Neuron Previews

Melting the Iceberg: Contrast Invariance in Visual Cortex

Matteo Carandini^{1,*}

¹The Smith-Kettlewell Eye Research Institute, San Francisco, CA 94115, USA *Correspondence: matteo@ski.org DOI 10.1016/j.neuron.2007.03.019

Neurons in visual cortex maintain their selectivity for stimulus orientation despite wide variations in stimulus contrast. Achieving this invariance is a challenge because of the iceberg effect: subthreshold responses are broader in selectivity than firing responses, and increasing contrast would bring them above threshold. An article in this issue of *Neuron* by Finn et al. explains how neurons solve this problem using simple mechanisms: contrast-gain control, additive noise, and firing threshold.

A major question in neuroscience concerns the circuits of the cerebral cortex and specifically the function of these circuits in terms of computation. There is great hope and reasonable expectation that the fundamental circuits are modular, i.e., they are repeated across cortical areas to apply similar computations to different purposes. If so, our best bet to understand them might be to study the primary visual cortex (V1). Area V1 is arguably the "giant squid axon" of cortical neurophysiology: we can control its sensory inputs with exquisite accuracy, we know quite a bit about is circuitry, we understand the basic response properties of its neurons, and we have a fairly clear idea of the computations that it performs. These computations are complex enough to be interesting and yet simple enough to study in detail.

The property of V1 neurons that has captured the most attention is their selectivity for orientation. Orientation selectivity was discovered by Hubel and Wiesel, who proposed that it arises through summation of inputs from the lateral geniculate nucleus (LGN). If a cortical cell summed the outputs of LGN neurons whose receptive fields are aligned, its own receptive field would be selective for orientation (Figure 1A). In the half-century that followed this classic model, mountains of evidence have accumulated both to support it and to show its limitations, resulting in considerable controversy (Ferster and Miller, 2000).

A key limitation of the classic model is a behavior known since at least the

1970s as the "iceberg effect" (e.g., Rose and Blakemore, 1974). The classical model relies on the spike threshold to hide the depolarizations caused by stimuli having the wrong orientation. This threshold needs to be high to cope with stimuli of high contrast (green in Figures 1B and 1C), but then it completely hides the responses to low contrasts (red in Figures 1B and 1C). A lower threshold would cope with the lower contrasts, but it would cause large, barely tuned firing responses at high contrast. The classic model, therefore, predicts that orientation selectivity should broaden markedly with increasing contrast, much as an iceberg could be made wider by raising it further out of the sea. This prediction is wrong: the selectivity of firing-rate responses in V1 is remarkably invariant with contrast; increasing contrast does make the responses get larger, but it does not broaden their selectivity (Skottun et al., 1987).

An elegant study in this issue of *Neuron* (Finn et al., 2007) explains how V1 neurons solve this problem. By recording intracellularly from V1 simple cells, Finn, Priebe, and Ferster were able to isolate few simple mechanisms that establish contrast-invariance to the orientation selectivity of firing-rate responses. They show that these mechanisms are sufficient to explain the phenomenon, and they clarify the role of each of them.

To understand their results, let's first revisit the problem in some detail (Figures 1A–1C). The classic model by Hubel and Wiesel is based entirely on excitation. The ON and OFF subregions of a simple cell's receptive field come respectively from summing the outputs of appropriate ON-center and OFF-center LGN neurons (Figure 1A). This model ensures that the cell is more depolarized by stimuli of preferred orientation than by stimuli of other orientations (green in Figure 1B). However, it results in some depolarization at all orientations, including the orthogonal one. This baseline depolarization does not cause firing because it lies below threshold (Carandini and Ferster, 2000), and the result is a nicely selective firing-rate response (green in Figure 1C). Alas, the high threshold that one needs to achieve this selectivity is too high for responses to stimuli of low contrast. Such stimuli produce depolarizations that are smaller both in peak and in baseline (red in Figure 1B); they do not reach threshold, so the resulting firing rate is zero (red in Figure 1C).

It would seem, therefore, that the main problem that the neurons face in achieving contrast invariance is the baseline depolarization, which grows with contrast but is not selective for orientation. Earlier proposals for how contrast invariance is achieved, indeed, proposed that this baseline depolarization might be removed through intracortical inhibition. This inhibition could predominate at nonpreferred orientations (Ben-Yishai et al., 1995; Somers et al., 1995), be specific to the preferred orientation (Troyer et al., 1998, 2002), or be insensitive to orientation (Lauritzen and Miller, 2003).

Neuron Previews



Figure 1. Achieving Contrast Invariance in Visual Cortex

(A) The classical model of orientation selectivity. The ON and OFF subregions of the receptive field of a V1 simple cell are obtained by summing appropriately aligned LGN inputs (circles).
(B) The orientation selectivity of membrane potential in response to stimuli of 50% contrast (green) and 5% contrast (red). The dashed line indicates the firing threshold (V_{thresh}). V_{rest} is the resting potential.

(C) The orientation selectivity of the corresponding firing-rate responses.

(D) The gain of responses in the visual system decreases with contrast, emphasizing responses to low-contrast stimuli relative to high-contrast stimuli.

(E and F) As in (B) and (C), after incoming signals have been modified by contrast-gain control. (G) The membrane potential of V1 neurons fluctuates around the mean visually driven value (V_{mean}). This noise causes the potential to cross threshold, occasionally even $V_{mean} < V_{thresh}$. (H and I) As in (E) and (F), after the addition of noise. Noise has higher variance at low contrast than at high contrast. The resulting firing rates (I) are contrast invariant.

Finn, Priebe, and Ferster, therefore, started by measuring the baseline depolarization in their simple cells and found that it is largest in those neurons that receive most of their drive from thalamus. To assess the cortical contribution to the visual responses, they measured visual responses both in normal conditions and after the cortex was locally inactivated (Chung and Ferster, 1998); the difference between the two indicates the cortical contribution. In neurons in which this contribution was large, the subthreshold membrane potential was much more tuned than expected from summation of thalamic inputs. This is reassuring, because those neurons receive most of their inputs from the firing of other V1 neurons, which have already solved the iceberg problem. The real issue, therefore, is with the neurons that receive most of their inputs from thalamus. As predicted by the classical model, for these neurons the authors found that the tuning subthreshold showed a sizeable baseline depolarization. This result indicates that inhibition does not suppress the baseline depolarization as had been proposed.

Next, Finn et al. (2007) measured the effect of changing contrast, and in particular the impact of contrast-gain control (Figures 1D-1F). The gain of visual responses is not constant but rather decreases with increasing contrast (Figure 1D). This effect becomes stronger at each stage of the visual system from retina to extrastriate cortex and is quite developed in area V1 (see Carandini [2004b] for a review). A consequence of contrast-gain control is that multiplying contrast by 10 (as in Figure 1) increases the membrane potential responses by much less than a factor of 10. Therefore, the membrane potential responses to the high and low contrast are much more similar to each other (Figure 1E) than in the absence of gain control (Figure 1B). It now becomes possible to set a threshold that yields firing responses for both contrast levels (Figure 1F).

This is a major step in the right direction, but it is not yet an entire solution to the iceberg problem. Specifically, the tuning curves for firing rate (Figure 1F) suffer from two problems. First, the tuning curve at high contrast is wider than that at low contrast. Second, both curves show unrealistically sharp transitions as they emerge from the floor.

To find the solution to these two problems, Finn et al. (2007) took the final step and measured the impact of noise fluctuations in membrane potential (Figures 1G-1I). Previous work from the Ferster laboratory had demonstrated that these fluctuations smooth the relationship between the visually driven membrane potential and the resulting firing rate, helping to achieve the contrast-invariance of orientation selectivity (Anderson et al., 2000). Fluctuations in membrane potential are approximately Gaussian (Carandini, 2004a); if a stimulus drives the mean membrane potential to a value that is sufficiently close to threshold, the tail of this Gaussian will reach above threshold and cause spikes (Figure 1G). In these conditions, the relationship between the mean (visually driven) membrane potential

Neuron **Previews**

and firing rate becomes a power function (see Anderson et al. [2000] and references therein). Adding noise fluctuations to the visually driven tuning curves (Figure 1H) solved the two problems mentioned above: the firing-rate responses obtained at the two contrasts resemble each other in all but a scaling factor (Figure 1H).

Intriguingly, the noise that was measured in the membrane potential was not fixed in amplitude but rather decreased with stimulus contrast (compare green and red bands in Figure 1H). This is in fact necessary to achieve contrast invariance, because the enhancement in the responses to suboptimal orientations must be larger at low contrasts than at high contrasts (otherwise, once again the responses at high contrast will be broader than those at low contrast). This contrast dependence of membrane potential noise does not seem to have an immediately obvious explanation. Presumably the noise originates from the firing rates of afferent LGN and V1 neurons, and these might be expected to become more variable as contrast is increased, not less variable. Indeed, increasing contrast increases firing rate, and noise is generally thought to grow with firing rate. Moreover, noise in membrane potential is independent of the mean depolarization (Carandini, 2004a), so it comes as an intriguing surprise to see that it decreases with contrast.

In summary, the work of Finn, Priebe, and Ferster provides a compelling explanation for how simple cells solve the iceberg problem. This explanation centers on two mechanisms, neither of which knows anything about stimulus orientation. The first is contrast-gain control, which increases the responses of stimuli of low contrast relative to the responses to stimuli of high contrast (Figures 1D– 1F). The second is a power law in the relationship between visually driven (mean) potential and firing rate, which is achieved through a combination of noise in membrane potential and a hard threshold for firing (Figures 1G–1I).

These results are in excellent agreement with the proposals of Heeger in the early 1990s, which identified two key mechanisms in the operation of V1 neurons: contrast normalization, which is contrast-gain control implemented through a divisive term (Heeger, 1992b), and half-squaring, which is a threshold followed by a power law with an exponent of approximately two (Heeger, 1992a). At the time of these proposals, it was not known how neurons could achieve squaring or division. Squaring is now explained through a combination of noise and a threshold. Divisive contrast-gain control, instead, might not have a single explanation. It operates at all stages of the early visual system, and the component of it that is provided by V1 might rely on more than one mechanism: synaptic inhibition from neurons that are not tuned for orientation, or from a pool of neurons tuned for many orientations, and depression at the thalamocortical synapse (see Carandini [2004b] for a review). A possible additional mechanism is the one discovered by Finn et al. (2007): decreased noise at high contrasts, which makes the cells less responsive because their potential is less likely to cross threshold.

With this study, contrast invariance of orientation selectivity goes to join a number of other complex phenomena that are explained by simple mechanisms of gain control, noise, and threshold. If our hopes and expectations are correct and the cortex indeed makes use of a limited set of tools, then these mechanisms might provide a general explanation for how cortical neurons maintain their selectivity in the face of wide variations in the strength of their incoming signals.

REFERENCES

Anderson, J.S., Lampl, I., Gillespie, D.C., and Ferster, D. (2000). Science 290, 1968–1972.

Ben-Yishai, R., Lev Bar Or, R., and Sompolinsky, H. (1995). Proc. Natl. Acad. Sci. USA *92*, 3844–3848.

Carandini, M. (2004a). PLoS Biol. 2, e264. 10.1371/journal.pbio.0020264.

Carandini, M. (2004b). The Cognitive Neurosciences III, M.S. Gazzaniga, ed. (Cambridge, MA: MIT Press), pp. 313–326.

Carandini, M., and Ferster, D. (2000). J. Neurosci. 20, 470–484.

Chung, S., and Ferster, D. (1998). Neuron 20, 1177–1189.

Ferster, D., and Miller, K.D. (2000). Annu. Rev. Neurosci. 23, 441–471.

Finn, I.M., Priebe, N.J., and Ferster, D. (2007). Neuron *54*, this issue, 137–152.

Heeger, D.J. (1992a). Vis. Neurosci. 9, 427–443.

Heeger, D.J. (1992b). Vis. Neurosci. 9, 181–197.

Lauritzen, T.Z., and Miller, K.D. (2003). J. Neurosci. 23, 10201–10213.

Rose, D., and Blakemore, C. (1974). Nature 249, 375-377.

Skottun, B.C., Bradley, A., Sclar, G., Ohzawa, I., and Freeman, R. (1987). J. Neurophysiol. 57, 773–786.

Somers, D.C., Nelson, S.B., and Sur, M. (1995). J. Neurosci. *15*, 5448–5465.

Troyer, T.W., Krukowski, A.E., Priebe, N.J., and Miller, K.D. (1998). J. Neurosci. *18*, 5908–5927.

Troyer, T.W., Krukowski, A.E., and Miller, K.D. (2002). J. Neurophysiol. 87, 2741–2752.