

28 Neural Computation in Sensory Systems

XAQ PITKOW AND MARKUS MEISTER

ABSTRACT An animal's brain must extract from the onslaught of raw sense data the few bits of information that actually matter for the guidance of behavior. Accordingly, one finds neurons in the upper echelons of a sensory system that report very selectively on high-level features of the stimulus while remaining invariant to many low-level perturbations. For example, "face cells" in the primate visual cortex respond selectively to one person's face regardless of the view angle, scale, or illumination. The emergence of such complex pattern detectors is one of the stunning phenomena in sensory neuroscience. In studies that trace signal flow through the associated neural circuits, the most common tool is the measurement of sensory receptive fields. Here we review the tenets and basic results of receptive field analysis and place it in a common framework with the task of selective feature detection. We show that classic receptive field measurement—with its focus on the linear summation of stimulus variables—has little to contribute on the subject of feature computation. On the other hand, a cascade of simple nonlinear operations can indeed account for high-level pattern detector neurons, and hints of such organization are found in the brains of diverse animals.

Consider a human typist copying a handwritten manuscript. His brain converts visual signals that enter the eye into motor signals that drive the fingers. How much information is involved in this task? At the input, the typist's eye can receive about 10^9 bits per second of raw visual information.¹ At the output, the typist generates about 20 bits per second of movement information.² This modest number is in line with other estimates for the information rate of human behavior (Eriksson, 1996; Pierce & Karlin, 1957). So the task of the typist's sensory system is to extract, every second from a data deluge of 10^9 bits, the 20 bits that really matter for visually driven behavior. Of course, those few important bits are hidden in the raw sensory input in convoluted ways. The same applies to all other cases of sensory-driven behaviors. This process of highly selective extraction is appropriately called computation; indeed, many of the tasks that we hand to our man-made computing machines involve similar challenges of pattern extraction.

¹Our back of the envelope reads: Six million cone photoreceptors, each modulated at up to 25 Hz with a signal-to-noise ratio of 100.

²One hundred words per minute, six letters per word, two bits of entropy per letter in the English language.

How does the brain attack this problem? At the very input, it converts photons into electrical signals in nerve cells. At the output, it converts electrical signals of nerve cells into movement. In fact, all the world events with which we interact are represented by neuronal membrane potentials: sight, sound, smell, touch, speech, movement, sweating, and internal phenomena like thought, dreams, and emotions. This is a truly remarkable step of abstraction, by which phenomena of entirely different physical nature are encoded with the same symbol set of membrane voltages. From that point on the brain can use general purpose devices—namely, neurons and synapses—to create connections between these world events; for example, combining sights and sounds to cause thoughts and speech. We sometimes forget the immense power of this abstraction, but it offers another parallel to man-made electronic computers, which similarly represent all world events by voltage signals.

In this chapter, we consider how these sensory computations happen; namely, how the neurons and synapses are arranged to accomplish the selective extraction of relevant bits from the sensory stream. We start with some illustrative examples of selective computation. In tracing this process through the associated neural circuitry, the primary tool has been measurement of receptive fields, and we discuss and critique this approach. How complex neuronal-response properties come about through the actual signal flow in neural circuits presents a more daunting challenge, and we cite some examples of progress in this area. Finally, we consider several theories of brain function that purport to explain why sensory computation is organized the way it is. Throughout, we make liberal use of examples from the visual system, partly because of our own exposure bias, but the chapter focuses on principles thought to apply in the other senses as well.

Selectivity and invariance

Continuing our observation of the typist: note how his right thumb flicks downward only when he needs to type a space. So there is a motor neuron innervating his thumb muscle that has become a perfectly selective

visual pattern detector. It fires reliably and only in response to visual images that look like the space between two words. Of course, these images can take on many different forms, given the varieties of handwriting, types of paper and ink, and illumination. The motor neuron is invariant to this enormous variation of the raw sensory stimulus. After work, the typist may ride home on a bicycle and his right thumb is used to flick a small bell on the handlebar. Now that same motor neuron has become a selective detector of pedestrians. We draw two lessons from these simple observations. First, the brain can indeed construct circuits that make single neurons perfectly selective for a high-level abstract feature of the sensory input. Second, this mapping is flexible and changes dramatically when the brain engages a different task or context.

Given that the brain can produce motor neurons selective for high-level features, one suspects that it may do so already within the sensory system, namely, prior to the commitment to a particular motor output: once individual neurons have extracted a high-level concept, their firing could easily be mapped into different behavioral outputs, or associated with other concepts for the purpose of learning, memory retrieval, or thinking (Barlow, 1972, 2009; Quiroga, 2012). Indeed, this is the case. A compelling example can be found in the “face cells” of the macaque inferotemporal cortex (Freiwald & Tsao, 2010; Gross, 1992; Rolls, 1992; Tsao & Livingstone, 2008). When the animal is presented with a broad variety of visual images, these neurons remain silent, except on presentation of a face. Some neurons appear selective for the faces of specific individuals. Yet their response to that face is largely invariant under widely different view angles, distances, or illumination conditions that result during natural interactions. The macaque cortex has several small regions in which the great majority of neurons are specialized for faces, to the extent where these regions can even be resolved by functional MRI (fMRI) imaging (Tsao, Freiwald, Knutsen, Mandeville, & Tootell, 2003). In turn, fMRI of human subjects suggests that they have similar regions dominated by face-selective neurons (Kanwisher, McDermott, & Chun, 1997). Recordings from multimodal areas of the human brain have suggested specialist neurons with even broader invariances, such as the “Halle Berry cell” that responded selectively both to the name of the actress and her image (Quiroga, Reddy, Kreiman, Koch, & Fried, 2005).

Selective pattern detectors have been described in other species, often in the context of sensory tasks that are an essential part of an animal’s behavior. For example, owls can localize prey on the ground based on sound cues alone. The information about azimuthal

direction of the sound source comes from the time delay between sounds arriving at the two ears. And in the higher stations of the auditory system, one indeed finds neurons whose response is selective for a particular interaural time delay, but invariant to many other features of the sound waveform (Carr & Konishi, 1990; Konishi, 2003). Another instance arises in weakly electric fish that sense their surroundings by producing an oscillating electric field in the water. When two fish with similar oscillation frequencies approach too close, their signals threaten to interfere, and they shift the two frequencies further apart. Among the many stations of the electrosensory nervous system, one finds neurons whose response reports selectively the frequency difference of the other fish, invariant to the location or orientation of the fish or even the absolute field frequencies (Heiligenberg, 1991; Rose, Kawasaki, & Heiligenberg, 1988). These two examples are instructive, because much has been discovered about the circuits that generate the selectivity (Konishi, 2006).

We will use these high-level pattern-detector neurons as a guide to understanding sensory computation. Obviously, this is not the only purpose of a sensory system. But in exploring how a system works, it is useful to have some concrete phenomenon in mind whose explanation is likely to reveal something fundamental. Face cells—and their analogs in other species—impress us by the combination of selectivity and invariance. Their response is selective for faces, but not just one picture of a face; that could easily be accomplished by a template match to the target image. Instead, they respond equally to images of that same face rotated, translated, and illuminated in all different ways. How does that arise?

First of all, one must recognize that selectivity and invariance per se are not unusual: every neuron has them. A photoreceptor is selective for light that falls on one point in the image, and invariant to the pattern of light on all the other photoreceptors. Further into the sensory system, each neuron receives a synaptic input current that represents a single scalar function of the sensory stimulus, and that defines its selectivity. The neuron’s response is invariant under all the myriad stimuli that leave this scalar function unchanged. So one is led to a more subtle assessment: the “face cell” is remarkable because its selectivity and invariance have been shaped so exquisitely to match a specific region in the space of all images that corresponds to different views of the grandmother’s face. Furthermore, its response is invariant under image transformations that profoundly alter the signals of each sensory receptor and most neurons in the early visual system. To

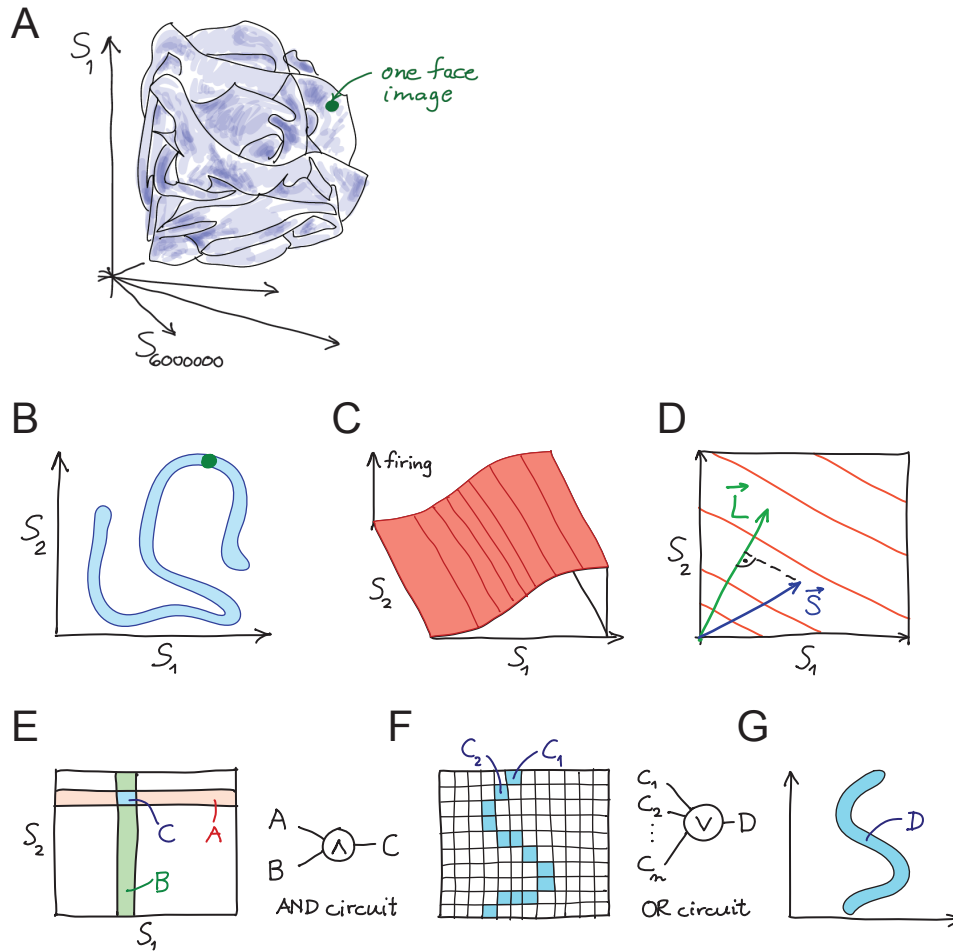


FIGURE 28.1 A geometric view of sensory computations. (A) Each stimulus is a vector in a high-dimensional space, whose axes are defined by the excitation of all the sensory receptors, for example, the intensity on the 6 million cones in a human eye. One retinal image of a face is a point in that space. The same face seen under different views and illuminations defines a high-dimensional surface (DiCarlo & Cox, 2007). (B) Two-dimensional stimulus space. A “face cell” responds only when the stimulus is in a thin, convoluted, shoelace-shaped region. (C) The response of a receptor as a

function of the stimulus variables. Note the contour lines are straight and orthogonal to the receptor’s axis of sensitivity (here the S_1 axis). (D) A receptive field is represented by a vector L in stimulus space. Integration of the stimulus S by this receptive field is the projection of S onto L . The response of a linear-nonlinear (LN) neuron depends only on this projection: the contour lines of its response function are straight and orthogonal to L . (E–G) How to create a shoelace-shaped response region in three simple steps. See text for details. (See color plate 24.)

appreciate this, it is helpful to think about the space of stimuli in geometric terms.

Your eye has about 6 million cone photoreceptors. For daylight vision, this is the number of pixels of the raw visual stimulus. Suppose the image of a face appears on the retina: like all other retinal images, it can be represented as a point in an abstract space that has 6 million dimensions (DiCarlo & Cox, 2007). Each coordinate of that space represents the intensity on one photoreceptor (figure 28.1A). Now rotate the face about one axis: this produces a series of stimulus points that trace out a curved line in the stimulus space. Rotate it also about another axis: now one has a two-dimensional surface. Add several other transformations,

like scaling, translation, illumination by a point source at various angles, background illumination, diverse facial expressions. Together these define a high-dimensional and highly convoluted surface in the stimulus space (figure 28.1A). The face cell fires only if the stimulus is on or near that surface.

For those few readers with difficulty imagining a 6-million-dimensional space, let us consider a rudimentary sensory system with just two receptors. The space of all stimuli is a two-dimensional plane (figure 28.1B). Any given image is a point on that surface. The activity of those two receptors can be represented with two response functions on that surface, whose contour lines run straight and parallel to the axes (figure 28.1C). But

a hypothetical “face cell” fires only in a thin stimulus region that looks like a convoluted shoelace (figure 28.1B). The challenge for the brain (or more aptly for us trying to understand the brain) is how the shoelace-shaped selectivity region can be built from the straight-line selectivities of the receptors. Of course, in the 6-million-dimensional space of real face detection, this problem is greatly magnified.

Receptive fields

How does one go about exploring these sensory circuits? The primary approach has been to record the signals of single neurons along the pathway from the primary receptors to the high-level feature detectors. At each stage, one measures the properties of neuronal responses to sensory stimuli, in the hope that they will gradually approach the selectivity and invariance of the “face cell.” In characterizing the function of a sensory neuron, one would ideally like to report its response to any relevant sensory stimulus. Given limited experimental time, one can test only a small number of stimuli, and then the results must be extrapolated in a way that extends to stimuli outside that limited set.

The most common form of reporting such a summary of response properties is the “receptive field.” The meaning of this term has undergone some evolution over the years, adapting to increased needs and powers of sensory physiology. Originally applied to entire reflex arcs (Sherrington, 1906), the receptive field meant the area of a dog’s skin in which touch would trigger a scratch reflex. Later it was applied more specifically to characterize neural responses, but generalized to other sensory surfaces like the retina (figure 28.2A–C), and even to abstract surfaces like an auditory spectrogram (figure 28.2E). Then the concept was extended to allow for both excitation and inhibition, meaning that area of the sensory field in which point stimuli directly increase or decrease the neuron’s firing rate. Finally, in today’s usage, the receptive field also conveys quantitative information about how much the firing rate is altered and with what time course.

In the visual system, for example, the spatiotemporal receptive field $L(x, y, t)$ is a function of space and time (sometimes also wavelength) that spells out how much an impulse of light at location (x, y) in the visual field alters the firing a time t later. An arbitrary stimulus can be constructed from a superposition of short flashes of light at different locations and times (Rodieck & Stone, 1965). Thus a neuron’s response summary should include a description of how these many point-like stimuli at different locations and times combine to generate the response of the neuron. When an author

quotes just the receptive field, the best (only) guess for such integration is a linear summation: at any given time, the system sums the intensity of the point-like stimuli in the recent past, weighted by the amplitude of the receptive field. Formally, this linear summation is expressed as

$$g(t) = \iiint_{x, y, t'=-\infty}^t S(x, y, t-t')L(x, y, t')dt'dydx, \quad (1)$$

where $S(x, y, t)$ represents the stimulus at location (x, y) and time t and $L(x, y, t)$ is the receptive field. In practice, one always deals with a finite number of discrete locations \bar{x}_i and times discretized to multiples of some small interval Δt . So the stimulus can also be represented as a vector that lists the intensities at those points in the recent past,

$$\mathbf{S} = \{S_{ij}\} = \{S(\bar{x}_i, t - j\Delta t)\}, \quad (2)$$

and similarly the receptive field is a vector

$$\mathbf{L} = \{L_{ij}\} = \{L(\bar{x}_i, j\Delta t)\}. \quad (3)$$

In this vector representation, the linear summation of stimuli becomes a simple dot product

$$g = \mathbf{S} \bullet \mathbf{L} \quad (4)$$

The effect of any given stimulus is determined entirely by the projection of the stimulus vector onto the special direction indicated by the receptive field vector \mathbf{L} . In this way the receptive field of a sensory neuron—at least in this modern incarnation—identifies its axis of selectivity in the abstract space of stimuli. The neuron’s response varies only along that axis, and remains invariant along all orthogonal directions (figure 28.1D).

How does the response vary along the special direction? Whereas the above expression for g can take on any value, a neuron’s response is generally limited to a finite range. For example, firing rates cannot be negative or exceed a biophysical limit from membrane refractoriness. Thus a more realistic model of sensory responses includes a nonlinear distortion function by which g turns into the neuron’s response:

$$r = N(g) = N(\mathbf{S} \bullet \mathbf{L}). \quad (5)$$

Generally, $N(g)$ has a sigmoid shape, with zero or constant response for very small values of g and saturation to a maximal response at very large values. This expression for a neuron’s response to stimuli is often called the “LN model”: a linear summation of stimuli followed by a nonlinear response function. One can measure the nonlinearity N during the same experiment that measures the receptive field L (Chichilnisky, 2001; Ringach, 2004). A complete report that includes both receptive field and nonlinearity allows the reader to predict the neuron’s response to any desired stimulus.

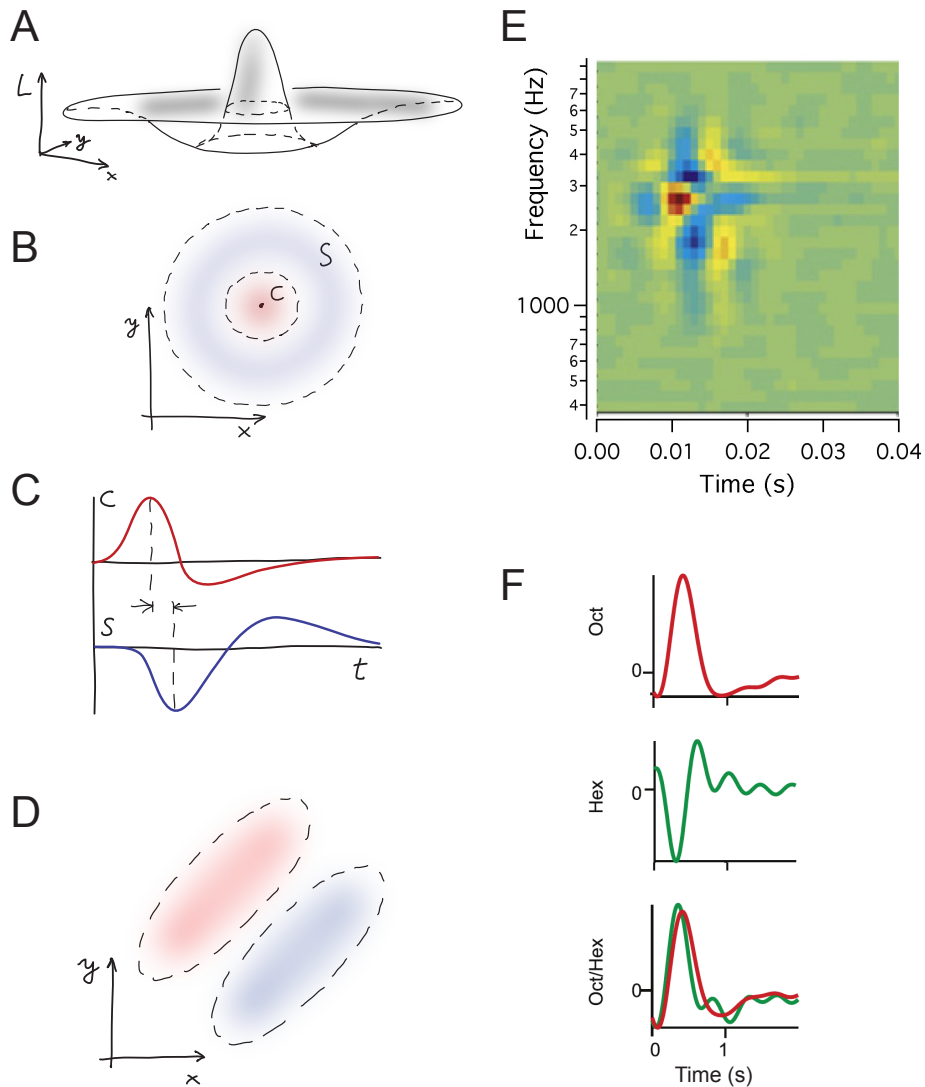


FIGURE 28.2 Receptive fields in various sensory systems. (A–C) Spatiotemporal receptive fields in the early visual system. Here the components of the stimulus vector \mathbf{S} are the light intensity at different locations and time points. Any given stimulus vector consists of a short “movie” $S(x, y, t)$. The receptive field $L(x, y, t)$ can similarly be viewed as a movie. For displays on paper, one often shows a single frame $L(x, y, t_0)$ of this movie (panels A–B) along with the time course $L(x_0, y_0, t)$ at a particular location (panel C). Note that L has both positive and negative components, meaning that light is excitatory (red) at some locations and times, and inhibitory (blue) at others. For retinal bipolar cells, ganglion cells, and thalamic relay cells, the receptive fields look very similar except for some scaling along the spatial and temporal axes. The spatial profile (B) contains a center region (C) and a surround region (S) in which light has opposite action. When displayed in three dimensions (A), this profile has the appearance of a “Mexican hat.” The time course of the receptive field (C) is biphasic, both in the center and the surround. Generally, the surround response is slightly delayed relative to the center (C). (D) For a simple cell in primary visual cortex, the receptive field profile typically shows elongated side-by-side regions in which light has opposite action. The

time course again is biphasic, as in panel C. (E) Spectrotemporal receptive field of a neuron in the auditory area field L of the zebra finch (Nagel & Doupe, 2008). Here the stimulus $S(f, t)$ is represented as a spectrogram, plotting sound power as a function of time and frequency. The receptive field of this cell shows excitatory (red) and inhibitory (blue) regions. Along the spectral dimension one finds a “Mexican hat” profile, with a central region of excitatory frequencies and adjacent inhibitory regions. The action of the surround is delayed relative to the center. In the time domain, the receptive field is strongly biphasic or even triphasic. This is only one of many receptive field shapes encountered in this brain region. (F) Osmotemporal receptive field of a neuron in the olfactory pathway of the locust (Geffen et al., 2009). Here the stimulus $S(o, t)$ is the time-varying concentration of two discrete odors, octanone or hexanal. The receptive field $L(o, t)$ specifies the contribution to the response of each odor at different times. *Top and middle:* Receptive field derived from experiments in which only one odor was presented at a time. *Bottom:* Receptive field measured when both odors were varied simultaneously. Note the neuron’s response to hexanal changes polarity under this condition. (See color plate 25.)

Figure 28.2A–D presents typical spatiotemporal receptive fields encountered along the mammalian visual pathway. Beginning in retinal bipolar cells, the spatial receptive field profile shows what is called “center-surround antagonism” (Asari & Meister, 2012; Burkhardt, Bartoletti, & Thoreson, 2011; Hare & Owen, 1990). This means that light in a small central region near the cell body has opposite effects from light in a larger surrounding region (figure 28.2A–B). The temporal part of the receptive field shows antagonism in time: light in the recent past has opposite effects from light in the more distant past (Baccus, Ölveczky, Manu, & Meister, 2008). This is reflected in the biphasic time course of the receptive field (figure 28.2C). Generally this time course is slightly different in the center and the surround, with a notable delay in the response to the surround (Baccus et al., 2008; Fahey & Burkhardt, 2003). These same features dominate the receptive fields reported for downstream neurons, like retinal ganglion cells (Benardete & Kaplan, 1997; Enroth-Cugell, Robson, Schweitzer-Tong, & Watson, 1983; Kuffler, 1953; Meister & Berry, 1999) and thalamic relay neurons (Hubel & Wiesel, 1961; Reid & Shapley, 2002). One synapse further, in the primary visual cortex, the receptive fields begin to change (DeAngelis, Ohzawa, & Freeman, 1995; Hubel & Wiesel, 1962; Martinez et al., 2005; Ringach, 2004): for so-called simple cells, the canonical shape consists of elongated side-by-side regions of opposite polarity (figure 28.2D). The temporal component of the receptive field again has a biphasic time course (figure 28.2D). The same receptive field formalism has been used for other sensory modalities, and we illustrate examples from the auditory (figure 28.2E) and olfactory systems (figure 28.2F). Note that in all these cases the dynamics look biphasic. Fundamentally, this means that the neurons emphasize changes in the sensory input.³

Why does the brain use neurons with these particular receptive fields? Attempts to explain the receptive

field shapes have been particularly effective in the early visual system. One theory postulates that the retina is concerned primarily with transmitting visual information efficiently through the bottleneck of the optic nerve (Atick & Redlich, 1992; Doi et al., 2012; Graham, Chandler, & Field, 2006; Srinivasan, Laughlin, & Dubs, 1982; van Hateren, 1992). In particular, this means avoiding redundancy among the signals of different retinal ganglion cells. The images from the natural world contain a great deal of correlation, because nearby points in space and time tend to have very similar intensity (Field, 1987). In this view, the retinal circuitry filters the incoming movies in space and time to remove those correlations. Formalizing this principle leads to a theory that predicts image filters with center-surround antagonism in space and a biphasic time course, bearing close resemblance to actual receptive fields measured in the retina. Furthermore, the same theory appears to explain a list of performance features of the early visual system (van Hateren, 1993). This has led to widespread acceptance of the efficient coding idea, although some skeptics remain, including its earliest proponent (Barlow, 2001). For example, it appears that lateral inhibition in the early visual system performs much less redundancy reduction than should be possible (Pitkow & Meister, 2012), and there are proposals that it serves a different purpose (Balboa & Grzywacz, 2000).

Deeper into the visual system, once the coding constraints of the optic nerve have fallen away, receptive fields have been interpreted under a different perspective. Here it has been suggested that the brain seeks to represent the stimulus using a large population of neurons that are each active only rarely and are statistically independent of each other (Olshausen & Field, 2004). Such a code can help to highlight “suspicious coincidences” (Barlow, 1994) of firings in the population, which are the hallmarks of higher-level stimulus structures like objects. When analyzing natural visual scenes by this criterion, one predicts receptive fields that respond selectively to edges, similar to those seen to excite neurons in the early visual cortex (figure 28.2D; Bell & Sejnowski, 1997; Olshausen & Field, 1996). A similar correspondence has been noted in the auditory system (Smith & Lewicki, 2006). On the one hand, it is quite remarkable that an *ab initio* theory based only on concepts of signal coding and the statistics of natural stimuli can predict with some success the sensory responses in living brains. On the other hand, the particular response features that it explains, namely receptive fields, do not tell the whole story of visual processing.

³Incidentally, the biphasic time course of receptive fields extends even to other phylogenetic kingdoms. The dynamics of human photoreceptors (Schnapf, Nunn, Meister, & Baylor, 1990), bacterial chemosensory responses (Block, Segall, & Berg, 1982), and light responses of fungi (Lipson, 1975) all follow the same time course, except for a scaling of the time axis. It is tempting to postulate this as a universal law, presumably related to the need of all living things to detect changes in their environment. Someday, life will be discovered on other planets. While those creatures may or may not be based on carbon chemistry, it is a safe bet that they have biphasic sensory responses.

Limits of receptive field analysis

For all its sophistication, the method of receptive field measurement—even in its modern embodiment by the LN model—ultimately delivers only partial insights into sensory computations. We focus here on four commonly noted limitations.

First, the LN model for a given neuron often accounts for only a portion of its neural response. This can be tested directly by comparing the firing rate predicted from the receptive field and nonlinearity (Eq. 5) with the actual firing rate. Often this prediction works well when one uses simple artificial stimuli, with limited or tightly defined spatial and temporal structure (Chichilnisky, 2001; Keat, Reinagel, Reid, & Meister, 2001). But when natural stimuli are used as inputs, the simple receptive field models often fail dramatically (David, Mesgarani, Fritz, & Shamma, 2009; Machens, Wehr, & Zador, 2004; Mante, Bonin, & Carandini, 2008; Sharpee, Miller, & Stryker, 2008). Such complex stimuli seem to engage pathways that go beyond linear summation. As an example, consider the “object motion sensitive” ganglion cells found in the vertebrate retina (Baccus et al., 2008; Ölveczky et al., 2003; van Wyk, Taylor, & Vaney, 2006; Zhang, Kim, Sanes, & Meister, 2012). These neurons fire selectively when an object moves with a trajectory different from that of the background. This response is largely invariant to the pattern on the object. But the cells remain silent when the entire retinal image moves in concert. This complex response function is of obvious utility for detecting prey or predator within the visual scene. But it cannot be formulated in terms of an LN model.

Second, receptive fields are fickle and can change on a moment’s notice. One finds very commonly that a simple change in the stimulus environment—going from dim to bright lights, or from soft to loud sounds, or from one to two odors—elicits substantial changes in the receptive field (Bair, 2005; Geffen, Broome, Laurent, & Meister, 2009; Geffen, de Vries, & Meister, 2007; Gilbert & Li, 2013; Mante et al., 2008; Nagel & Doupe, 2006; Theunissen, Sen, & Doupe, 2000; figure 28.2F). The LN model may well produce acceptable fits to the response under one stimulus environment, but changing the environment leads to entirely different results for L and N . Sometimes these changes can be interpreted as “adaptation” to the new environment (Hosoya, Baccus, & Meister, 2005; Shapley & Enroth-Cugell, 1984; Wark, Lundstrom, & Fairhall, 2007), but they are a clear indicator that the actual computation performed by the system is more intricate than the LN model. If one knew the true nature of that computation, there would be no need to invoke “adaptation”

(Borst, Flanagin, & Sompolinsky, 2005; Garvert & Gollisch, 2013; Ozuysal & Baccus, 2012).

Third, the sensory computation expressed by the LN model is rather primitive. It seems implausible a priori that neurons many synapses into the brain are still concerned with a simple weighted summation of stimuli. For example, the receptive fields quoted for neurons along the mammalian visual pathway are remarkably similar. There is virtually no change from retinal bipolar cells to thalamic relay neurons: they all have spatial profiles shaped like “Mexican hats” with opposite effects from light in the center and surround, and a biphasic temporal profile (figure 28.2A–C). Some change in the spatial profile appears at the level of primary visual cortex, where the antagonistic regions tend to lie side by side rather than concentrically (figure 28.2D). This reflects a rather modest computational accomplishment, given the enormous amount of neural machinery that is engaged, including no fewer than 70 different neuron types in the retina alone. One gets the suspicion that these receptive field measurements have very little to do with what these circuits really accomplish. In the vertebrate retina, we now know that many of the approximately 20 types of retinal ganglion cell perform quite sophisticated image computations (Field & Chichilnisky, 2007; Gollisch & Meister, 2010), which become apparent under more ecologically relevant stimuli. Yet when probed with the white-noise flicker stimulus that is popular for receptive field measurements, these RGCs will always hide their true personality under a modest sombrero.

Finally, and perhaps most fundamentally, receptive field analysis of this type simply does not help much in understanding interesting computations, such as those leading to face cells. This is because the LN neuron has a trivially simple invariance. Its response varies along a single direction in stimulus space (the vector L in Eq. 5), and it remains invariant to all the directions orthogonal to it. This basic limitation is independent of the form of the nonlinearity N . In our reduced two-receptor system (figure 28.1D), the response function of any LN neuron has straight and parallel contour lines. Given such a straight-laced neuron, we are no closer to constructing the curly response space of a face cell (figure 28.1B) than we were with the original sensory receptors. At best, such a neuron can offer a local tangent to the desired response region. But within the receptive field formalism, one cannot explain why the region should bend around from one tangent to the next. So from the standpoint of stimulus geometry, the receptive field of any LN neuron has the same complexity as that of a sensory receptor, and is no more helpful in explaining the advanced stimulus

selectivities one encounters at the higher echelons of sensory systems.

Beyond receptive fields

Experimentally, there have long been indications that the receptive field alone cannot account for sensory neuron responses, even early in the sensory circuits. For example, stimuli that on their own do not elicit a response—and thus fall outside the receptive field—can often modulate the responses in powerful ways. When a simple cell in primary visual cortex is probed with short line segments as stimuli, it has a compact receptive field with elongated response regions (figure 28.2D). But if one adds another line segment outside this receptive field, the response changes dramatically depending on whether the segments are collinear or not (Kapadia, Ito, Gilbert, & Westheimer, 1995). Similarly, the response of a retinal ganglion cell can be powerfully suppressed by stimulus movement in distant regions on the retina, far outside the receptive field (Baccus et al., 2008; Roska & Werblin, 2003). These and similar response components have been described as the “extraclassical” receptive field, or “surround effects,” or “contextual influences” (Fitzpatrick, 2000). Presumably these effects contribute to bending a neuron’s response surface. This gets us closer to explaining sensory computation, and indeed for that goal the extraclassical effects are more important than the neuron’s receptive field alone.

Given the importance of characterizing a neuron’s response beyond a single receptive field, how can those additional components be treated in a formal analysis?

1. One approach is to experiment in a stimulus space of reduced dimensionality, for example, with just two independent stimulus components (figure 28.1B). Within that space, one can then map out a neuron’s response region in complete detail (Bolinger & Gollisch, 2012; Gollisch, Schutze, Benda, & Herz, 2002). A recent study on retinal ganglion cells showed that their response regions are often highly curved, and that different cell types produce opposite curvatures in the same part of stimulus space (Bolinger & Gollisch, 2012). These could serve as primitive parts for constructing a more complex response region.

2. Alternatively, one can maintain the full dimensionality of the stimulus space, but ask whether there exists a subset of these directions \mathbf{L}_i that can affect the neuron’s response. A formal framework for such responses is the multi-LN model:

$$r = N(g_1, \dots, g_n) = N(\mathbf{S} \cdot \mathbf{L}_1, \dots, \mathbf{S} \cdot \mathbf{L}_n). \quad (6)$$

Here the stimulus \mathbf{S} is projected onto n different vectors \mathbf{L}_i , producing n scalars g_i . Then the response is computed as a nonlinear function of all these variables (figure 28.3A). There are principled methods for discovering the special vectors \mathbf{L}_i (Marmarelis & Orme, 1993; Schwartz, Pillow, Rust, & Simoncelli, 2006), and indeed many sensory neurons probed this way have more than one special direction in stimulus space (Fairhall et al., 2006; Maravall, Petersen, Fairhall, Arabzadeh, & Diamond, 2007; Rust, Schwartz, Movshon, & Simoncelli, 2005; Slee, Higgs, Fairhall, & Spain, 2005; Touryan, Lau, & Dan, 2002). Already the simple process of spike generation, which turns a membrane current into action potentials, involves two special directions (Agüera y Arcas, Fairhall, & Bialek, 2003). These methods can serve to identify a relevant subspace in which the neuron’s computation seems to take place. However, mapping out the nonlinearity $N(\dots)$ in the response function is a challenge when it depends on more than two variables. Moreover, if it is simple enough to characterize this way, the mathematical form of this response function is still rather restrictive, and implausible a priori for neurons deep into a sensory system.

3. A third approach is to model the neural system as a cascade of LN stages leading up to the sensory neuron in question (French & Korenberg, 1989; Mante et al., 2008; Shapley & Victor, 1981; van Hateren, Ruttiger, Sun, & Lee, 2002). The structure of such a model can be inspired by what is known about the anatomy of the circuits. For example, in modeling retinal ganglion cell responses, one would include circuit elements like bipolar cells and amacrine cells with their known synaptic relationships (Baccus et al., 2008; Chen et al., 2013; Gollisch & Meister, 2008; Greschner, Thiel, Kretzberg, & Ammermüller, 2006; figure 28.3B). Each circuit element performs a simple LN operation on its inputs. Yet a cascade of such simple units—including divergence, convergence, and feedback of signals—can in principle compute arbitrarily complex functions of the stimulus (Cybenko, 1989). This modeling approach has been successful in capturing many “extraclassical” response features (Gollisch & Meister, 2010). It also offers a useful linkage to the actual neural circuits that are the biophysical substrate of all these computations.

Expanded nonlinear representations

With these insights, we can now return to the simple two-receptor system considered above, and ask: how could one construct a neuron with a curved and convoluted response region? One very simple approach is to first construct a population of neurons that have small, local, almost point-like response regions in

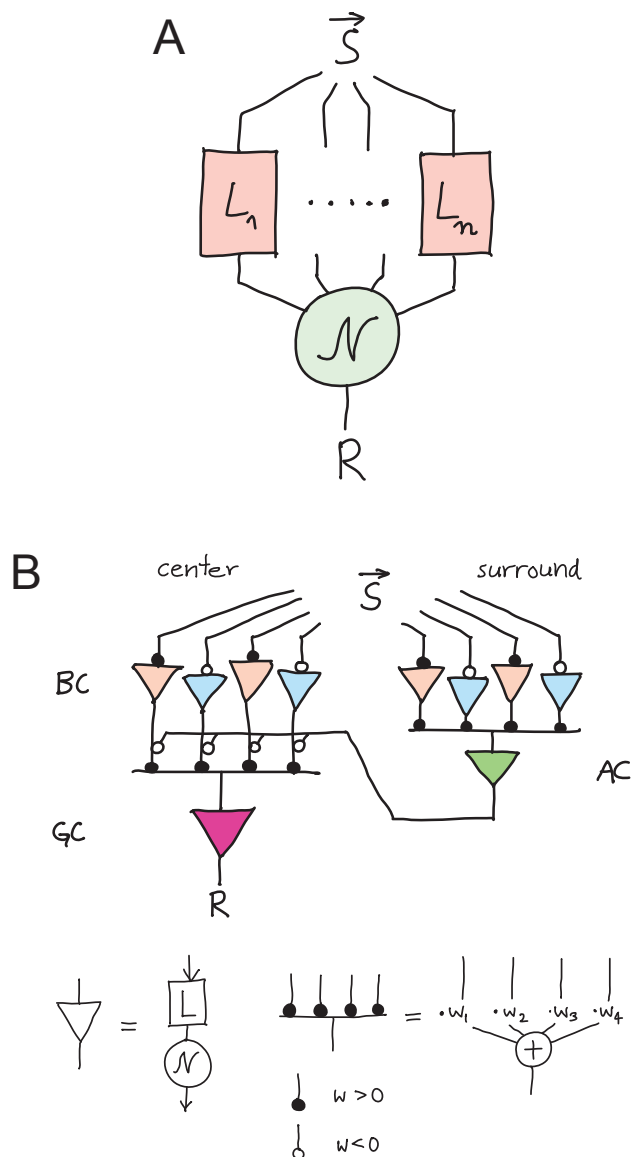


FIGURE 28.3 Response models beyond the receptive field. (A) Multi-LN model. Here the stimulus, S , is processed linearly through multiple parallel filters. Their outputs are then combined in a single static nonlinearity, yielding the response, R . (B) Circuit model with a cascade of LN stages that may be arranged in series or parallel, including feedforward and feedback pathways. This example takes inspiration from retinal circuitry, and each participating neuron type (On and Off bipolar cells, BC; amacrine cells, AC; ganglion cells, GC) is represented by a simple LN model: a weighted summation of inputs (weights w_i) followed with a temporal filter and a nonlinear response function. With appropriate choice of those parameters, the circuit makes quantitatively accurate predictions for the response of the so-called object motion sensitive ganglion cells that sense differential motion between foreground and background (Ólveczky et al., 2003). (See color plate 26.)

stimulus space. Then one can combine a suitable collection of these to create a neuron with an arbitrary response region.

In principle, this can be achieved in three steps (figure 28.1E–G):

1. For each of the two stimulus axes, create n neurons that respond to only a small range along that axis (e.g., neurons A and B in figure 28.1E).

2. Combine such neurons from each axis with a logical AND to create a large population of n^2 neurons that each respond in a tiny region of the stimulus space (neurons C_i in figure 28.1F; for example, neuron C_1 responds only when stimulus intensity S_1 is medium and intensity S_2 is maximal).

3. Combine many of these by a logical OR to create a neuron with a thin and convoluted response region (neuron D in figure 28.1G).

Note that the neurons from step 1 can still be described by an LN model—though with an unconventional hump-shaped nonlinearity—and thus have a defined receptive field, whereas this is no longer the case for the neurons in step 2. A key ingredient of this scheme is the dramatic expansion of the neural population. Starting with just two receptor neurons, after step 2 the stimuli are represented by n^2 neurons. This is followed by an equally dramatic contraction in step 3, which yields a neuron with complex response function.

Indeed, one can find such extreme expansions and contractions in several sensory systems. For example, in the retina, one cone photoreceptor connects to ~ 10 bipolar cells (Wässle, Puller, Muller, & Haverkamp, 2009), and each bipolar cell in turn has many synaptic terminals. Each of these terminals may receive a different set of amacrine cell inputs, resulting in different response properties (Asari & Meister, 2012; Baden, Berens, Bethge, & Euler, 2013). Finally, a retinal ganglion cell pools over a select subset of all these bipolar cell terminals to construct its specific response function (Gollisch & Meister, 2010). From the retina to primary visual cortex, one finds another expansion in the neuron number by a factor of ~ 100 . At the same time, the activity in these cortical populations becomes more and more sparse (Barth & Poulet, 2012; Hromadka & Zador, 2009; Isaacson, 2010; Olshausen & Field, 2004; Wolfe, Houweling, & Brecht, 2010): whereas natural stimulation will drive early sensory neurons to fire much of the time, neurons in the cortex are active more rarely, with an average spike rate estimated at only one to three spikes per second (Attwell & Laughlin, 2001; Lennie, 2003). Of course, each cortical neuron pools several thousand of such inputs to construct its own response space. A similar expansion into a sparsely

active population has been demonstrated in the insect olfactory system: From the second-order to third-order neurons in this circuit, the number of distinct response types in the population increases, and the activity within each type decreases dramatically (Laurent, 2002).

As illustrated above, the purpose of such a nonlinear expansion may be to allow for construction of complex response functions in the subsequent pooling stage. Interestingly, this does not require a careful design of response functions in the expanded representation. Even random combinations of features are sufficient to construct very complex response spaces (Rigotti, Ben Dayan Rubin, Wang, & Fusi, 2010). However, the nonlinearity is essential: only by bending the input into higher dimensions can linear projections produce a different outcome than on the original space. Surprisingly, only one layer of nonlinear neurons is required to generate any arbitrary function (Cybenko, 1989), although this may require a huge expansion and delicate cancellations in the subsequent convergence. A much better use of neural resources appears to be a series of nonlinear expansions and linear projections, an architecture that is generally described as a “deep network” (Bengio, Courville, & Vincent, 2013).

This kind of architecture forms the basis of an influential model of sensory processing leading to object detection (Mel, 1997; Riesenhuber & Poggio, 1999; Serre, Oliva, & Poggio, 2007). According to this model, the brain constructs complex percepts by a sequence of linear and nonlinear operations. The linear operations reduce the dimensionality of the input to encode only the presence of particular patterns, generating selectivity. The nonlinear operations compute the maximum over similar patterns that differ only by a transformation like shifting or rotation, thereby generating invariance to that transformation. A cascade of these operations can develop sensitivity to complex combinations of object features, while retaining invariance to changes in viewpoint, scale, or illumination (see chapter 30 by Nicole Rust, this volume).

Thus we have arrived at a plausible model to explain the puzzling phenomenon of “face cells” that motivated our foray, though it should be said that the computational models don’t yet match human performance on such tasks (DiCarlo, Zoccolan, & Rust, 2012). Along the way we learned that the measurement of receptive fields—perhaps the single most common activity in sensory neuroscience—is of limited use for understanding how sensory computations arise. It needs to be extended with more intricate and flexible models of neural responses. Ultimately, the most effective modeling strategy will be to actually understand the neural circuits that underlie the responses: the arrangement

of physiological cell types, their connectivity, synaptic integration, and the patterns of signal flow. This parallel pursuit of structural and functional information has greatly helped in understanding peripheral sensory circuits, like retina and olfactory bulb, and will ultimately be essential to cracking the mysteries of cortical circuits as well.

REFERENCES

- AGÜERA Y ARCAS, B., FAIRHALL, A. L., & BIALEK, W. (2003). Computation in a single neuron: Hodgkin and Huxley revisited. *Neural Comput*, *15*, 1715–1749.
- ASARI, H., & MEISTER, M. (2012). Divergence of visual channels in the inner retina. *Nat Neurosci*, *15*, 1581–1589.
- ATICK, J. J., & REDLICH, A. N. (1992). What does the retina know about natural scenes? *Neural Comput*, *4*, 196–210.
- ATTWELL, D., & LAUGHLIN, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab*, *21*, 1133–1145.
- BACCUS, S. A., ÖLVEÇZY, B. P., MANU, M., & MEISTER, M. (2008). A retinal circuit that computes object motion. *J Neurosci*, *28*, 6807–6817.
- BADEN, T., BERENS, P., BETHGE, M., & EULER, T. (2013). Spikes in mammalian bipolar cells support temporal layering of the inner retina. *Curr Biol*, *23*, 48–52.
- BAIR, W. (2005). Visual receptive field organization. *Curr Opin Neurobiol*, *15*, 459–464.
- BALBOA, R. M., & GRZYWACZ, N. M. (2000). The role of early retinal lateral inhibition: More than maximizing luminance information. *Vis Neurosci*, *17*, 77–89.
- BARLOW, H. B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, *1*, 371–394.
- BARLOW, H. B. (1994). What is the computational goal of the neocortex? In C. Koch & J. L. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 1–22). Cambridge, MA: MIT Press.
- BARLOW, H. B. (2001). Redundancy reduction revisited. *Network*, *12*, 241–253.
- BARLOW, H. B. (2009). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, *38*, 795–798.
- BARTH, A. L., & POULET, J. F. (2012). Experimental evidence for sparse firing in the neocortex. *Trends Neurosci*, *35*, 345–355.
- BELL, A. J., & SEJNOWSKI, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vis Res*, *37*, 3327–3338.
- BENARDETE, E. A., & KAPLAN, E. (1997). The receptive field of the primate P retinal ganglion cell, I: Linear dynamics. *Vis Neurosci*, *14*, 169–185.
- BENGIO, Y., COURVILLE, A., & VINCENT, P. (2013). Representation learning: A review and new perspectives. *IEEE Trans Pattern Anal Mach Intell*, *35*, 1798–1828.
- BLOCK, S. M., SEGALL, J. E., & BERG, H. C. (1982). Impulse responses in bacterial chemotaxis. *Cell*, *31*, 215–226.
- BOLINGER, D., & GOLLISCH, T. (2012). Closed-loop measurements of iso-response stimuli reveal dynamic nonlinear stimulus integration in the retina. *Neuron*, *73*, 333–346.
- BORST, A., FLANAGIN, V. L., & SOMPOLINSKY, H. (2005). Adaptation without parameter change: Dynamic gain

- control in motion detection. *Proc Natl Acad Sci USA*, *102*, 6172–6176.
- BURKHARDT, D. A., BARTOLETTI, T. M., & THORESON, W. B. (2011). Center/surround organization of retinal bipolar cells: High correlation of fundamental responses of center and surround to sinusoidal contrasts. *Vis Neurosci*, *28*, 183–192.
- CARR, C. E., & KONISHI, M. (1990). A circuit for detection of interaural time differences in the brain stem of the barn owl. *J Neurosci*, *10*, 3227–3246.
- CHEN, E. Y., MARRE, O., FISHER, C., SCHWARTZ, G., LEVY, J., DA SILVEIRA, R. A., & BERRY, M. J. 2ND. (2013). Alert response to motion onset in the retina. *J Neurosci*, *33*, 120–132.
- CHICHILNISKY, E. J. (2001). A simple white noise analysis of neuronal light responses. *Network*, *12*, 199–213.
- CYBENKO, G. (1989). Approximation by superpositions of a sigmoidal function. *Math Control Signal*, *2*, 303–314.
- DAVID, S. V., MESGARANI, N., FRITZ, J. B., & SHAMMA, S. A. (2009). Rapid synaptic depression explains nonlinear modulation of spectro-temporal tuning in primary auditory cortex by natural stimuli. *J Neurosci*, *29*, 3374–3386.
- DEANGELIS, G. C., OHZAWA, I., & FREEMAN, R. D. (1995). Receptive-field dynamics in the central visual pathways. *Trends Neurosci*, *18*, 451–458.
- DICARLO, J. J., & COX, D. D. (2007). Untangling invariant object recognition. *Trends Cogn Sci*, *11*, 333–341.
- DICARLO, J. J., ZOCOLAN, D., & RUST, N. C. (2012). How does the brain solve visual object recognition? *Neuron*, *73*, 415–434.
- DOI, E., GAUTHIER, J. L., FIELD, G. D., SHLENS, J., SHER, A., GRESCHNER, M., ... SIMONCELLI, E. P. (2012). Efficient coding of spatial information in the primate retina. *J Neurosci*, *32*, 16256–16264.
- ENROTH-CUGELL, C., ROBSON, J. G., SCHWEITZER-TONG, D. E., & WATSON, A. B. (1983). Spatio-temporal interactions in cat retinal ganglion cells showing linear spatial summation. *J Physiol*, *341*, 279–307.
- ERIKSSON, J. T. (1996). Impact of information compression on intellectual activities in the brain. *Int J Neural Syst*, *7*, 543–550.
- FAHEY, P. K., & BURKHARDT, D. A. (2003). Center-surround organization in bipolar cells: Symmetry for opposing contrasts. *Vis Neurosci*, *20*, 1–10.
- FAIRHALL, A. L., BURLINGAME, C. A., NARASIMHAN, R., HARRIS, R. A., PUCHALLA, J. L., & BERRY, M. J. 2ND. (2006). Selectivity for multiple stimulus features in retinal ganglion cells. *J Neurophysiol*, *96*, 2724–2738.
- FIELD, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A*, *4*, 2379–2394.
- FIELD, G. D., & CHICHILNISKY, E. J. (2007). Information processing in the primate retina: Circuitry and coding. *Annu Rev Neurosci*, *30*, 1–30.
- FITZPATRICK, D. (2000). Seeing beyond the receptive field in primary visual cortex. *Curr Opin Neurobiol*, *10*, 438–443.
- FREIWALD, W. A., & TSAO, D. Y. (2010). Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science*, *330*, 845–851.
- FRENCH, A. S., & KORENBERG, M. J. (1989). A nonlinear cascade model for action potential encoding in an insect sensory neuron. *Biophys J*, *55*, 655–661.
- GARVERT, M. M., & GOLLISCH, T. (2013). Local and global contrast adaptation in retinal ganglion cells. *Neuron*, *77*, 915–928.
- GEFFEN, M. N., BROOME, B. M., LAURENT, G., & MEISTER, M. (2009). Neural encoding of rapidly fluctuating odors. *Neuron*, *61*, 570–586.
- GEFFEN, M. N., DE VRIES, S. E., & MEISTER, M. (2007). Retinal ganglion cells can rapidly change polarity from Off to On. *PLoS Biol*, *5*, e65.
- GILBERT, C. D., & LI, W. (2013). Top-down influences on visual processing. *Nat Rev Neurosci*, *14*, 350–363.
- GOLLISCH, T., & MEISTER, M. (2008). Rapid neural coding in the retina with relative spike latencies. *Science*, *319*, 1108–1111.
- GOLLISCH, T., & MEISTER, M. (2010). Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron*, *65*, 150–164.
- GOLLISCH, T., SCHUTZE, H., BENDA, J., & HERZ, A. V. (2002). Energy integration describes sound-intensity coding in an insect auditory system. *J Neurosci*, *22*, 10434–10448.
- GRAHAM, D. J., CHANDLER, D. M., & FIELD, D. J. (2006). Can the theory of “whitening” explain the center-surround properties of retinal ganglion cell receptive fields? *Vis Res*, *46*, 2901–2913.
- GRESCHNER, M., THIEL, A., KRETZBERG, J., & AMMERMÜLLER, J. (2006). Complex spike-event pattern of transient ON-OFF retinal ganglion cells. *J Neurophysiol*, *96*, 2845–2856.
- GROSS, C. G. (1992). Representation of visual stimuli in inferior temporal cortex. *Philos Trans R Soc Lond B Biol Sci*, *335*, 3–10.
- HARE, W. A., & OWEN, W. G. (1990). Spatial organization of the bipolar cell’s receptive field in the retina of the tiger salamander. *J Physiol*, *421*, 223–245.
- HEILIGENBERG, W. (1991). *Neural nets in electric fish*. Cambridge, MA: MIT Press.
- HOSOYA, T., BACCUS, S. A., & MEISTER, M. (2005). Dynamic predictive coding by the retina. *Nature*, *436*, 71–77.
- HROMADKA, T., & ZADOR, A. M. (2009). Representations in auditory cortex. *Curr Opin Neurobiol*, *19*, 430–433.
- HUBEL, D. H., & WIESEL, T. N. (1961). Integrative action in the cat’s lateral geniculate body. *J Physiol*, *155*, 385–398.
- HUBEL, D. H., & WIESEL, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J Physiol*, *160*, 106–154.
- ISAACSON, J. S. (2010). Odor representations in mammalian cortical circuits. *Curr Opin Neurobiol*, *20*, 328–331.
- KANWISHER, N., McDERMOTT, J., & CHUN, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *J Neurosci*, *17*, 4302–4311.
- KAPADIA, M. K., ITO, M., GILBERT, C. D., & WESTHEIMER, G. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron*, *15*, 843–856.
- KEAT, J., REINAGEL, P., REID, R. C., & MEISTER, M. (2001). Predicting every spike: A model for the responses of visual neurons. *Neuron*, *30*, 803–817.
- KONISHI, M. (2003). Coding of auditory space. *Annu Rev Neurosci*, *26*, 31–55.
- KONISHI, M. (2006). Behavioral guides for sensory neurophysiology. *J Comp Physiol A*, *192*, 671–676.
- KUFFLER, S. W. (1953). Discharge patterns and functional organization of mammalian retina. *J Neurophysiol*, *16*, 37–68.

- LAURENT, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nat Rev Neurosci*, 3, 884–895.
- LENNIE, P. (2003). The cost of cortical computation. *Curr Biol*, 13, 493–497.
- LIPSON, E. D. (1975). White noise analysis of *Phycomyces* light growth response system. I. Normal intensity range. *Biophys J*, 15, 989–1011.
- MACHENS, C. K., WEHR, M. S., & ZADOR, A. M. (2004). Linearity of cortical receptive fields measured with natural sounds. *J Neurosci*, 24, 1089–1100.
- MANTE, V., BONIN, V., & CARANDINI, M. (2008). Functional mechanisms shaping lateral geniculate responses to artificial and natural stimuli. *Neuron*, 58, 625–638.
- MARAVALL, M., PETERSEN, R. S., FAIRHALL, A. L., ARABZADEH, E., & DIAMOND, M. E. (2007). Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biol*, 5, e19.
- MARMARELIS, V. Z., & ORME, M. E. (1993). Modeling of neural systems by use of neuronal modes. *IEEE Trans Biomed Eng*, 40, 1149–1158.
- MARTINEZ, L. M., WANG, Q., REID, R. C., PILLAI, C., ALONSO, J. M., SOMMER, F. T., & HIRSCH, J. A. (2005). Receptive field structure varies with layer in the primary visual cortex. *Nat Neurosci*, 8, 372–379.
- MEISTER, M., & BERRY, M. J. (1999). The neural code of the retina. *Neuron*, 22, 435–450.
- MEL, B. W. (1997). SEEMORE: Combining color, shape, and texture histogramming in a neurally inspired approach to visual object recognition. *Neural Comput*, 9, 777–804.
- NAGEL, K. I., & DOUPE, A. J. (2006). Temporal processing and adaptation in the songbird auditory forebrain. *Neuron*, 51, 845–859.
- NAGEL, K. I., & DOUPE, A. J. (2008). Organizing principles of spectro-temporal encoding in the avian primary auditory area field L. *Neuron*, 58, 938–955.
- OLSHAUSEN, B. A., & FIELD, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–609.
- OLSHAUSEN, B. A., & FIELD, D. J. (2004). Sparse coding of sensory inputs. *Curr Opin Neurobiol*, 14, 481–487.
- ÖLVEČKY, B. P., BACCUS, S. A., & MEISTER, M. (2003). Segregation of object and background motion in the retina. *Nature*, 423, 401–408.
- OZUYSAL, Y., & BACCUS, S. A. (2012). Linking the computational structure of variance adaptation to biophysical mechanisms. *Neuron*, 73, 1002–1015.
- PIERCE, J. R., & KARLIN, J. E. (1957). Reading rates and the information rate of a human channel. *Bell Syst Tech J*, 36, 497–516.
- PITKOW, X., & MEISTER, M. (2012). Decorrelation and efficient coding by retinal ganglion cells. *Nat Neurosci*, 15, 628–635.
- QUIROGA, R. Q. (2012). Concept cells: The building blocks of declarative memory functions. *Nat Rev Neurosci*, 13, 587–597.
- QUIROGA, R. Q., REDDY, L., KREIMAN, G., KOCH, C., & FRIED, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, 435, 1102–1107.
- REID, R. C., & SHAPLEY, R. M. (2002). Space and time maps of cone photoreceptor signals in macaque lateral geniculate nucleus. *J Neurosci*, 22, 6158–6175.
- RIESENHUBER, M., & POGGIO, T. (1999). Hierarchical models of object recognition in cortex. *Nat Neurosci*, 2, 1019–1025.
- RIGOTTI, M., BEN DAYAN RUBIN, D., WANG, X. J., & FUSI, S. (2010). Internal representation of task rules by recurrent dynamics: The importance of the diversity of neural responses. *Front Comput Neurosci*, 4(24).
- RINGACH, D. L. (2004). Mapping receptive fields in primary visual cortex. *J Physiol*, 558, 717–728.
- RODIECK, R. W., & STONE, J. (1965). Analysis of receptive fields of cat retinal ganglion cells. *J Neurophysiol*, 28, 832–849.
- ROLLS, E. T. (1992). Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Philos Trans R Soc Lond B Biol Sci*, 335, 11–20.
- ROSE, G. J., KAWASAKI, M., & HEILIGENBERG, W. (1988). “Recognition units” at the top of a neuronal hierarchy? Prepacemaker neurons in *Eigenmannia* code the sign of frequency differences unambiguously. *J Comp Physiol A*, 162, 759–772.
- ROSKA, B., & WERBLIN, F. (2003). Rapid global shifts in natural scenes block spiking in specific ganglion cell types. *Nat Neurosci*, 6, 600–608.
- RUST, N. C., SCHWARTZ, O., MOVSHON, J. A., & SIMONCELLI, E. P. (2005). Spatiotemporal elements of macaque V1 receptive fields. *Neuron*, 46, 945–956.
- SCHNAPF, J. L., NUNN, B. J., MEISTER, M., & BAYLOR, D. A. (1990). Visual transduction in cones of the monkey *Macaca fascicularis*. *J Physiol*, 427, 681–713.
- SCHWARTZ, O., PILLOW, J. W., RUST, N. C., & SIMONCELLI, E. P. (2006). Spike-triggered neural characterization. *J Vis*, 6, 484–507.
- SERRE, T., OLIVA, A., & POGGIO, T. (2007). A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci USA*, 104, 6424–6429.
- SHAPLEY, R., & ENROTH-CUGELL, C. (1984). Visual adaptation and retinal gain controls. *Prog Retin Eye Res*, 3, 263–346.
- SHAPLEY, R. M., & VICTOR, J. D. (1981). How the contrast gain control modifies the frequency responses of cat retinal ganglion cells. *J Physiol*, 318, 161–79.
- SHARPEE, T. O., MILLER, K. D., & STRYKER, M. P. (2008). On the importance of static nonlinearity in estimating spatiotemporal neural filters with natural stimuli. *J Neurophysiol*, 99, 2496–2509.
- SHERRINGTON, C. S. (1906). Observations on the scratch-reflex in the spinal dog. *J Physiol*, 34, 1–50.
- SLEE, S. J., HIGGS, M. H., FAIRHALL, A. L., & SPAIN, W. J. (2005). Two-dimensional time coding in the auditory brainstem. *J Neurosci*, 25, 9978–9988.
- SMITH, E. C., & LEWICKI, M. S. (2006). Efficient auditory coding. *Nature*, 439, 978–982.
- SRINIVASAN, M. V., LAUGHLIN, S. B., & DUBS, A. (1982). Predictive coding: A fresh view of inhibition in the retina. *Proc R Soc Lond B Biol Sci*, 216, 427–459.
- THEUNISSEN, F. E., SEN, K., & DOUPE, A. J. (2000). Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *J Neurosci*, 20, 2315–2331.
- TOURYAN, J., LAU, B., & DAN, Y. (2002). Isolation of relevant visual features from random stimuli for cortical complex cells. *J Neurosci*, 22, 10811–10818.
- TSAO, D. Y., FREIWALD, W. A., KNUTSEN, T. A., MANDEVILLE, J. B., & TOOTELL, R. B. (2003). Faces and objects in macaque cerebral cortex. *Nat Neurosci*, 6, 989–995.

- TSAO, D. Y., & LIVINGSTONE, M. S. (2008). Mechanisms of face perception. *Annu Rev Neurosci*, *31*, 411–437.
- VAN HATEREN, J. H. (1992). A theory of maximizing sensory information. *Biol Cybern*, *68*, 23–29.
- VAN HATEREN, J. H. (1993). Spatiotemporal contrast sensitivity of early vision. *Vis Res*, *33*, 257–267.
- VAN HATEREN, J. H., RUTTIGER, L., SUN, H., & LEE, B. B. (2002). Processing of natural temporal stimuli by macaque retinal ganglion cells. *J Neurosci*, *22*, 9945–9960.
- VAN WYK, M., TAYLOR, W. R., & VANEY, D. I. (2006). Local edge detectors: A substrate for fine spatial vision at low temporal frequencies in rabbit retina. *J Neurosci*, *26*, 13250–13263.
- WARK, B., LUNDSTROM, B. N., & FAIRHALL, A. (2007). Sensory adaptation. *Curr Opin Neurobiol*, *17*, 423–429.
- WÄSSLE, H., PULLER, C., MULLER, F., & HAVERKAMP, S. (2009). Cone contacts, mosaics, and territories of bipolar cells in the mouse retina. *J Neurosci*, *29*, 106–117.
- WOLFE, J., HOUWELING, A. R., & BRECHT, M. (2010). Sparse and powerful cortical spikes. *Curr Opin Neurobiol*, *20*, 306–312.
- ZHANG, Y., KIM, I. J., SANES, J. R., & MEISTER, M. (2012). The most numerous ganglion cell type of the mouse retina is a selective feature detector. *Proc Natl Acad Sci USA*, *109*, 2391–2398.