

Teaching the Brain to See

How do some people identify a wine by year and vineyard, while others struggle to discriminate Pinot Gris, Chablis, and Chardonnay? How do experts learn to see meaning in brain scans and X rays, while novices see incomprehensible splotches of black, white, and gray? Practice makes perfect—or, at least, practice makes much better. In recent years, behavioral and cognitive neuroscientists have become increasingly interested in the effects of practice and learning on brain and behavior, and some of the most exciting research has been done in the area of perceptual learning.

It is now well established that the ability to discriminate sensory stimuli improves with practice, and that such improvements are often highly specific to trained stimuli and maintained over quite long periods of time (Sagi and Tanne, 1994). It is also well established that neuronal responses change with learning, at the level of individual cells, populations of cells within one cortical area, and functional connections across cortical areas (Buonomano and Merzenich, 1998). For example, after learning, ~25% of neurons in inferotemporal cortex show object selectivity for previously novel visual objects (Logothetis et al., 1995; Kobatake et al., 1998). Such a result suggests that neurons become more selectively tuned to objects through experience. However, learning is also frequently accompanied by a decline of neuronal responsiveness. Although, at first glance, this result may seem counterintuitive, researchers have suggested that such a decline may reflect an increased

efficiency of processing stimuli with increased familiarity. However, little, if any, direct support for this theory has been provided.

In an article in the current issue of *Neuron*, Rainer and Miller (2000) describe an elegant study designed to fill that gap. Rainer and Miller trained monkeys to discriminate particular sets of objects, presented in varying amounts of visual noise (or degradation). The effect of learning was apparent in terms of both behavioral and physiological responses. When the stimuli were degraded by noise, monkeys performed more accurately with familiar stimuli than with novel stimuli. In other words, degraded familiar objects could be discriminated from one another better than degraded unfamiliar objects. Such behavioral effects of learning were correlated with several physiological changes in the prefrontal (PF) cortex. First, the responsiveness of PF neurons declined as a function of learning, in terms of both maximum firing rate and the number of selectively tuned neurons for nondegraded stimuli. Second, although fewer PF neurons responded selectively to familiar stimuli, those neurons were slightly, but significantly, more narrowly tuned to particular nondegraded stimuli than were neurons that responded to unfamiliar stimuli. Third, some degree of object selectivity was maintained in PF neurons over larger amounts of stimulus degradation for familiar objects than for unfamiliar objects.

These results are exciting because they provide the best evidence to date that decreased neuronal activity is actually accompanied by a greater selectivity for learned stimuli. The results also show that the largest behavioral and physiological benefits of learning actually occur with degraded stimuli. Furthermore, across all levels of degradation, there is a striking and suggestive correlation between the extent to which learning improves overall performance of the animal and the extent to which learning improves PF neurons' ability to discriminate stimuli. The present study also represents significant methodological advances over previous research. In particular, Rainer and Miller use a within-animal design, enabling the direct comparison of various stimulus conditions and ensuring that comparable regions of PF were tested for novel and familiar stimuli. In addition, great care was taken in designing stimulus sets that were equated for global properties, such as spatial frequency, luminance, and color.

Of course, as with most important research, the results of this study lead to more questions than answers. For example, why does learning reduce neural activity? What aspect of information is altered by learning? And how are the physiological changes in PF related to changes in other brain regions?

In what sense does the reduction in neural activity increase processing efficiency? Perhaps, prior to learning, many of the so-called selective neurons were in fact broadly tuned to objects, and learning reduces the responsiveness of these neurons because they carry relatively little information. However, even weakly selective cells potentially can provide some information that leads to better object discrimination, so it is not obvious why removing these responses ought to improve discrimination. Instead, one might have expected learning to result in better utilization of the information carried

by these neurons. One possible explanation for the reduction is that object recognition is based on interactions among PF cells, and that there are constraints on the number of interactions that can be processed optimally. In other words, the costs of including additional weakly selective cells may outweigh the benefits. Under these conditions, pruning the neural representation to a small set of informative cells might significantly improve behavioral performance. Another possibility is that the initial neural representation of a set of unfamiliar objects consists of many cells whose responses are highly correlated with each other, and that learning removes neuronal responses to reduce redundancy and improve efficiency (Barlow, 1990). Unfortunately, almost nothing is known about how learning affects correlations among PF cells, or how the responses are converted into a perceptual decision (but see Kim and Shadlen, 1999). It will take very clever experiments to determine if either of these factors contributes to the physiological and behavioral changes observed during learning.

Although Rainer and Miller's results indicate that more information is transmitted by PF cells after learning, their data do not specify how this increase occurs. From an information-processing perspective, there are two possibilities. First, learning could alter the features that cells represent, so that only highly informative features are encoded at the end of learning. Second, the features represented by PF cells might remain the same, but variability in the neural responses could decrease with learning. Gold et al. (1999) used psychophysical techniques to examine which of these factors contribute to perceptual learning in human observers. Like Rainer and Miller, Gold et al. found that learning improves the recognition of complex patterns embedded in varying amounts of visual noise. In a critical experiment, Gold et al. presented identical stimulus and noise combinations on more than one trial and recorded the consistency of each observers' responses. This measure of response consistency can be used to estimate the internal noise that affects an observer's decision. For example, a system that contained no internal noise would respond with perfect consistency. Of course, the human visual system is noisy, and so human observers are not perfectly consistent. The important result from Gold et al. was that response consistency did not vary as a function of learning, suggesting that internal noise remained constant, even when sensitivity increased by up to 400%. Thus, the results of Gold et al. imply that learning affects the signal but not noise.

One obvious way to increase the discriminative signal is to encode informative features. Assume for the moment that observers encode patterns with linear templates or receptive fields. Ideal observer theory can be used to derive the templates that maximize performance, and learning might serve to bring the templates used by human observers closer to the ideal ones. Although this idea is plausible, until recently it has been difficult to test because there were no satisfactory methods for inferring the templates used by human observers. Recently, though, researchers have used reverse correlation techniques to map templates, or "behavioral receptive fields," used in visual discrimination tasks (Neri et al., 1999; Gold et al., 2000). Although these techniques have not yet been adapted to study learning,

they potentially could enable us to see how receptive fields—of single cells and of entire observers—change with learning.

Rainer and Miller suggest that learning-related alterations in physiological responses of PF neurons may be due to selective strengthening of connections between PF neurons and neurons earlier in the visual system (e.g., area IT). Much of the electrophysiological and neuroimaging work to date has focused on activation in isolated regions of the visual system, not on interactions across brain regions. It is becoming increasingly clear that systems of interactive analyses are critical for a complete understanding of perceptual and cognitive processes. An example of one such analysis is seen within the context of neuroimaging. McIntosh and colleagues have asked not just what part of the brain “lights up,” but how different brain regions interact with one another, and how those interactions relate to performance (Horwitz et al., 1999). Although the spatial resolution of positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) necessitate that questions are addressed at a different level of analysis than with single-cell electrophysiology, recent technological advances are enabling the combination of electrophysiology and neuroimaging. Thus, one can imagine combining the best of both worlds—fine grained analyses of neuronal responses within particular cortical regions, and investigations of the functional interactions across those regions.

A thorough understanding of perceptual learning will require the integration of the best aspects of psychophysics and physiology. Rainer and Miller’s study provides an excellent example of what can be accomplished when one takes this approach. Their results provide a significant advance in our understanding of the neural mechanisms associated with perceptual learning and suggest a possible physiological basis for learning-related behavioral changes.

Allison B. Sekuler and Patrick J. Bennett

Department of Psychology
University of Toronto
Toronto, Ontario, M5S 3G3
Canada

Selected Reading

- Barlow, H. (1990). *Vision Res.* 30, 1561–1571.
- Buonomano, D., and Merzenich, M. (1998). *Annu. Rev. Neurosci.* 21, 149–186.
- Gold, J., Bennett, P.J., and Sekuler, A.B. (1999). *Nature* 402, 176–178.
- Gold, J.M., Murray, R.F., Bennett, P.J., and Sekuler, A.B. (2000). *Curr. Biol.* 10, 663–666.
- Horwitz, B., Tagamets, M.-A., and McIntosh, A.R. (1999). *Trends Cogn. Sci.* 3, 91–98.
- Kim, J.N., and Shadlen, M.N. (1999). *Nat. Neurosci.* 2, 176–185.
- Kobatake, E., Wang, G., and Tanaka, K. (1998). *J. Neurophysiol.* 80, 324–330.
- Logothetis, N.K., Pauls, J., and Poggio, T. (1995). *Curr. Biol.* 5, 552–563.
- Neri, P., Parker, A.J., and Blakemore, C. (1999). *Nature* 401, 695–698.
- Rainer, G., and Miller, E.K. (2000). *Neuron* 27, this issue, 179–189.
- Sagi, D., and Tanne, D. (1994). *Curr. Opin. Neurobiol.* 4, 195–199.