Paul Miller: Research Statement

Memory Systems in the Brain

Several questions are important for all types of memory system. 1) Is the memory encoded in a discrete, even binary manner, or is it analog? 2) How is the stability of the memory maintained for a timescale much greater than that of the underlying processes? 3) How is the memory encoded and later decoded? My research addresses these questions through mathematical and computational modeling. I am investigating short-term memory systems, based on cortical activity, and long-term memory systems, based on molecular switches within the postsynaptic density.

Cortical models of decision-making based on working memory.

Introduction

Most decisions that we make are based on earlier information, so require memory. My research focuses on tasks that require such a decision based on prior information, or context. The prior information is recent, typically a few seconds old, so can be stored in working memory, which is maintained by neuronal activity. The tasks that I model have three phases. First, an initial stimulus or context that must be encoded. Second, the initial encoding should persist in some form during a delay period. Third, a second stimulus results in a motor response. The important feature of these tasks is that the motor response depends on the first stimulus, so requires memory. I have collaborated with and have access to data from the group of Ranulfo Romo, a world leader in this line of experimental research.

Sequential discrimination by integral feedback control.

Ranulfo Romo's group have carried out a set of experiments where they have recorded activity of neurons that respond to a discrimination task (see e.g. Romo et al. 1999, Brody et al. 2002, Romo et al. 2004). In the task, a monkey experiences a vibrational stimulus to its index finger, at a frequency, f1, followed by a delay of three to six seconds, then a second vibrational stimulus at a frequency, f2. The monkey must indicate by moving a lever to the left or right, whether the second frequency is higher or lower than the first. Romo's group have recorded from neurons in many areas of the brain, such as somatosensory cortex, prefrontal cortex, premotor areas and motor cortex. They see neurons in the sensory areas whose activity correlates with the stimulus, but only while the stimulus is present. Neurons in motor areas indicate the decision, or response of the monkey at the end of the second stimulus. Prefrontal, premotor and supplementary motor neurons can show stimulus-dependent activity, delay activity and decision-making activity. Hence from the behavior of the different neurons involved, it is possible to verify models of the neuronal circuitry used to complete the task. I am constructing models and simulating networks of spiking neurons to perform the task.

The models I have produced to carry out the sequential discrimination task use integral feedback control (Yi et al. 2000). Two groups of neurons are required, which can be found in the prefrontal cortex. The first group receives stimulus-dependent input that is time-locked to the stimulus, from somatosensory areas. The second group integrates the activity of the first group, and can maintain its firing rate once the first group is quiescent (Seung et al. 2000, Miller et al. 2003). The hallmark of integral feedback control is the presence of inhibitory feedback from the integrator to its input. This type of network has the important feature that the amplitude of the stimulus can be encoded, whereas an integrator alone would encode the product of amplitude and duration of a stimulus. I have simulated this model, and variations of it using a network of integrate-and-fire neurons (Miller et al. 2004b). Many of the features of neurons in the network are observable in the experimental data.

The ability of humans and monkeys to carry out this type of discrimination depends on the two frequencies and grows with the difference in frequencies. Using noisy spiking network models, with trial to trial variations, I will calculate neurometric functions, which measure whether the information contained in the spikes of a single neuron is sufficient to discriminate f2>f1 versus f2<f1 on a trial by trial basis. I will compare these functions with those measured experimentally, and with the psychometric curves evaluated on the basis of performance. In particular, it will be informative to see how the threshold for discrimination (that is the difference between f2 and f1 needed for 75% correct responses) varies with the absolute frequencies.

Future projects: Reading out the position of a population 'bump' of activity.

Networks capable of supporting a 'bump' attractor have achieved much success in modeling spatial working memory (Compte et al. 2000). In spatial working memory tasks, the position of a spot on a screen must be remembered for a few seconds, after which the animal makes an eye movement to the prior location of the spot. In the experiments, the location of the spot is along a ring, so in models of the task, neurons are labeled by their preferred location as an angle around the ring. A 'bump' attractor is a localized peak in activity of neurons, triggered by the stimulus (the spot on the screen). The 'bump' persists once the stimulus is removed, so can maintain the memory of the earlier stimulus location. In computer models, the position of the center of the bump is given by the population vector (the sum of vectors, one for every neuron, in the direction of each neuron's preferred angle, with amplitude of the firing rate of that neuron). However, it is far from clear how any motor response could be based on the information contained within a population vector.

I will formulate and compare methods for making a motor response based on the activity of a 'bump' attractor. The motor response is an eye saccade. The direction of saccade is the vector sum of components for vertical and horizontal movement, as these are the directions muscles are able to move the eye. Hence a coordinate transformation is necessary, from polar to rectangular, between neurons encoding angle and the motor neurons. The loss of angular symmetry entailed may explain observed biases in saccades — errors occur along preferred directions.

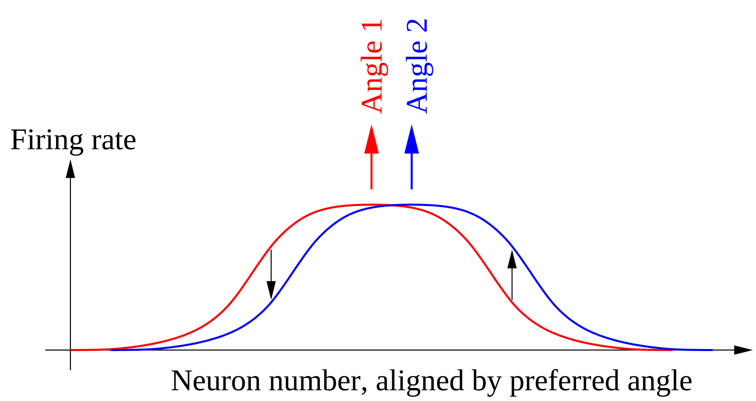

Figure 1 A bump attractor encodes the memory of spatial location. If a bump initially encodes location at angle 1 (left curve) then moves to angle 2 (right curve), the neurons whose activity changes the most are those on the flanks of the bump (marked by arrows between curves). Neurons near the center of the bump change little in their activity.

The neurons which most sensitively encode the position of the bump are not those at the center of the bump, but those at its sides, as indicated in Figure 1. This is because if the bump moves a little, neurons at the center do not change their firing rate by much, as they sit on a plateau of activity. It is the neurons on the side which rapidly change their firing rate when the bump moves. Some methods for stabilizing the position of a bump result in the profile of the bump changing. In particular, the bump may get wider while its center maintains a constant position. It will be important to test whether different readout mechanisms are sensitive to such changes in the width of a bump.

The required connection strengths between orientation-sensitive neurons in the bump attractor and horizontal/vertical motor output neurons can be learned. I will investigate the use of reward-based Hebbian learning rules to generate the appropriate connections strengths. Reward-based learning results in the potentiation of synapses if there is presynaptic activity, postsynaptic activity and a correct response. It is likely that the criterion for correct response, in terms of direction of eye saccade, will initially need a broad window in order to obtain some correct responses. As learning proceeds and the saccades become more accurate, the allowed error for a correct, rewarded response, can be narrowed. Hence the transformation from memory to motor output will allow me to test methods of reward-based learning.

Stochastic modeling of molecular processes underlying long-term memory.

Introduction

Chemical reactions sustain all biological systems. However, intuition based on chemistry in the test tube, to all intents and purposes a macroscopic system, may not be appropriate for understanding the chemistry of living cells (Halling 1989). Biological reactions are often confined to small compartments within a cell and may involve only tens of molecules for any given reactant. Hence stochastic effects, particularly shot noise, become important.

Shot noise arises when reactions are considered as discrete events. The law of mass action is appropriate for a macroscopic system, where it is meaningful to talk of a rate of reaction that can vary continuously with the concentration of reactants. However, when only a small number of molecules of each reactant are available, the concentrations and hence the rates vary discretely as a result of each individual reaction. Reaction steps occur probabilistically (Gillespie 1977), so large fluctuations about the mean rate can occur. Such fluctuations are known as shot noise. A particular effect of shot noise is a reduction in stability of otherwise stable states. Hence we investigate the stable lifetime of memory systems composed of a small number of molecules.

Bistability in the CaMKII-Phosphatase system and synaptic plasticity.

Synapses are the connections between neurons, whose strengths determine the ability of one neuron to excite another. Each neuron receives thousands of inputs through its synapses, each of which can become stronger or weaker following appropriate stimulation. The strengthening and weakening of synapses affects the network's responses to external inputs and is believed to play a key role in long-term memory (Morris 2003).

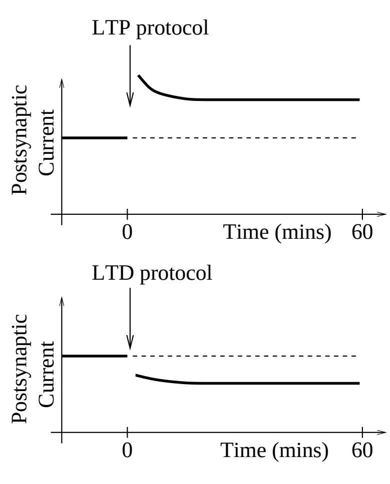

Figure 2 Experimental protocols lead to LTP or LTD. The postsynaptic current is measured in response to presynaptic action potentials. A specific pattern of applied stimuli are given at the time indicated by the vertical arrows. The stimuli cause a change in the postsynaptic current in response to later presynaptic action potentials. If the change lasts for more than tens of minutes, such strengthening is known as long-term potentiation (LTP) and weakening as long-term depression (LTD). The general rule is that a moderate increase in free calcium ion concentration in the post-synaptic neuron leads to LTD, whereas a larger increase leads to LTP — though it is now becoming accepted that multiple mechanisms exist for LTP across different synapses and different parts of the brain.

Experimental protocols can lead to a strengthening (LTP) or weakening (LTD) of synapses (see Figure 2). Activation of calcium-dependent protein kinase II (CaMKII) is known to be essential for many forms of LTP. A key issue is whether the ten to twenty holoenzymes of CaMKII present post-synaptically are capable of maintaining the activation that may underly memory storage. I have addressed this question using Monte Carlo computer simulations of a model for a molecular switch (Lisman and Zhabotinsky 2001) and using stochastic analysis in terms of estimating the hopping frequency over barriers. The key result is that fifteen to twenty holoenzymes can be sufficient to maintain stable states for as long as a human lifetime (Miller et al. 2004a).

Future directions — potentiation versus depression.

The system studied so far possesses bistability. A state with primarily unphosphorylated CaMKII is stable at low calcium concentrations, a state of highly phosphorylation is stable at high calcium and both can be stable over an intermediate range of calcium. In the model system, any increase in the concentration in calcium leads to a greater probability for phosphorylation, which correlates with synaptic strength. However, experiments show that a small increase in calcium concentration can lead to a reduction in synaptic strength (LTD). Therefore I will extend the system of CaMKII and phosphatase, in cooperation with appropriate inhibitors, to model LTD as well as LTP. At a minimum, the system must possess (1) bistability at low calcium concentrations, (2) a state with low activity that is stable alone at intermediate calcium concentrations and (3) a state with high activity that is stable alone at high concentrations.

Once a scheme for bidirectional plasticity is established, I will address key features of the phenomenon computationally. The first issue is to relate the changes in phosphorylation to the time course of a transient calcium signal, as the system may not reach a steady state during the transient period of elevated calcium. Second, it will be important to model a range of synapses with a distribution of volumes and concentrations of reacting species. This is important, as while changes in a single synapse may be all or none, total synaptic strength changes in a graded manner when measured across a population. The ease of LTP is history-dependent. This can be explained with a population of synapses that range from difficult to easy to potentiate. Once the easy ones are potentiated, it becomes more difficult to observe further potentiation. The corollary is that as more synapses are potentiated, the easier it is to observe depression, as there are more 'UP' to knock 'DOWN'. I will investigate these ideas with the bidirectional model.

Funding agencies and requirements.

In my first year as Assistant Professor, I will apply for funding to the National Institutes of Health (from whom I have already obtained a five-year grant), the National Science Foundation (in particular for a CRCNS grant, aimed at computational modeling in neuroscience), the David and Lucile Packard Foundation and the Searle Scholars Program.

I require access to computational resources in order to be successful in these projects. I have experience at running parallel code on national supercomputers, and in managing a large in-house Beowulf cluster. I will use start-up money to set up an initial computer cluster if possible, which can be expanded as further funds become available. If I need additional computational power at any time, I will apply for use of national resources, such as those at the Pittsburgh Supercomputing Center.

References

Brody CD, Hernandez A, Zainos A, Lemus L, Romo R (2002) Analysing neuronal correlates of the comparison of two sequentially presented sensory stimuli. Philos Trans R Soc Lond B Biol Sci 357:1843-50

Compte A, Brunel N, Goldman-Rakic PS, Wang X-J (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb. Cortex 10:910-923.

Gillespie DT (1977) Exact stochastic simulation of coupled chemical reactions. J Phys Chem 81:2340-2361.

Halling PJ (1989) Do the laws of chemistry apply to living cells? Trends Biochem Sci 14:317-319.

Landauer R (1962) Fluctuations in bistable tunnel diode circuits. J Appl Phys 33:2209-2216.

Lisman JE, Zhabotinsky AM (2001) A model of synaptic memory: A CaMKII/PP1 switch that potentiates transmission by organizing an AMPA receptor anchoring assembly. Neuron 31:191-201.

Miller P, Brody CD, Romo R, Wang X-J (2003) A recurrent network model of somatosensory parametric working memory in the prefrontal cortex. Cerebral Cortex 13:1208-1218.

Miller P, Zhabotinsky AM, Lisman JE, Wang X-J (2004a) The stability of a stochastic CaMKII switch: dependence on the number of enzyme molecules and protein turnover. Submitted to PLoS Biology.

Miller P, Romo R, Wang X-J (2004b) Discrimination of temporally separated stimuli by integral feedback control. *In preparation.*

Morris RG (2003) Long-term potentiation and memory. Philos Trans R Soc Lond B Biol Sci 358:643-647.

Romo R, Brody CD, Hernàndez A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. Nature 339:470-473.

Romo R, Hernandez A, Zainos A (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. Neuron 41:165-173.

Seung HS, Lee DD, Reis BY, Tank DW (2000) Stability of the memory of eye position in a recurrent network of conductance-based model neurons. Neuron 26:259-271.

Yi TM, Huang Y, Simon MI, Doyle J (2000) Robust perfect adaptation in bacterial chemotaxis through integral feedback control. Proc Natl Acad Sci U S A 97:4649-4653.

Paul Miller: Statement of Teaching Philosophy and Interests.

Teaching Philosophy

As a teacher I see my main goal is to enable understanding of concepts, while students gain mastery of problem-solving techniques. I believe it is important to differentiate the physical, chemical or biological processes involved from the mathematical techniques that are best used to investigate them. In this way, I aim to teach students to gain an intuition of what the answer to a problem should be, or could be, so that when solving it they are not just 'turning the handle'.

Teaching Diverse students:
The ability to grasp new concepts differs from the ability to manipulate mathematical equations, so students will often excel in one ability to a greater extent than the other. By clarifying the concepts using different analogies and models, the student who has difficulty with the mathematical formalism still gets to appreciate the key idea or process. Such appreciation encourages the student to work on mathematical skills with a clear purpose and motivation. On the other hand, a student who excels more at the mathematical method is forced to think more deeply about the subject and to connect different ideas. Students who simply crank through equations to get the correct answers, with little understanding of what they are doing, will be challenged, as they should be if they are to develop into critical thinkers and good scientists.

The questions I ask in class, homeworks and exams are diverse in style and difficulty. It is important that all students are encouraged by what they have learned, but also that the strong students are stretched mentally. Similarly, I ask intuitive and visual questions, to ensure understanding, as well as requiring strong problem-solving skills.

Using Diverse methods:
I think it is important to encourage full class involvement. While a certain amount of teaching from the front is necessary, I usually see if the logical steps in an argument or calculation can come from others in the classroom. Since not all students volunteer answers, at certain points in the class I will ask questions whose answers should be written down by everyone. In smaller classes I check their responses before continuing by giving the 'official' version. At the end of class I usually give a two minute test on the main point of that class. The tests count in a small way toward the final course grade. The idea is to make the one or two questions simple enough that all students who were attentive in the class answer correctly. Hence attendance and attentiveness are encouraged, and the students leave class with a reminder of the one or two main points, or main skills they have acquired.

An important skill in the classroom is to encourage questions without getting side-tracked from the goals of the lesson, and to be welcoming of all attempts at answers, while clearly distinguishing correct responses from any errors. I think I succeed at this by being empathetic, often being able to address the cause of any mistake rather than just saying 'wrong', by being approachable, with a friendly demeanor, but being firm when necessary.

In summary, I understand that students are far from clones of myself, but each bring their own unique perspective to the classroom. To be a good teacher, the ability to listen is as important as the ability to lecture, because it is only by listening that the teacher learns the students way of thinking. I have the breadth of experience to be able to address students using a number of approaches, so am able to communicate effectively with a wide range of students.

Teaching Experience

My teaching experiences are varied and extensive. I have taught on three different continents (Europe, Africa and America); I have taught students ranging from high-school freshmen to college seniors; and I have taught physics to students in non-scientific disciplines, in the life sciences and in the physical sciences.

As a doctoral student I was a teaching assistant in computing classes and laboratory classes. I learned the importance of clarifying the assumptions of an experiment, and distinguishing those assumptions from the conclusions.

After completing my Ph.D., I spent 16 months teaching Physical Science, Mathematics and English at high-school level, in Malawi, Africa. I was motivated by my belief that education is our most valuable asset, and should be available to everyone in the world. While there, I had to develop various skills at classroom management, made especially difficult by the compulsory use of English in school classes, when a significant fraction of students were far from proficient in the language. I learned how to communicate clearly and simply, in particular the importance of finding relevant analogies and examples, when explaining ideas in an unfamiliar culture. I saw that gaining understanding and knowledge, at whatever the level, can bring joy.

At Georgetown in the fall of 1999, I co-taught the course 'The Quantum World Around Us' which was an undergraduate course on quantum mechanics for non-scientists. Instead of requiring full mathematical calculations, we treated the wave function as an arrow (a vector with amplitude and phase) and used the sum over trajectories approach to explain effects ranging from interference to discretization of bound states. My reviews for this class were good (available verbatim on request). The course showed me that non-scientists, while often lacking in mathematical skills, can otherwise be just as astute at tackling scientific problems as a trained scientist. However, I did gain an idea of the importance of teaching basic mathematical skills, particularly probability, to students who will one day be making business or policy decisions.

Also while at Georgetown I taught mathematics to three high-school students, in a voluntary role at a local school.

At Brandeis I taught the sophomore physics course, 'Oscillations and Waves' in the Fall of 2000. The course had the reputation of being tough and had a notoriously high drop-out rate, dissuading many students from majoring in Physics. I am happy to say that nobody dropped the class when I taught it, and that the course evaluation received the highest score in the department that semester (records are attached, reviews available verbatim on request). I also filled in for two weeks while a professor was away, teaching a total of six classes in 'Thermodynamics and Statistical Physics', in the Fall of 2003. Teaching a complete course requires a great deal of organization and preparation. I am now fully aware of the key decisions that need to be made in the months before the course starts, of the importance of early feedback to both me and the students, and the necessity of a reliable routine to maintain momentum through the semester.

My experience indicates that I have interdisciplinary expertise. I am able to teach a wide range of courses and I am able to adapt the way I teach a course to address the needs of the students.

Teaching Interests

I will enjoy teaching any class that involves applied mathematics, or the introduction of new concepts to the students. My particular areas of expertise span from quantum mechanics to computer modeling to noise in chemical and biological systems to introductory neuroscience.