the figure (panel A) describes the minimum requirements for the DS response, which are a spatial asymmetry and a time delay between the excitatory and inhibitory synapses. Yoshida et al. (2001) demonstrate that the starburst cell may contribute to both of these minimum requirements. However, one can predict that we have to understand much more of the molecular and synaptic details of the starburst cell circuitry before directional selectivity, this classical problem of computational neuroscience, is finally understood.

Heinz Wässle
Max-Planck-Institut für Hirnforschung
Deutschordenstrasse 46
D-60528 Frankfurt/M.
Germany

Selected Reading

Ariel, M., and Daw, N.W. (1982). J. Physiol. (Lond.) *324*, 161–185.

Barlow, H.B., and Levick, W.R. (1965). J. Physiol. (Lond.) *178*, 477–504.

Borg-Graham, L. (2001). Nat. Neurosci. *4*, 176–183.

Brandstätter, J.H., Greferath, U., Euler, T., and Wässle, H. (1995). Vis. Neurosci. *12*, 345–358.

Famiglietti, E.V. (1991). J. Comp. Neurol. *309*, 40–70.

He, S., and Masland, R.H. (1997). Nature *389*, 378–382.

Koulen, P., Malitschek, B., Kuhn, R., Wässle, H., and Brandstätter, J.H. (1996). Eur. J. Neurosci. *8*, 2177–2187.

O'Malley, D.M., Sandell, J.H., and Masland, R.H. (1992). J. Neurosci. *12*, 1394–1408.

Taylor, W.R., He, S., Levick, W.R., and Vaney, D.I. (2000). Science *289*, 2347–2350.

Vaney, D.I. (1991). Progr. Ret. Res. *9*, 1–28.

Wässle, H., and Boycott, B.B. (1991). Physiol. Rev. *71*, 447–480.

Yoshida, K., Watanabe, D., Ishikane, H., Tachibana, M., Pastan, I., and Nakanishi, S. (2001). Neuron *30*, this issue, 771–780.

Reverse Spikeology: Predicting Single Spikes

Neural models that simulate single spike trains can help us understand the basic principles of neural coding in vision. Keat et al. (2001) develop a hybrid model that combines spatiotemporal filtering with nonlinear spike generation. The model does a good job of predicting the responses of single retinal ganglion cells and thalamic relay neurons.

When a digital camera is used to capture an image, the camera converts the analog luminance signal into a binary series of zeros and ones. Given only this binary signal and no information about the coding process, it would be difficult for an observer to reconstruct the original image. A similar coding process happens in the brain when analog signals in the retina are recoded by retinal ganglion cells into action potentials. From that point onward, neurons throughout the visual system represent the world in terms of time varying spike trains. Understanding these neural codes is one of the central goals of systems neuroscience. The report by Keat et al. (2001 [this issue of *Neuron*]) addresses this issue by modeling how early visual neurons translate stimuli into spikes.

To investigate neural coding, many researchers have used computational tools adopted from linear systems analysis (Marmarelis and Marmarelis, 1978). One of the most common approaches is to treat sensory neurons as filters whose outputs correspond to specific stimulus attributes in a scene. For example, neurons in primary visual cortex are often modeled as filters tuned to specific orientations and spatial frequencies (DeValois and DeValois, 1990). Often such models consider only a few critical stimulus attributes and ignore others. For computational simplicity, they often assume that stimuli are fixed and unchanging. However, even the simplest static visual stimulus can elicit a dynamic series of action potentials from visual neurons, and these responses can develop and persist over time. The simplest filter models fail to account for these complex temporal response dynamics.

Recent efforts have focused more closely on modeling responses to time-varying visual stimuli. One simple coding assumption that might reasonably account for the data is that neurons transmit information about time-varying stimuli by means of a temporally modulated rate code (for discussion of coding schemes, see Rieke et al. [1997]). According to this view, information is encoded in the instantaneous firing rate, and downstream neurons must estimate the instantaneous rate to decode the spike train. This process cannot be modeled in terms of a static filtering operation but requires a spatiotemporal filter sensitive to time-varying signals (Bialek et al., 1991; Theunissen et al., 2001; Vinje and Gallant, 1998).

While spatiotemporal filtering models have clear advantages over static filters, they also have limitations. Because real neurons transmit information via spike trains, they must simultaneously estimate the instantaneous firing rate of their inputs from a discrete series of spikes and translate their own time-varying state into a series of spikes for output to other cells. To resolve these decoding and encoding problems, modelers must consider the relationship between spatiotemporal filtering and the mechanisms governing spike generation. A model that produces realistic spike trains in response to a time-varying visual stimulus has the potential to substantially improve our understanding of neural coding.

Unfortunately, several practical problems have dissuaded researchers from constructing such models. Most importantly, many neurophysiological experiments have reported that action potentials are elicited unreliably and that their timing accuracy is poor (Shadlen and Newsome, 1998). In fact, these experiments have long been used to argue that single spikes are irrelevant and to justify simplifying assumptions (such as the mean-rate-code assumption described above). An additional complication stems from the spatiotemporal filtering framework itself: filtering is usually implemented as a quasilinear operation, but spiking is an inherently nonlinear process that requires a different modeling approach. Computational models for spike generation have been proposed (Gabbiani and Koch, 1998), but there has been relatively little work aimed at developing hybrid models that integrate a filter-based input stage with a realistic

spike generator (but see Chance et al., 1998; Worgotter and Koch, 1991).

In this issue of *Neuron*, Keat et al. (2001) address this problem head on. Their stated goal is to create a model that predicts individual spikes elicited in response to complex visual stimulation, focusing on retinal ganglion cells and neurons in the lateral geniculate nucleus. These cells generate relatively reliable and accurate action potentials (Berry et al., 1997; Kara et al., 2000), which is a necessary prerequisite for predicting single spikes solely from retinal input. They employ a standard linear filtering approach to account for the basic stimulus selectivity of their neurons. Additional mechanisms are then invoked to control spike generation and timing. Spike generation is governed by thresholding the filter output, and a feedback mechanism is proposed to account for both the refractory period and the observed bursting patterns of spikes.

Unfortunately, even peripheral sensory neurons and thalamic relay cells are neither perfectly accurate nor completely reliable. When stimuli are repeated precisely, the resulting spike trains are still somewhat variable. This apparently random trial-to-trial variability cannot be accounted for by a deterministic model. Keat et al. approach this problem by considering responses at two levels. First, they model the occurrence of separate spike clusters (which they call "events") that can be identified in the records of these cells. Event occurrence is governed by the estimated filter and the threshold, while variability in event timing is modeled by adding stochastic noise to the threshold. Second, they model the statistical distribution (mean number and variance) of spikes within each event. This is accomplished by adding a stochastic noise term to the feedback loop. Unfortunately, the use of stochastic variables eliminates any chance of predicting when every single spike occurs. However, it still allows them to account for the variability of event timing and numbers of spikes within each event. The resulting model is very compact. Although it requires 20 parameters, 15 of these are used to estimate the filter, and only 5 are required to control spike generation and stochastic variability.

Keat et al. recorded spike trains from several preparations in response to a complex temporal input. Recordings were made from retinal ganglion cells in salamander, rabbit, and cat, and also from cat lateral geniculate nucleus. Certain aspects of their model evaluation procedures are particularly noteworthy. Keat et al. fit their model separately to data acquired from each cell. In all cases, model parameters are fit using only a portion of the data and the model is evaluated in terms of its ability to predict other data not used in the initial fit. This clean separation of exploratory and confirmatory data sets ensures that the results are not biased by overfitting.

Remarkably, Keat et al. report good performance of their model in all three preparations, even though the spike trains encountered in each case had very different statistical characteristics. In many cases, the model predicts the occurrence and timing variability of events quite well. The model does not predict single spikes precisely, but this is to be expected due to the stochastic nature of the variables governing spike generation. The model also tends to fail when events are not elicited reliably or when the timing variability of events exceeds their duration.

By attempting to predict spike trains of individual peripheral sensory and thalamic relay neurons, Keat et al. demonstrate the practical utility of hybrid single-cell models that combine spatiotemporal filtering with nonlinear spike generation. The Keat et al. model makes a clear distinction between deterministic and stochastic aspects of the neural response. Future studies might reveal whether the apparently random stochastic features are due to processes within the cell (which might be measurable using intracellular techniques) or whether they reflect instead the influence of the local neural network. Because this model explicitly addresses spike generation while retaining the power and flexibility of the filtering approach, it might also serve as a starting point for modeling more central sensory neurons. Quantitative models such as this allow us both to predict neural responses at a fine time scale and to understand the limits of such predictions. They clearly demonstrate the growing power of computational approaches for understanding how sensory neurons encode information.

Jack L. Gallant[1] and William E. Vinje[2]
[1] Department of Psychology and
[2] Department of Molecular and Cellular Biology
University of California, Berkeley
Berkeley, California 94720

Selected Reading

Berry, M.J., Warland, D.K., and Meister, M. (1997). Proc. Natl. Acad. Sci. USA *94*, 5411–5416.

Bialek, W., Rieke, F., de Ruyter van Steveninck, R.R., and Warland, D. (1991). Science *252*, 1854–1857.

Chance, F.S., Nelson, S.B., and Abbott, L.F. (1998). J. Neurosci. *18*, 4785–4799.

DeValois, R.L., and DeValois, K.K. (1990). Spatial Vision (New York, NY: Oxford).

Gabbiani, F., and Koch, C. (1998). In Methods in Neural Modeling, C. Koch and I. Segev, eds. (Cambridge, MA: MIT Press), pp. 313–360.

Kara, P., Reinagel, P., and Reid, R.C. (2000). Neuron *27*, 635–646.

Keat, J., Reinagel, P., Reid, R.C., and Meister, M. (2001). Neuron *30*, this issue, 803–817.

Marmarelis, P.Z., and Marmarelis, V.Z. (1978). Analysis of Physiological Systems: The White Noise Approach (New York, NY: Plenum).

Rieke, F., Warland, D., de Ruyter van Steveninck, R., and Bialek, W. (1997). Spikes: Exploring the Neural Code (Cambridge, MA: MIT Press).

Shadlen, M.N., and Newsome, W.T. (1998). J. Neurosci. *18*, 3870–3896.

Theunissen, F.E., David, S.V., Singh, N.C., Hsu, A., Vinje, W.E., and Gallant, J.L. (2001). Network: Computation in Neural Systems, in press.

Vinje, W.E., and Gallant, J.L. (1998). In Advances in Neural Information Processing Systems 10, M.I. Jordan, M.J. Kearns, and S.A. Solla, eds. (Canbridge, MA: MIT Press), pp. 236–242.

Worgotter, F., and Koch, C. (1991). J. Neurosci. *11*, 1959–1979.