

Toward Functional Classification of Neuronal Types

Tatyana O. Sharpee^{1,*}

¹Computational Neurobiology Laboratory, Salk Institute for Biological Studies, La Jolla, CA 92037, USA

*Correspondence: sharpee@salk.edu

<http://dx.doi.org/10.1016/j.neuron.2014.08.040>

How many types of neurons are there in the brain? This basic neuroscience question remains unsettled despite many decades of research. Classification schemes have been proposed based on anatomical, electrophysiological, or molecular properties. However, different schemes do not always agree with each other. This raises the question of whether one can classify neurons based on their function directly. For example, among sensory neurons, can a classification scheme be devised that is based on their role in encoding sensory stimuli? Here, theoretical arguments are outlined for how this can be achieved using information theory by looking at optimal numbers of cell types and paying attention to two key properties: correlations between inputs and noise in neural responses. This theoretical framework could help to map the hierarchical tree relating different neuronal classes within and across species.

As other publications in this issue of *Neuron* describe, we now stand at the threshold of being able to measure the anatomy and connections within neural circuits at an exquisite spatial resolution. The increasing availability of large-scale connectomics data sets holds the promises of answering one of the most basic and long-standing questions in neuroscience—how many different types of neurons are there in the brain (Jonas and Kording, 2014). Yet, the large size of these data sets also presents a set of challenges. To tackle these challenges, it would be advantageous to focus any analysis on those statistical properties of neural circuits that are likely to generalize across samples, brain regions, and perhaps species. In this respect, theoretical frameworks based on optimization principles can be very helpful, in particular in making predictions as to how many different neuron types one can expect to find in any given brain region and pinpointing experiments that make differences among neural types more apparent. In this article, I review a set of previous results and outline directions for future research on how we might be able to understand the relationship between neuronal types across species and brain regions, from the operation of identified neurons in invertebrate circuits to large neural populations in mammals.

Neurons can be classified based on anatomical (Boycott and Wässle, 1991; Davis and Sterling, 1979; Masland, 2004), molecular (Baraban and Tallent, 2004; Kodama et al., 2012; Nelson et al., 2006; Trimarchi et al., 2007) or electrophysiological properties (Markram et al., 2004). Yet, definitive mapping between these classification schema has proved difficult (Mott and Dingledine, 2003) because some of the diversity is due to morphological differences and some to differences in intrinsic firing or synaptic properties (Nelson et al., 2006). To complicate the matter further, sometimes vastly different molecular characteristics yield the same electrophysiological properties (Schulz et al., 2006). When such discrepancies arise, perhaps an insight based on neuronal function within the circuits could help. Functional classifications typically rely, in the case of sensory neurons, on differences in the types of input features encoded by different neurons (Balasubramanian and Sterling, 2009; Garrigan et al., 2010; Segev et al., 2006). (Analogous definitions are possible

in motor circuits but will not be discussed here). The relevant input features represent the stimulus components that can modulate neural responses most strongly. (The strength of this modulation can be quantified using either mutual information [Adelman et al., 2003; Rajan and Bialek, 2013; Sharpee et al., 2004] or by the amount of change in the mean or variance between distributions of inputs that elicit and do not elicit spikes [Bialek and de Ruyter van Steveninck, 2005; Schwartz et al., 2006].) Examples of classification based on functional properties include the separation of neurons based on their chromatic preferences (Chichilnisky and Wandell, 1999; Derrington et al., 1984; Wiesel and Hubel, 1966), into On and Off types of visual neurons (Chichilnisky and Kalmar, 2002; Masland, 1986; Werblin and Dowling, 1969), with further subdivisions into parvo- and magnocellular visual neurons (Sincich and Horton, 2005) distinguished according to different preferred spatial and temporal frequencies. These functional distinctions are often reflected in anatomical differences. For example, the dendrites of On and Off neurons segregate in different sublaminae of the inner plexiform layer (Masland, 1986).

Less commonly, neurons can sometimes be classified based on their nonlinear properties. Two of the most famous examples include the distinction between X and Y cells in the retina (Hochstein and Shapley, 1976) and between simple and complex cells in the primary visual cortex (Hubel and Wiesel, 1962, 1968). While the existence of separate classes of simple and complex classes has been questioned (Chance et al., 1999; Mechler and Ringach, 2002), recent studies provide other examples of where neurons are primarily distinguished based on their nonlinear properties (Kastner and Baccus, 2011). For example, the so-called adapting and sensitizing cells differ primarily in terms of their thresholds but not their preferred visual features (Kastner and Baccus, 2011). These observations call attention to the need to understand the circumstances under which neuron types can be defined primarily based on features, nonlinearities, or both.

In prior work, the application of information theoretic principles has been very useful in accounting for the relative proportions of cells tasked with encoding different stimulus features (Balasubramanian and Sterling, 2009; Ratliff et al., 2010). Further, searching for an optimal solution could account for spatial

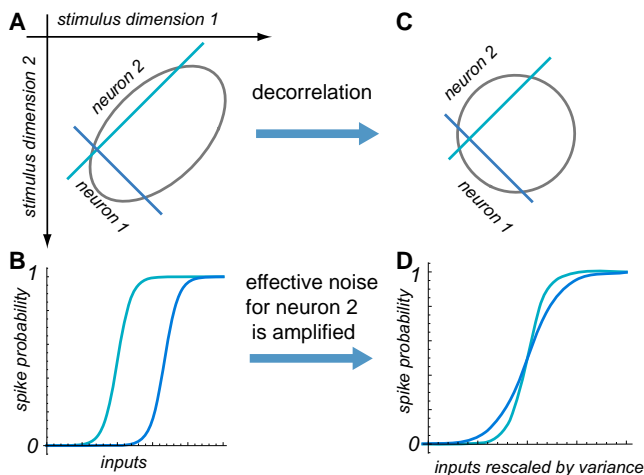


Figure 1. Decorrelation Amplifies Effective Noise in Neural Coding
 (A) One of the possible optimal configurations of neural thresholds in the presence of correlations. Neuron 1 encodes the input value along the direction of largest variance, whereas neuron 2 encodes the input value along the other dimension.
 (B) The thresholds values of the two neurons are proportional to the input variances along their respective relevant dimensions.
 (C and D) Equivalent to (A) and (B) except that the inputs are rescaled by their respective variances. This representation emphasizes that decorrelation increases the effective noise for neuron 2.

and temporal filtering properties of retinal neurons (Atick and Redlich, 1990, 1992; Bialek, 1987; Bialek and Owen, 1990; Rieke et al., 1991) as well as for changes in the filtering properties as a function of signal to noise ratio (Doi et al., 2012; Potters and Bialek, 1994; Srinivasan et al., 1982). Building on these previous successes, I will examine here the optimal solutions that maximize mutual information assuming that the responses of each neuron are binary and are described by logistic tuning curves (Pitkow and Meister, 2012). Logistic tuning curves correspond to the least constrained model nonlinearity that can be made consistent with measured input-output correlations (Fitzgerald et al., 2011; Globerson et al., 2009). It is likely that similar conclusions will be obtained if one were to consider Fisher information (Dayan and Abbott, 2001) instead of the Shannon mutual information (Brunel and Nadal, 1998; Harper and McAlpine, 2004; Salinas, 2006) and to allow for graded neural responses. This is because using Fisher information one can construct a bound for the mutual information and at least in the case of low noise and large neural ensembles, map a population of binary neurons to single-neuron encoding stimuli using graded responses (Brunel and Nadal, 1998).

A key feature of the environment that neural circuits have to operate in is that the signals are not only multidimensional but also highly correlated. This is uniformly true across sensory modalities (Ruderman and Bialek, 1994; Simoncelli and Olshausen, 2001) and motor outputs (Todorov, 2004). For example, the inputs to the fly visual system derive from 10^6 photoreceptors each of which conveys a slightly distinct message from its neighbors, but overall correlations between the inputs are strong. For the sake of simplicity, I will approximate input signals as correlated Gaussian, ignoring for now the (fairly substantial) non-

Gaussian effects in natural inputs (Ruderman and Bialek, 1994; Simoncelli and Olshausen, 2001; Singh and Theunissen, 2003). It should be noted that non-Gaussian effects can change the maximally informative dimensions within the input ensemble (Bell and Sejnowski, 1997; Karklin and Simoncelli, 2011; Olshausen and Field, 1996). Yet many insights can be obtained within the Gaussian approximation on which I will focus here. In the Gaussian approximation, the contours of equal probability are ellipses. The fact that inputs are correlated means that these ellipses are stretched in certain directions (Figure 1). The stronger the correlations, the more elongated the equal probability contours become. The fact that natural signals are strongly correlated means that the variance along certain stimulus dimensions in natural signals can be orders of magnitude larger than in other dimensions (Aljadeff et al., 2013; van Hateren and Ruderman, 1998). These observations on the properties of natural sensory signals have led Horace Barlow to propose the idea of a redundancy reduction, also known as “decorrelation” (Barlow, 1961, 2001; Barlow and Foldiak, 1989). Indeed, in the high signal-to-noise regime, it would make sense to first encode the input value along the dimension that has the largest variance and to allocate subsequent neurons to encode dimensions of progressively smaller variance. The high signal-to-noise regime presupposes that the ratio of noise to input variance is small even for dimensions along which input variance is small. With Gaussian inputs, one of the optimal configurations corresponds to sets of neurons that each encode inputs along orthogonal dimensions (Fitzgerald and Sharpee, 2009). For non-Gaussian inputs, optimal solutions specify that the sensitivity of single neurons to multiple inputs simultaneously (Horwitz and Hass, 2012; Sharpee and Bialek, 2007) with orthogonal dimensions substituted by sets of independent components (Bell and Sejnowski, 1997; Olshausen and Field, 1996; van Hateren and Ruderman, 1998) in the natural input ensemble.

The decorrelation comes with its own cost, however, in terms of its potential for noise amplification. Such problems are well known in the field of receptive field estimation where the goal is to estimate the relevant features from the neural responses (Sharpee, 2013; Theunissen et al., 2001; Sharpee et al., 2004). However, neural representations also have to solve the same issues. To illustrate the main ideas I will consider correlations between two input dimensions and how they might affect optimal encoding by two neurons. I will discuss along the way how these results might generalize to the higher dimensionality of the input space and neural populations. When two signals are correlated with strength ρ , the ratio of variances along these dimensions is given by $(1 + \rho)/(1 - \rho)$, which diverges as the correlation coefficient ρ increases toward 1. Assuming that input signals can be measured with a certain precision ν that is the same for all input dimensions, this neural noise ν will have a greater impact on encoding the dimension of smaller input variance. Specifically, the effective noise along the secondary dimension will be $(1 + \rho)/(1 - \rho)$ larger than the effective noise along the first dimension (Figure 1B). In the case of natural stimuli where variances along different input dimensions differ by orders of magnitude (Aljadeff et al., 2013; van Hateren and Ruderman, 1998), the effective noise along certain dimensions will be greatly amplified. In this situation, it might be better to allocate both neurons to

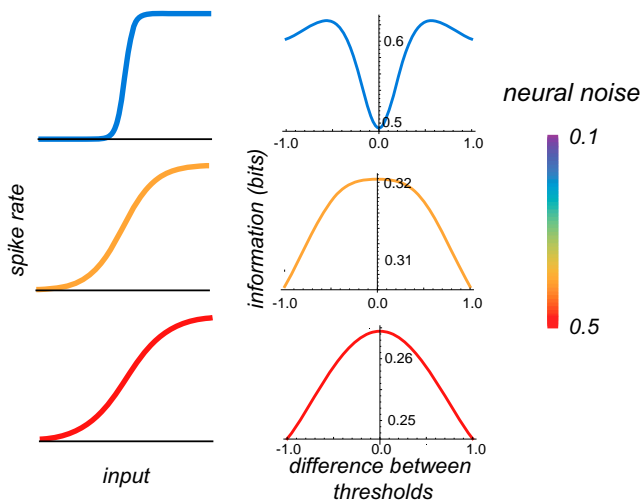


Figure 2. Separate Classes of Neurons Encoding the Same Input Dimension Can Appear with Decreasing Neural Noise

The noise values measured relative to the input SD are $\nu = 0.1, 0.3, 0.5$, top to bottom row. The probability of a spike from a neuron with a threshold μ and noise value ν is modeled using logistic function $P(\text{spike}|x) = [1 + \exp((\mu - x)/\nu)]^{-1}$, with x representing a Gaussian input of unit variance. For two neurons, there are four different response patterns (00, 01, 10, 11 where each neuron is assigned response 0 or 1 if it produces a spike). Neurons are assumed to be conditionally independent given x . With these assumptions, it is possible to evaluate mutual information as the difference between response entropy $R = -p_{00}\log_2 p_{00} - p_{01}\log_2 p_{01} - p_{10}\log_2 p_{10} - p_{11}\log_2 p_{11}$ and noise entropy

$$N = \int dx \frac{P(x)}{1 + \exp\left(\frac{\mu_1 - x}{\nu}\right)} \log_2 \left[1 + \exp\left(\frac{\mu_1 - x}{\nu}\right) \right] + \int dx \frac{P(x)}{1 + \exp\left(\frac{x - \mu_1}{\nu}\right)} \log_2 \left[1 + \exp\left(\frac{x - \mu_1}{\nu}\right) \right] + \int dx \frac{P(x)}{1 + \exp\left(\frac{\mu_2 - x}{\nu}\right)} \log_2 \left[1 + \exp\left(\frac{\mu_2 - x}{\nu}\right) \right] + \int dx \frac{P(x)}{1 + \exp\left(\frac{x - \mu_2}{\nu}\right)} \log_2 \left[1 + \exp\left(\frac{x - \mu_2}{\nu}\right) \right]$$

where μ_1 and μ_2 are the thresholds for the neurons 1 and 2, respectively. The mutual information $I = R - N$ is plotted here as a function of the difference between thresholds $\mu_1 - \mu_2$. For each value of $\mu_1 - \mu_2$ and ν , the average threshold between the two neurons is adjusted to ensure that the spike rate remains constant for all information values shown. In these simulation, the average spike rate across all inputs x and summed for neurons 1 and 2 was set to 0.2. With these parameter values, the transition from redundant coding observed for large noise values to coding based on different thresholds (observed at low noise values) occurs at $\nu \sim 0.4$.

encoding signal values along the primary dimension of variance as opposed to tasking them with encoding both input dimensions. In other words, when the amount of variance along the secondary dimension becomes smaller or comparable to neural noise, signals become effectively 1D.

Neurons that have the same threshold and encode the same dimension would then functionally be classified as belonging to

one class. Thus, the amount of neural noise together with correlations between inputs may determine the optimal number of neural classes. When the ratio of the neural noise to the smallest variance in input dimensions remains small, it will be beneficial to have as many neuronal classes as the input dimensions. Given that the noise in neural responses is substantial and can further enhance correlations between afferent neurons beyond correlations due to inputs (Ala-Laurila et al., 2011), multiple neurons might be assigned to encode the same input dimension, thereby reducing the number of distinct neuronal classes. Here, however, an interesting complication arises. It turns out that when multiple neurons encode the same input dimension, it might be optimal (in the sense of maximizing the Shannon mutual information about this input dimension) to assign these neurons different thresholds (McDonnell et al., 2006; Nikitin et al., 2009; Kastner et al., 2014). In this case, the neurons that encode the same dimensions would belong to separate classes, as has been observed in the retina (Kastner and Baccus, 2011). Whether or not this happens, i.e., whether mutual information is maximized when neurons have the same or different thresholds, again depends on the amount of neural noise (Figure 2). When noise in neural responses is large, information is maximal when the thresholds are the same (Figure 2A). In this regime, which might be termed redundant coding (Potters and Bialek, 1994; Tkacik et al., 2008), both neurons belong to the same functional class. As the noise in neural responses decreases, it may become beneficial to assign different thresholds (Figure 2) to the two neurons (Kastner et al., 2014).

It is also worth pointing out that in some cases correlations between inputs can be substantially increased beyond those present in the environment by noise fluctuations in afferent inputs (Ala-Laurila et al., 2011). The presence of these so-called noise correlations (Gawne and Richmond, 1993) could make encoding of stimulus values along the dimension of largest variance less important. In this case, multiple neurons could be assigned to encode stimulus values along dimensions of smaller variance.

In a more general case where more than two neurons are encoding the same input dimension, one would expect to find a series of bifurcations with decreasing noise level (McDonnell et al., 2006). When the noise is large, all neurons have the same thresholds and therefore form a single neuronal class. As the noise decreases, neurons might first split up into two groups, and if noise can be reduced further, subsequent neurons will separate from either of the two initial classes to form new classes. Thus, multiple neuronal classes may emerge even among neurons that encode the same input dimension and one could envision making a hierarchical tree of how different neuronal classes relate to each other.

Considering the effects due to input correlations and neural noise together, we observe that in some cases separate neural classes appear because neurons are encoding different input dimensions whereas in other cases neural classes appear because they have different thresholds along the same input dimension. For example, in the hypothetical case where correlations between inputs are zero, information is maximized by encoding separate input dimensions (Figure 3), regardless of neural noise level. Although this case is unrealistic, encoding along the




Figure 3. Crossover between Different Encoding Schemes

In scheme 1, neuron classes are defined based on different input features. Here, the mutual information depends on both correlation between inputs and the neural noise level. In scheme 2, neuron classes are defined based on differences in thresholds for the same input feature. The mutual information is independent of stimulus correlations and is shown here with lines parallel to the x axis. Color denotes the noise level relative to the SD along the dimension of largest variance. All simulations have been matched to have the same average spike rate across inputs (values are as in Figure 2). Suboptimal solutions (after the crossover) are shown using dashed and dotted lines, respectively for the case of parallel and orthogonal solutions.

orthogonal features will continue to be optimal provided that the correlations between inputs do not exceed a certain critical value. This critical value increases as the neural noise decreases (Figure 3). More generally, encoding along the orthogonal features corresponds to encoding along the so-called independent components of the input distribution, taking into account non-Gaussian characteristics of input statistics. However, the main observation that neurons in this regime will be distinguished based on their most relevant input feature is still likely to remain valid.

In the opposite regime of strong input correlations and small neural noise, neurons are distinguished based on their nonlinear properties. As the neural noise increases and correlations remain strong, optimal solutions specify neurons to have both the same relevant features and the same thresholds. Overall, the ratio between input correlations and neural noise emerges as the main parameter controlling the separation of neurons into classes based on either their relevant features or nonlinearities.

Could neurons be segregated by both their relevant features and their nonlinearities? The answer is likely no, or at least the differences in thresholds for neurons tuned to different features are likely to be small. To see why this is the case, we recall that independent solutions are optimal only when correlations are weak, meaning along dimensions with similar variance. In this case, the effective noise along these dimensions will be comparable. Thus, the difference in noise level will not drive neural thresholds to different positions relative to the peak of the distribution. At the same time the mutual information (Cover and Thomas, 1991) is maximized when the spikes from two neurons are used equally often (this maximizes the so-called response entropy) (Brenner et al., 2000). Thus, the thresholds will be positioned at similar distances from the peak of the input distribution

for two neurons. These arguments carry over to the case of multiple neurons encoding multiple input dimensions.

It is important to distinguish the effects of noise that drive the separation of neurons into two classes from the effects that the noise may have on changes in neural encoding within a given class. For example, changes in both the environmental conditions, such as the average light intensity and the metabolic state of the animal, may lead to an overall amount of noise across different inputs and neurons. Reducing the average light intensity will reduce correlations between photoreceptor outputs and change the direction of the maximal variance in the input distribution. Correspondingly, the optimal features that neurons encode will (and real features are known to) change (Atick and Redlich, 1990, 1992; Potters and Bialek, 1994; Srinivasan et al., 1982). However, provided the correlation between inputs stays above the critical value (marked by points in Figure 3 for different noise values), the number of neuron classes will not change. Neurons will just change their relevant features (the stimulus features will “rotate” within the input space). Decreasing the signal-to-noise ratio can also make input signals more Gaussian, which would initiate the change from the relevant features that are aligned with independent components to features that are aligned with principal components (Karklin and Simoncelli, 2011). Upon a further increase in the noise level, neurons might temporarily switch their features to encode the same input dimension. Depending on the noise level, their thresholds might temporarily merge, but that does not mean that they belong to the same class, provided environmental conditions exist where the neural responses differ qualitatively. Thus, the number of neuronal classes is determined by the smallest noise level that the circuit may experience in different environmental conditions.

Overall, the information-theoretic arguments discussed above highlight the difficulties and opportunities for creating a taxonomy of neuronal classes that spans across brain regions and species. The difficulties stem from the fact that both features and nonlinearities are context-dependent and might change depending on experimental conditions. Functional differences among neurons might be prominent in some experiments and absent in other. A theoretical framework describing how the features and thresholds should change could help pinpoint critical experiments that can reveal differences between classes of neurons that may have been missed up to now. In turn, a more detailed physiological description could help bridge differences in classification based on electrophysiological (Markram et al., 2004), molecular (Trimarchi et al., 2007), and anatomical parameters (Masland, 2004). Notable in that regard is the recent computational effort to characterize neuronal types based on statistical analysis of connectomics data sets (Jonas and Kording, 2014). At the same time, many opportunities exist for creating a catalog of optimal neural classes that could help relate differences and commonalities that exist across species. For example, in the case of retinal processing, where the neural classes are perhaps known in the most exquisite detail (Masland, 2001), noticeable differences between neural types in the rabbit, mouse, or primate exist. While these differences have largely been left unexplained so far, perhaps a theoretical framework that takes into account differences in environmental niches could offer clues. More

generally, it now seems feasible to obtain a unified framework that can bridge the divide between encoding using identified neurons in invertebrate circuits to encoding based on neural populations in mammalian and other species. Assuming that the intrinsic noise level is similar across species, one could potentially find a mapping between single identified neurons within the smaller invertebrate neural circuit with the corresponding populations of neurons in larger mammalian circuits.

ACKNOWLEDGMENTS

The author thanks David Kastner, Stephen Baccus, and John Berkowitz for many helpful conversations. The work was supported by the National Science Foundation (NSF) Career Award 1254123 and NIH grant R01EY019493.

REFERENCES

- Adelman, T.L., Bialek, W., and Olberg, R.M. (2003). The information content of receptive fields. *Neuron* 40, 823–833.
- Ala-Laurila, P., Greschner, M., Chichilnisky, E.J., and Rieke, F. (2011). Cone photoreceptor contributions to noise and correlations in the retinal output. *Nat. Neurosci.* 14, 1309–1316.
- Aljadeff, J., Segev, R., Berry, M.J., 2nd, and Sharpee, T.O. (2013). Spike triggered covariance in strongly correlated gaussian stimuli. *PLoS Comput. Biol.* 9, e1003206.
- Atick, J.J., and Redlich, A.N. (1990). Towards a theory of early visual processing. *Neural Comput.* 2, 308–320.
- Atick, J.J., and Redlich, A.N. (1992). What does the retina know about natural scenes? *Neural Comput.* 4, 196–210.
- Balasubramanian, V., and Sterling, P. (2009). Receptive fields and functional architecture in the retina. *J. Physiol.* 587, 2753–2767.
- Baraban, S.C., and Tallent, M.K. (2004). Interneuron diversity series: interneuronal neuropeptides—endogenous regulators of neuronal excitability. *Trends Neurosci.* 27, 135–142.
- Barlow, H.B. (1961). Possible principles underlying the transformation of sensory messages. In *Sensory Communication*, W.A. Rosenblith, ed. (Cambridge, MA: MIT Press), pp. 217–234.
- Barlow, H. (2001). Redundancy reduction revisited. *Network* 12, 241–253.
- Barlow, H., and Foldiak, P. (1989). Adaptation and decorrelation in the cortex. In *The Computing Neuron*, R. Durbin, C. Miall, and G. Mitchinson, eds. (New York: Addison-Wesley).
- Bell, A.J., and Sejnowski, T.J. (1997). The “independent components” of natural scenes are edge filters. *Vision Res.* 37, 3327–3338.
- Bialek, W. (1987). Physical limits to sensation and perception. *Annu. Rev. Biophys. Biophys. Chem.* 16, 455–478.
- Bialek, W., and de Ruyter van Steveninck, R.R. (2005). Features and dimensions: motion estimation in fly vision. In arxiv.org/abs/q-bio.NC/0505003.
- Bialek, W., and Owen, W.G. (1990). Temporal filtering in retinal bipolar cells. Elements of an optimal computation? *Biophys. J.* 58, 1227–1233.
- Boycott, B.B., and Wässle, H. (1991). Morphological classification of bipolar cells of the primate retina. *Eur. J. Neurosci.* 3, 1069–1088.
- Brenner, N., Strong, S.P., Koberle, R., Bialek, W., and de Ruyter van Steveninck, R.R. (2000). Synergy in a neural code. *Neural Comput.* 12, 1531–1552.
- Brunel, N., and Nadal, J.P. (1998). Mutual information, Fisher information, and population coding. *Neural Comput.* 10, 1731–1757.
- Chance, F.S., Nelson, S.B., and Abbott, L.F. (1999). Complex cells as cortically amplified simple cells. *Nat. Neurosci.* 2, 277–282.
- Chichilnisky, E.J., and Wandell, B.A. (1999). Trichromatic opponent color classification. *Vision Res.* 39, 3444–3458.
- Chichilnisky, E.J., and Kalmar, R.S. (2002). Functional asymmetries in ON and OFF ganglion cells of primate retina. *J. Neurosci.* 22, 2737–2747.
- Cover, T.M., and Thomas, J.A. (1991). *Elements of Information Theory*. (New York: Wiley-Interscience).
- Davis, T.L., and Sterling, P. (1979). Microcircuitry of cat visual cortex: Classification of neurons in layer (IV) of area 17, and identification of the patterns of lateral geniculate input. *J. Comp. Neurol.* 188, 599–628.
- Dayan, P., and Abbott, L.F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. (Cambridge, Massachusetts: MIT Press).
- Derrington, A.M., Krauskopf, J., and Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *J. Physiol.* 357, 241–265.
- Doi, E., Gauthier, J.L., Field, G.D., Shlens, J., Sher, A., Greschner, M., Machado, T.A., Jepsen, L.H., Mathieson, K., Gunning, D.E., et al. (2012). Efficient coding of spatial information in the primate retina. *J. Neurosci.* 32, 16256–16264.
- Fitzgerald, J.D., and Sharpee, T.O. (2009). Maximally informative pairwise interactions in networks. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 80, 031914.
- Fitzgerald, J.D., Sincich, L.C., and Sharpee, T.O. (2011). Minimal models of multidimensional computations. *PLoS Comput. Biol.* 7, e1001111.
- Garrigan, P., Ratliff, C.P., Klein, J.M., Sterling, P., Brainard, D.H., and Balasubramanian, V. (2010). Design of a trichromatic cone array. *PLoS Comput. Biol.* 6, e1000677.
- Gawne, T.J., and Richmond, B.J. (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.* 13, 2758–2771.
- Globerson, A., Stark, E., Vaadia, E., and Tishby, N. (2009). The minimum information principle and its application to neural code analysis. *Proc. Natl. Acad. Sci. USA* 106, 3490–3495.
- Harper, N.S., and McAlpine, D. (2004). Optimal neural population coding of an auditory spatial cue. *Nature* 430, 682–686.
- Hochstein, S., and Shapley, R.M. (1976). Quantitative analysis of retinal ganglion cell classifications. *J. Physiol.* 262, 237–264.
- Horwitz, G.D., and Hass, C.A. (2012). Nonlinear analysis of macaque V1 color tuning reveals cardinal directions for cortical color processing. *Nat. Neurosci.* 15, 913–919.
- Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J. Physiol.* 160, 106–154.
- Hubel, D.H., and Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* 195, 215–243.
- Jonas, E., and Kording, K. (2014). Automatic discovery of cell types and microcircuitry from neural connectomics. In [arXiv:1407.4137 \[q-bio.NC\]](https://arxiv.org/abs/1407.4137).
- Karklin, Y., and Simoncelli, E.P. (2011). Efficient coding of natural images with a population of noisy Linear-Nonlinear neurons. *NIPS*, 999–1007.
- Kastner, D.B., and Baccus, S.A. (2011). Coordinated dynamic encoding in the retina using opposing forms of plasticity. *Nat. Neurosci.* 14, 1317–1322.
- Kastner, D.B., Baccus, S.A., and Sharpee, T.O. (2014). Critical and maximally informative encoding between neural populations in the retina. In [arXiv:1409.2604 \[q-bio.NC\]](https://arxiv.org/abs/1409.2604).
- Kodama, T., Guerrero, S., Shin, M., Moghadam, S., Faulstich, M., and du Lac, S. (2012). Neuronal classification and marker gene identification via single-cell expression profiling of brainstem vestibular neurons subserving cerebellar learning. *J. Neurosci.* 32, 7819–7831.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nat. Rev. Neurosci.* 5, 793–807.
- Masland, R.H. (1986). The functional architecture of the retina. *Sci. Am.* 255, 102–111.

- Masland, R.H. (2001). The fundamental plan of the retina. *Nat. Neurosci.* 4, 877–886.
- Masland, R.H. (2004). Neuronal cell types. *Curr. Biol.* 14, R497–R500.
- McDonnell, M.D., Stocks, N.G., Pearce, C.E.M., and Abbott, D. (2006). Optimal information transmission in nonlinear arrays through suprathreshold stochastic resonance. *Phys. Lett.* 352, 183–189.
- Mechler, F., and Ringach, D.L. (2002). On the classification of simple and complex cells. *Vision Res.* 42, 1017–1033.
- Mott, D.D., and Dingledine, R. (2003). Interneuron Diversity series: Interneuron research—challenges and strategies. *Trends Neurosci.* 26, 484–488.
- Nelson, S.B., Sugino, K., and Hempel, C.M. (2006). The problem of neuronal cell types: a physiological genomics approach. *Trends Neurosci.* 29, 339–345.
- Nikitin, A.P., Stocks, N.G., Morse, R.P., and McDonnell, M.D. (2009). Neural population coding is optimized by discrete tuning curves. *Phys. Rev. Lett.* 103, 138101.
- Olshausen, B.A., and Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607–609.
- Pitkow, X., and Meister, M. (2012). Decorrelation and efficient coding by retinal ganglion cells. *Nat. Neurosci.* 15, 628–635.
- Potters, M., and Bialek, W. (1994). Statistical mechanics and visual signal processing. *J. Phys. I France* 4, 1755.
- Rajan, K., and Bialek, W. (2013). Maximally informative “stimulus energies” in the analysis of neural responses to natural signals. *PLoS ONE* 8, e71959.
- Ratliff, C.P., Borghuis, B.G., Kao, Y.H., Sterling, P., and Balasubramanian, V. (2010). Retina is structured to process an excess of darkness in natural scenes. *Proc. Natl. Acad. Sci. USA* 107, 17368–17373.
- Rieke, F., Owen, W.G., and Bialek, W. (1991). Advances in Neural Information Processing 3, R. Lippman, J. Moody, and D. Touretzky, eds. (San Mateo, CA: Morgan Kaufmann), pp. 377–383.
- Ruderman, D.L., and Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Phys. Rev. Lett.* 73, 814–817.
- Salinas, E. (2006). How behavioral constraints may determine optimal sensory representations. *PLoS Biol.* 4, e387.
- Schulz, D.J., Goallard, J.M., and Marder, E. (2006). Variable channel expression in identified single and electrically coupled neurons in different animals. *Nat. Neurosci.* 9, 356–362.
- Schwartz, O., Pillow, J.W., Rust, N.C., and Simoncelli, E.P. (2006). Spike-triggered neural characterization. *J. Vis.* 6, 484–507.
- Segev, R., Puchalla, J., and Berry, M.J., 2nd. (2006). Functional organization of ganglion cells in the salamander retina. *J. Neurophysiol.* 95, 2277–2292.
- Sharpee, T.O. (2013). Computational identification of receptive fields. *Annu. Rev. Neurosci.* 36, 103–120.
- Sharpee, T., and Bialek, W. (2007). Neural decision boundaries for maximal information transmission. *PLoS ONE* 2, e646.
- Sharpee, T., Rust, N.C., and Bialek, W. (2004). Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput.* 16, 223–250.
- Simoncelli, E.P., and Olshausen, B.A. (2001). Natural image statistics and neural representation. *Annu. Rev. Neurosci.* 24, 1193–1216.
- Sincich, L.C., and Horton, J.C. (2005). The circuitry of V1 and V2: integration of color, form, and motion. *Annu. Rev. Neurosci.* 28, 303–326.
- Singh, N.C., and Theunissen, F.E. (2003). Modulation spectra of natural sounds and ethological theories of auditory processing. *J. Acoust. Soc. Am.* 114, 3394–3411.
- Srinivasan, M.V., Laughlin, S.B., and 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., David, S.V., Singh, N.C., Hsu, A., Vinje, W.E., and Gallant, J.L. (2001). Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 12, 289–316.
- Tkacik, G., Callan, C.G., Jr., and Bialek, W. (2008). Information flow and optimization in transcriptional regulation. *Proc. Natl. Acad. Sci. USA* 105, 12265–12270.
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nat. Neurosci.* 7, 907–915.
- Trimarchi, J.M., Stadler, M.B., Roska, B., Billings, N., Sun, B., Bartch, B., and Cepko, C.L. (2007). Molecular heterogeneity of developing retinal ganglion and amacrine cells revealed through single cell gene expression profiling. *J. Comp. Neurol.* 502, 1047–1065.
- van Hateren, J.H., and Ruderman, D.L. (1998). Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc. Biol. Sci.* 265, 2315–2320.
- Werblin, F.S., and Dowling, J.E. (1969). Organization of the retina of the mudpuppy, *Necturus maculosus*. II. Intracellular recording. *J. Neurophysiol.* 32, 339–355.
- Wiesel, T.N., and Hubel, D.H. (1966). Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *J. Neurophysiol.* 29, 1115–1156.