Carnegie Mellon University

Annual Progress Report: 2008 Formula Grant

Reporting Period

July 1, 2009 – June 30, 2010

Formula Grant Overview

The Carnegie Mellon University received $747,818 in formula funds for the grant award period January 1, 2009 through December 31, 2011. Accomplishments for the reporting period are described below.

Research Infrastructure Project 1: Project Title and Purpose

Mellon Institute Vivarium – Research Infrastructure - Many faculty in the Mellon College of Science, which includes the departments of Biological Sciences, Chemistry, Mathematics and Physics, as well as several associated interdisciplinary centers, conduct research that uses vertebrate animals. This infrastructure project will contribute to the design and construction of a new vivarium, the facility in which those animals are kept. It is located in the historic Mellon Institute building.

Anticipated Duration of Project

1/1/2009 – 9/30/2011

Project Overview

The new Mellon Institute Centralized Vivarium (MICV) will be used by a variety of Carnegie Mellon University (CMU) investigators for projects that involve the use of rodents for biomedical research. The primary purpose of the facility will be to facilitate the housing and breeding of mice and rats, including unique strains of genetically modified animals.

Specifically, this project will provide partial support for the vivarium during construction, Phase 1, and initial operation, Phase 2. Through an open bid process, Carnegie Mellon has engaged the services of CUH2A, a firm with specialized knowledge of vivarium construction and operations. CUH2A has assisted in the planning and design phases of the project and is now engaged in preparing architectural drawings to be used to seek competitive bids on this construction project. Bids are expected to be completed in early April with construction beginning in May 2009. Construction should be completed in March 2010. In Phase 2, once construction is completed, the vivarium will require several months to become operational. The initial population must be introduced in a controlled manner which includes sentinel testing. This project will assist in funding construction and initial operations.

Principal Investigator

Fred Gilman, PhD
Dean, Mellon College of Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Other Participating Researchers

Nathan Urban, PhD, Max Dorosa, Sharon McCarl - employed by Carnegie Mellon University

Expected Research Outcomes and Benefits

The new facility will offer a high level of protection and isolation so that unique transgenic strains of animals will be protected from pathogens and will offer much improved containment of biohazardous animals. The MICV also will be equipped with areas designed for mouse and rat surgery, behavioral testing, in vivo physiological recording as well as whole animal and in vivo brain imaging.

A major goal of the facility is to reduce the risk of infection in the animals housed at CMU so that we are better able to preserve precious animals generated and used at CMU. To this end, separate areas for quarantine and housing of animals that have been manipulated in labs outside the facility will be built and maintained.

A contemporary animal facility is essential to keep our life sciences researchers competitive in grant proposals to NIH and other funding agencies. An important benefit of centralizing the vivarium is that it will foster interactions among research teams that inevitably lead to new collaborative projects.

Summary of Research Completed

In the past year, Carnegie Mellon University has made substantial progress on the design and construction on the Mellon Institute Centralized Vivarium (MICV). This main facility is being built in space located on the second floor of the nine-story Mellon Institute (MI) building. The MICV project will provide newly renovated space for dedicated vivarium, laboratory and cage processing, HVAC supply and exhaust systems, and fixed equipment. The project area includes renovated spaces on the second floor, a dedicated mechanical room on the third floor, new dedicated exhaust risers within the existing interior courtyard light well, and roof area for ductwork and fans to convey and exhaust air out of the building.

Partial selective demolition was performed beginning in 2008, leaving the space as vacant, shell space. In 2009, Mosites Construction Company was hired as a construction management firm and Burt Hill Associates as a commissioning agent. Construction has been ongoing since. This work has included selection of subcontractors, coordination of design and selected equipment and substantial amounts of construction. Currently we are nearing the completion of all work

outside the second floor vivarium space, including the third floor mechanical room space and also space on the second floor adjacent to the main vivarium space. A major event in this construction will be the installation of a 72" diameter exhaust riser in the Mellon Institute courtyard that will span all none floors of the building.

For this progress report we summarize below the status of the major components/systems of the project.

HVAC – The HVAC system is the single largest component of this project and its proper function is critical to the proper operation of this facility. Two main air handling units for the facility have been fabricated, installed and tested. These units will provide 50% redundant capacity to this space. The exhaust riser duct, which will be installed in the Mellon Institute courtyard, is nearly completely fabricated and will be put in place on the 1st of August . Following this placement, this exhaust riser will be anchored, fans will be installed and this riser will then be connected to the HVAC system of the MICV space. Ductwork within the space is ~ 80% complete including valves and controls. We anticipate that HVAC systems will be essentially complete and will undergo preliminary balancing in early September. This preliminary balancing is critical to ensuring that the system performs as per design specifications. Installation of ceilings and other finishes will not be performed until after the results of this balancing are known so that any problems identified can be easily fixed while access to ductwork values and controls is optimal.

Plumbing – Plumbing for lab water, lab gasses, animal watering system, drains is ~ 80% complete. Floor penetrations to spaces above and below are complete, rough-ins for equipment, sinks, etc are in progress.

Electrical- Electrical systems including lighting, power for equipment etc is ~ 60% complete. Lighting installation will await completion of ceilings. Substantial effort was required to relocate building data cables that run through the vivarium space and which will be sealed into the vivarium ceilings.

Carpentry – Stud walls are mostly (>80%) framed. Door frames are being installed now. Drywall and ceiling installation will begin in the next few weeks.

Finishes and Flooring – Floor areas were cut to accommodate large equipment including bulk autoclave and cage/rack washer. Flooring, ceilings and wall finishes have been chosen and installation will begin as soon as possible after preliminary balancing of HVAC system is complete.

Fire protection/Security – Permitting and code issues have been resolved and fire protection subcontractor has been involved in coordinating plumbing and electrical work.

Equipment – Most (75%) equipment has been ordered and delivery times are being coordinated between vendors and construction schedule. Cage systems, rack and cage washer, biosafety cabinets and autoclaves are the biggest items and all these have been ordered and are in production now.

The current schedule calls for substantial completion of construction and equipment installation by the first week in January 2011. This will be followed by about two months of commissioning and testing of the facility at which point it will be ready for occupancy.

Research Project 2: Project Title and Purpose

Computational and Neural Basis of Visual Inference - The purpose of this project is to understand how neurons in the visual cortex work together to resolve ambiguity during object perception from both a computational and a neurophysiological perspective. The basic idea is that prior visual experiences, encoded in connections among neurons, constrain the visual computations underlying the interpretations of visual scenes. We will test this idea by monitoring the activities of neurons in the earliest visual area (V1), an intermediate visual area (V4), and the highest visual area (IT) of the primate visual system. The computational studies will assess how robust visual inference can be accomplished in the framework of hierarchical Bayesian inference. The electrophysiological studies will study the dynamics of neuronal activities during visual processing to test the computational hypotheses.

Anticipated Duration of Project

1/1/2009 – 6/30/2011

Project Overview

This project combines computational and neurophysiological approaches to attack a core problem in neuroscience: how networks of neurons work together to resolve ambiguity during perceptual inference. The investigators propose to study this question in the context of object perception in the visual cortex. Computational studies on figure-ground segregation (Aim 1) and object recognition (Aim 2) will be carried out to elucidate the computational principles and mechanisms underlying object perception. Large scale neuronal recording using single electrodes and multi-electrode arrays will be carried out to understand how objects are represented and processed in neuronal population activities in the primary visual cortex V1 and in areas V4 and IT (Aim 3). The computational studies will be carried out in the framework of hierarchical Bayesian inference using factor-graph and Bayesian belief propagation to explore the specific hypothesis that robust visual processing can be achieved with *analysis by synthesis.* The neurophysiological studies will test this hypothesis by examining the temporal dynamics of V1 and IT activities to determine the nature of bottom-up and top-down interaction during visual processing. This effort will provide valuable new information into the computational and representational roles of V1, V4 and IT in object perception. More generally, it will provide valuable insights into how neurons represent statistical information and how neuronal circuits perform probabilistic reasoning.

Principal Investigator

Carl R. Olson, PhD
Professor
CNBC

Mellon Institute 115
4400 Fifth Avenue
Pittsburgh, PA 15213

Other Participating Researchers

Tai Sing Lee, PhD - employed by Carnegie Mellon University

Expected Research Outcomes and Benefits

These experiments will cast light on how neurons in the primate visual systems perform inference. The information will be directly relevant to understanding several debilitating disorders of vision from which patients with congenital or traumatic brain pathology suffer, including topographic amnesia (an inability to recognize formerly familiar places) and prosopagnosia (an inability to recognize formerly familiar faces). Advances in theories of visual inference will provide new methods of processing visual information. In addition, progress in these areas will lead to the design of better algorithms for computer vision which may one day allow computers and robots to interact with the natural sensory environment the way we do. Finally, by better understanding the properties of visual representations and the computational principles of the visual system, we will gain deeper insight into the neural basis of object perception.

Summary of Research Completed

Computational study: development of object representation

To understand better how higher order visual concepts can be learned statistically from natural scenes and serve as a prior for probabilistic perceptual inference, Tai Sing Lee and his colleagues have studied a class of Bayesian non-parametric methods for unsupervised hierarchical learning of invariant visual concept. Our working hypothesis is that higher extra-striate cortical neurons like V2, V4, and IT encode more abstract visual concepts with successively higher degree of invariance. At a higher level, a visual concept encompasses its multiple attributes or realizations. By visual concept, we could mean an object, a shape or an object part. Each concept defines a probabilistic distribution of its realization, for example, an object can be defined as a linked network of many of its views under different viewing angles or lighting conditions, similar (but much denser) in a way to the classical aspect graph in computer vision. We hypothesize each visual concept can have an indefinite number of states and each state can explain a certain observation or realization, and the transition between these states are learned from temporal continuity or similarity of the observations. A natural but relatively novel way to model such a concept is to use a hidden Markov model (HMM), with different visual concepts modeled by distinct HMMs. States within a concept can transit to one another, but not to states in another concept. Thus, the HMM transition matrix exhibits a block diagonal structure (see Figure 1). Given the number of states and the number of concepts possible is unknown in advance or infinite in theory. An appropriate framework suited for learning such structures is the recently developed infinite HMM (iHMM), a realization of discrete-state hidden Markov models based on hierarchical Dirichlet process (HDP). The iHMM is like an ordinary HMM, but the procedure

that "fits" it to data allows it to allocate a variable number of states to describe the data. The existing HDP framework allows an infinite number of states, but there is no built-in mechanism to guide or encourage states to form clusters or concepts. We develop a new framework as a generalization of the iHMM to allow flexible introduction of reasonable priors to guide the unsupervised learning or development of concepts, each of which is a meaningful cluster of states learned based on various Gestalt principles such as spatiotemporal continuity, similarity in various visual cues and attributes (Stepleton et al. 2009). At some level, this is a clustering technique, an alternative to standard techniques like normalized cut or K-mean, but potentially more flexible and powerful. Furthermore, this is a generative approach, allowing top-down influence in a hierarchy that might be more appropriate for modeling the visual cortex from our perspective (Lee and Mumford 2003).

Our generalization of the infinite HMM (Stepleton et al. 2009) offers several new capabilities: (1) it can isolate distinct behavioral regimes in the dynamics of the temporal processes, i.e. it can identify sub-behaviors such as recurring motifs in music, repeated gestures in a dance, common words and phrases in human speech, and different views of an object; (2) it can partition data sequences into segments corresponding to the times when these different sub-behaviors are executed. This model can be used to model the development of neural structures for representing different visual concepts, with different observations of the same concept linked together, but observations or realizations of different concepts, despite their apparent similarity in other cue dimensions, could be kept apart. Figure 2 illustrates the learning of different gestures in a video game (batting, boxing, pitching and tennis swing gestures) in which each gesture can be considered a concept. Each point in the graph is a low-dimensional projection of a particular view. Views that are very dissimilar in appearance can be linked together by virtue of the temporal continuity of their appearances and deemed to belong to the same concept, while views that are only slightly different can be kept distinct as different concepts because they have not been found to transit into each other in our dynamic visual experience.

This approach is potentially very powerful with a wide range of applications in many fields, not just for an innovative conceptualization of object representation and learning in the brain. The several applications that we have considered are simply scratching the surface of its potential.

This work resulted in a presentation in the *International Conference on Artificial Intelligence and Statistics,* and a journal publication in the *Journal of Machine Learning Research.*

- Stepleton, T., Ghahramani, Z., Gordon G., Lee, T.S. (2009) The Block Diagonal Infinite Hidden Markov Model. *Journal of Machine Learning Research: 5: 544-551.*

Neurophysiological studies: statistical learning in IT

The general idea behind this project is that functional connections among neurons in inferotemporal cortex (IT) change according to the statistics of visual experience so as to maximize the efficiency with which images are processed. The specific idea we tested in experiments during the preceding year was that if an image is highly predictable then IT neurons will cease to represent it strongly because it conveys no new information. We explored this possibility by exposing monkeys repeatedly to images presented in a certain sequence, so that the first became a strong predictor of the second, and then asking whether IT neurons responded less

strongly to an image when it was predicted than when it was not. The answer was in the affirmative. We have submitted this study for publication.

We began by exposing two monkeys repeatedly to six pairs of images (Fig. 3a). On each trial, while the monkey looked at the center of the screen, two images appeared in succession at fixation for half a second each. The members of each pair were always presented in the same order so that the leading image became a strong predictor of the trailing image. Training involved viewing each sequential pair of images over 800 times. The training runs were distributed across three months in monkey 1 and one month in monkey 2. Following training, we began to collect data from single ITC neurons in microelectrode recording sessions. The procedure employed during testing was identical to the training procedure with the exception that trials involving untrained sequences were interleaved among the trials involving trained sequences. An untrained sequence consisted of a leading image and a trailing image that belonged to different training pairs. So as to minimize any attenuation of the effects of training, we adopted a design in which each of the six trained sequences was presented eight times in a full run whereas each of the 30 untrained sequences was presented only once.

Using this procedure, we collected data from 81 visually responsive ITC neurons (46 in monkey 1 and 35 in monkey 2). ITC neurons responded more strongly when the trailing image followed another image's predictor than when it followed its own (Fig. 3). This effect achieved significance in 33/81 individual neurons (ANOVA with prediction status and stimulus identity as factors and firing rate 50-500 ms after stimulus onset as the dependent variable, $\alpha = 0.05$). The effect emerged clearly at the level of population activity. The firing rate elicited by the trailing image was greater on average when it was unpredicted (mean = 16.8 Hz) than when it was predicted (mean = 13.0 Hz) and the difference was highly significant (paired t-test, $n = 81$, $p = 1.5e^{-11}$; $p < 0.0005$ in each monkey). Thus neurons in IT robustly signal when a prediction based on the statistics of past experience has been violated.

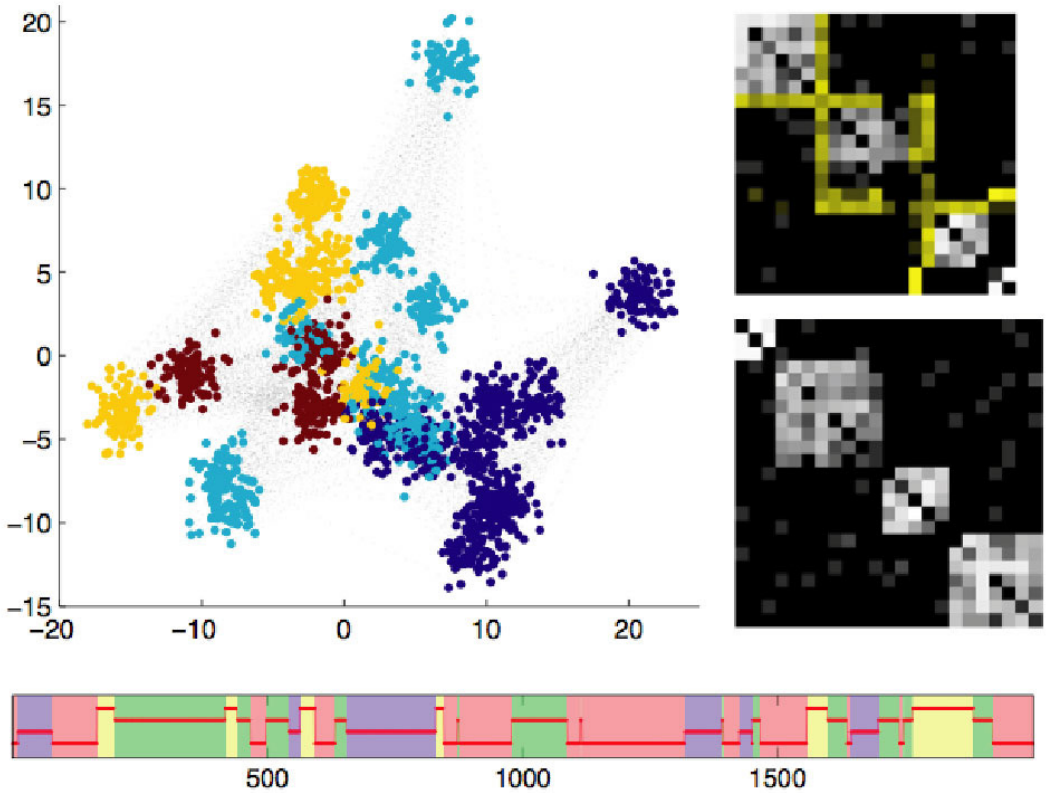

Figure 1: Top left: synthetic data for the BD-IHMM inference task in §6.2.1; colors mark the four sub-behaviors. At right, visualizations of the transition count matrices inferred for this data by the IHMM (top) and the BD-IHMM (bottom)—by recognizing sub-behaviors, the BD-IHMM correctly identifies separate states that the IHMM conflates (yellow). Below, a timeline relating the sub-behavior inferred for the training set sequence (red line) with the ground truth (background colors).

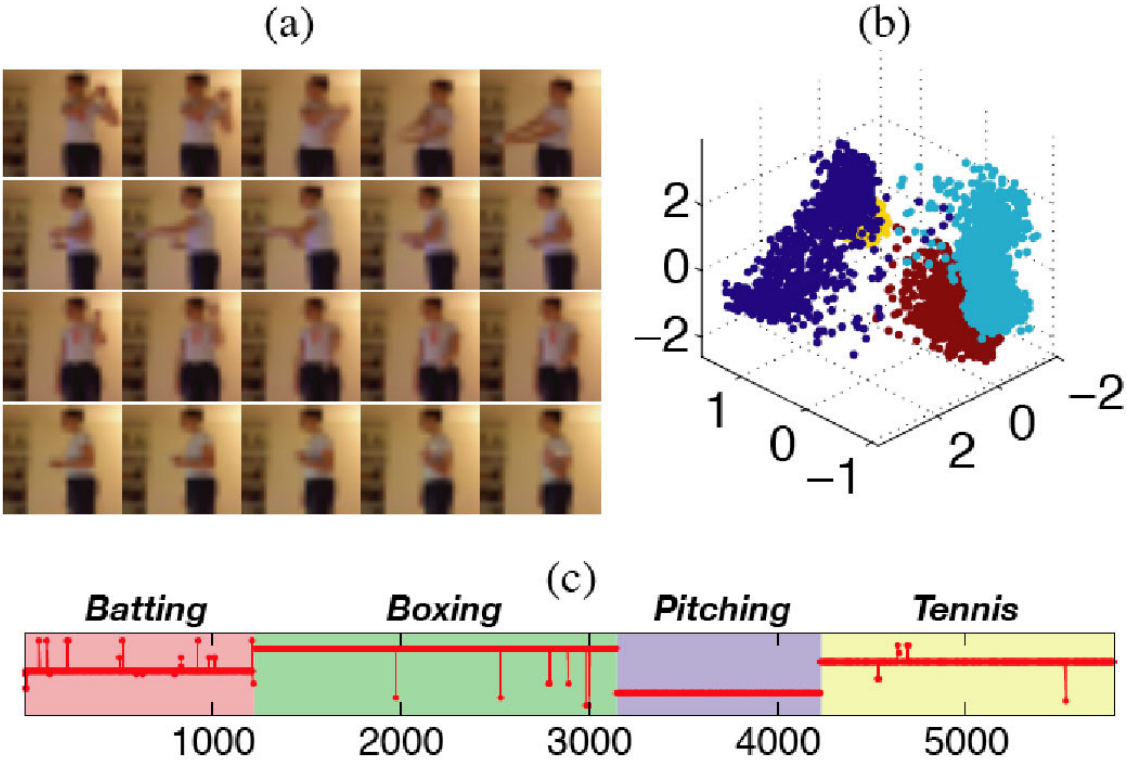

Figure 2: Top left: selected downscaled video frames for (top to bottom) batting, boxing, pitching, and tennis swing gestures. Top right: first three dimensions of video frames' PCA projections, colored by gesture type. Below: a timeline relating the sub-behavior inferred for the training set sequence (red line) with the ground truth (background colors). Labeled sub-behavior regions actually comprise numerous video clips of the same gesture.

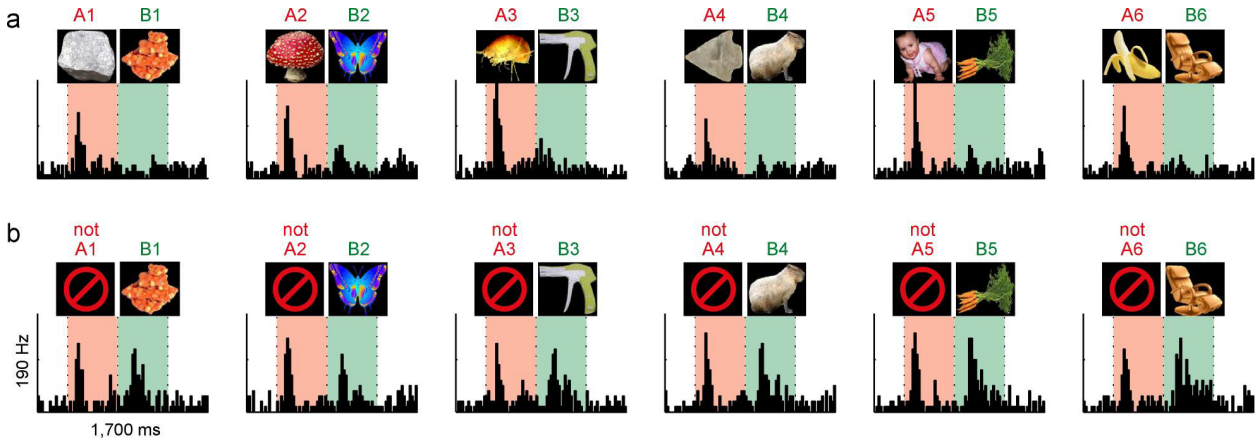

Figure 3. Visual prediction error signals in a representative neuron. a, Predicted images elicit weak responses. Each histogram represents mean firing rate as a function of time across eight trials in which the leading image (A_m) and the trailing image (B_n) were paired as they had been during training ($m = n$) with the consequence that the trailing image was predicted. The leading and trailing images were visible during periods demarcated by hatched lines and shaded with red and green respectively. b, Unpredicted images elicit strong responses. Each histogram represents firing on trials in which the leading image (A_m) and the trailing image (B_n) were not paired as they had been during training ($m \neq n$) with the consequence that the trailing image was unpredicted. Each histogram represents data from five trials beginning with presentation of the five leading images not paired with the trailing image during training.

Research Project 3: Project Title and Purpose

Memory-Based Neural Activity in the Hippocampus - The hippocampus, located in the temporal lobe of the brain, is known to be involved in both memory formation and the representation of spatial information. Early studies in rats showed that the activity pattern of hippocampal cells correlates strongly with the animal's current location in its environment. But later experiments have revealed that hippocampal activity can be much more dynamic, replaying trajectories the animal experienced previously, or producing trajectories that correspond to a path it is anticipating taking. This project will analyze data on hippocampal trajectory encodings and explore computational mechanisms to account for the phenomenon.

Anticipated Duration of Project

1/1/2009 – 12/31/2010

Project Overview

Understanding how cognitive processes arise from neural tissue involves, in part, determining the computational roles played by various brain structures. The hippocampus is one of the major structures of the brain, receiving inputs from much of the cortex and from several important subcortical areas, such as the basal ganglia and amygdala. It is believed to be important for both memory formation and spatial representation. In the rodent, hippocampal pyramidal cells have compact spatial firing fields, called "place fields," suggesting that the hippocampus might be specialized for encoding the rodent's location on a "cognitive map".

Recent experiments have shown that the hippocampus encodes more than just the animal's current location. Under some circumstances, the hippocampus "replays" a trajectory the animal recently experienced; and in a process called "vicarious trial and error", which can occur when the animal pauses at a choice point in a T-maze, the hippocampus can "play" a trajectory corresponding to travel down one of the two maze arms. The goal of this project is to study these trajectory events in order to arrive at a better understanding of hippocampal information processing.

Two approaches will be used. The first is to apply statistical analysis techniques to multiunit recording data (obtained from other laboratories) to reconstruct and characterize trajectories. The second is to develop computational models to try to replicate and thus explain the experimental data. Attractor-based neural network models have been used to explain many aspects of hippocampal function and are the leading candidate for modeling these trajectory phenomena.

Principal Investigator

David S. Touretzky, PhD
Research Professor
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

Other Participating Researchers

Anoopum S. Gupta, BS – employed by Carnegie Mellon University

Expected Research Outcomes and Benefits

When the brain temporarily loses its oxygen supply due to a stroke, heart stoppage, drowning, etc., the hippocampus is among the first regions to be affected. Hippocampal deterioration is also a major factor in Alzheimer's disease. Persons with hippocampal damage exhibit marked cognitive effects, including memory deficits and impaired learning. A better understanding of what the hippocampus does will improve the prospects for aiding patients with hippocampal damage. Understanding hippocampal function will also contribute to the fundamental scientific problem of how brains learn and remember things.

The data analysis and computational modeling work to be performed in this research project will focus on how the rodent hippocampus represents trajectories through a familiar environment. While there has been much work on hippocampal encoding of individual spatial locations, trajectory data, in which the hippocampus replays a sequence of locations, has only recently become available.

Summary of Research Completed

We completed revisions to the manuscript on hippocampal replay of learned and novel

trajectories that was mentioned in last year's progress report. The article has now been published:

> Gupta, A., van der Meer, M. A. A., Touretzky, D. S., and Redish, A. D. (2010) Hippocampal replay is not a simple function of experience. *Neuron*, 65(5):695–705.

Based in part on this work, Anoopum Gupta, an MD/PhD student in the Robotics doctoral program, was awarded an NIH (National Institutes of Health) F30 graduate training fellowship that will begin July 1, 2010.

We have turned our attention from data analysis to the computational modeling component of the project, and have a significant new result: a model of backward sequence learning. When a rat is in an awake but paused and inattentive state, the hippocampus switches from theta rhythm to a state known as LIA (Large-amplitude Irregular Activity), punctuated by "sharp waves" in which many hippocampal neurons fire together. Analysis of sharp wave activity patterns reveals neural firing sequences representing trajectories the rat has experienced (forward replay). But there are also sequences corresponding to the reverse of these trajectories (backward replay). Sequences are assumed to result from a chain of synaptic connections, e.g., observing the sequence A-B-C would suggest that there is an excitatory connection from neuron A to neuron B, and a similar connection from neuron B to neuron C. The synaptic learning rule for potentiating a synapse from A to B requires that neuron A fire before neuron B. Since the rat only experiences the trajectory in forward order, how could the brain learn the connections from C to B and from B to A that would allow it to replay the sequence in reverse order?

Our proposed solution to this puzzle is to exploit the recently reported theta rhythm phase gradient along the longitudinal (septo-temporal) axis of the hippocampus (Lubenov and Siapas, 2009). While cells at all levels are firing in forward sequence as the rat experiences a trajectory, the firing of more temporal cells can be delayed by as much as 120 ms relative to cells at the dorsal pole. Since pyramidal cells in hippocampal area CA3 have wide-ranging projections, it is possible for dorsal cells representing B to acquire excitatory projections to more temporal cells representing A, as shown in the figure below (see red connection from B_2 to A_6). Forward connections can be learned at the same time (e.g., green connection from A_6 to B_6).

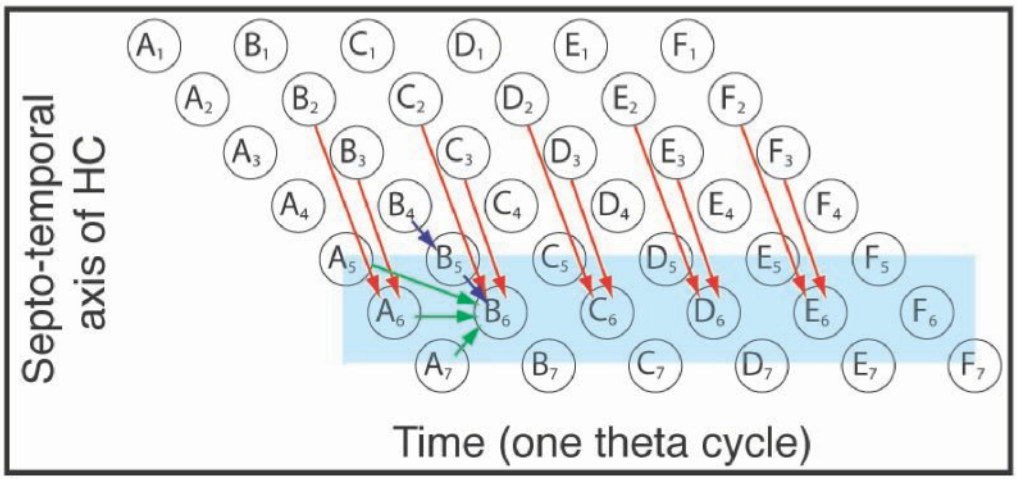

We developed a computer simulation of hippocampal learning that includes the theta rhythm phase gradient and longitudinal connections, and verified that both forward and backward connections are learned when the model experiences only forward trajectories. The figure below shows the learned connection strengths between a representative cell at each of 10 septo-temporal levels and all other cells in model. Color distinguishes forward from backward connections. Connections in both directions are evident at all levels.

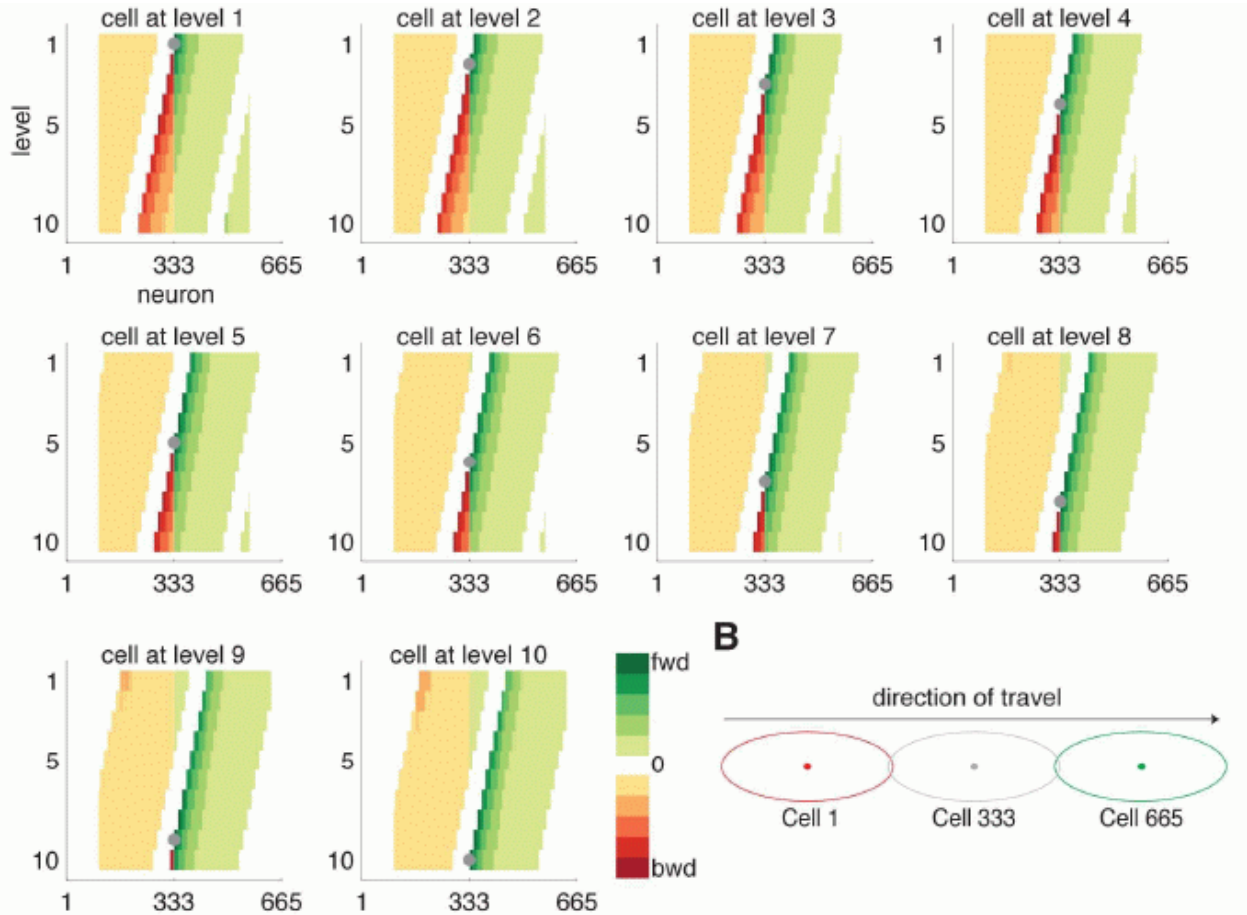

We then developed a separate model of hippocampal sharp wave replay, and showed that if the connection weights followed the pattern constructed by the first model, sequences could be replayed in either forward or backward order.

A manuscript describing this novel result is nearly ready for submission.