Szentagóthai, J. (1978). Proc. R. Soc. Lond. B. Biol. Sci. *201*, 219–248.

Toledo-Rodriguez, M., Blumenfeld, B., Wu, C., Luo, J., Attali, B., Goodman, P., and Markram, H. (2004). Cereb. Cortex *14*, 1310–1327.

Wichterle, H., Garcia-Verdugo, J., Herrera, D., and Alvarez-Buylla, A. (1999). Nat. Neurosci. *2*, 461–466.

Xu, Q., Cobos, I., De La Cruz, E., Rubenstein, J., and Anderson, S. (2004). J. Neurosci. *24*, 2612–2622.

DOI 10.1016/j.neuron.2005.11.012

A Rose by Any Other Code

In this issue of *Neuron*, Mazor and Laurent demonstrate that the internal representation of an odor in the antennal lobe of locusts is broadly distributed across the population of projection neurons and is formatted in a manner that requires deciphering of response transients rather than steady-state activity patterns.

How is the odor of a rose represented by neural activity patterns within the first few stages of our olfactory system? What are the neural mechanisms through which we can distinguish that rose's odor from one of a different variety? Issues related to olfactory coding have inspired a great deal of recent experimental and theoretical research and have yielded a host of significant insights into mechanisms underlying population coding of complex stimuli. However, several fundamental questions related to olfactory coding have recently been raised that are the focus of considerable debate. Two questions of particular interest relate to the density and the dynamics of the olfactory representation. How broadly distributed is the trace of a particular odor across the population of principal neurons of the first olfactory processing stage? Can the identity of an odor be "read out" from a static steady-state activity pattern that develops across this population during a "sniff," or is the information about the odor formatted into some aspect of the population's dynamic activity patterns during the sniff? Mazor and Laurent report the results of a spectacular set of experiments that provide definitive, quantitative answers to these questions: they demonstrate that the odor representations are very broadly distributed in the first processing stage and that they are, indeed, "dynamically formatted" (Mazor and Laurent, 2005).

The experiments were carried out on the locust, which is one of several invertebrate and vertebrate preparations Laurent and colleagues (and researchers at several other institutions) have been studying over the past decade. The specific targets of the study were the projection neurons (PNs) in the animal's antennal lobe (AL), which functions as an "encoding machine" to transform the olfactory input in a manner that enables the formation of olfactory memories within the next processing stage (the mushroom body). (The PNs and AL are functionally equivalent to the mitral cells and olfactory bulb in vertebrates.) There are only about 830 PNs in the locust's entire antennal lobe. These PNs receive direct input from the animal's array of olfactory receptors and interact with one another through a group of local interneurons in the AL. The focus of this study was on a complex and unresolved issue: what is the precise nature of the coding scheme established at these early stages? Specifically, the goals were to determine (1) how many of the 830 PNs were activated over one single oscillation cycle of a long-lasting response, (2) how reliably the spikes were produced by individual PNs, (3) how rapidly the representations (i.e., the population activity patterns) for single odors evolved, (4) if these representations eventually stabilized to a fixed pattern, and if so, (5) whether or not those stabilized patterns were optimally discriminable.

The experiments designed to answer these questions were very straightforward, though technically a tour de force: different odors were presented to a group of test animals, and the fully resolved spike-train responses of 99 out of the 830 PNs were recorded. This represents 12% of the entire PN population. For some of the experiments, local field potentials were recorded simultaneously. Although this sample of 99 PNs was drawn from 10 different animals, population data could be assembled by combining sets of simultaneously recorded PNs across experiments, as described in earlier experiments (Stopfer et al., 2003). The authors characterized the population's responses to five different odors, presented for durations ranging from 0.3 to 10 s, with enough repetitions of all stimuli to get quantitative statistical assessments of pattern discriminability.

The results, in a nutshell, were as follows. In the resting initial state, in the absence of any odor, only 1% of the PNs were highly active, 23% were silent, and the remaining 76% were "flickering" on and off. Upon presentation of the odor, the PN population activity became much more structured: about 10% of the PNs became reliably active and highly correlated with one another in their firing, 60% became inactive, and the remaining 30% "flickered" unreliably (i.e., without any significant correlation to the stimulus). After approximately 50 ms, this 10-60-30 pattern was re-established, but with a somewhat different set of PN cells in each of the three activity categories. (Each one of these 50 ms epochs partially correlated activity corresponded to one of cycle of the 20 Hz LFP oscillations.) During maintained odor presentations, this characteristic decorrelation and recorrelation of different sets of PN cells continued for 1–2 s, after which the pattern stabilized to a fixed set of active, inactive, and unreliable PNs. By the end of the 1–2 s period of transient dynamical activity, approximately half of the 830 PN cells had participated in one of the "10%" (reliably active and highly correlated) groups.

A simple analogy might help illustrate this complex scenario. Imagine that 832 PN cells are configured as a low-resolution Palm-Pilot-like screen, having only 26 × 32 pixels. Before the odor presentation, eight to ten of the pixels scattered throughout the array are latched up to full brightness, about 190 are blacked out, and the remaining 630 are flickering on and off like "snow" on a video monitor. The overwhelming impression would be of totally unpatterned activity. Upon odor presentation, however, a spatio-temporal pattern emerges: 500 of the pixels go black, and about 80 go bright at nearly the same instant. The remaining 250 pixels flicker on and off with no apparent correlation to either the bright or dark pixel groups. The screen image then starts to throb at 20 Hz: upon every cycle, a new pattern is established

that has the same basic numbers of "on," "off," and "flickering" pixels. Each successive pattern shares some of the same pixels, but not all of them: each pattern is slightly decorrelated with respect to the preceding patterns. As the sequence progresses, patterns separated in time by six or more of the 20 Hz cycles have no significant correlation. This general activity proceeds for a second or two, dims a bit in intensity, and then settles into a pattern that is nearly static with respect to the location of the "on" and "off" cells. When the odor is removed, a new pattern flashes onto the screen similar in statistical characteristics, and the whole process reverses itself: a sequence of patterns ensues at the 20 Hz frequency until the initial (prestimulus) pattern is re-established, and the cycles die out. Upon presentation of different odors, you find that each odor causes a different sequence of patterns for the "on" and "off" transients; but strangely enough, the steady-state "fixed point" patterns achieved after 1–2 s are somewhat similar across the different odors.

So what aspects of these sequences encode the information about the different odors? Through the application of elegant analytical approaches, Mazor and Laurent demonstrate that optimal stimulus separation occurs during the "on" and "off" transient periods, rather than during the fixed point patterns. In other words, most of the information about odor identity is dynamically formatted into the complex sequence of activity patterns. One surprise noted by the authors is that the observed evolution of the dynamic patterns to a fixed point is at odds, in several respects, from the "winnerless competition" model that had been proposed in earlier studies (Rabinovich et al., 2001, Laurent et al., 2001, Afraimovich et al., 2004).

Rather than stopping here, Mazor and Laurent go on to analyze the relative significance of the PN cells' transient and set-point response segments at the level of their postsynaptic targets: the Kenyon cells in the mushroom bodies. They demonstrate that these target neurons are the least responsive when their PN population input is stabilized at the fixed point and most responsive during the transient periods after stimulus presentation and termination. Thus, all PN spikes are not equally "meaningful." The implications are sobering: a full understanding of the neural representation scheme in the AL may require the evaluation of correlations between neurons in the PN population within the context of the decoding operation implemented at the subsequent processing stage in the mushroom body.

All of these results support the general working hypothesis that has guided recent research in Laurent's group: that the AL's principal functions are to (1) create a large coding space for odor representations and (2) optimize the distribution of the odor memories within that space. It does so, they demonstrate here, through the use of time as a computational variable. In other words, time is used not just to increase accuracy through temporal integration or signal averaging, but for the actual encoding of information in patterns that evolve through a sequence of informative states.

The large size of the representation space is a direct consequence of the temporal dynamics: the number of spatio-temporal trajectories through a neural population is much larger than the number of static spatial patterns that can be supported on that population. To extend the pixilated-screen analogy (with reckless exuberance), if each cycle of the 20 Hz pattern can be thought of as a "word," then the entire sequence of patterns might be analogous to a "sentence." Given any finite number of words in a dictionary (i.e., the number of possible pixel patterns on the screen), you can distinguish between many more categories with descriptive sentences than with single-word descriptors. (Ask your local wine snob to describe two different cabernets….) Laurent points out (Laurent, 2002), however, that an aspect of this scheme that is much more important than storing a huge number of olfactory memories is to increase the efficiency and accuracy of storing and recalling a smaller number of memories that the animal will need to recognize during its lifetime. To this end, a second consequence of the oscillatory temporal dynamics is the increased stability of each odor representation in the presence of noise and the optimal distribution of the odor codes within the representation space: the odor codes are maximally "spread out" in the available representation space (i.e., they are decorrelated), in a manner that enables their robust discrimination in a very short time after stimulus presentation, long before a set point is reached.

These results will very likely be generalizable to other olfactory systems, including those in vertebrates (Dulac, 1997; Abraham et al., 2004; Uchida and Mainen, 2003; Lin et al., 2005). Indeed, the results may generalize to other sensory modalities and cortical systems (see, for example, Vinje and Gallant [2000]). These experiments also offer several important lessons to those of us attempting to understand coding schemes in other sensory and cortical systems. First, measures of neural activity that have low temporal resolution would not have detected the transient dynamic patterns after odor presentation, which are precisely the aspects of the stimulus-evoked responses that carry the most information about the odor. (Imagine the consequences of trying to discriminate two different odors from a long-time-frame, coarse-spatial-scale average of the Palm Pilot screen in the earlier analogy.) Second, even with the high temporal precision achieved here, the overall coding scheme could not have been realized if the authors had not recorded such a large statistical sample of fully resolved spike trains from such a large proportion of the neurons in this population. (Again, imagine trying to discriminate the odors from the screen images if all but a few screen pixels were obscured.) Finally, and perhaps the most subtle and problematic, it appears that a researcher might need to know how subsequent processing stages interpret individual action potentials in order to measure the relevant information content of a spike train. Yikes!

John P. Miller
Center for Computational Biology
Montana State University
Bozeman, Montana 59717

Selected Reading

Abraham, N., Spors, H., Carleton, A., Margrie, T., Kuner, T., and Schaefer, A. (2004). Neuron *44*, 744–747.

Afraimovich, V., Zhigulin, V., and Rabinovich, M. (2004). Chaos *14*, 1123–1129.

Dulac, C. (1997). Neuron *19*, 477–480.

Laurent, G. (2002). Nat. Rev. Neurosci. *3*, 884–895.

Laurent, G., Stopfer, M., Friedrich, R.W., Rabinovich, M.I., Volkovskii, A., and Abarbanel, H.D. (2001). Annu. Rev. Neurosci. *24*, 263–297.

Lin, D., Zhang, S.-Z., Bolck, E., and Katz, L. (2005). Nature *434*, 470–477.

Mazor, O., and Laurent, G. (2005). Neuron, this issue, 661–673.

Rabinovich, M., Volkovskii, A., Lecanda, P., Huerta, R., Abarbanel, H.D., and Laurent, G. (2001). Phys. Rev. Lett. *87*, 068102. 10.1103/PhysRevLett.87.068102.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Neuron *39*, 991–1004.

Uchida, N., and Mainen, Z.F. (2003). Nat. Neurosci. *6*, 1224–1229.

Vinje, W.E., and Gallant, J.L. (2000). Science *287*, 1273–1276.

DOI 10.1016/j.neuron.2005.11.006

Cortical Map Reorganization without Cholinergic Modulation

Acetylcholine has been shown to modulate many forms of cortical plasticity. New evidence indicates that reorganization of adult primary auditory cortex is still possible after removal of cholinergic inputs. This finding suggests that acetylcholine may act less as a gate and more as a gain control on cortical plasticity.

Cholinergic modulation of learning and memory formation was first implicated a century ago when it was recognized that cholinergic antagonists can cause amnesia in human subjects. Substantial evidence now supports the hypothesis that acetylcholine regulates neural plasticity, learning, and memory (Hasselmo, 1995). The nucleus basalis (NB) provides the major source of cholinergic input to the cerebral cortex, and NB lesions in rats are known to impair motor learning, sensory discrimination, spatial navigation, short-term memory, and recovery from brain damage (McGaughy et al., 2000; Conner et al., 2005). NB lesions also block cortical plasticity, including the robust topographic reorganizations that follow motor skill learning, peripheral denervation, and digit amputation (Juliano et al., 1991; Webster et al., 1991; Conner et al., 2005). NB neurons respond strongly to behaviorally arousing stimuli, whether aversive or rewarding (Richardson and DeLong, 1991). Electrical stimulation of NB is sufficient to generate profound and long-lasting changes in receptive field organization within the primary sensory cortex. These findings suggest that acetylcholine released by NB neurons serves as a reinforcement signal to guide cortical plasticity (Weinberger, 2003; Kilgard et al., 2002).

It is not yet clear whether it is more accurate to describe cholinergic modulation as a gate on cortical plasticity or as a gain control mechanism. In this issue of *Neuron*, Kamke and colleagues demonstrate that cholinergic input is not required for lesion-induced plasticity in adult cortex (Kamke et al., 2005). While several other studies have reported that cortical plasticity is eliminated by NB lesions, their authors could not rule out the possibility that plasticity is only delayed. In this study, the authors used high-density microelectrode mapping techniques to document cortical map plasticity after damage to NB and to the high-frequency region of the cochlea. Several weeks after the cochlea was damaged, A1 neurons that had previously responded exclusively to high-frequency tones were found to respond to mid-frequency tones. This plasticity is analogous to that seen after digit amputation and results in a prominent over-representation of the highest frequency spared by the lesion. Nearly identical map reorganization was also observed in cats with near-complete destruction of cholinergic neurons projecting to the cortex.

Several aspects of the experimental design resulted in particularly clear results. The authors were able to confidently document the reorganization of inputs from the lesioned ear by comparing the frequency tuning for tones delivered to the lesioned ear and those delivered to the unlesioned (ipsilateral) ear. Responses to the unlesioned ear exhibited typical topographic organization, which presumably indicates the frequency tuning prior to cochlear damage. The difference between the frequency tuning for tones delivered to each ear reveals the degree of receptive field plasticity at each site without the need for multiple mapping surgeries that might alter cortical plasticity and topography. By selecting a toxin that specifically targets neurons expressing the low-affinity NGF receptor, the authors were able to precisely eliminate the subset of NB neurons that are cholinergic and project to the cortex, without damaging other basal forebrain neurons. Although the precision of the experimental manipulations and measurements strongly supports their conclusion that map plasticity is possible without cholinergic modulation, we are left wondering why NB lesions eliminate map reorganization after peripheral damage to the somatosensory system.

While it is possible that species- or modality-specific differences underlie this discrepancy, many other studies have reported more similarities than differences in the influence of acetylcholine across species and modalities. Differences in the experimental time course, lesioning technique, and even housing conditions could account for the apparent inconsistency and shed new light on the function of cholinergic modulation in the cortex.

While the studies of map plasticity in somatosensory cortex were conducted either 4 days or 4 weeks after the lesions were made, the current study documents plasticity that occurred at least 7 weeks after cochlear damage and 9 weeks after NB lesioning. The greater interval may provide time for recovery of cortical plasticity via upregulation of biochemical pathways or functional substitution by another system. Alternatively, it may be that acetylcholine is never really obligatory, but functions to increase the rate of cortical plasticity.

Another difference between this and earlier studies is the great precision with which cholinergic neurons that project to the cortex were destroyed. NB neurons that project to the amygdala and reticular thalamus were likely spared, as were the majority of NB projection neurons that are noncholinergic. Although there is currently little evidence that these pathways are critical for lesion-induced plasticity, it remains a real possibility (Sarter and Bruno, 2002). A third possibility is that some aspect