Carnegie Mellon University

Annual Progress Report: 2009 Formula Grant

Reporting Period

July 1, 2010 – June 30, 2011

Formula Grant Overview

The Carnegie Mellon University received $910,547 in formula funds for the grant award period January 1, 2010 through December 31, 2012. 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.

Anticipated Duration of Project

1/1/2010 – 12/31/2012

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

Over the past year we have been using novel neuroimaging methods to investigate the structural connectivity of face-selective cortical regions in humans. In particular, because it has been extensively studied and provides robust category-selective neural responses, face processing can serve as a paradigmatic model for studying more general object processing in the human ventral-temporal cortex. We have been exploring the connectivity of the multiple brain areas that are recruited during the perception of faces. Although it is clear that these regions reflect the neural machinery used for this high-level perceptual task, the precise architecture of the structural connections in the "face network" remains largely unknown.

We have used whole-brain diffusion spectrum imaging (DSI) in conjunction with fMRI to better elucidate connectivity within the face network within individual subjects. To this end, the same subjects were scanned using DSI and an fMRI face "localizer" that compares the neural activation for faces versus common objects. High-definition fiber-tracking (HDFT) was applied to the DSI data and connectivity between face-selective brain areas was determined on a pair-wise basis using two functional ROIs (as defined by the functional localizers) at a time. ROIs were expanded two voxels into white matter for tracking purposes. The occipital and temporal lobes were seeded 1,000 times in each voxel (34-50 million total seeds) and fibers intersecting the two ROIs were selected. After tracking, any fibers that did not connect and terminate in the ROIs were removed.

These analyses (which are ongoing) have produced three significant results.

First, in exploring connectivity within the "core" face network (IOG/OFA, FFA, STS, and aIT), we, surprisingly, find *no* connectivity between these face-selective regions and the STS. That is, there is strong connectivity between the IOG and FFA, and between FFA and aIT, but none of these regions has any significant white matter fiber tracts connecting to STS (Figure 1). Recent control analyses reveal that STS does have significant white matter connectivity to other category-selective regions within the ventral-temporal cortex (e.g., PPA); thus, the lack of connectivity we see with respect to the core face network provides new constraints for how we understand mechanisms of face perception.

Second, across a wide range of neuroimaging studies examining face-selectivity in early visual areas, it has been observed that the location of the face selective area referred to as the 'Occipital Face Area' or OFA varies between studies. Areas labeled as OFA have been located on IOG, in ITS on the posterior fusiform, and more dorsally in the vicinity of MTG. Subjects often show two or more of these areas as selectively active for faces. Critically, using a new state-of-the-art, Siemens Verio with a 32-channel head coil, our face localizer fMRI scans consistently reveal multiple "OFA's" within subjects. Figure 2 shows the connectivity of these multiple OFAs to

FFA for two sample subjects. Notably, all three of these early, face-selective regions show connectivity to FFA, yet control regions with the occipital cortex do not. Thus, it appears as if there are three distinct regions within occipital cortex that are functionally significant for face processing. We are now developing fMRI studies to titrate out the specific computational role(s) of these brain regions.

Third, and quite preliminarily, we have some evidence that the fiber endpoints from face-selective regions terminating within FFA are located in a small region of white matter *directly above* the FFA functionally-defined voxels showing the highest neural activity. Although this result requires a great deal more analysis and follow-up to be conclusive, we speculate that the higher neural responses observed within category-selective regions may not necessarily be the result of greater neural selectivity for particular stimulus categories. Rather, it is possible that peaks in the BOLD response actually reflect a nexus of information flow, whereby the summed activation of the neural units (e.g., voxels) in this local region tend to be larger than in surrounding units more distant from the nexus. We plan to explore this result using a variety of methods, including applying linear classifiers that are less sensitive to neural summation than are standard GLM contrasts. However, if this result stands up to careful scrutiny, it may necessitate some rethinking about what we may infer from category-selectivity as observed in fMRI.

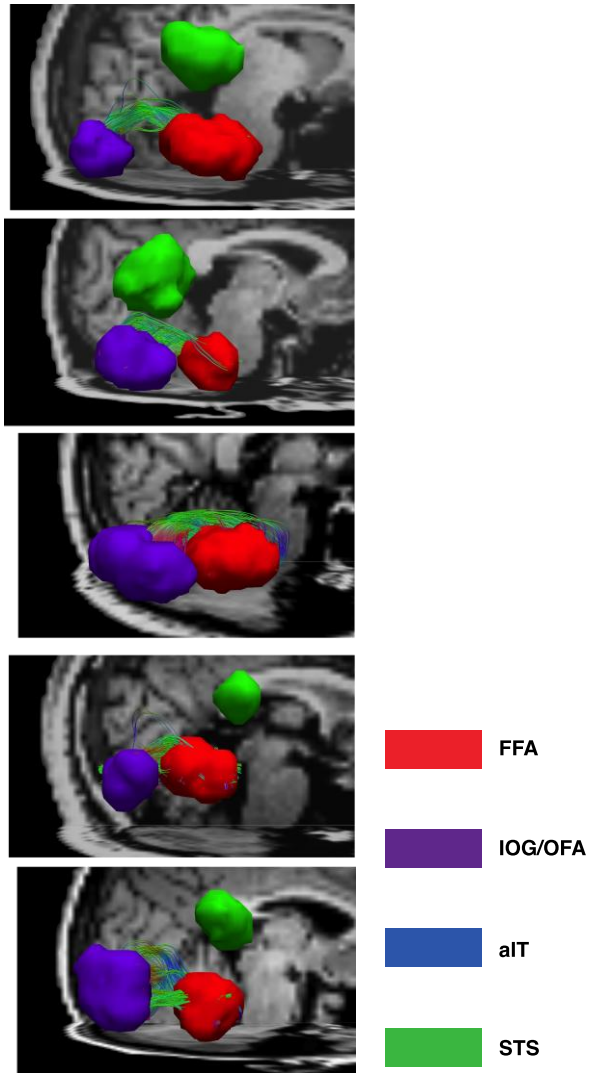

Figure 1. No fibers from STS to FFA or IOG. Strong connectivity between FFA and IOG. Data is shown individually for 5 subjects.

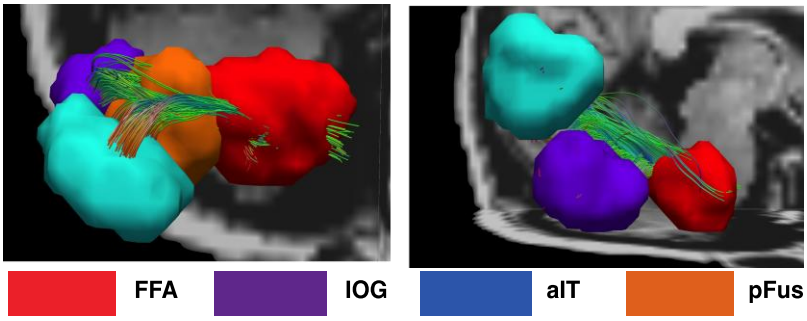

Figure 2. Connectivity of occipital "OFA" face-selective areas to FFA.

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.

Anticipated Duration of Project

1/1/2010 – 12/31/2012

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

The positive outcomes of our work to date have been to clarify the differences between alternative causal models of AVH and to initiate a characterization of how to experimentally distinguish between them.

Current accounts of the positive symptoms of schizophrenia are dominated by self-monitoring accounts. These argue that the basis of these symptoms result from the patient's failure to monitor that they are themselves the source of those episodes. For example, they may "talk in their heads" (inner speech) and fail to recognize that they are the one's conducting inner speech, leading to the experience of that speech as externally generated and hence, auditory verbal hallucination (AVH).

In recent work with Raymond Cho, we have undertaken to examine the assumption of such self-monitoring accounts so as to assess their empirical viability. We have discovered that one hindrance to progress is the lack of detailed mechanisms that allow for clear empirical interventions. Moreover, we believe that there is a plausible alternative that has received less attention. To make progress, however, we must clarify the substance and content of the two theories.

A first step is to provide more mechanistic detail in all alternatives. We have achieved a simple division between two alternatives. Briefly, our current understanding is that self-monitoring involves a "top-down" mechanism whose disruption is necessary for the positive symptoms in schizophrenia. In contrast, our favored alternative account, which is less prevalent in the literature, is "bottom-up", namely the overactivity of sensory areas in the brain. We propose that in auditory hallucination, the overactivation of auditory perceptual areas and auditory memory would be the primary cause of AVH.

The general distinction between top-down and bottom-up mechanisms is captured by the theoretical distinction in psychology between a controlled and automatic process for which one of us has now provided a rigorous characterization in two publications.

The next two steps are (1) to fill in the mechanistic details for each model and (2) to consider the types of experiments that might decide between them. We have made some initial progress on the second project.

We have investigated at a theoretical level how to distinguish these two models. One result of our work has been to recognize that while much empirical work has emphasized that patients with schizophrenia show clear deficits in the self-monitoring mechanism, this often reported result *does not distinguish* between the two mechanisms we have just noted. Thus, while these results support a general defect in self-monitoring, they do not establish causality and thus do not provide evidence against the bottom-up model.

The reason this evidence is not relevant to deciding between the two alternatives is that *both* predict that a failure of self-monitoring activity is a signature of AVH. In the case of our preferred, bottom-up account, we make this prediction simply because the basic bottom-up mechanism does not rely on self-monitoring. Thus, it follows that one should not see normal self-monitoring activity in patients with AVH.

We have begun to assess potential experimental tests that can decide between or give evidence in favor of one of the two models. As the problem is rarely formulated in this way, we see our conceptual work as a needed step in advancing our understanding of the etiology of AVH. We have discovered that the experimental possibilities are fairly restricted.

In healthy subjects, two obvious interventions could be undertaken, but they illustrate specific challenges for empirically dissecting the issues.

On the one hand, as disruption of self-monitoring is taken to be at least necessary for AVH, one could attempt to induce AVH by disrupting self-monitoring with Transcranial Magnetic Stimulation (TMS) during auditory imagery or inner speech. A simple prediction of the self-monitoring account would be the induction of AVH in conditions when the subject was performing inner speech or auditory imagination of speech. To our knowledge, this precise test has not yet been done. A clear null effect would seem to be evidence *against* self-monitoring accounts, but even if one induced AVH like symptoms, the question remains open whether this is the mechanism generating AVH in patient's with schizophrenia.

On the other hand, we might consider stimulating auditory cortices to test the bottom-up model. Again this would show that such activity would be sufficient for inducing auditory hallucination, but in fact, we know *a priori* that some form of auditory cortical stimulation must be sufficient for hallucinatory auditory experience. Moreover, Wilder Penfield showed in the 1950s that stimulating cortical areas of awake patients during neurosurgery could induce complex hallucinations.

We can, however, couple these two ideas to test the bottom-up account: stimulation of auditory cortical areas, say those identified as active in the normal course of AVH, should induce AVH in patients and disruption of those same areas during AVH should correlate with reduction of AVH. A similar paradigm applied to the top-down model also applies: restoration of self-monitoring should ameliorate AVH while disruption induce or exacerbate TMS.

Progress in understanding the mechanisms of AVH requires that we know what the alternatives are and begin to map how to experimentally test these models. We have made progress in understanding the theoretical space wherein further experimental progress is then possible and to provide an initial map of the experiments that need to be undertaken.

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.

Anticipated Duration of Project

1/1/2010 – 6/30/2012

Project Overview

This project seeks to elucidate immunological patterns and mechanisms in pre-clinical composite tissue and solid organ transplants under various conditions of rejection, inflammation and tolerance. To achieve this, machine learning, computational linguistics, agent based modeling, and differential equation based methods of analysis will be applied to proteomic and genetic data. The raw data is being gathered through ongoing collaborative projects with the University of Pittsburgh Departments of Immunology, Plastic and Reconstructive Surgery, The McGowan Institute for Regenerative Medicine, The Pittsburgh Nuclear Magnetic Resonance Center, and the University of Innsbruck School of Medicine.

The objective of the proteomic analysis is a set of models with predictive power for rejection state (and possibly grade), time point, and graft type. Models will be built for specific tissues including skin, muscle, heart, lung, and other tissues for which data may be available including the liver, kidney, pancreas, lymph node, and serum. Protein data is provided in concentration values (pg/ml) as read by a Luminex 100-IS machine at the University of Pittsburgh Department of Immunology.

Genetic analysis will focus on the skin and muscle components, utilizing gene array data provided by the composite tissue transplant program at the University of Pittsburgh Department of Plastic and Reconstructive Surgery as well as the University of Innsbruck School of Medicine. The objectives of this analysis include development of models with predictive power for rejection state, identification of rejection specific gene expression and enumeration of these targets for verification by collaborating researchers, and corroboration of gene-protein expression under various rejection and inflammation conditions. Time and resources permitting, we will also investigate the feasibility of models with predictive power for time point and physiological process.

Principal Investigator

Jaime Carbonell, PhD
Director, Language Technology Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Other Participating Researchers

Ravi Starzl – employed by Carnegie Mellon University

Expected Research Outcomes and Benefits

Both solid organ and composite tissue transplantation has traditionally required extensive life-long immunosuppression therapy to prevent rejection of the graft. It is this long-term high-dose immunosuppression regimen that is the limiting factor for both the quality of organs that can be used for a solid organ transplant and for the widespread adoption of composite tissue

transplantation. Further, the present clinical method for detecting rejection in solid organ transplant is fluctuation of organ function (caused by damage from rejection), or skin discoloration in the case of composite tissue rejection.

The recent paradigm of transplant tolerance, or chimerism, has revealed a conceptual framework where it may be possible to detect the onset of rejection before tissue damage occurs, adjust immunosuppression dosage to optimal levels, and potentially wean patients to very low or no maintenance doses of immunosuppression.

This study will generate novel quantitative models, methods and protocols that will allow the toleragenic paradigm to be personalized and adapted to each individual patient. This would allow better allograft management and the treatment of rejection prior to the accumulation of damage that affects organ function. With better proactive management, a larger pool of organs that could be used in transplant would be made available, helping to alleviate the worldwide shortage of kidney, liver, pancreas and other organs. Additionally, this study will enumerate genetic markers of susceptibility to rejection, as well as targets and strategies for the development of novel non-toxic immunoregulatory pharmaceuticals.

Summary of Research Completed

Introduction and Significance

The primary focus of this project's research is the elucidation of immune signaling patterns and potential mechanisms in a pre-clinical context, utilizing data from ongoing experiments in composite-tissue allograft (CTA) transplantation, solid-organ transplantation, as well as wound-healing and regenerative medicine, at collaborating institutions. The primary platforms for immunological signaling in the immune system are cytokines and chemokines, which affect cell differentiation, migration, degranulation, and a host of other essential immune activities. We therefore seek to elucidate the profile of 14 of the most prominent inflammation-associated cytokines and chemokines to generate a cytokine network profile at multiple timepoints. This profile provides a multi-dimensional view of the inflammation-associated immune signaling occurring within a given sample over time.

Understanding the patterns of immune signaling in depth within certain contexts will lead to advanced diagnostics that are able to detect the onset of rejection before it is visible to histological analysis and far before organ function is affected. In a clinical context, this would allow pro-active adjustment of immunosuppression treatment to maximize therapeutic effect, maximize pro-toleragenic practices, and minimize potent side effects.

Synopsis of Previous Research Activity

This project has already reported on the application of logistic classifier methods to raw luminex data, utilizing feature selection methods such as Least Absolute Selected and Shrinkage Operator (LASSO). These straightforward methods were able to achieve high classification accuracy rate in differentially diagnosing rejection from pure trauma, with 10-fold cross validation. In the current research period, we build upon these promising results by extending into more informative discriminative models and begin defining methods by which the parameters of generative models can be discerned.

Groups and Data Sets:
In addition to the previously reported Cardiac Transplant (Heart and Lung), and CTA transplant data sets, surgeries for three additional data sets have been completed or are nearing completion by collaborating surgeons and the Luminex readings are currently being analyzed. These data sets are:

- Naive Wound data set, which will provide additional information about the inflammation profiles surrounding surgical trauma and reperfusion injury
- Full-Thickness Skin Transplant, which will provide information about the inflammation profile associated with skin-only transplant. This model is expected to show the most vigorous rejection response, as the skin is highly antigenic and there are no "protective" tissues such as muscle or bone that may ameliorate the acute immune response.
- Skin-Muscle-Bone (SMB) transplant, which will provide information about the inflammation profile associated with transplants that incorporate a skin, muscle, and bone component. The vigour of rejection in this model is expected to be somewhere between Full-Thickness Skin (most vigorous) and CTA (least vigorous)

In these groups, new time points for sampling were devised based upon information gathered from previous experiments that showed a significant amount of predictive immune activity occurred within the first 24 hours post transplant.

Summary of Findings
During this research period, several important findings have been made that are paving the way to a better understanding of immune signaling, a highly accurate diagnostic for rejection before it is visible in histology, and that may lead to novel methods of modulating the immune response in rejection as well as in other forms of inflammation.

We have developed a preliminary method by which we can infer the functional cellular population demographics based on the cytokine mélange as measured by Luminex. This allows us to produce a new type of graph we have termed the "cellulograph." This is a snapshot of the functional cell population as it probably exists in order to have produced the cytokine profile observed (Figure 1).

To create this snapshot, the individual cytokine activity percentages are calculated and each of the 14 measured cytokines are associated with the specific cell type(s) that are known to emit them. For cases where more than one emitting cell type is known, we assume independent identical distribution (IID) of emission, except in cases where rate of emission actually is known. This method produces the cellulograph panel.

We have identified three apparent stages of rejection, labeled here as innate response, adaptive response, and end-stage response (Figure 2). These stages appear in both CTA and solid organ contexts. Importantly, the earliest phases of immune response appear to be driven by cytokines that are produced predominantly by macrophages or other monocytes, indicated the innate response has a key role in the downstream adaptive immune response, although the exact nature of this interaction is not fully understood yet.

It is also evident from the data that there are tissue-specific immune signaling profiles at each stage of rejection. In Figure 3, data from the cardiac transplant model is shown to elucidate this point, because it is the model with tissue samples from the widest range of anatomical locations. However, the finding that each tissue has distinct cytokine profiles holds for all tissue analyzed, in all models. Additionally, patterns of cytokine expression can be found within each tissue type.

To further elucidate the underlying drivers of the observed patterns in immune signaling, and to enhance discriminative classifier performance, feature transformation methods such as Principle Components Analysis were applied to the data and analyzed. As a result, key immune signaling molecules are being identified in each tissue and model that seem to govern the local immune response (Figures 4, 5). For brevity, PCA plots from the cardiac transplant model and the skin wounding model are shown here. Each model and tissue reveal a unique set of governing cytokines.

PCA allows the cytokines to be identified that are the most informative to the immune process. Further, we can understand what immune signaling components are exerting influence on other components, and how strong this influence is. By combining this information with a more detailed analysis of the underlying biological mechanisms of each of these cytokines, and the cells that produce them, we can derive a narrative that explains why the local cytokine milieu takes on the particular form observed in each tissue and immune response. This additional step we term biological function analysis, and is a method of disambiguating cytokine function in context. It also plays a role in identifying parameters to incorporate into future generative models.

Conclusion

During this research period, we have uncovered a wealth of information about the distinct immune response patterns present in different tissues at different time points. We now seek to unify this information under a single analytic framework that will allow accurate tissue specific discriminative models that are capable of rendering predictions with data about local cell populations as well as local cytokine concentrations.

Additionally, the PCA and other feature transformation methods have allowed us to begin defining parameters and modeling rules for preliminary generative models. These generative models will provide a clinically relevant diagnostic tool, and may enable in-silico simulations of immune response modulation under various projected treatments, thus providing truly personalized immune therapy guidance. These models may also find use as tools for discovery of novel, non-toxic methods of immune modulation.

Table 1: Listing of pro-inflammatory cytokines measured by Luminex.

IFNy	IL-1a	IL-1b	IL-2	IL-4	IL-5	IL-6
IL-10	IL-12p70	IL-18	MCP-1	GRO/KC	TNFa	GM-CSF

Table 2: Rat surgical models and tissue types that this project is analyzing

Surgical Model	*Tissues Sampled*	*Time Points Sampled*
Hind Limb CTA	Skin, Muscle, Serum	POD 3,5,7,9,11
Heart and Lung Transplant	native heart, graft heart, native lung, graft lung, lymphnode, spleen, kidney, liver and serum	5 hours, POD 1,3,5,7,9
Naïve Wound	Skin, Muscle	N/A
Full-Thickness Skin Transplant	Skin, Muscle, Serum	Hours 6, 12, 24 POD 3, 21, 45
SMB Transplant	Skin, Muscle, Serum	Hours 6, 12, 24 POD 3, 21, 45

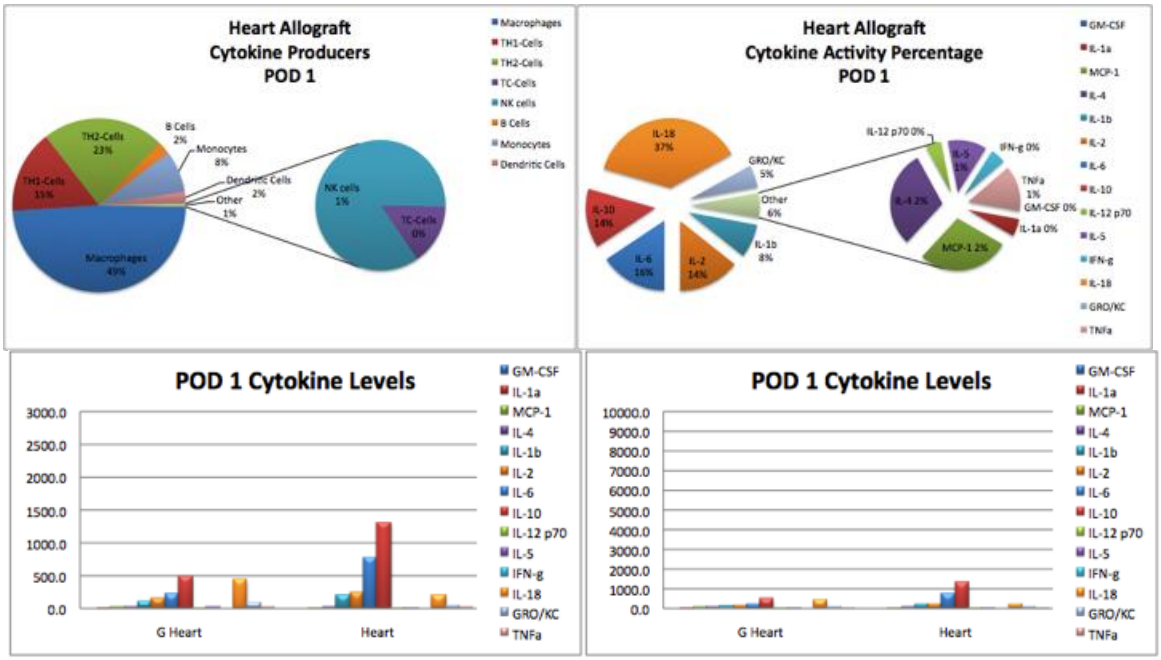

Figure 1: Cellulograph panel inferring cell population from heart allograft cytokine levels

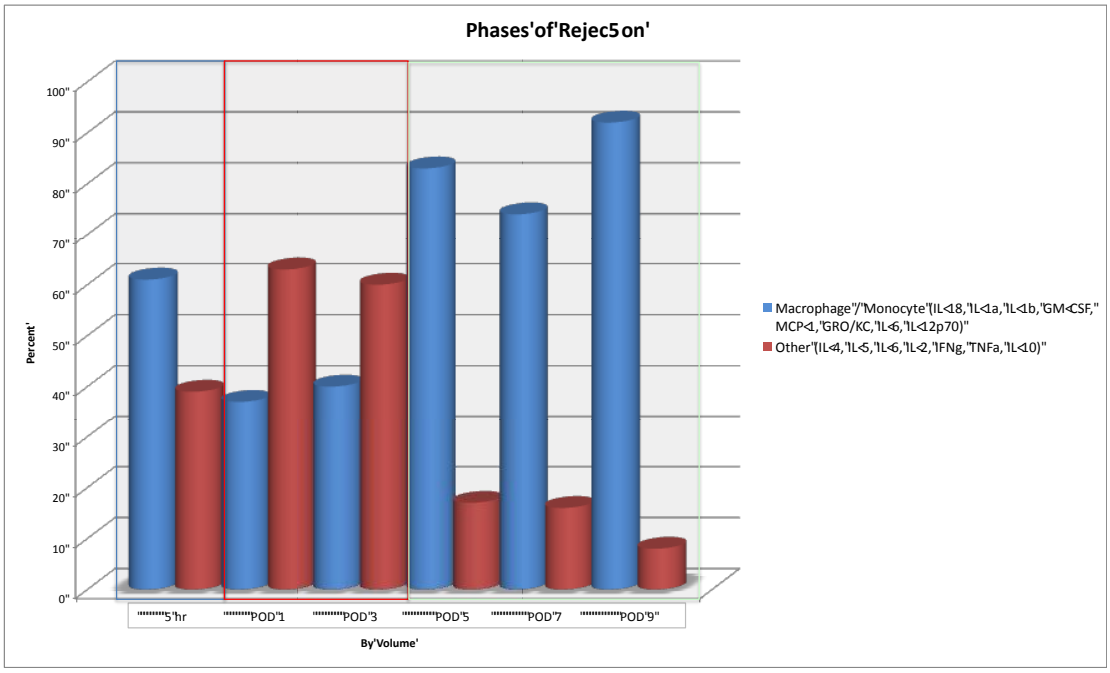

Figure 2: There appear to be at least three phases in the course of rejection. This graph is of data from the cardiac transplant series, however the same apparent "phases" are evident in all data sets analyzed.

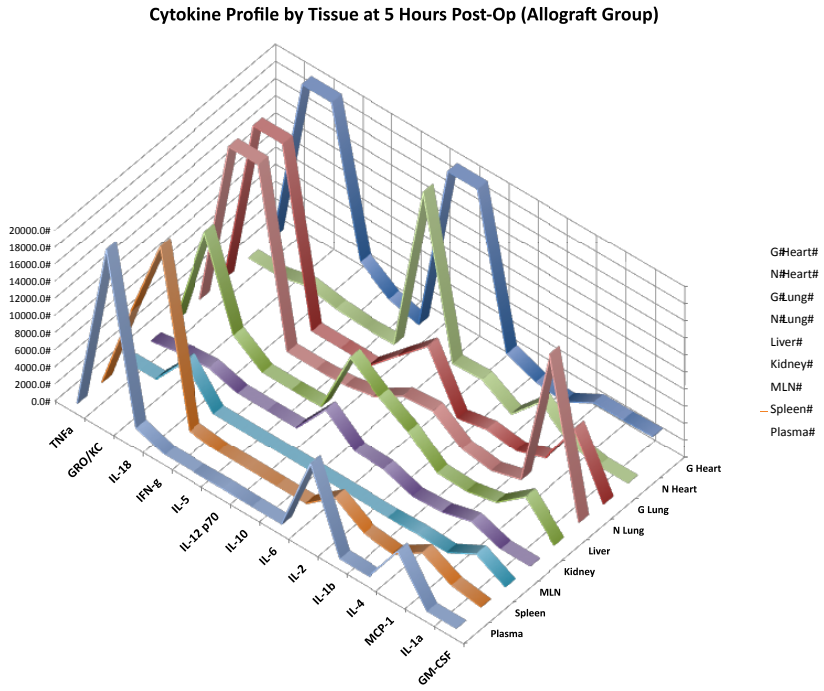

Figure 3: Distinct cytokine network profiles are evident for each tissue in this graph of the cardiac transplant model at 5 hours post op.

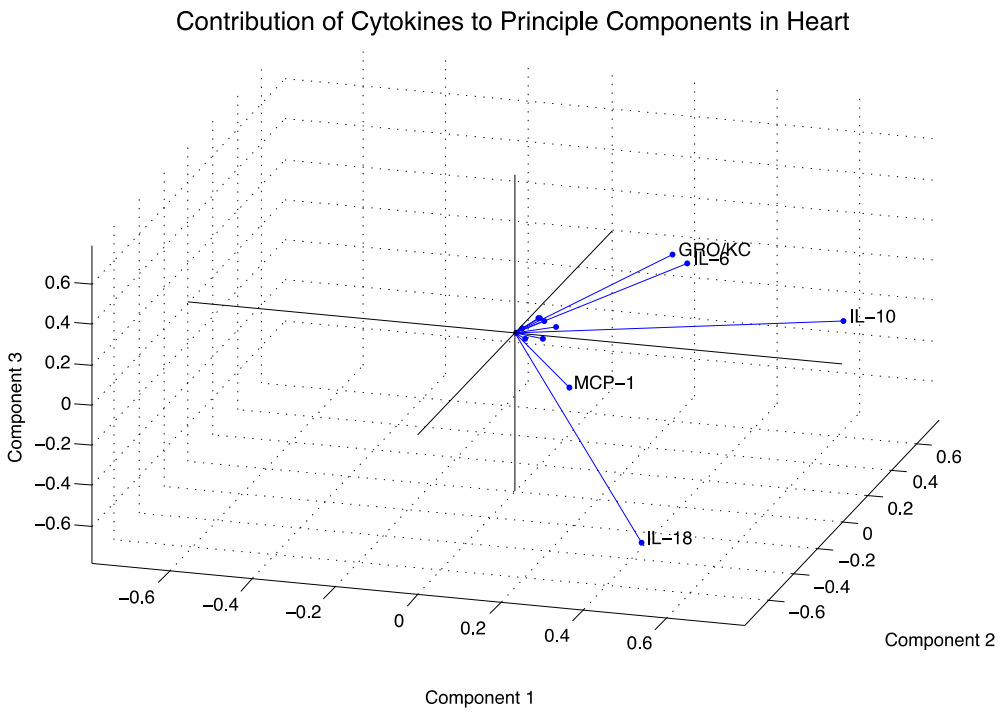

Figure 4: Contribution of cytokines to each of the three principle components in the graft heart from the cardiac transplant model. Three principle components explain >95% of variance. Length and direction of vectors indicates contribution to principle components.

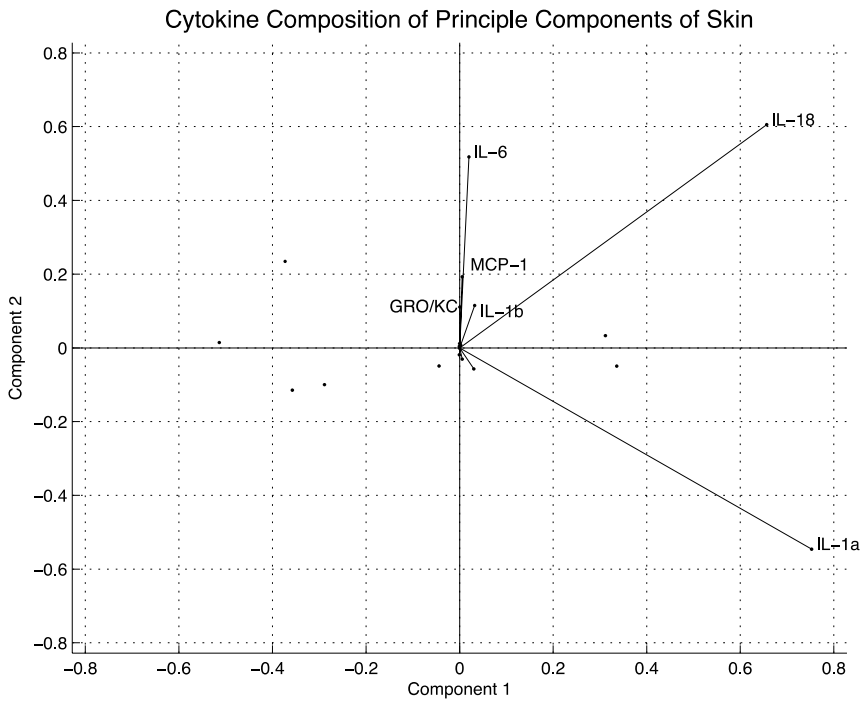

Figure 5: Contribution of cytokines to each of the two principle components in the graft skin from the skin wound model. Two principle components explain >95% of variance. Length and direction of vectors indicates contribution to principle components.