

motor response that leads to the same reinforcement for the mouse. We speculate that there may be an added savings in the GNG task that makes the discrimination and response easier. The mechanisms may involve some of the areas that are connected in feedforward or feedback arrangements with the olfactory bulb or cortex (Shiple and Adamek, 1984; Sobel et al., 1998). For most of the mice in the Rinberg et al. ZAC, while the animals enjoy better accuracy with longer sampling, the accuracy levels for the most difficult task are not as high as asymptote as those for the easiest task. Furthermore, sampling time increases are larger than in the prior study (on the order of 200–500 ms). In the Abraham et al. GNG task, accuracy is close to 100% with 70–100 ms longer sampling times for the more difficult discriminations.

One issue that the increase in sampling times for harder odor discriminations addresses is time evolution of spatial and temporal structure of odorant representations (Friedrich, 2006; Spors and Grinvald, 2002). In zebrafish these patterns evolve into more specific representations by about 400 ms and can persist for up to 1.5 s, but the question has arisen that if an animal can do the discrimination in a very brief time, what are these extended temporal codes used for? They could be used for representation of more specific odor features or even behavioral elements associated with the odorants, as it appears that mice do gain something with extra time. It has been shown that disruption of fast oscillatory coordination among insect antennal lobe projection neurons impairs fine odor discrimination, and, inversely, enhancement of fast oscillatory coordination among mammalian olfactory bulb mitral cells improves fine odor discrimination (Kay and Stopfer, 2006). What is needed is to examine the system physiology during the two different tasks in intact animals to determine whether and what central structures might be differentially involved in each task. Are sustained oscillations produced in response to difficult-to-discriminate odorants during the extra sampling time, similar to those seen in locusts, honeybees, and zebrafish (Kay and Stopfer, 2006; Laurent, 2002)? Owing to the disparity in time increases for the two tasks, we ask whether the system is doing the same thing during the extra 200–500 ms in the ZAC as it is doing during the 70–100 ms in the GNG task. We suspect that the performance improvement in the GNG task does take large advantage of the inputs to the olfactory bulb from higher-order reward and cognitive circuits, as we have shown previously (Kay and Freeman, 1998; Martin et al., 2006). In this case, the stimulus would not be processed separately from its meaning, and its neural representation would involve a network extending beyond the olfactory structures.

Finally, since we know that mice and rats can be very good at discriminating very small differences among odorants, we ask whether the tasks that we neurophysiologists use are the best ones for evaluating performance ability in these animals. The tasks certainly help us to uncover the behavioral and physiological mechanisms that *can* operate in this system, but *what else* can the system do? Is it that mice only strive for adequacy, or that our means of assessing their abilities only encourage them to rise to the level of good enough?

Leslie M. Kay,^{1,2} Jennifer Beshel,^{1,2} and Claire Martin²

¹Department of Psychology

²Institute for Mind and Biology

The University of Chicago

Chicago, Illinois 60637

Selected Reading

Abraham, N.M., Spors, H., Carleton, A., Margrie, T.W., Kuner, T., and Schaefer, A.T. (2004). *Neuron* 44, 865–876.

Friedrich, R.W. (2006). *Trends Neurosci.* 29, 40–47.

Kay, L.M., and Freeman, W.J. (1998). *Behav. Neurosci.* 112, 541–553.

Kay, L.M., and Stopfer, M. (2006). *Semin. Cell Dev. Biol.*, in press. Published online May 5, 2006. 10.1016/j.semcdb.2006.04.012.

Khan, R.M., and Sobel, N. (2004). *Neuron* 44, 744–747.

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

Martin, C., Gervais, R., Messaoudi, B., and Ravel, N. (2006). *Eur. J. Neurosci.* 23, 1801–1810.

Rinberg, D., Koulakov, A., and Gelperin, A. (2006). *Neuron* 51, this issue, 351–358.

Shiple, M.T., and Adamek, G.D. (1984). *Brain Res. Bull.* 12, 669–688.

Sobel, N., Prabhakaran, V., Desmond, J.E., Glover, G.H., Goode, R.L., Sullivan, E.V., and Gabrieli, J.D. (1998). *Nature* 392, 282–286.

Spors, H., and Grinvald, A. (2002). *Neuron* 34, 301–315.

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

DOI 10.1016/j.neuron.2006.07.015

Auditory Filters, Features, and Redundant Representations

Responses in auditory cortex tend to be weaker, more phasic, and noisier than those of auditory brainstem and midbrain nuclei. Is the activity in cortex therefore merely a “degraded echo” of lower-level neural representations? In this issue of *Neuron*, Chechik and colleagues show that, while cortical responses indeed convey less sensory information than auditory midbrain neurons, their responses are also much less redundant.

Recent years have seen a steady and sustained increase in the number of studies aiming to understand the workings of auditory cortex, but despite many ingenious experiments and meticulous studies, auditory cortex remains much less well understood than its visual counterpart. It is as if each new experiment uncovers a new and sometimes unexpected piece of a great puzzle, but so far the various pieces steadfastly refuse to fall together to form a coherent picture. Thus, central auditory processing seems a great deal more complicated (or, as auditory researchers would put it, a great deal more “interesting”) than its visual counterpart. The gaps in our understanding of the auditory system are particularly marked in the field of auditory object recognition. Visual objects lend themselves to stylized representations, for example through line drawings, and ever since Hubel and Wiesel’s classic work, line segments and edges have been recognized as “archetypal” low-level visual features. They are easily detected by simple cells in primary visual cortex while complex cells signal their

presence in a somewhat position invariant manner. As one ascends the visual processing stream toward infra-temporal cortex, one finds neurons which are selective to specific, complex, and increasingly abstracted combinations of such low-level features. In this manner, these neurons become detectors of essential feature combinations which characterize particular visual objects irrespective of the precise position, size, or contrast of the visual stimulus. These transformations can be modeled in a biologically plausible manner, as has been done, for example, in the VisNet model developed by Rolls and colleagues (Elliffe et al., 2002; Rolls and Treves, 1998), and while it might be premature to conclude that vision is “solved,” there are at least very good indications as to how visual object recognition might work.

One might think that auditory object recognition could similarly be done through a processing hierarchy wherein neurons at successively higher levels become responsive to increasingly complex and diverse combinations of low-level features. But auditory neuroscientists are not yet in a position to confirm this, and their efforts are at least in part hampered by the fact that we still have no suitable set of low-level auditory features that might be considered the auditory equivalent of the line segment. Certainly, Fourier analysis allows us to decompose any arbitrary sound into pure tone components, but in a natural sound the temporal interrelations of acoustic features are all important, and Fourier analysis fails to capture these in an intuitive manner. It is therefore not surprising that auditory neuroscientists are increasingly turning away from pure tones and use either artificial stimuli, rich in temporal modulations (Escabi et al., 2003; Garcia-Lazaro et al., 2006; Klein et al., 2000), or recordings of natural sounds, possibly after some suitable manipulation (Nelken et al., 1999; Schnupp et al., 2006; Wang and Kadia, 2001).

One interesting new study which embraces this later approach is published in this issue of *Neuron* (Chechik et al., 2006). In this study, Chechik and colleagues recorded the responses of single neurons to a battery of “raw” and manipulated bird song recordings at three key stages of the cat central auditory pathway, namely the inferior colliculus (IC), the medial geniculate nucleus (MGN), and the primary auditory cortex (A1). All three structures are known to have a tonotopic organization, and the traditional, although somewhat simplistic, view is that each of these stages can therefore be thought of as a set of frequency-tuned neural filter banks. Naturally, if the neurons are merely frequency filters, then neurons with similar frequency tuning should give essentially similar responses. Frequency tuning curves in the central auditory pathways are often broad and overlapping, and one would therefore expect to see a fair amount of “redundancy” in the neural population response, as a number of similarly tuned neurons should tell us very similar things about any given stimulus. The results presented by Chechik et al. (2006) indicate that this expectation is largely borne out at the level of the IC, but, surprisingly, appears not to apply at the level of the auditory cortex.

Of course, whether the responses in two simultaneously recorded neurons appear “similar” and therefore “redundant” may depend critically on which mea-

sure of the neural response one considers important. For example, two neural responses may be very similar in overall spike count, yet exhibit crucial differences in their temporal discharge pattern (Schnupp et al., 2006). To achieve a measure of redundancy which is largely independent of assumptions about the relevant response metric, Chechik and colleagues resorted to an information theoretic analysis of a variety of different neural response metrics, but obtained the same result in each case.

Information theory views neural spike trains and sensory stimuli as random variables and quantifies their interdependence by a measure known as mutual information. If two random events x and y are entirely independent of each other, then observing x cannot tell you anything about y , and the probability of x and y occurring together $p(x,y)$ would equal the product of their individual probabilities $p(x) \times p(y)$. Information theory therefore quantifies the mutual information between x and y as a departure from independence, and quantifies this as the logarithm of the ratio $p(x,y)/(p(x) \times p(y))$. For independent x and y this ratio evaluates to one and the log of one is zero, i.e., x carries no information about y . The advantage of this approach is that it allows us to quantify the information content of a neural spike train without having to make any assumptions about how the spike train ought to be decoded. However, in order to calculate the information content of a neural response precisely, one needs to know the value of the joint probability $p(x,y)$ of any possible spike train x to occur together with any possible stimulus y . While getting very accurate estimates for these probability distributions can be very difficult in practice, there are nevertheless good reasons to believe that reasonably accurate estimates of the information content of the response of an auditory neuron can be obtained (Nelken et al., 2005). From determining the information content of individual neural responses it is a relatively small step to quantifying the informational redundancy of pairs of neural responses. For independently firing neurons, the information conveyed by each should sum, and redundancy can be quantified as a departure from this expectation of additive information.

In the IC, Chechik et al. (2006) found neural responses to be highly redundant, a fact that they attribute largely to stimulus-induced correlations in the neural firing patterns. Responses of neurons in auditory cortex carried on average less information per response than those in IC, but the informational redundancy between pairs of cortical neurons was also much lower. The decline in the information content of individual responses in cortex relative to IC was to be expected. It is in large part attributable to the lower mean response firing rates seen in cortex, and to a lesser extent an inevitable consequence of the data processing inequality, a physical law which states that subsequent levels of any information processing system, unable to invent new information, can at best do a tolerably good job at preserving the information passed up from previous levels. Consequently, the often divergent connections in sensory pathways lead to the seemingly peculiar situation that the activity of increasingly larger numbers of neurons represents ever smaller total amounts of sensory information.

But while the decline in information content per response is therefore inevitable, the sharp decline in informational redundancy is not, and it is very interesting to consider what the reasons and the purpose behind this redundancy reduction might be. Ed Rolls' VisNet model of cortical visual object processing also exhibits a redundancy reduction at successively higher levels of the processing stream. There the redundancy reduction arises as higher levels of the network become sensitive to increasingly abstract feature combinations. In the VisNet model, redundancy reduction is therefore a hallmark of a transition from a "feature-based" to an "object-based" representation, and it is intriguing to speculate that the redundancy reduction in the auditory pathway described by Chechik et al. (2006) might similarly be interpreted as the fingerprint of a transition from an acoustic-based feature toward a more "auditory object-based" representation.

The VisNet model is not "born" with a low-redundancy representation of its stimuli in its top layers. The redundancy reduction only arises after a competitive learning process in which higher layers become sensitive to specific feature combinations (Rolls and Treves, 1998; Rolls, 1995). One might predict that the low-redundancy representations observed at the levels of the auditory cortex may similarly be the result of developmental or learning mechanisms which decorrelate the responses of individual cortical neurons. This could easily be tested by measuring redundancy in the cortex of young, naive animals with little auditory experience. Achieving a low-redundancy representation in auditory cortex could well be an important part of learning how to hear.

Jan Schnupp¹

¹Department of Physiology, Anatomy, and Genetics
Sherrington Building
University of Oxford
Oxford OX1 3PT
United Kingdom

Selected Reading

- Chechik, G., Anderson, M.J., Bar-Yosef, O., Young, E.D., Tishby, N., and Nelken, I. (2006). *Neuron* 51, this issue, 359–368.
- Elliffe, M.C., Rolls, E.T., and Stringer, S.M. (2002). *Biol. Cybern.* 86, 59–71.
- Escabi, M.A., Miller, L.M., Read, H.L., and Schreiner, C.E. (2003). *J. Neurosci.* 23, 11489–11504.
- Garcia-Lazaro, J.A., Ahmed, B., and Schnupp, J.W. (2006). *Curr. Biol.* 16, 264–271.
- Klein, D.J., Depireux, D.A., Simon, J.Z., and Shamma, S.A. (2000). *J. Comput. Neurosci.* 9, 85–111.
- Nelken, I., Rotman, Y., and Bar Yosef, O. (1999). *Nature* 397, 154–157.
- Nelken, I., Chechik, G., Mscic-Flogel, T.D., King, A.J., and Schnupp, J.W. (2005). *J. Comput. Neurosci.* 19, 199–221.
- Rolls, E.T. (1995). *Behav. Brain Res.* 66, 177–185.
- Rolls, E., and Treves, A. (1998). *Neural Networks and Brain Function* (Oxford, England: Oxford University Press).
- Schnupp, J.W., Hall, T.M., Kokelaar, R.F., and Ahmed, B. (2006). *J. Neurosci.* 26, 4785–4795.
- Wang, X., and Kadia, S.C. (2001). *J. Neurophysiol.* 86, 2616–2620.

The Lure of the Unknown

Using event-related fMRI, Bunzeck and Düzel show that midbrain regions putatively housing dopamine cell bodies activate more for novel pictures than for negative pictures, pictures requiring a motor response, or repeated pictures. These findings indicate that midbrain regions preferentially respond to novelty and suggest that novelty can serve as its own reward.

Meriwether Lewis and William Clark spent years working at it, Edmund Hillary and Tenzing Norgay climbed Mt. Everest for it, Neil Armstrong flew into space for it, and Robert Falcon Scott died for it—a chance to discover something never before seen. A long tradition of human exploration testifies to the motivating force of novelty. Evolutionary biologists have argued that in order to flourish, all foraging species must have a drive to explore the unknown (Panksepp, 1998). How such a drive manifests in the brain, though, has remained unclear. In this issue of *Neuron*, for the first time, Bunzeck and Düzel (2006) show that midbrain regions that putatively house dopamine neurons preferentially respond to novel rather than rare, arousing, or behaviorally relevant stimuli (Bunzeck and Düzel, 2006).

From the outside, the ventral tegmental area (VTA) and substantia nigra (SN) are easy to miss. Nestled deep in a bend of the brainstem, these nuclei house the bodies of most of the dopamine neurons that innervate the striatum and prefrontal cortex. Tract tracing studies indicate that while the VTA projects to more ventral regions of the striatum and prefrontal cortex, the SN projects to more dorsal and lateral regions of the striatum and prefrontal cortex. Though small, these nuclei are in a position to exert widespread influence. Indeed, from the inside, life without these midbrain neurons is far from easy. For instance, both organic lesions (due to Parkinson's disease) and synthetic lesions (due to improperly manufactured drugs) of the SN/VTA lead to mental and physical immobility.

While lesion studies suggest that dorsal pathways innervated by the SN play a role in movement, ventral pathways innervated by the VTA play a less-understood role in motivation (Haber and Fudge, 1997). Some prominent theories hypothesize that activity in this ventral pathway confers "salience" upon stimuli (Berridge and Robinson, 1993). However, theorists have defined salience differently, confounding empirical attempts to isolate the function of these nuclei. For instance, some definitions of salience invoke novelty, others invoke behavioral relevance, and still others invoke arousal.

Here, Bunzeck and Düzel operationally define "salience" in four different ways. During acquisition of event-related fMRI, the investigators showed subjects pictures of faces or outdoor scenes embodying different attributes of salience and then measured the SN/VTA response to these different stimuli. A first group of pictures was novel, or never seen before. A second group of pictures was behaviorally relevant, requiring a button press. A third group of pictures was negative and thus presumed to be arousing (i.e., a negative expression in the case of faces, or a car accident in the case of scenes).