Retinal Coding of Visual Scenes—Repetitive and Redundant Too?

Visual information reaches the brain by way of a fine cable, the optic nerve. The million or so axons in the optic nerve represent an information bottleneck in the visual pathway—where the fewest number of neurons convey the visual scene. It has long been thought that to make the most of the optic nerve's limited capacity the retina may encode visual information in an optimally efficient manner. In this issue of *Neuron*, Puchalla et al. report a test of this hypothesis using multielectrode recordings from retinal ganglion cells stimulated with movies of natural scenes. The authors find substantial redundancy in the retinal code and estimate that there is an ~10-fold overrepresentation of visual information.

In the traditional view of retinal function, each ganglion cell axon represents an independent channel of information from eye to brain. The bundle of such axons comprising the optic nerve forms a set of labeled lines, each conveying the presence of a specific visual feature. Given the limited capacity of the optic nerve, the idea that each of these features should be nearly unique has long been attractive (reviewed by Barlow, 2001). However, mounting evidence suggests that the traditional view of retinal coding is inadequate in several respects. Far from just detecting a fixed set of visual elements, the retina enjoys significant computational power and shapes its code in response to contrast changes (Smirnakis et al., 1997), object motion (Berry et al., 1999; Olveczky et al., 2003), and spatial scale (Smirnakis et al., 1997) within the visual scene. Other results challenge the view of ganglion cells as independent encoders of information. Patterns of concerted spiking among pairs (Mastronarde, 1989; Meister et al., 1995) and even larger groups of retinal ganglion cells (Schnitzer and Meister, 2003) convey spatial information that is distinct from that conveyed by activity from the same cells individually. Puchalla et al. again challenge our understanding of retinal coding by using an information theoretic analysis to test whether each ganglion cell actually transmits unique visual messages (Puchalla et al., 2005). The findings are largely inconsistent with uniqueness and suggest that there is considerable overlap in ganglion cells' visual messages. On average, the ganglion cell population appears to transmit roughly ten "copies" of each piece of visual information. To appreciate this result, it is helpful to know a little bit about information theory.

The Bell Labs mathematician Claude Shannon developed information theory as a means of quantifying the capacity of a communication channel (Shannon and Weaver, 1963). The theory is general and applies not just to phone lines but to communication media of all forms. The measure of information Shannon introduced, the bit, is now used to describe the storage capacity of computer disks, the size of digital images, and even the communication rate of humpback whale songs. Likewise, the bit can equally well quantify the amount of visual information transmitted down the optic nerve. An important aspect of Shannon's theory is that information is measured without reference to semantic meaning. Instead, information can be quantified through statistical examination of the symbol set used to represent the messages. Both the frequency of and correlations between symbols are important. For instance, rare letters in the English language, such as "*x*," carry more information than common letters, such as "*e*," because the "*x*" restricts the word being conveyed to a far smaller set of possibilities. Correlations between symbols decrease the efficiency of the symbolic representation. For example, the "*u*" in the obligate pairing "*qu*" is redundant in that it provides no additional information in English beyond that of the "*q*" alone.

In the study of retinal coding, the question of whether ganglion cells' visual messages are unique can be addressed by examining statistics of the appropriate symbol set, the spikes, which represent the visual stimuli. Thanks to Shannon and to more recent work on how spiking patterns can properly be analyzed as a discrete symbol set (Strong et al., 1998), this approach to quantifying transmitted information works, even though our ability to interpret the retinal code remains incomplete. Here the relevant statistics concern spike frequencies and correlations in ganglion cell activity patterns. By tabulating spike train statistics for individual and pairs of ganglion cells, Puchalla et al. studied whether two cells jointly convey more (synergistic coding), the same (independent coding), or less (redundant coding) information than the sum of what the two cells convey individually. With redundant coding, in spite of the diminished signaling efficiency, there is often increased tolerance to coding errors, much as with a backup copy of a computer file. Redundant coding can be achieved in many ways, for example by repeating every word to to ty guald guard guard against ageinst against mixspelling misspelling misspelling. Alternatively, redundancy can reflect a common noise source, and so need not provide error protection. Redundancy in the retinal code could reflect correlations in the visual scene or intrinsic aspects of retinal circuitry. To address these issues, one needs access to spiking patterns of ganglion cell populations.

Planar electrode arrays that record activity from numerous ganglion cells have made it possible to obtain large data sets about the retinal code. An earlier multielectrode array study examined responses to spatially uniform flicker and found that with this highly structured stimulus ganglion cells of the same functional type provided redundant information, but cell pairs of different types provided independent information (Warland et al., 1997). Puchalla et al. examined such issues by stimulating the salamander retina with movies of natural scenes and recording spiking responses from tens of cells simultaneously. In order to translate continuous spike trains into discrete symbols, the authors binned the spikes in 10 ms intervals and counted the frequencies of all spiking patterns extending over four intervals (1 1 1 1, 1 1 1 0, etc.), for individual and pairs of cells. This allowed the researchers to measure the likelihoods of the various spiking patterns in response to specific portions of the movie and thus to quantify the visual information transmitted via the spikes.

Interestingly, the authors find a broad distribution of values for the redundancy between cell pairs, ranging

from near zero to 50%. Among redundant cell pairs, the average redundancy was ~15% for close neighbors, lower for cell pairs further apart. At first blush this level of redundancy may seem modest. But keep in mind that each ganglion cell has a lot of neighbors—over a thousand neighbors within a radius of 500 μm—and so this degree of overrepresentation can add up quickly. In the absence of direct redundancy measurements on the entire cell population, the authors performed an extrapolation by assuming that the information conveyed redundantly is uniformly distributed across the cell population. In essence, the assumption is that the portion of cell A's messages encoded redundantly by cell B is set independently from the portion of A's messages encoded redundantly by cell C. Under this reasonable working hypothesis, the authors estimated that with natural scene stimulation there is ~10-fold overrepresentation of visual information across the ganglion cell population. Redundant and repetitive too! Of note, significant redundancy persists under visual stimulation with a randomly flickering checkerboard, which lacks the spatial structure present in natural scenes or uniform flicker. This indicates much of the redundancy likely originates within retinal circuitry.

How might retinal circuitry give rise to redundancy? Much ganglion cell activity occurs in the form of synchronous spiking by groups of multiple cells (Mastronarde, 1989; Meister et al., 1995; Schnitzer and Meister, 2003). Analysis of multielectrode data indicates more than ~50% of all spikes recorded from the retina may be of this form (Schnitzer and Meister, 2003). Such concerted spiking creates strong correlations between ganglion cell spike trains, which lower the efficiency of the representation of visual scenes. On the other hand, synchronous spiking groups also convey distinct visual messages from those conveyed by individual ganglion cells: the receptive fields of concerted spiking patterns are more sharply localized than those of single cells (Meister et al., 1995; Schnitzer and Meister, 2003). In principle, communication of this more precise spatial information through synchronous spiking might enable synergistic coding in the retina. In practice, receptive fields of the ganglion cell groups are only ~30% smaller in area than those of individual cells (Schnitzer and Meister, 2003), so any weak synergy in the transmission of spatial information is likely more than offset on average by the strong redundancy due to correlated spiking. Thus, the unexpectedly large values of redundancy found by Puchalla et al. may be closely related to the surprisingly large proportion of synchronous activity in the retina's outputs. Existing ideas about the physiological mechanisms underlying concerted spiking therefore lead to some potential explanations of redundancy.

Physiological and pharmacological data indicate that a significant portion of concerted ganglion cell spiking probably arises from common gap junction-mediated inputs from interneurons in the inner retina that either spike or exhibit fast voltage transients (Figure 1) (Brivanlou et al., 1998). These interneurons may be amacrine cells, or conceivably ganglion cells, that drive multiple ganglion cells to spike synchronously via common electrical inputs. Thus, the receptive field of an individual ganglion cell might be built from more localized receptive fields of the interneurons that induce the ganglion cell to spike (Figure 1). In such a way, the ganglion cells might multiplex messages from interneuron inputs, allowing the retina to convey a greater number of distinct messages than there are fibers in the optic nerve. Because these fibers are considerably fewer in number than the photoreceptors, correlated ganglion cell spiking might be used to overcome the limited thickness of the optic nerve and to transmit spatial information near the resolution limit set by photoreceptor density. In this view of retinal coding, it is the interneurons—i.e., the concerted spiking patterns—that should transmit nearly unique visual messages, not the ganglion cells. However, messages about fine spatial details come at the cost of extra spikes and redundancy in the retinal code.

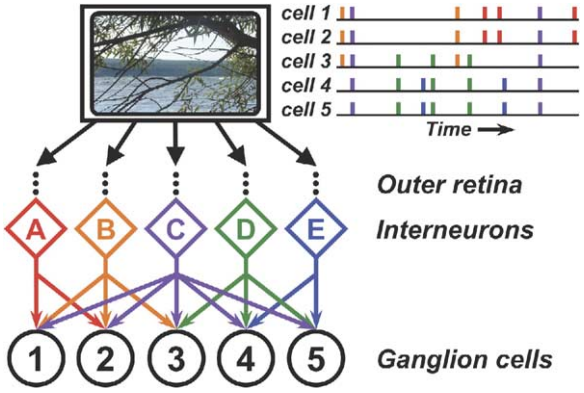

Figure 1. Concerted Ganglion Cell Spiking Arising Due to Shared Electrical Inputs from Interneurons in the Inner Retina

A portion of synchronous ganglion cell spiking activity is thought to arise via common gap junction-mediated inputs from interneurons in the inner retina (Brivanlou et al., 1998). Visual stimuli (such as movies of natural scenes) are detected by photoreceptors and processed in the outer retina (ellipses). The resulting signals lead to interneuron activity (lettered cells), which can drive concerted spiking in multiple ganglion cells (numbered cells) receiving common electrical inputs (colored arrows). Thus, the receptive fields of the interneurons are also those of the resulting concerted spiking patterns. The ganglion cell spike trains (upper right) exhibit synchronous activity in various patterns (colored spikes), reflecting input from different interneurons. In this way, ganglion cell receptive fields may be built from the more spatially localized receptive fields of concerted spiking patterns (Schnitzer and Meister, 2003).

This hypothesis concerning the role of gap junction inputs in creating redundancy should be testable. One might carry out an analytical test, by performing the analysis of Puchalla et al. on trains of concerted spikes and asking whether these exhibit redundancy in the transmitted visual information. Alternatively, one might examine coding redundancy in transgenic mice that lack specific classes of connexin proteins forming neuronal gap junctions. Although rearrangements of retinal circuitry during development might alter results from such mice, one might expect to see diminished levels of concerted spiking and redundancy in the ganglion cell code. Such experiments might lead to new understanding of how to interpret the retina's messages, bringing us from spike statistics to semantics.

Eran A. Mukamel[1] and Mark J. Schnitzer[2]
[1]Department of Physics
James H. Clark Center for
Biomedical Engineering & Sciences
Stanford University
Stanford, California 94305
[2]Department of Biological Sciences
Department of Applied Physics
James H. Clark Center for
Biomedical Engineering & Sciences
Stanford University
Stanford, California 94305

Selected Reading

Barlow, H. (2001). Network *12*, 241–253.

Berry, M.J., Brivanlou, I.H., Jordan, T.A., and Meister, M. (1999). Nature *398*, 334–338.

Brivanlou, I.H., Warland, D.K., and Meister, M. (1998). Neuron *20*, 527–539.

Mastronarde, D.N. (1989). Trends Neurosci. *12*, 75–80.

Meister, M., Lagnado, L., and Baylor, D.A. (1995). Science *270*, 1207–1210.

Olveczky, B.P., Baccus, S.A., and Meister, M. (2003). Nature *423*, 401–408.

Puchalla, J.L., Schneidman, E., Harris, R.A., and Berry, M.J. (2005). Neuron *46*, 493–504.

Schnitzer, M.J., and Meister, M. (2003). Neuron *37*, 499–511.

Shannon, C.E., and Weaver, W. (1963). The Mathematical Theory of Communication, Second Edition (Chicago: University of Illinois Press).

Smirnakis, S.M., Berry, M.J., Warland, D.K., Bialek, W., and Meister, M. (1997). Nature *386*, 69–73.

Strong, S.P., Koberle, R., de Ruyter van Steveninck, R.R., and Bialek, W. (1998). Phys. Rev. Lett. *80*, 197–200.

Warland, D.K., Reinagel, P., and Meister, M. (1997). J. Neurophysiol. *78*, 2336–2350.

DOI 10.1016/j.neuron.2005.04.018