Carnegie Mellon University

Annual Progress Report: 2008 Formula Grant

Reporting Period

July 1, 2010 – June 30, 2011

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

The Mellon Institute Centralized Vivarium (MICV) is nearly complete. We began moving in sentinel animals, and are currently waiting the results of sentinel testing before finalizing operating procedures and introducing the remaining animals. For this progress report we summarize below the progress made on the major components/systems of the project.

HVAC – System installation was completed and the system was balanced.

Plumbing – Sinks and wash basins were installed.

Carpentry – Stud framing was completed, followed by door frames. Dry wall and ceilings were installed.

Electrical – Data cables running through the space were re-routed. Lighting was installed.

Flooring – All flooring was installed

Equipment – The remaining (25%) equipment was been ordered and delivered. All equipment, including cage systems, rack and cage washer, biosafety cabinets and autoclaves, and benches were installed.

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

1. Learning nonlinear feature mapping using Generative and Discriminative models.
We have developed a new approach to learn features for classification using a combination of generative models and discrimination models. Perceptual inference, such as recognizing objects or 3D surface geometry of objects, requires features that are both specific enough to allow for discrimination of different classes of objects, or different surface geometries, and yet able to tolerate image variations due to changes in viewpoint and illumination, as well as variations and deformation of shapes and texture within each class. Learning the appropriate features that are both specific and invariant for object discrimination has been a difficult problem. There have been proposals that perceptual inference in the brain is accomplished by the interaction of a bottom-up feed-forward discriminative process that *extracts* features for classification and a top-down generative process that produce "perceptual proposals" to *explain away* the features. An important question is to what extent the learning of the features used in the discriminative models are themselves influenced by the generative feedback process. In this study, we explored the use of generative models to derive a nonlinear mapping that map image input to a high dimensional nonlinear feature space based on the principle of maximizing the divergence of the posterior probabilities of the different classes. The framework builds feature mapping based on 3 complementary measures (corresponding to the three components of the computational model):

(1) how much a sample image affects the model, (2) how well the model explains the sample, (3) how uncertain is the model's explanation of the sample. The resulting "features" when fed into the standard classifier such as support vector machine produce superior performance in classification performance in two classical vision benchmark applications: (1) scene recognition, (2) face vs non-face discrimination. This work is published in [Li 2011] IEEE Conference in Computer Vision and Pattern Recognition (CVPR) in 2011, a prestigious computer vision conference. Establishing the computational competence of this method is important not just for computer vision, but also for neuroscience, as it is possible the brain might likewise utilize a similar hybrid scheme of discriminative and generative models to learn features for classification. This class of hybrid methods for feature learning can provide concrete hypotheses of feature representation for testing neurons in the higher order visual cortex.

Figure 1 illustrates how our method (PD posterior divergence) performs in scene classification (blue line) relative to other state of the art methods as a function of number of scene categories. It demonstrates that our method has the best performance up to 50 classes, and remains close to the top in recognition rate beyond 50 classes (top). Figure 1 bottom panel illustrates the contribution to performance by the different components of our model. Figure 2 illustrates how our method (blue line) performs in face/non-face classification relative to other state-of-art methods (top) and the contribution to performance of the different components of our model.

2. Elucidation of cortical network using Generalized Linear models.

Perceptual inference relies on statistical priors we learned from the natural environments. These priors allow surround contextual information to help disambiguate local inference computation. In computer vision, network encoding statistical priors are formulated in terms of Markov random field or factor-graph, which are nodes with horizontal interactions. We hypothesized the local functional circuitries in the primary visual cortex can be modeled in terms of Markov random field. To evaluate these conceptualizations, we performed multi-electrode array recording studies in the primary visual cortex of monkeys to look into functional circuitry and coding of the neurons. We found that (1) neurons coding for depth interact in a way that is broadly consistent with the statistical constraints (specifically the continuity constraint and the uniqueness constraint) proposed by David Marr 30 years ago [Samonds 2009]. That is, neurons with similar disparity tunings cooperate across space, and neurons dissimilar in disparity tunings compete at each spatial location; (2) a neuron's activities can be predicted better by knowing the other neurons' activities than by the visual input in its receptive field, indicating extensive network interaction. Both of these studies support the notion of extensive horizontal interaction between neurons in the visual cortex to mediate perceptual computation. Definitely elucidating the network structures and comparing them with constraints learned from natural scenes is a long-term goal of our laboratory.

Figure 3 shows an example of a cell from the array population analyzed using Generalized Linear Models (GLM). The key finding is that knowing other neurons' responses will allow the model to predict a neuron's response much better than knowing the input stimulus alone. This suggests neurons tend to function in a network, rather than as an independent agent. Figure 4 (left) illustrates, with the example of one neuron's activity, how a full model (that includes network interaction) can predict the PSTH (peri-stimulus time histogram) of the neuron's response better than a model that is simply based on stimulus input. Figure 4 (right panel) shows

statistics of the improved performance of the full model over the stimulus-only model for the population of neurons.

3. Statistical learning in IT.

We have continued to explore in experiments based on single-neuron recording in monkeys an extremely robust phenomenon, involving visual inference, that we discovered in the course of this project. If a monkey is exposed repeatedly to a pair of images presented in immediate succession so that the leading image becomes a strong predictor for the trailing image, then the functional properties of neurons in inferotemporal cortex (IT) change. They virtually cease to respond to the predicted trailing image when it is presented in the trained context, which is to say, when it is predicted, but they continue to respond strongly to the image when it is presented outside the trained context, which is to say, when it is not predicted. The obvious interpretation of this effect is that visual circuitry leading up to IT is capable of drawing inferences from visual events (inferring from the occurrence of a given leading image that a given trailing image will follow) and uses this ability to suppress the representation of events that, because they are predictable, convey no new information. During the year under report, we carried out further analyses on data from initial experiments. Two key conclusions arising from these analyses are: (a) that the effect depends on the order in which the two stimuli are presented and (b) that strong responses to an unpredicted stimulus depend simply on its being unpredicted and not on the degree to which it deviates in nature from the predicted event. In addition, we launched new recording experiments designed to cast light on the effect by characterizing its spatial selectivity (must the two images occur at the same location) and its temporal selectivity (is it specific to the interval between leading and trailing stimulus that was imposed during training).

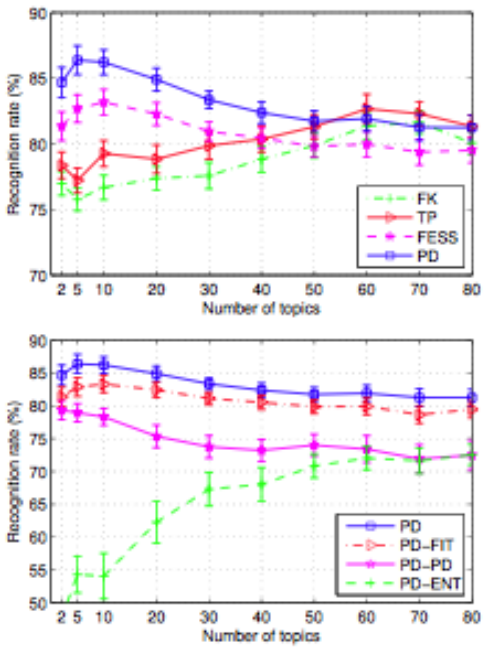

Figure 1. Performance comparisons on the PLSA model for scene recognition. The number of topics is variable. Four methods (Top) and the three components of PD (Bottom) are evaluated.

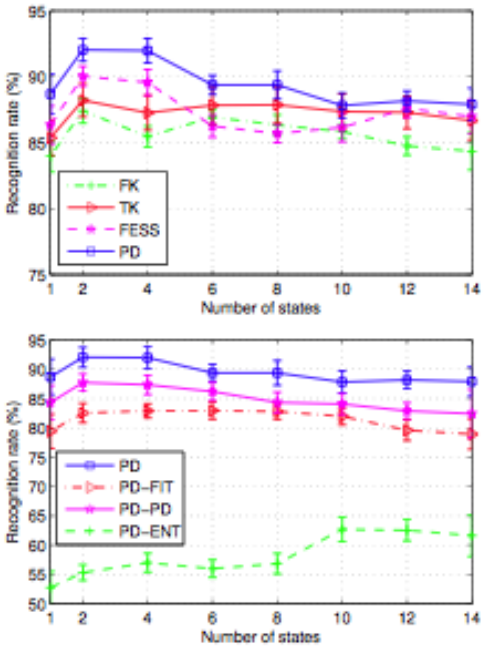

Figure 2. Performance comparisons of face and non-face classification. The number of states is variable. Four methods (Top) and the three components of PD (Bottom) are evaluated.

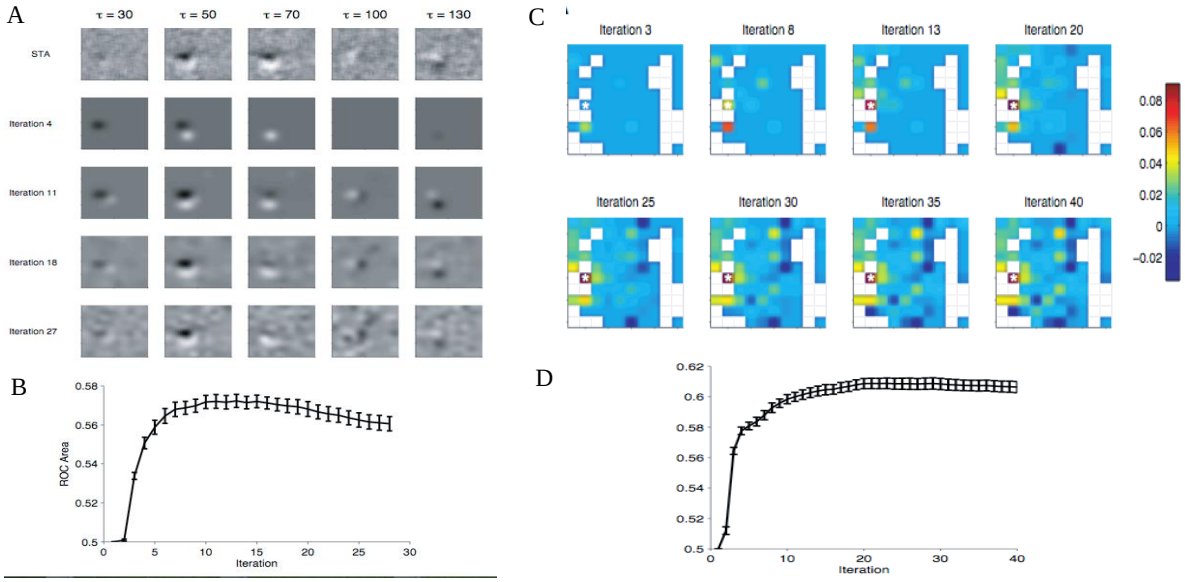

Figure 3: an example of a cell from the array population analyzed using GLM. A: Five time points of the spatiotemporal receptive field at four different iterations along the L1 regularization path. The spike triggered average result of the receptive field is shown on the top row for comparison. B: The 10-fold cross validation assessment of the regularization path. The best fit is obtained around iteration 11, Lambda (a regularization parameter) = 50 as determined using the cross-validation procedure, and measured in terms of area under the curve of a ROC. C: the same cell fitted with coupling terms instead of a spatiotemporal receptive field. Coupling of the surrounding neuron with the neuron being modeled (indicated by an asterisk *) is shown for 8 particular iterations alone the regularization path. These plots are in the electrode space - the colors indicate the value of the parameters lambda chosen for the corresponding coupling terms. D: The 10-fold cross validation assessment of the regularization path. The function hits its maximum at 20.

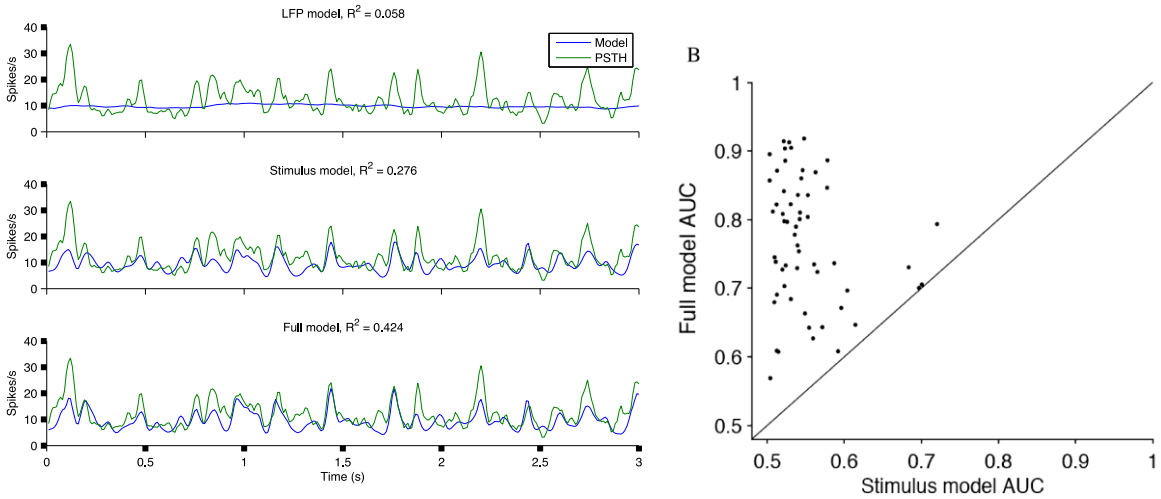

Figure 4: Left Panel: Estimation of the PSTH for a single cell. The reconstructions with 3 difference models are shown. The full model, capturing trial-by-trial variations in addition to stimulus effects, models the PSTH most faithfully. Right Panel: A comparison between the stimulus-only model and the full model. In most cells the full model outperforms the stimulus-only model.

Publications acknowledging the CURE grant:

[Li 2011] Li, X, Lee, T.S., Liu (2011) Hybrid generative-discriminative classifiers using posterior divergence. IEEE Conference in computer vision and pattern recognition (CVPR). 2713-2720.

[Samonds 2009] Samonds, J.M., Potetz, B., Lee, T.S., (2009) Cooperative and competitive interactions facilitate stereo computations in macaque primary visual cortex J. Neuroscience 29(50):15780-15795.

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 – 12/31/2011

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 continued work on a model of sequence learning in the hippocampus that could account for rats being able to mentally replay sequences in reverse order that they had only experienced in forward order. For example, if a rat performs a trajectory that causes a sequence of place cells A-B-C-D-E to fire, it may later, during a pause in its activities, replay that sequence in the order E-D-C-B-A. Since the neural learning phenomenon known as LTP (long term potentiation) requires the presynaptic cell to fire *before* the postsynaptic cell in order for their connection to be strengthened, this raises the question of how a connection from E to D could be potentiated by experiencing the sequence A-B-C-D-E.

Our original model, described in last year's report, relied on the reported existence of a phase gradient along the longitudinal axis of hippocampus, plus the known presence of long-distance longitudinal connections, to establish that some place cells coding for location E would actually project to and fire before some cells coding for location D. Thus, at least some E-D connections could be potentiated by the LTP learning rule when the animal experienced the sequence in forward order. But the fact that only these longer-range connections would be potentiated was problematic, because it complicated the issue of how sequences could be played back at every level of the hippocampus simultaneously.

We have since verified that a simpler solution can account for the learning of reversed sequences. Active place cells fire in a pattern that repeats with each theta cycle. (The theta rhythm is roughly 7-12 Hz in the rodent hippocampus.) Thus, the actual firing pattern the animal experiences is A-B-C-D-E--A-B-C-D-E--A-B-C-D-E--... While there is some question whether LTP can occur across and not just within theta cycles, if we assume that this is possible, then

there will be instances of cell E firing before (but not too far before) an earlier cell, say B, and this is enough for the hippocampus to be able to replay sequences bi-directionally. The coexistence of forward and backward connections allows sequences to be replayed in either order; attractor dynamics and a brief refractory period ensure that replay runs in only one direction at a time. We used a computer simulation to verify this result over a range of parameter values.

In the second half of the reporting period we turned our attention to a more detailed analysis of hippocampal activity sequences recorded in the Redish lab at the University of Minnesota. Animals ran through a two-T maze as shown in Figure 1A. At the choice point at the top of the maze (labeled T2 in the figure), they had to decide whether to turn left or right. Pairs of feeder cups (labeled F1 and F2) along the return arms provided a food reward if the animal had made the correct choice. As the animal moved through the maze, place cells fired, forming sequences of activity that repeated on each theta cycle. The distribution of these sequences was the focus of our study.

Theta cycle activity begins with the firing of place cells whose fields are centered slightly behind the rat's current position, and ends with cells whose place field centers are slightly ahead of it. Thus, the pattern of activity within a theta cycle sweeps along a segment of the path the rat is traveling.

We found that as velocity increased, these sequences represented longer paths in the environment (Figure 1B). When the animal was accelerating, the prospective (Pro) portion of paths included more space ahead of the animal; when it was decelerating, the retrospective (Ret) portion of paths included more space behind the animal (Figure 2). Furthermore, as path length increased, there were linear increases in both theta period and number of gamma cycles per theta cycle.

Finally, we observed that neural activity represented the environment in segments, corresponding with physical landmarks such as the corners, turn points, and feeder cups. Since the animals tended to accelerate as they left a landmark and decelerate as they approached the next one, they produced an alternating pattern of prospective and retrospective paths, resulting in a distribution of activity sequences that preferentially represented the segments between landmarks. We suggest that this can be regarded as the animal segmenting the environment into logical "chunks". These chunks could help the animal organize its memory and enable a more abstract encoding of the structure of the task it was performing.

These results have been written up as a journal article and submitted for publication.

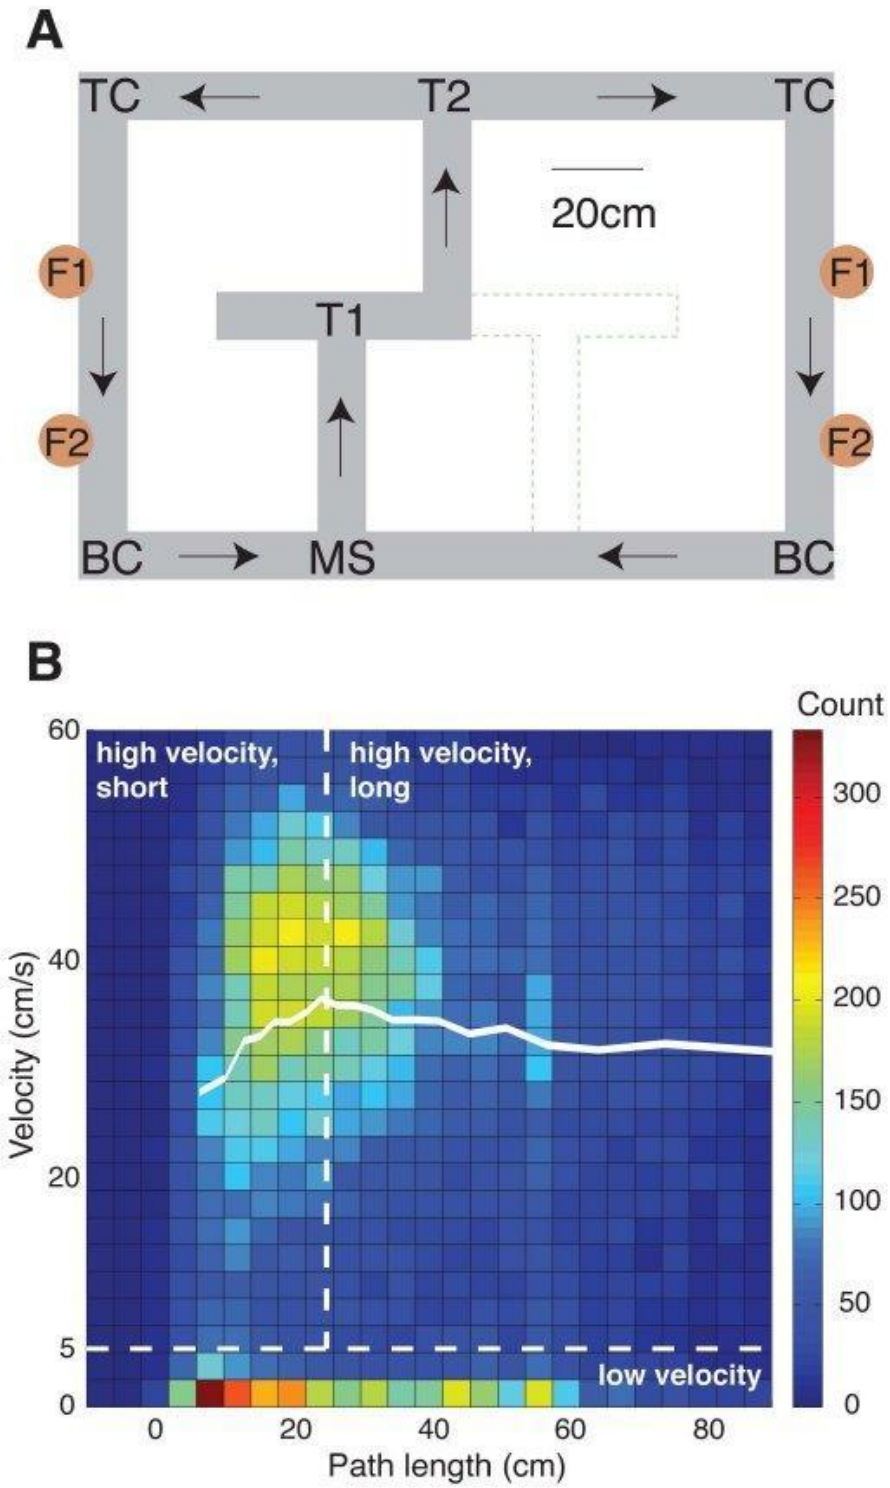

Figure 1: (A) The two-T maze, with prominent landmarks: maze start (MS), turns 1 and 2 (T1 and T2), top corners (TC), feeder cups (F1 and F2), and bottom corners (BC). (B) Histogram of paths (recorded sequences of place cell activations) broken down by velocity and length; for short path, high velocity sequences, path length increases with velocity.

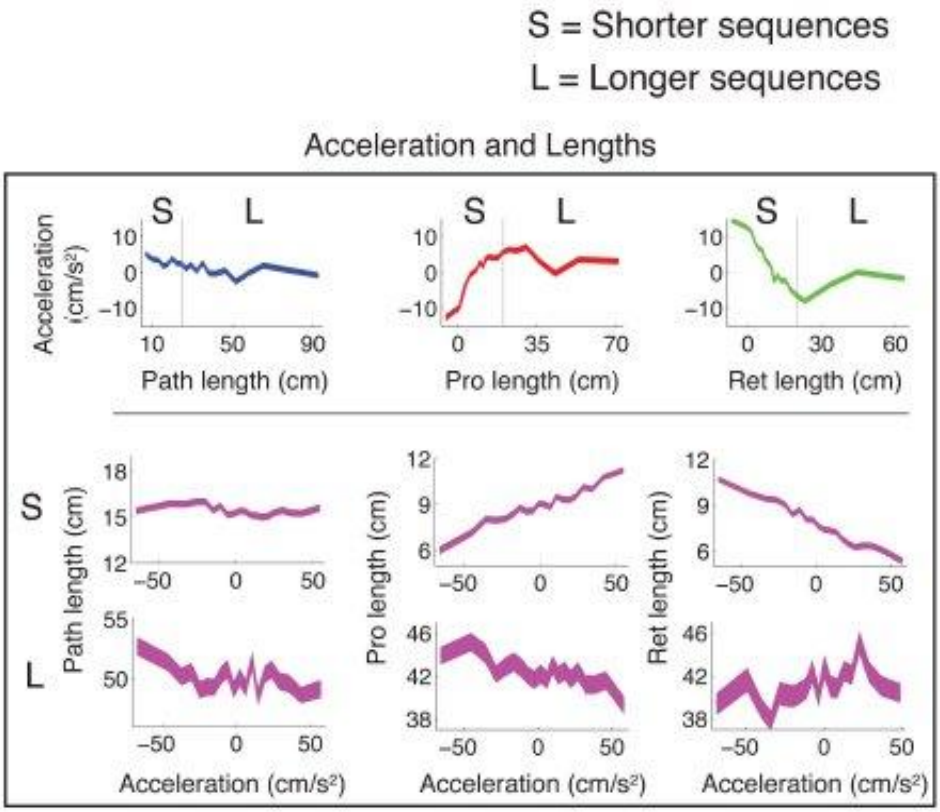

Figure 2: Relationship between acceleration and path length. “Path” refers to total path length, “Pro” to prospective length (paths from the rat’s current position extending forward), and “Ret” to retrospective length (from the start of the sequence to the rat’s current position). The top row shows accleration as a function of Path, Pro and Ret length. The next two rows show path length as a function of acceleration, plotted separately for short and long paths. For short paths, Pro length increases with acceleration, whereas Ret length decreases with acceleration.