneous activity accounts. These tables provide a logical framework for use by AVH researchers to test models of the underlying mechanisms of AVH, making the space of experimental possibilities much clearer than they were before [Tables are adapted from Cho and Wu (2013)].

Cho and Wu (2013) "Mechanisms of Auditory Verbal Hallucination in Schizophrenia" *Frontiers in Psychiatry: Schizophrenia*

Table 1. AVH Mechanisms: Proof of concept studies

Intervention	Possible Outcomes	Commentary
Disrupt self-monitoring during inner speech	AVH of one's own inner voice (does not reproduce typical schizophrenia AVH)	To date, there have been no reports in the literature of such experiments for self-monitoring.
Disrupt self-monitoring during auditory imagery	AVH of imagined voices	As above.
Stimulate auditory cortices	AVH with typical phenomenology	There are relevant cases for stimulation of auditory cortices, e.g. neurosurgical studies and studies in epilepsy.

Table 2. AVH Mechanisms: Testing Self-monitoring Account in Schizophrenia

Intervention	Possible Outcomes	Commentary
Disrupt self-monitoring	Worsening of existing AVH Increase in types of experienced AVH (e.g. inner speech).	As self-monitoring posits loss of function, it is unclear that disruption of self-monitoring during discrete AVH episodes provides relevant outcomes. Further disruption of self-monitoring should lead to increase in AVH, but with predicted expansion of the type of AVH experienced (e.g. inner speech).
Remediation of monitoring deficits	Decrease in frequency of AVH	An ideal intervention of targeted remediation of self-monitoring during discrete AVH episodes in patients is currently unavailable.

Table 3. AVH Mechanisms: Testing Spontaneous Activity Account in Schizophrenia

Intervention	Possible Outcomes	Commentary
Inhibition of sensory cortex during discrete AVH episodes	AVH decrease	Demonstrates that a modulation of auditory cortical activity correspondingly modulates AVH severity
During quiescent non-AVH period, stimulation of same cortical region shown to be sensitive for inhibition of AVH	AVH with complete phenomenology	As above.
Generalized inhibition of sensory cortex	Decrease in AVH frequency	Ralph Hoffman (Yale University) and co-workers have implemented generalized inhibition, demonstrating overall decrease in frequency of AVH

Research Project 3: Project Title and Purpose

Computational Immunology for Toleragenic Composite Tissue and Solid-Organ Transplantation
We propose to elucidate the tissue specific biological and informational principles of complex immunological mechanisms and systems that drive rejection and inflammation in solid or composite tissue transplants. Through advanced computational methods, machine learning and other high dimensional analytic techniques, a new understanding of the patterns and governing principles of the immune system will be formulated, and will pave the way for novel clinically relevant toleragenic transplant protocols, therapeutics and long term patient care strategies. The longer term objective is to improve transplant tolerance and reduce side effects from indiscriminant immunosuppression.

Duration of Project

1/1/2010 – 12/31/2012

Summary of Research Completed

This project ended during a prior state fiscal year. For additional information, please refer to the Commonwealth Universal Research Enhancement C.U.R.E. Annual Reports on the Department's Tobacco Settlement/Act 77 web page at http://www.health.state.pa.us/cure.