

Edvard I. Moser
Centre for the Biology of Memory
NTNU
NO-7489 Trondheim
Norway

Selected Reading

- Bailey, C.H., Giustetto, M., Huang, Y.Y., Hawkins, R.D., and Kandel, E.R. (2000). *Nat. Rev. Neurosci.* 7, 11–20.
- Bostock, E., Muller, R.U., and Kubie, J.L. (1991). *Hippocampus* 1, 193–205.
- Carli, M., Robbins, T.W., Evenden, J.L., and Everitt, B.J. (1983). *Behav. Brain Res.* 9, 361–380.
- Desimone, R., and Duncan, J. (1995). *Annu. Rev. Neurosci.* 18, 193–222.
- Granon, S., Passetti, F., Thomas, K.L., Dalley, J.W., Everitt, B.J., and Robbins, T.W. (2000). *J. Neurosci.* 20, 1208–1215.
- Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). *Neuron* 42, this issue, 283–295.
- Markus, E.J., Qin, Y.L., Leonard, B., Skaggs, W.E., McNaughton, B.L., and Barnes, C.A. (1995). *J. Neurosci.* 15, 7079–7094.
- McGaughy, J., and Sarter, M. (1995). *Psychopharmacology (Berl.)* 117, 340–357.
- Miller, E.K., and Cohen, J.D. (2001). *Annu. Rev. Neurosci.* 24, 167–202.
- O’Keefe, J., and Conway, D.H. (1978). *Exp. Brain Res.* 31, 573–590.
- O’Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Oxford: Clarendon Press).
- Quirk, G.J., Muller, R.U., and Kubie, J.L. (1990). *J. Neurosci.* 10, 2008–2017.

Taking Strategies to Task

In this issue of *Neuron*, Uka and DeAngelis report a neural signature of the strategy that monkeys’ brains developed to solve a visual discrimination task by training them on one version of the task and testing them on another. Extensive training on one version caused decision networks in the animals’ brains to ignore certain classes of neurons whose signals would have been useful on the modified version of the task used to test them.

If you need advice on how to fix a leaky faucet, you call a plumber rather than a florist. This is common sense. But how do you know that it is a *plumber* that possesses this expertise, and how do you locate the plumber? Again, this seems easy, but only because brains and societies have developed effective strategies to solve these problems. In fact, how various pockets of expertise are discovered, categorized, and stored for future use is a rather challenging problem.

How do populations of neurons interact to solve this kind of problem? Consider the following task: you are given a noisy visual display that contains moving dots, and you have to report in which direction it is moving: right or left. Now, imagine you are a neuron or a small circuit of neurons, a few in a sea of billions, and you must make the final choice as to whether to look to the

right or the left. Somewhere in this sea are a few neurons whose signals are particularly relevant to this task. Most of the others are useless. We know that the brain of the monkey contains such neurons, many of the most informative in a few square millimeters of cortex referred to as “MT,” and they are informative by virtue of the fact that they respond vigorously to some directions of motion but not to others. In addition, we know that monkeys can indeed learn to perform this task, ultimately as well as humans do.

But learning this task takes the monkey a rather long time—several months if he is to learn to discriminate arbitrary direction pairs presented anywhere in his visual field. Interestingly, however, the direction selectivity of the neurons in MT doesn’t change noticeably during the training—though it hasn’t been directly tested, most MT-ologists would say that the neurons are every bit as informative about the stimulus on day 1, when the monkey is clueless, as they are on day 100, when the monkey is an expert at the task. It appears that the stimulus-derived information the monkey needs to solve the task is present in his brain from the beginning. Thus, much of what must be occurring in his brain during those long weeks of training is the decision circuits “figuring out” which MT neurons to listen to for a given configuration of the task—i.e., changes in connections. How does the monkey’s brain do this?

Answering this question is key to understanding one of the most interesting things that brains do: making decisions. For the student of brain decisions, it requires first learning which particular sensory neurons are tapped for the decisions so that he or she may then ask how the sensory signals are combined and how they come to efficiently communicate with the decision circuits.

For a typical experiment of the sort described above, the decision scientist compares three different kinds of information: (1) the visual *stimulus* displayed, (2) the action potentials (*spikes*) produced by one or more neurons recorded during the task, and (3) the animal’s *choices* (e.g., “right” or “left”). By analyzing the relationship between visual stimulus and spikes, one can compute a measure of how informative a given neuron is about some critical feature of the stimulus. For trials on which the stimulus is weak or ambiguous, one can analyze the relationship between spikes and choices to learn how informative the neuron’s spikes are about the monkey’s decision. The latter is often referred to as “choice probability,” and it can be thought of as a measure of the extent to which the decision networks in the monkey’s brain are using that neuron’s spikes—or others with which that neuron is correlated—to arrive at the final choice. This is a potent metric, since knowing what kinds of neurons the decision networks are “listening to” can be used to infer the nature of the strategy the brain is using to solve the task.

This approach was pioneered by Newsome and his colleagues at Stanford (Newsome et al., 1989; Britten et al., 1992, 1996). DeAngelis, one of Newsome’s scientific progeny, now has applied a clever variation on this theme to suggest an interesting approach to studying the brain circuitry underlying task strategy. Uka and DeAngelis (this issue of *Neuron*) trained monkeys on a similar two-choice discrimination task in which the ani-



Figure 1. Task Strategy for a Two-Alternative Forced Choice Depth Discrimination

mal had to decide whether a field of dots was nearer to or farther than a reference point. The only available visual cues that the monkey had to make his choices were the relative positions in the two eyes of some of the dots in the stimulus—a cue referred to as “binocular disparity.” Previously, the DeAngelis lab had shown that MT neurons are both sensitive to such cues and that they are used by monkeys to solve the task (DeAngelis et al., 1998; Uka and DeAngelis, 2003). But which neurons and under what circumstances?

It turns out that different MT neurons have qualitatively different response profiles—referred to as “tuning curves”—to visual stimuli containing binocular disparity information. Some have tuning curves that are symmetrical with respect to the plane of fixation. For example, a neuron of the kind known as “tuned excitatory” responds maximally to stimuli at the same depth as the point of fixation but respond progressively less well to stimuli that are either nearer to or farther than this point. Other neurons, however, have very asymmetrical tuning curves and might respond vigorously to stimuli anywhere behind the fixation plane but very poorly to those anywhere in front of it. This would be classified as a “far” neuron, and its inverse as a “near” neuron.

In the study reported in this issue of *Neuron*, Uka and DeAngelis (2004) trained the monkeys initially using *only* symmetrical pairs of disparity stimuli. That is, the monkey always had to discriminate between two stimuli that were nearer or farther *by an equal amount*. Over several months, they became experts at solving this symmetrical task. It is quite easy to see that a “tuned excitatory” neuron would not be very useful for solving such a task, since it will give roughly equal firing rates to the two possible visual stimuli. A more sensible strategy would be to compare the outputs of the *near* and *far* neurons and to ignore the symmetrically tuned neurons altogether (Figure 1).

When it was time to examine responses of MT neurons

during this task, however, the authors modified the task slightly so that it was optimized for the neuron from which they were recording. Thus, if they happened to be recording from a tuned excitatory neuron, they chose one stimulus to be a very good stimulus for the neuron—say, just on the near side of zero disparity—and one to be a poor stimulus, such as a very far disparity. In this case, the responses of the neuron were highly informative about the two possible stimuli. Common sense would dictate that the signals be used, and previous studies have found that monkey decision circuits do indeed tend to give greater weight to the most highly informative neurons (Britten et al., 1996; Parker et al., 2002). In this case, however, Uka and DeAngelis found that neurons with symmetrical disparity tuning curves were pretty much ignored—i.e., they had low choice probabilities—regardless of their relevance to the task at hand. It was as if, during all those months of training, the monkey’s brain had learned to ignore neurons with symmetrical tuning curves, and he could not change this strategy on the fly.

This observation provokes a number of questions. Was the monkey’s performance notably worse for the new, asymmetrical task? Uka and DeAngelis found that this was not the case, presumably because the original strategy of comparing pools of *far* and *near* tuned neurons still worked. But at some point this strategy could be made to break down. One wonders how their monkeys would have performed if both stimuli were made to lie on the *same* side of the fixation plane—a case for which the *near* versus *far* strategy would fail. Assuming this would make the monkey’s performance worse, at least initially, if the potentially highly informative neuron could be recorded from while the monkey mastered this new task, would the neuron’s choice probability increase over time? And, if so, over what length of time—hours? days? weeks? Presumably, the answers to such questions would shed light on the plasticity mechanisms

that underlie the acquisition of new strategies. With such a hook into the system, one could then imagine using techniques such as antidromic activation to monitor the underlying changes in connectivity that must be occurring.

Another interesting feature of Uka and DeAngelis's result is that it argues for a "bottom-up" implementation of task strategy—as has been implicit in most of the preceding discussion—rather than a "top-down" influence of feedback from the decision networks themselves. They found that subtle variations in the animal's performance over time were correlated with similar fluctuations in the sensitivity of single neurons and that this coupling was strongest for those neurons with the highest choice probability. In other words, the monkey's decision networks, through training, had become more or less hardwired to a particular group of low-level neurons, and as these neurons performed, so performed the monkey. As noted above, this hints strongly at where in the visual pathways one might start to look for such changes in wiring.

The results of Uka and DeAngelis have indicated just how potentially powerful measurements of choice probability can be about how a monkey's brain is solving a particular task. And they point the way to potentially more incisive experiments in which the experimenter monitors dynamic changes in choice probability while animals are confronted with specific variations on previously learned tasks. While the ultimate elucidation of the means by which animal brains form neural strategies will require many more experiments, the findings of Uka and DeAngelis suggest an interesting strategy for this task.

Richard T. Born

Department of Neurobiology
Harvard Medical School
Boston, Massachusetts 02115

Selected Reading

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). *J. Neurosci.* 12, 4745–4765.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). *Vis. Neurosci.* 13, 87–100.

DeAngelis, G.C., Cumming, B.G., and Newsome, W.T. (1998). *Nature* 394, 677–680.

Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). *Nature* 341, 52–54.

Parker, A.J., Krug, K., and Cumming, B.G. (2002). *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 357, 1053–1062.

Uka, T., and DeAngelis, G.C. (2003). *J. Neurosci.* 23, 3515–3530.

Uka, T., and DeAngelis, G.C. (2004). *Neuron* 42, this issue, 297–310.