

Interacting Linear and Nonlinear Characteristics Produce Population Coding Asymmetries between ON and OFF Cells in the Retina

Zachary Nichols,¹ Sheila Nirenberg,^{1,2} and Jonathan Victor^{2,3}

¹Department of Physiology and Biophysics, ²Institute for Computational Biology, and ³Brain and Mind Research Institute, Weill Cornell Medical College, New York, New York 10065

The early visual system is a model for understanding the roles of cell populations in parallel processing. Cells in this system can be classified according to their responsiveness to different stimuli; a prominent example is the division between cells that respond to stimuli of opposite contrasts (ON vs OFF cells). These two cell classes display many asymmetries in their physiological characteristics (including temporal characteristics, spatial characteristics, and nonlinear characteristics) that, individually, are known to have important roles in population coding. Here we describe a novel distinction between the information that ON and OFF ganglion cell populations carry in mouse—that OFF cells are able to signal motion information about both light and dark objects, while ON cells have a selective deficit at signaling the motion of dark objects. We found that none of the previously reported asymmetries in physiological characteristics could account for this distinction. We therefore analyzed its basis via a recently developed linear–nonlinear–Poisson model that faithfully captures input/output relationships for a broad range of stimuli (Bomash et al., 2013). While the coding differences between ON and OFF cell populations could not be ascribed to the linear or nonlinear components of the model individually, they had a simple explanation in the way that these components interact. Sensory transformations in other systems can likewise be described by these models, and thus our findings suggest that similar interactions between component properties may help account for the roles of cell classes in population coding more generally.

Introduction

The structure of visual system is a prime example of parallel organization in the brain (Masland, 2001; Wässle, 2004). At multiple levels within this system, information is processed simultaneously in different cell populations. A canonical case of this parallel processing is the separation of ON and OFF responses (Hartline, 1938), which first occurs at the bipolar cell synapse (Werblin and Dowling, 1969) and continues into the brain. The utility of this separation is indicated by its conservation across the retinas of vertebrates, from cartilaginous fishes (Dowling and Ripps, 1970) to amphibians (Hartline, 1938; Schwartz, 1974) to mammals (Kuffler, 1953; for review, see Schiller, 2010). But despite its ubiquity and presumed selective advantage, the functional implications of this separation are incompletely understood.

An important aspect of this incomplete understanding is the fact that ON and OFF pathways are not simply equal and opposite. Asymmetries begin at the retinal level and include spatial filtering properties (Chichilnisky and Kalmar, 2002; Balasubramanian and Sterling, 2009), temporal filtering properties (Chichilnisky and Kalmar, 2002; Sagdullaev and McCall, 2005; Pandarinath et al., 2010), and nonlinear properties (Chichilnisky and Kalmar, 2002; Zaghoul et al., 2003; Molnar et al., 2009). Asymmetries also continue downstream, where circuitry devotes unequal resources to processing lights and darks (Zemon et al., 1988; Jin et al., 2008; Yeh et al., 2009).

These asymmetries contribute to the challenge of understanding the roles of the ON and OFF channels for two reasons. First, they complicate approaches that rely on the design of stimuli that selectively activate one or another of the channels. But more importantly, these asymmetries raise the possibility that the functional roles of the two classes are not restricted to a simple partitioning of scenes into light and dark components, since the two cell classes also have different spatial and temporal characteristics.

Here we used a data-driven computational approach—the virtual retina (Bomash et al., 2013)—that addresses both of these issues. First, it allows for clean isolation of the information carried by ON and OFF ganglion cell populations, by reconstructing or decoding the responses of just one population. Second, as presented by Bomash et al. (2013), it allows for rapid *in silico* screening of hypotheses concerning the functional roles of ON

Received March 4, 2013; revised Aug. 5, 2013; accepted Aug. 9, 2013.

Author contributions: Z.N., S.N., and J.V. designed research; Z.N. performed research; Z.N. analyzed data; Z.N., S.N., and J.V. wrote the paper.

This work was supported by NIH Grants R01 EY12978 (S.N.) and R01 EY07977 (J.V.), the Tri-Institutional Training Program in Computational Biology and Medicine (Z.N., NIH Grant T32GM083937), and the Tri-Institutional Training Program in Vision Research (NIH Grant 5T32EY007138-19). We thank Illya Bomash and Chethan Pandarinath for helpful comments on this manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to Jonathan Victor, Brain and Mind Research Institute, Weill Cornell Medical College, New York, NY 10065. E-mail: jdvicto@med.cornell.edu.

DOI:10.1523/JNEUROSCI.1004-13.2013

Copyright © 2013 the authors 0270-6474/13/3314958-16\$15.00/0

and OFF populations, so that physiological experiments can be focused on ones that are viable.

Using this approach, we identified an unexpected selective deficit for motion processing in ON cells and analyzed its physiological basis. In particular, we first found that model-based stimulus reconstruction experiments suggest that OFF populations are able to transmit information about the motion of both light and dark objects, while ON populations have a deficit in transmitting information about the motion of dark objects. We then designed a motion-decoding task that allowed us to confirm this difference with electrophysiological recordings directly, independently of models. Finally, we analyzed the source of this difference and found that it results from an interaction between asymmetries that involve the linear and nonlinear components of ganglion cell processing.

Materials and Methods

Tissue preparation and recording. Electrophysiological recordings were obtained *in vitro* from the isolated retinas of C57BL/6 mice. All procedures were performed with approval of the Institutional Animal Care and Use Committee of Weill Cornell Medical College (protocol #0807-769A).

Central retinal ganglion cell (RGC) responses were recorded on a 64-channel multielectrode array using methods described previously (Pandarinath et al., 2010). Briefly, 7- to 9-week-old female mice were dark adapted for 1–3 h, after which they were killed and their retinas dissected under dim red light into oxygenated Ringer's solution. The central retina (adjacent to the optic nerve) was then isolated, yielding a piece 1.5 to 3 mm on a side, which was placed onto the multielectrode array for recording. Spike waveforms were amplified and digitized via a Plexon Instruments Multichannel Neuronal Acquisition Processor. A standard spike sorting method (Fee et al., 1996) was used to identify individual cells.

Only units that were well isolated, as measured by refractory period violations (fewer than 2% of the spikes occurring within 1.5 ms of the previous spike), and whose firing rate was stable over the course of the experiment were included in the analyses. (The mean drop in firing rate over the course of the experiment was 5% of the initial recorded firing rate; cells whose firing rate fell by >30% were discarded.) Overall, data were recorded from 58 retinas, yielding 512 ganglion cells that met these criteria.

Stimulation. Retinas were presented with spatiotemporally varying grayscale stimuli (mean luminance of $0.24 \mu\text{W}/\text{cm}^2$) using a Sony LCD computer monitor. This yielded photoisomerization rates in the central retina as described previously: $1.8 \times 10^3 R^*/\text{rod}/\text{s}$, $900 R^*/\text{M cone}/\text{s}$, and $40 R^*/\text{S cone}/\text{s}$ (Pandarinath et al., 2010). Note, however, that cones in the mouse central retina coexpress M and S opsins (Applebury et al., 2000; Nikonov et al., 2006), and that near the optic disk, the proportion varies substantially along the ventral to dorsal axis (Wang et al., 2011). Since the exact regions recorded in these experiments are unknown, the rates $900 R^*/\text{M cone}/\text{s}$, and $40 R^*/\text{S cone}/\text{s}$ constitute the range of cone photoisomerizations at this light level.

Stimulus frames were presented at 15 Hz; each frame consisted of 20×18 checks, each subending $80 \times 80 \mu\text{m}$ on the retina. This corresponds to a total stimulus area of $\sim 77.5 \times 70^\circ$ of visual angle, and $3.9 \times 3.9^\circ$ per check. The spatial and temporal resolution given by this system covers the response range for mouse ganglion cells under photopic conditions, as measured by their responses to spots and gratings (Stone and Pinto, 1993; Carciari et al., 2003; Umino et al., 2008; Pandarinath et al., 2010).

Three types of stimuli were used: binary spatiotemporal white noise (WN), natural scene (NS) movies, and coherent motion (CM) stimuli. WN and NS stimuli were used to build models of neuronal responses; CM stimuli were used to test the ability of populations to transmit infor-

mation about moving objects. All stimuli were presented at grayscale, with 8 bit resolution.

For the WN stimulus, checks were randomly and independently assigned to either of two luminances, 0.15 or $0.33 \mu\text{W}/\text{cm}^2$. This yielded a root-mean-squared (RMS) excursion of $0.087 \mu\text{W}/\text{cm}^2$ about the mean. For the NS stimulus, a movie filmed in Central Park of New York City [as in the study by Meytlis et al. (2012)] was digitized at the spatial and temporal resolution of the display by block averaging. The WN and NS movies were linearly scaled to have the same mean luminance and RMS excursion. The NS movie had a temporal power spectrum of $1/f^{2.04}$ (where f is temporal frequency) and a spatial power spectrum of $1/\omega^{2.09}$ (where ω is spatial frequency).

For the CM stimuli, two versions were constructed: dark objects moving on a lighter background and light objects moving on a darker background (see Fig. 2A). Each consisted of an array of 16 objects (disks with diameter of 6.5° of visual angle) centered on a 4×4 grid. The CM stimuli were presented at the same size and display resolution as the WN and NS stimuli. Light objects had a maximum luminance matching the light checks of the WN stimulus, and their background had the luminance of the dark checks; the reverse was true for the dark objects. To create coherent motion, a subset of these objects moved either left or right at a constant speed (17.5° of visual angle per second). Each motion stimulus lasted 1.4 s, so the total distance moved by each object was 24.5° , $\sim 140\%$ of the distance between object centers. Since the background luminances for the light-object and dark-object CM stimuli differed (0.15 vs $0.33 \mu\text{W}/\text{cm}^2$), the Weber contrasts differed as well, by a factor of 2.2:1 (the inverse of this ratio).

To use CM stimuli to test transmission of information about moving objects, they were constructed at several levels of difficulty. Difficulty was controlled by varying the number of nonmoving objects from 2 (easiest) to 10 (hardest). For each polarity (light and dark objects) and each level of difficulty (2, 4, 6, 8, and 10 nonmoving objects), we constructed a library containing 32 examples of left-moving stimuli and 32 examples of right-moving stimuli. Members of the library varied according to the (random) choice of the specific objects that moved. Stimuli were excluded from the library if the distribution of moving objects was highly nonuniform; our criterion was that the minimum and maximum number of moving objects within any 2×2 subsection of the grid differed by no more than 2.

Models. Models were constructed from neural responses to the WN and NS stimuli using a procedure that was shown to accurately reproduce neural responses to a wide variety of stimuli (Nirenberg and Pandarinath, 2012; Bomash et al., 2013). The models were composed of a linear–nonlinear–Poisson (LNP) cascade: a given spatiotemporal pattern of light was transformed by a linear filter (L); the output of the linear filter fed into a spline nonlinearity (N), and the output of the spline nonlinearity was the instantaneous firing rate of the model (for general reviews of LNP models, see Simoncelli et al., 2004; Pillow et al., 2008). This firing rate is then used as the generator for an inhomogeneous Poisson process, allowing the model to simulate spike trains in response to the given stimulus.

Mathematically, the instantaneous firing rate λ at the time t is represented as follows:

$$\lambda(t) = N[X * L(t)], \quad (1)$$

where $*$ denotes convolution, X is the stimulus, L is the cell's linear filter, and N is the cell's nonlinearity.

Each cell's linear filter consisted of the product of a spatial function specified on a grid of 10 by 10 checks and a temporal function specified on 18 time bins (of 67 ms each), examples of which are shown in Figure 8A. The temporal function's dimensionality was reduced by constraining it to a linear combination of 10 basis functions (raised cosines) as in the studies by Nirenberg and Pandarinath (2012) and Bomash et al. (2013), following Pillow et al. (2008). The nonlinearity was parameterized as a cubic spline with six knots, with knots covering the range of output values from the linear filter.

As in the studies by Nirenberg and Pandarinath (2012) and Bomash et al. (2013), parameters for the model were fit from neural responses

to WN and NS stimuli, using a maximum likelihood approach. The quantity maximized is the log-likelihood of the model given the response (Z):

$$Z = \sum_m \left[\sum_i \log[\lambda_m(\tau_m(i))] - \int_{t=0}^{\text{end}} \lambda_m(t) dt \right], \quad (2)$$

where λ_m is the firing rate, and τ_m is the spike train of cell m .

Each model cell was constructed from responses of a cell to 20 min of spatiotemporal stimulation (10 min of NS and 10 min of WN). As reported previously (Nirenberg and Pandarinath, 2012, their supporting information), the rationale for this strategy is that the WN and NS stimuli work in a complementary way (for example, the NS stimulus has more power at low spatial and temporal frequencies, while the WN stimulus has more power at high spatial and temporal frequencies), so that the combination constrains the model fitness landscape better than either stimulus alone.

Descriptive parameters. Response latency, response duration, and bias index were defined as in the study by Carciari et al. (2003). Response latency was quantified as the time to peak firing rate after presentation of optimal spot, and response duration was quantified as the time for this response to decay to half of its peak. The preference for light versus dark stimuli was quantified by the bias index: $BI = (R_{ON} - R_{OFF}) / (R_{ON} + R_{OFF})$, where R_{ON} is the peak response to an ON spot and R_{OFF} is the peak response to an OFF spot, with -1 corresponding to a pure OFF cell, 0 corresponding to an ON–OFF cell, and 1 corresponding to a pure ON cell. To determine these values, we simulated the responses to optimal spot stimuli [defined as in the study by Carciari et al. (2003) as spots with diameter yielding the highest peak response for each individual cell] according to Equation 1, using the model derived from WN and NS responses. The distributions obtained for these descriptive parameters based on the simulated response to an optimal spot responses match previously reported distributions (Carciari et al., 2003).

Receptive field size was quantified by fitting a 2D Gaussian to the spatial component of the model filter L . As in the study by Chichilnisky and Kalmar (2002), the size parameter is then given by the square root of the area of the elliptical contour at height $1/e$ relative to the peak.

The degree of nonlinearity [the nonlinearity index (NI)] was quantified by the extent to which the model spline nonlinearity N deviated from a straight line. Specifically, with x_5 , x_{50} , and x_{95} , respectively, denoting the 5, 50, and 95% quantiles of the linear stage's output, the NI was defined as the difference between the spline output at x_{50} and the straight line drawn from the end quantiles: $NI = N(x_{50}) + [(x_{50} - x_5) \times (N(x_{95}) - N(x_5)) / (x_{95} - x_5)] - N(x_{50})$. Responses to WN stimuli were used to determine x_5 , x_{50} , and x_{95} [to match the convention used by Chichilnisky and Kalmar (2002)]. For the case of an exponential nonlinearity [used by Chichilnisky and Kalmar (2002)] rather than a spline nonlinearity (used by these models), this index yields values roughly proportional to the indices used by Chichilnisky and Kalmar (2002) and Zaghoul et al. (2003), which were based on the log ratios of the slopes of the nonlinearity.

The biphasic index (BI) was calculated following Cai et al. (1997) and Jin et al. (2011b) as the ratio of the refractory amplitude to the peak response amplitude. Refractory and peak amplitudes were calculated using the temporal component of the model filter L as an estimate of the impulse response: the peak amplitude a was given by the height of the first (positive) lobe of the temporal filter, and the refractory amplitude b was given by the depth of the second (negative) lobe of the temporal filter, so that $BI = b/a$. To accurately estimate the height and depth values, the temporal filter was interpolated with cubic splines.

Direction selectivity was quantified by the direction selectivity index (DSI), calculated [following the study by Grzywacz and Amthor (2007)] as $DSI = (R_{\text{pref}} - R_{\text{null}}) / (R_{\text{pref}} + R_{\text{null}})$, where R_{pref} is the response (spike count) for movement in the preferred direction, and R_{null} is response for movement in the null direction. Responses used for this measure were taken from each segment of the CM stimuli in which a single object crosses the receptive field of the cell. Significance of the DSI measure was determined by an unpaired, two-sided t test comparing the trials where movement was in the preferred direction to trials in which movement was in the null direction, with cells whose DSI was >0.3 with 95% con-

fidence considered direction selective (DS). Since the CM stimuli present motion in only two directions and do not stimulate all recorded cells, this procedure does not identify all directionally selective cells, but it does identify the cells whose direction selectivity could impact the CM experiments (i.e., cells with a rightward vs leftward bias and whose receptive fields were in the path of a moving object).

Cell classification. Retinal ganglion cells were classified into four groups: long-latency cells, short-latency ON cells, short-latency ON–OFF cells, and short-latency OFF cells. Following the study by Carciari et al. (2003), cells were designated as long latency if their simulated peak response to an optimal spot occurred after 400 ms, and the remaining (short-latency) cells were classified into ON, ON–OFF, or OFF based on their BI (ON, $BI > 0.5$; ON–OFF, $0.5 \geq BI > -0.5$; OFF, $BI \leq -0.5$). Using this criteria, the 512 recorded RGCs were classified into 13 long-latency cells, 238 ON cells, 121 OFF cells, and 140 ON–OFF cells.

Decoding spike trains. To determine the extent to which responses to the CM stimuli conveyed information about leftward versus rightward object motion, they were analyzed using standard Bayesian decoding on binned responses. To do this, we determined the CM stimulus that was most likely to elicit each given response (out of all CM stimuli in that library, sharing the same luminance polarity and the same level of difficulty). These decoded responses were then tabulated into a confusion matrix, which were used to calculate the fraction of responses that yielded the correct direction (left vs right).

To determine the most likely CM stimulus given a multineuronal response \mathbf{r} , we applied Bayes' theorem:

$$p(s_i | \mathbf{r}) \propto p(\mathbf{r} | s_i) p(s_i), \quad (3)$$

where $p(s_i | \mathbf{r})$ is the a posteriori probability of a stimulus s_i given the response \mathbf{r} , $p(\mathbf{r} | s_i)$ is the probability of a response \mathbf{r} given a stimulus s_i , and $p(s_i)$ is the a priori probability of the stimulus s_i occurring. Since all $p(s_i)$ are uniformly equal (to $1/64$), finding the stimulus that maximizes the a posteriori probability $p(s_i | \mathbf{r})$ is equivalent to maximizing $p(\mathbf{r} | s_i)$.

To calculate $p(\mathbf{r} | s_i)$, each neuron's response was considered to be conditionally independent—which for natural scenes does not result in significant loss of information (Oizumi et al., 2010; Meytlis et al., 2012)—enabling the probability of a multineuronal response to be written as a product of the probabilities of the responses of the individual neurons. Each cell's responses were then treated as an inhomogeneous Poisson process, i.e., as a Poisson process whose rate could vary across time. To capture this time-dependent firing rate, the response period was divided into eight bins of 125 ms width (bin sizes both larger and smaller than 125 ms have been shown to produce similar decoding results for gratings and natural scenes; Jacobs et al., 2009; Pandarinath et al., 2010; Bomash et al., 2013), starting at 400 ms after stimulus onset (see Fig. 2A). Spikes occurring over the initial 400 ms of stimulus presentation were not included as part of the binned response, so that the previous stimulus would not influence the data used for decoding.

The response distribution $p(\mathbf{r} | s_i)$ for each stimulus was estimated by taking 20 randomly chosen trials as the training set and using the sample mean number of spikes per bin to construct the binwise Poisson distributions for each neuron. The remaining 20 trials were then used as a testing set for decoding.

Spatiotemporal reconstructions. For movie reconstructions, Bayesian decoding was also used (using Equation 3 as above), again using a uniformly flat prior (i.e., all possible spatiotemporal movies considered a priori equally likely) and again assuming that the response distribution was described by an inhomogeneous Poisson process whose instantaneous rate function depends on the stimulus. Since we required this rate function for all possible movies s (in contrast to the CM decoding described above, in which we required the rate function only for the 64 alternative stimuli), here it was necessary to use models of the cells to calculate $p(\mathbf{r} | s)$.

To obtain the spiking data required for the reconstructions, virtual populations were created (using the LNP model of Equation 1, obtained as described above). The density of model cells in the virtual populations was matched to reported ganglion cell density in the mouse retina (Jeon et al., 1998). Model cell parameters were drawn from previously recorded

experiments (Bomash et al., 2013) and were chosen to have approximately the same mean firing rate (~ 12 Hz). To keep the computation tractable, a small number of example cells were used (six total); these cells were replicated and positioned with their centers covering the stimulus in the desired density.

Reconstructions were performed using gradient ascent to find the maximum likelihood stimulus, given the simulated population responses [via Eq. 3, again setting $p(s)$ as uniform]. The likelihood that the observed response was elicited by a candidate movie, which is proportional to $p(r|s)$ in Equation 3, was calculated using Equation 2, where the instantaneous firing rate is again given by Equation 1. Maximization was performed using the scaled conjugate gradients library provided by SciPy (<http://www.scipy.org>). To keep the calculation tractable, reconstructions were performed independently in adjacent blocks of seven by seven checks (18.2° of visual angle) over 15 frames (1 s). To ensure that convergence was to a global maximum, multiple random seeds were used for each block.

The above calculations were performed in parallel on a 72 CPU cluster.

Organization of physiological experiments. We recorded the RGC responses from each retina to two groups of stimuli: modeling stimuli and decoding-task stimuli.

All recordings included the same modeling stimuli, consisting of 10 min of nonrepeating spatiotemporal WN and 10 min of nonrepeating spatiotemporal NS, together used for model fitting, and an out-of-sample 5 min sequence of a repeated NS stimulus (30 presentations of a 10 s scene) used for validation, as in the studies by Meytlis et al. (2012) and Bomash et al. (2013). To keep the recording time within practical limits, each recording included only a subset of the decoding-task stimuli: either the light or dark CM, and either difficulty levels 2, 6, and 10 or difficulty levels 4 and 8. Thus, the retinas (58 in total) were subdivided into four cohorts (12 to 19 per cohort): one subdivision based on whether a light or dark CM experiment was run and a second subdivision based on the two sets of difficulty levels (2, 6, and 10 or 4 and 8).

Results

To identify and characterize differences in information carried by ON and OFF cell populations, we proceeded in several steps. First, we used model-based reconstructions of spatiotemporal stimuli to quickly probe for differences between the two cell classes. We found that ON cell model populations display a selective deficit in carrying information about moving stimuli (a deficit that OFF cell populations do not exhibit). Next, we used ideal-observer decoding methods to confirm the presence of this deficit with *in vitro* microelectrode recordings. Finally, we showed that this information coding difference between ON and OFF cell populations is related to an interaction between their linear and nonlinear response characteristics.

ON and OFF cell populations differ in a reconstruction task

To survey for deficits in information carried by ON and OFF RGCs, we used their responses to movies consisting of moving objects. This was a three-step process (see Materials and Methods): first, several types of model populations were built, each containing either ON cells, OFF cells, or both ON and OFF cells together. Next, responses were simulated for each of the model populations to spatiotemporal movies. Finally, the responses were decoded to yield reconstructed spatiotemporal movies.

Fig. 1 shows the results of the reconstruction experiments for two types of stimuli, a field of moving light objects on a dark background (Fig. 1A) and a field of moving dark objects on a light background (Fig. 1B). For each type of stimulus (positive and negative contrast), objects of various sizes move toward the viewer (increasing in size as they do so). For each reconstruction experiment, we simulated responses with a ganglion cell density roughly matching that in the mouse retina (Jeon et al., 1998), corresponding to $\sim 20,000$ cells for a $120 \times 120^\circ$ field. Recon-

structions were then performed using gradient ascent to find the stimulus that most likely accounts for the responses (i.e., the maximum likelihood stimulus; see Materials and Methods).

Figure 1A shows the three reconstructions (mixed ON and OFF cells, only OFF cells, or only ON cells) for the moving light objects. These were of similar quality: for each of these reconstructions, despite the presence of noise, it is possible to accurately judge the presence and position of all but the smallest objects (Fig. 1A, compare bottom three rows).

For the reconstructions of moving dark objects in Figure 1B, this was not the case (note the difference between the bottom two rows). While the OFF RGC population (third row) yielded a reconstruction similar in quality to that of the mixed population (second row), the ON RGC population reconstruction (bottom row) contained many errors, including spurious dark regions and missing objects. Consequently, finding the location and speed of the original objects in the ON cell reconstruction is more difficult. The difference between the bottom two rows of Figure 1B, while subjective in that it relies on a visual interpretation of the reconstruction, led us to frame the hypothesis that ON RGCs have a deficit in signaling moving dark objects. This appeared to be a selective deficit and to represent an asymmetry between cell classes, as there was no corresponding difference in the OFF versus ON reconstructions for moving light objects (Fig. 1A, bottom two rows).

ON and OFF cell populations differ in an object motion discrimination task

To test the hypothesis that ON RGCs have a selective deficit in signaling moving dark objects, and to do so in a quantitative way that did not depend on inspection of model-based reconstructions, we designed a physiological experiment that allowed us to quantify the performance of recorded RGC populations in a motion decoding task.

The task was constructed by adapting the coherent dot-motion paradigm of Britten et al. (1992) to the needs of Bayesian decoding of experimentally recorded population responses. There were two key constraints. First, the two classes of stimuli to be distinguished (leftward vs rightward motion) needed to contain a sufficient number of examples to show that the decoded signal captured motion per se, and not idiosyncratic features of individual stimuli. Second, the number of distinct stimulus examples could not be too large, as it is necessary to record a sufficient number of responses to each stimulus to have an accurate empirical estimate of the stimulus–response relationship. To determine design parameters that were able to meet these constraints, we used the virtual retina (Bomash et al., 2013) approach. This led to the specific experimental design shown in Figure 2A, in which stimuli consisted of an array of 16 objects, a portion of which moved either left or right. Graded levels of difficulty were created by altering the number of objects that were not moving, over the range from 2 to 10. For each level of difficulty, 32 left-moving and 32 right-moving stimulus examples were used (for further details, see Materials and Methods, Stimulation and Decoding spike trains).

To visualize and quantify the reliability of the directional information carried by ON and OFF RGC populations to the CM stimuli, we calculated confusion matrices at each level of difficulty. Briefly (for details, see Materials and Methods), confusion matrices are a tally of the number of times that a stimulus is confused with another: an entry in row i and column j of the matrix corresponds to the number of trials in which stimulus i is presented and stimulus j is decoded from the response. Confu-

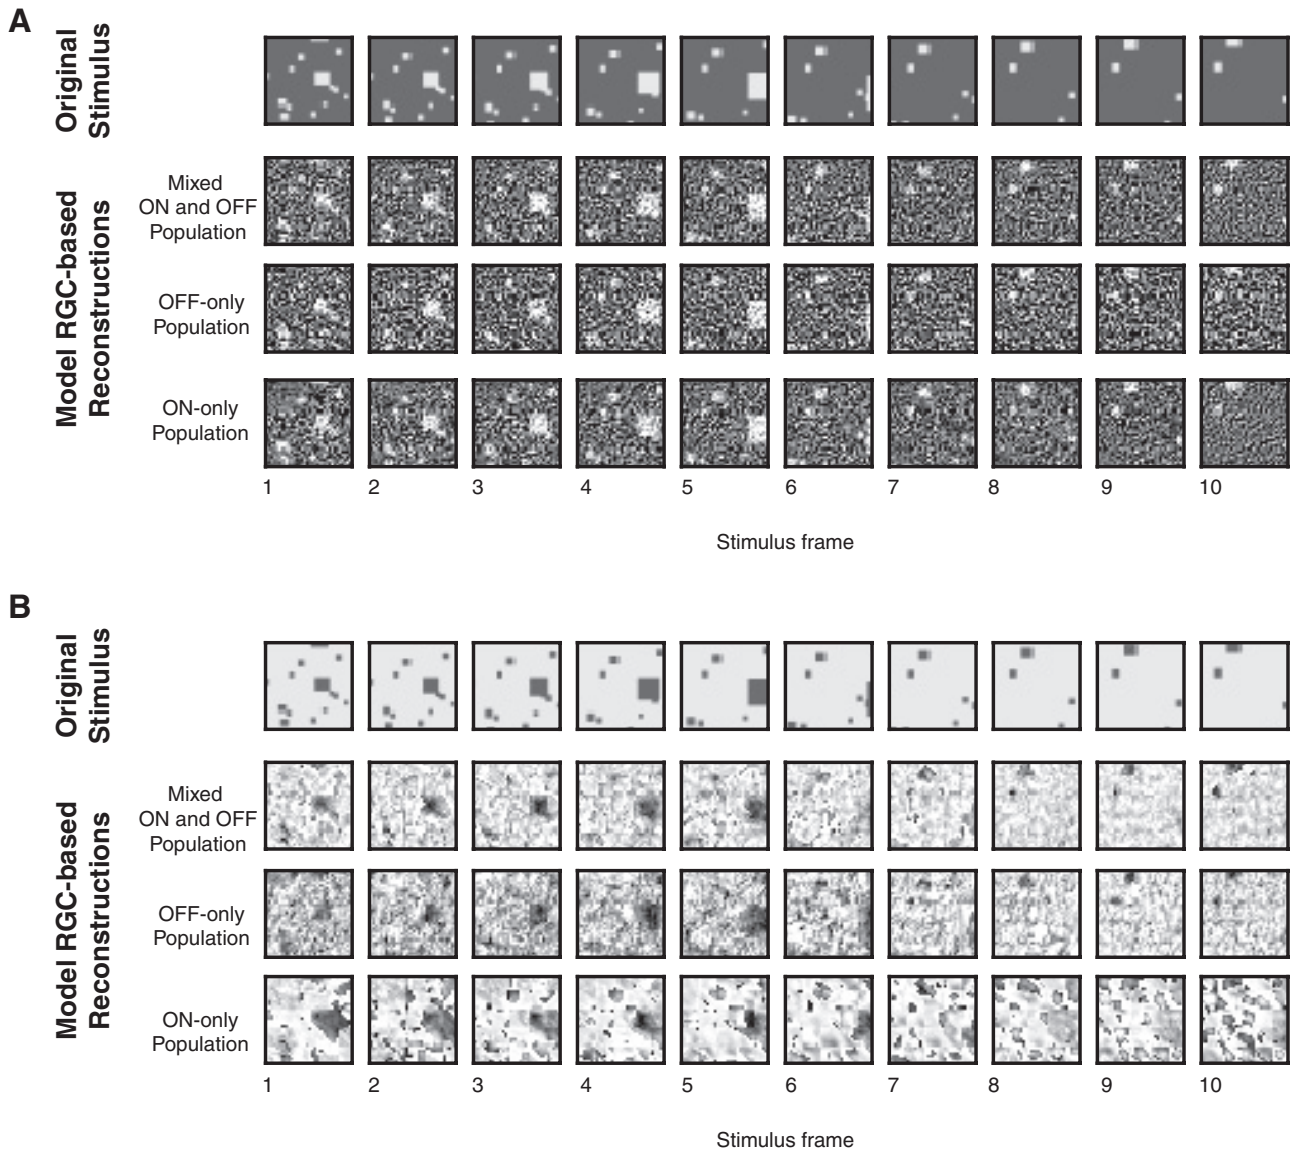


Figure 1. Examples of reconstructed spatiotemporal stimuli from model population responses. Each stimulus contains a set of objects moving toward the viewer and covers a visual angle of $\sim 120^\circ \times 120^\circ$. Reconstructions were performed with a ganglion cell density matching the mouse retina ($\sim 20,000$ model cells for a stimulus this size). **A**, The reconstructions of a stimulus consisting of moving light objects are largely similar for virtual retinas consisting of both ON and OFF RGCs (second row), only OFF RGCs (third row), and only ON RGCs (bottom row). **B**, The reconstruction for a stimulus consisting of moving dark objects is markedly altered when only ON RGCs are present: compared to the reconstruction based on the mixed population (second row) or the OFF RGCs only (third row), the reconstruction based on a population of only ON RGCs (bottom row) has several spurious objects.

sion matrices were ordered so that visualization of motion information was straightforward: left-moving stimuli are indexed in the first half of each row and each column, and right-moving stimuli are indexed in the second half. Entries off the diagonal indicate that one stimulus is confused for another, but if the correct motion direction can nevertheless be correctly decoded, this confusion will be confined to the top-left and bottom-right quadrants. Thus, the confusion matrices serve to quantify task performance, and also distinguish between performance that is driven by overall motion information (the quantity under study) and performance that is driven by the idiosyncratic position of individual objects in each stimulus (a potential confound): in the former case, there will be within-class confusion, and tallies will be spread across the top-left and bottom-right quadrants; in the latter, tallies will be confined to the main diagonal.

Example confusion matrices for single ON and OFF RGC populations in virtual experiments (using the same model pa-

rameters as the reconstructions) are shown in Figure 2*B*. Here, confusion matrices were computed using spike trains simulated from randomly chosen model populations of 20 ON RGCs or 20 OFF RGCs for a library of dark-object CM stimuli at difficulty level 4. Consistent with our hypothesis, the decoding errors for OFF cells are primarily confined to the top left and bottom right quadrants, while for ON cells, decoding errors are scattered throughout the matrix, covering all four quadrants; that is, for OFF cells but not for ON cells, directional information is preserved. Most correct tallies are not confined to the main diagonal (the percentage of tallies on the diagonal ranges from 5 to 25% across the confusion matrices), indicating that performance is indeed driven by overall motion direction, not the idiosyncrasies of specific stimulus examples.

Figure 2*C* shows the parallel results for recorded populations. The key features of the confusion matrices for the simulated re-

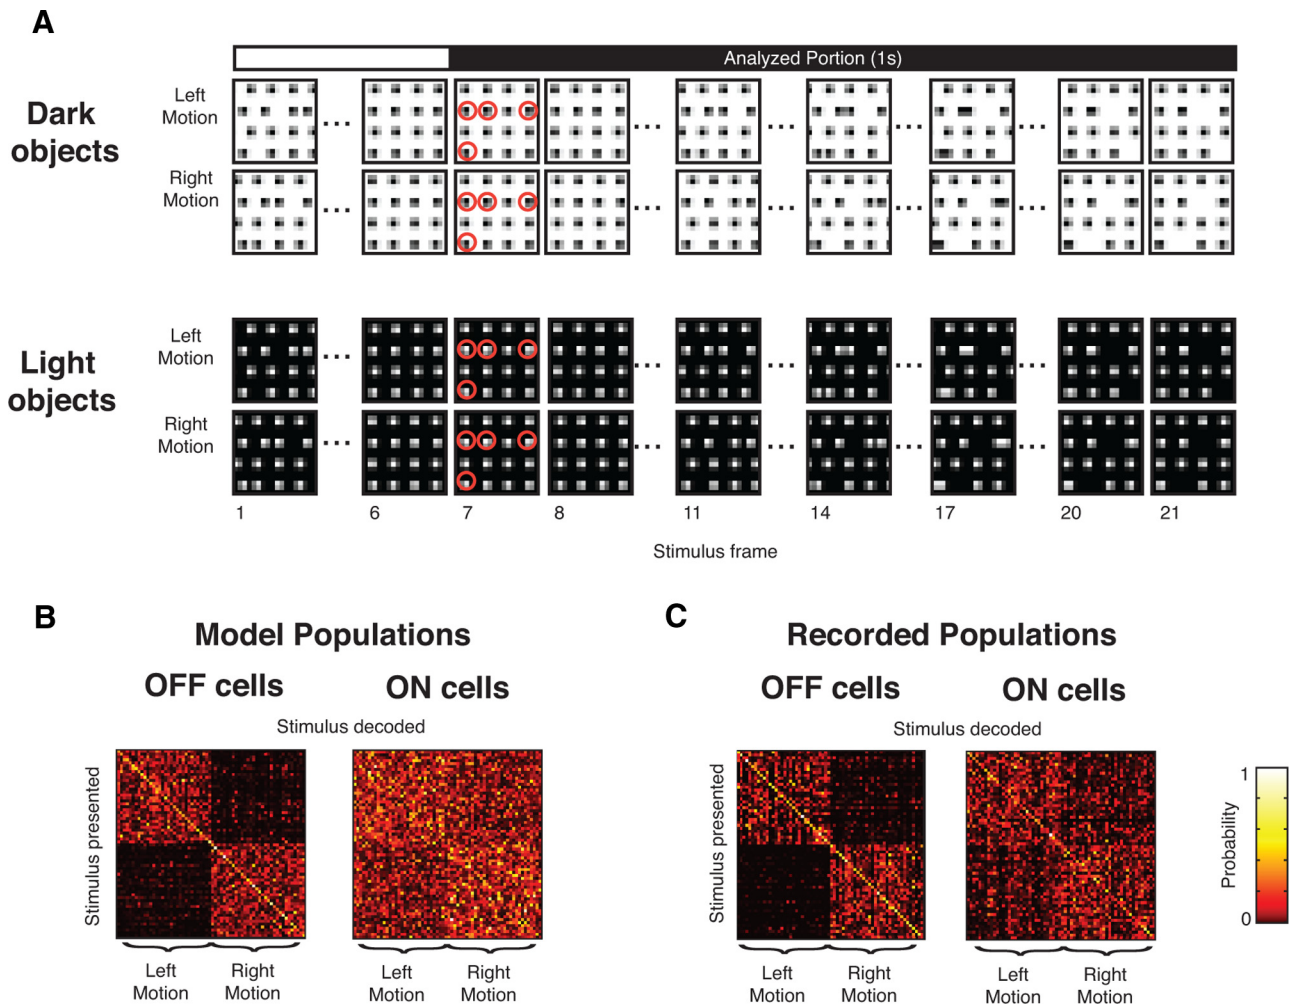


Figure 2. Responses to the CM stimuli reveal differences in the information carried by ON and OFF cells. **A**, Examples of the CM stimulus for dark objects (top) and light objects (bottom). Each stimulus is a movie lasting 21 frames (1.4 s), displaying a grid of 16 objects. A subset of the objects move either left or right; the remaining objects are stationary distractors (highlighted in red, frame 7). Only responses obtained during the final 1 s (the final 15 frames) are analyzed, to avoid transients due to stimulus onset and carryover from the previous stimulus (see Materials and Methods). **B**, **C**, Confusion matrices built from responses of populations of 20 OFF (left matrix) and ON (right matrix) RGCs to a library of dark-object CM stimuli with four stationary objects. For populations of model cells (**B**), OFF cells are able to transmit motion information (most of the tallies in the confusion matrices are confined to two quadrants), whereas ON cells are not (the tallies are scattered throughout the confusion matrices). The same pattern is seen in the populations of RGCs for which responses were directly recorded (**C**).

sponses are retained. Specifically, even though decoding the recorded OFF cell population responses often misidentifies the specific stimulus exemplar, directional information is preserved, as these errors are confined to the top left and bottom right quadrants. In contrast, the decoding errors for ON cells are common in all quadrants, indicating frequent confusion of the stimulus direction.

Figure 3 extends this analysis to all levels of task difficulty, and to the CM task for light objects. For each condition, performance is summarized by the fraction correct, i.e., the fraction of tallies in the confusion matrix that are located in the correct quadrant. Figure 3A shows the results for simulated virtual ON and OFF populations used to design the experiment. Error bars show the range of values encountered across 20 different simulations, each drawing 20 ON and OFF model cells with receptive fields placed in random locations on the stimulus (see Materials and Methods). For the dark-object stimuli, the model ON cell populations display a deficit relative to OFF cell populations, while for light-object stimuli, both ON and OFF populations perform well. This shows that this asymmetry appears to be robust, and not due to the idiosyncrasies of a particular spatial arrangement of cells.

The performance of the laboratory-recorded RGC populations at the motion decoding task is presented in Figure 3C. In parallel with the results in Figure 3A, the OFF populations outperform the ON populations at the dark-object motion decoding task, but not the light-object decoding task, across many levels of task difficulty. To establish significance (Fig. 3C, asterisks), *p* values were calculated with a one-tailed permutation test (Moore and McCabe, 2006), where the measured decoding performance difference was compared to the distribution of differences that were obtained by randomly assigning cells to either class (ON or OFF).

To determine whether the modest differences in Figure 3 between **A** and **C** (the convergence of the performance curves in **C** for dark objects and the overall more jagged appearance) are consistent with the effect of sampling populations from a finite pool of cells, we used a bootstrap technique. Specifically, we repeated the virtual experiment, this time drawing from model neurons in a finite pool, comparable to the size of a recorded data set. Model parameters used in this virtual experiment were drawn from the same cell recordings as the data in Figure 3C. The results of 25 such virtual experiments are shown in Figure 3B: each trace

corresponds to a simulated experiment in which 50 model cells of a given class were randomly chosen, and the decoding analysis was performed over draws (each without replacement) of 20 cells from this finite pool. Again it is clear that, on average (though individual traces vary from each other), OFF cell populations perform better than ON populations at the dark-object motion task, while both populations perform well at the light-object motion task. Moreover, the behavior in the original virtual retina analysis (Fig. 3A) was well within the range of the bootstrapped curves (Fig. 3B).

To ensure that details of the decoding measurements did not influence the result, we repeated the calculations in Figure 3C in several other ways, specifically, using smaller population sizes (Fig. 4A), using smaller bin sizes (Fig. 4B), and omitting the on-diagonal elements (Fig. 4C). In each case, the result remained the same: OFF populations were significantly better than ON cells at signaling moving dark objects, whereas there was no significant difference between ON and OFF cells for the light objects. We also performed the decoding under conditions that do not assume Poisson firing rates, using a multinomial decoding model, dividing rates into quartiles, and found that the results remained the same (OFF populations outperformed ON populations for four of five conditions on the dark-object task; no significant difference for the light-object task; data not shown).

To also ensure that speed and contrast of the stimulus were not critical for the observed asymmetries, we conducted virtual experiments using faster (140%) and slower (50%) speeds (but with the same stimulus duration; Fig. 4E), as well as lower (50%) contrast (Fig. 4F). In each case, the virtual experiments yielded the same asymmetries as they did for the original stimulus (Fig. 4D).

The calculations shown in Figure 4 also allow us to assess the possibility that the similarity between ON and OFF population performance for the light object condition is due to a saturation of performance (e.g., due to the differences in Weber contrast). For all analyses shown in Figure 4 (except for B, top), the performance is well below saturation, yet the similarity between the ON and OFF populations for decoding the light-object stimuli remains.

Task performance is driven by the population activity of non-DS cells

Since this motion decoding task has a strong directional component, it is natural to ask whether DS ganglion cells are playing a significant role in these results, or alternatively that the directional signal in this task is carried by the population activity of non-DS responses. We found that the latter is the case. To make this assessment, we first identified the cells that displayed direction-selective responses to the

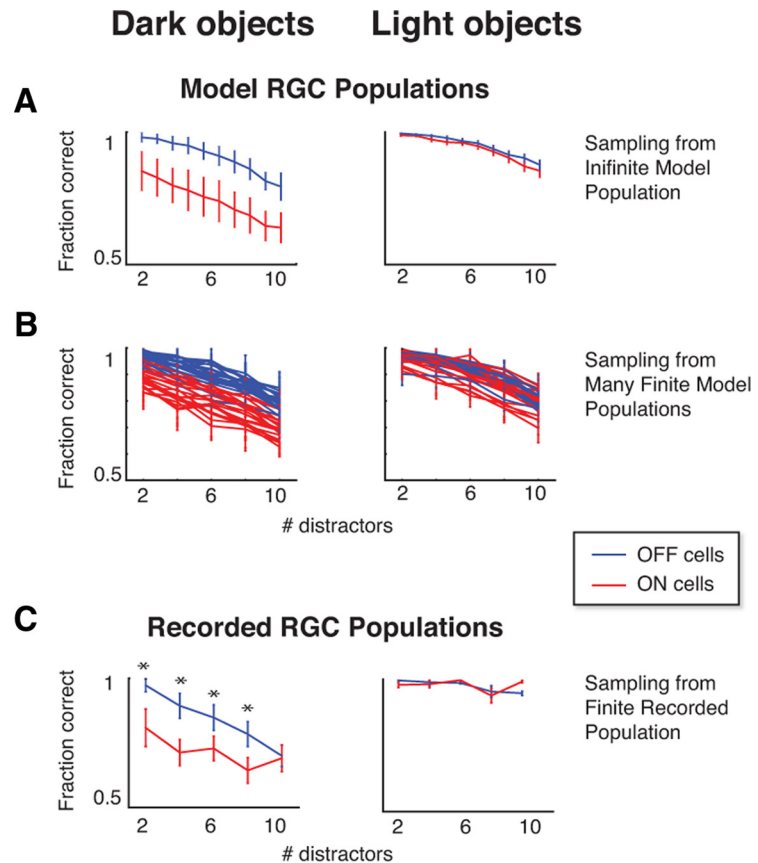


Figure 3. ON RGC populations discriminate motion direction poorly for dark objects, while OFF RGC populations perform well for both light and dark objects. Each panel shows the performance (fraction correct) of populations of 20 ON or 20 OFF RGCs in signaling the direction of coherent motion, across a range of task difficulties (the number of stationary distractors). In **A**, the populations are generated by drawing 20 neurons, without replacement, from an infinite pool of model neurons. Error bars indicate 1 SD, as estimated from 20 draws for each kind of population (see Results). In **B**, each curve represents analysis of a single simulated experiment in which populations were sampled (without replacement) from a finite pool of 50 ON and 50 OFF model RGCs. Error bars indicate 1 SD across the 20 populations in each simulated experiment. This indicates that recordings of 50 ON and 50 OFF cells should provide an adequate data set to identify a difference in performance in the CM task. **C** shows the decoding analysis of a population of electrophysiological recordings of ON and OFF RGCs roughly matching the pool sizes in **B**, demonstrating an ON cell deficit in the dark-object CM task and no difference between ON and OFF populations for the light-object CM task. Asterisks indicate performance differences with $p < 0.05$ (by a permutation test, see Results). Error bars indicate 1 SD across 20 resamplings of the populations used for decoding.

CM stimulus in our populations (DSI > 0.3 with 95% confidence; see Materials and Methods). Example responses from DS and non-DS cells are shown in Figure 5A; using this criterion, 28 of 359 ON and OFF cells were identified as directionally selective. (Note that this does not identify all DS cells, only those whose DS asymmetry could contribute to the task, because it is manifest at the orientation of the motion path that occurs in the CM stimulus.)

We next tested whether removal of these cells from the populations altered our results. As is shown in Figure 5B (confusion matrices for the same level of task difficulty as in Fig. 2B, C) and Figure 5C (a summary of performance across all levels of task difficulty), removal of DS cells had little overall effect and did not change the population coding asymmetries between ON and OFF cells. Since there is a subtle difference when DS cells are excluded (lowering the performance of ON and OFF populations for the light object stimuli; compare right panels of Figs. 3C, 5C), we cannot rule out a small contribution of DS activity to the decoding. However, these results show that the primary source of information for this task is distributed across the population of non-DS cells, rather than contained in the activity of the small

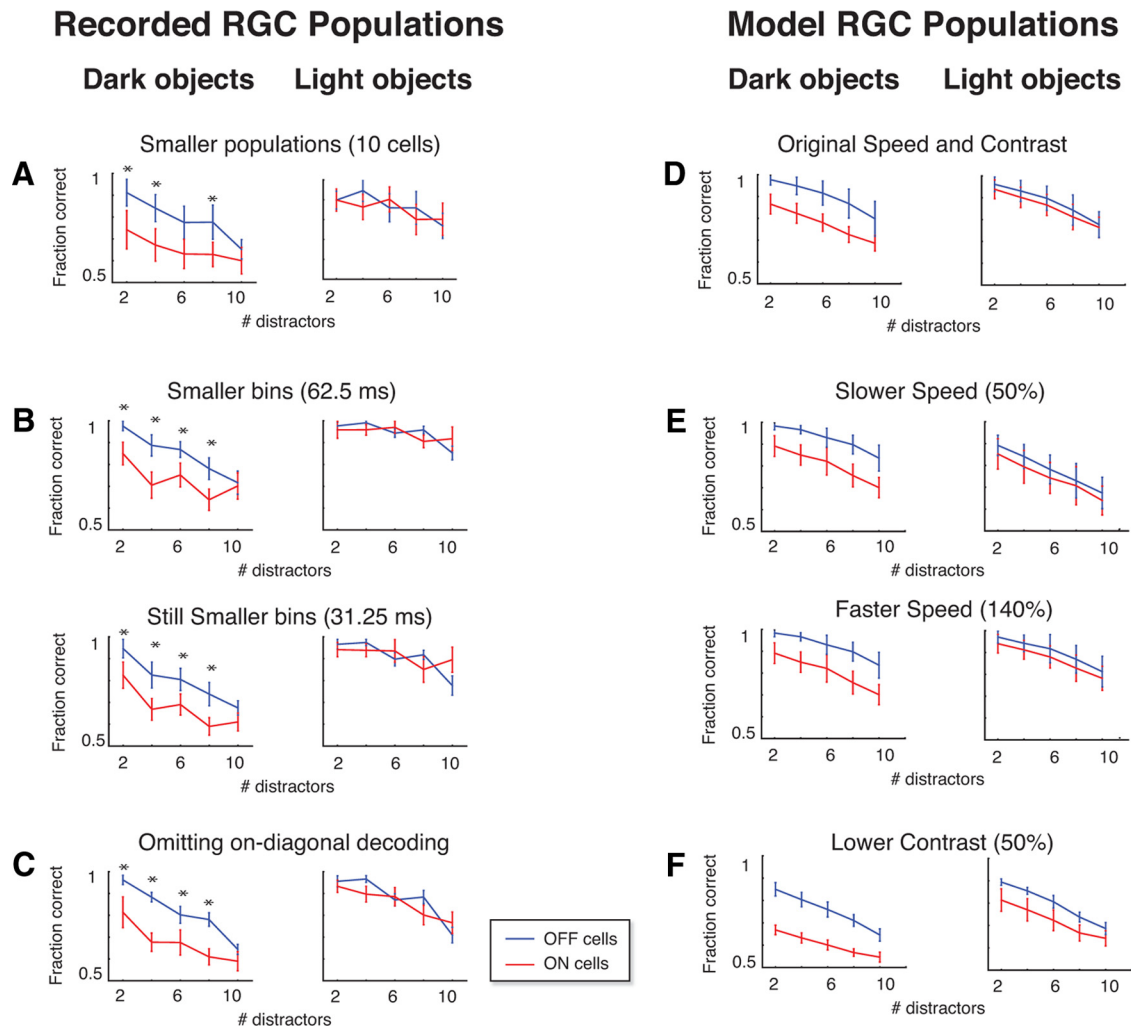


Figure 4. Coding asymmetries between ON and OFF populations are robust to details of the stimulus and of the decoding measurements. **A–C**, Decoding the recorded RGC responses in several different ways yields the same asymmetries as the main results (compare to Fig. 3C). **A**, Decoding with fewer cells in each population (10 vs the original 20). **B**, Decoding with finer bins (62.5 and 31.25 ms). **C**, Decoding while excluding on-diagonal elements from the confusion matrix (to rule out results driven by idiosyncrasies of the stimulus). For **A–C**, as in Figure 3C, asterisks indicate performance differences with $p < 0.05$ (by a permutation test; see Results), and error bars indicate 1 SD across 20 resamplings of the populations used for decoding. **D–F**, Virtual experiments reproduce the same asymmetries across different stimulus speeds and contrasts. **D**, Average model population performance for the original CM stimuli (averages taken across the traces in Fig. 3B). **E**, Average performance for faster (140%) and slower (50%) speed stimuli. **F**, Average performance for lower (50%) contrast. For **D–F**, error bars indicate 1 SD across 25 simulated experiments.

subset of cells that are directionally selective, and the DS cells do not contribute to the asymmetry in performance between ON and OFF populations for the dark-object CM stimuli.

The slight loss of performance when DS cells are removed again shows that the similar performance of ON and OFF cells for the light-object CM stimulus is not merely because performance is at saturation for both cell classes.

Differences in performance between ON and OFF cell populations are related to an interaction between linear and nonlinear characteristics

Having ruled out direction selectivity of individual cells as an explanatory factor, we next focused on which physiological differences between ON and OFF RGCs at the population level could be contributing to the result.

First, we characterized several properties of the cells in our data set that are known to vary between ON and OFF RGCs. These include peak firing rate, latency of response, duration of response, receptive field size, nonlinearity of the response, and the degree to which responses were biphasic (for how these are

calculated, see Materials and Methods). Figure 6, A and B, shows how the physiological measures differ between ON and OFF cells at the level of individual neurons: Figure 6A for the entire population, and Figure 6B for the cells from a single cohort of retinas (in this case, dark-object CM experiments with 2, 6, and 10 stationary distractors). Figure 6C shows, not surprisingly, that when populations of 20 cells are randomly drawn from these cohorts, the mean values of the parameters show the same kinds of asymmetries as do the individual neurons.

The asymmetries mean that while individual ON and OFF RGCs cover a broad distribution of values, and these distributions have substantial overlap, many of the differences between the two populations are significant (for a statistical summary of the data in Fig. 6A, see Table 1). For several physiological parameters (including receptive field size, peak firing rate, and the degree of nonlinearity), these differences are consistent with previous reports in a wide range of species (Chichilnisky and Kalmar, 2002; Carcieri et al., 2003; Zaghoul et al., 2003; Balasubramanian and Sterling, 2009; Molnar et al., 2009). Measurements relating to temporal kinetics (response latency, response duration, and bi-

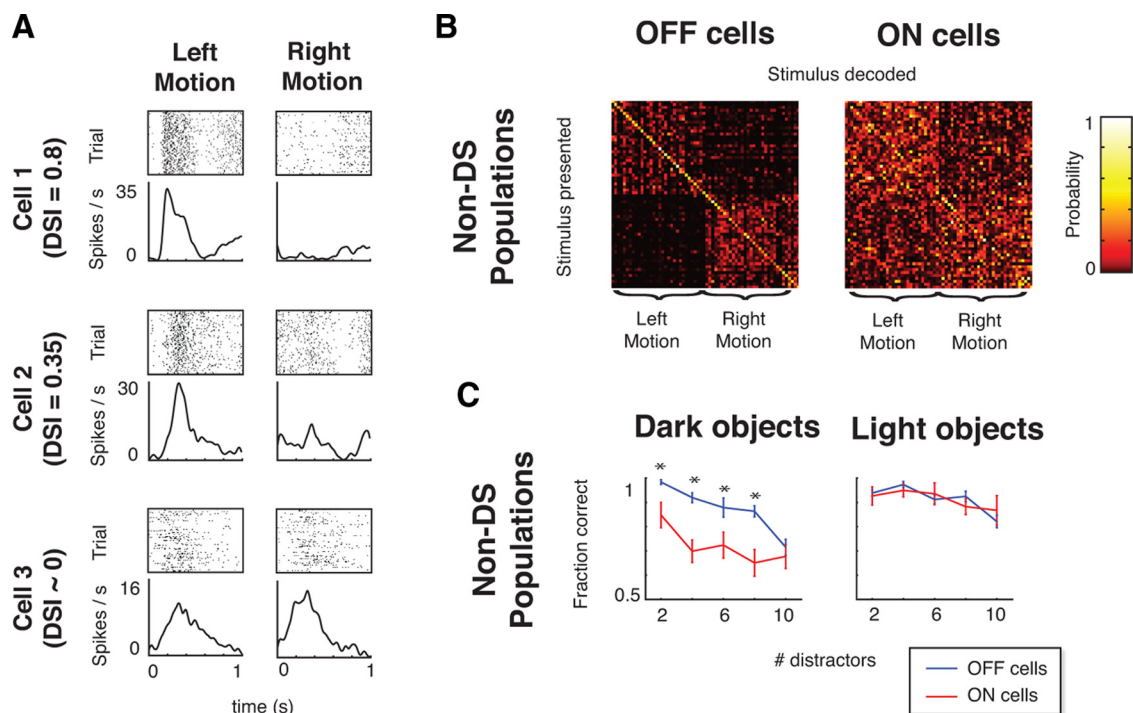


Figure 5. Performance in the CM task does not rely on DS responses of individual cells. **A**, Rasters (80 trials) and smoothed firing rates for example cells that are highly DS (top), mildly DS (middle), and non-DS (bottom) for leftward versus rightward motion. Data cover the interval over which direction selectivity was calculated, i.e., the period during which a single moving object passed over the cell's receptive field. **B**, Confusion matrices calculated for populations of non-DS ($DSI < 0.3$) RGCs (for dark CM stimuli with 4 stationary distractors) are similar to those calculated from the complete populations (compare Fig. 2C): OFF non-DS populations were able to signal motion direction, whereas ON non-DS cells were not. **C**, Non-DS population performance matched the complete population performance across task difficulties: ON non-DS populations performed poorly in the dark-object CM task, whereas OFF non-DS populations performed well in the light- and dark-object CM tasks (compare Fig. 3C). Asterisks (significance) and error bars were calculated as in Figure 3C.

phasic index) appear to be less consistent across species however: while the ON pathway was found to have faster responses than the OFF pathway in primate ganglion cells (Chichilnisky and Kalmar, 2002), the opposite was reported in salamander bipolar cells (Burkhardt, 2011), salamander ganglion cells (Burkhardt et al., 1998; Gollisch and Meister, 2008), and cat LGN cells (Jin et al., 2011b). We find that OFF RGCs have a shorter response latency and duration than ON RGCs in our data set, and that OFF cells have a more biphasic response than ON cells. We also find that the overall temporal kinetics and latency of RGC responses in our data set (consistent with previous reports in mouse; Carciari et al., 2003) are similar to the latencies also reported in salamander (e.g., Burkhardt et al., 1998), and are longer than latencies found in primate (Chichilnisky and Kalmar, 2002) or cat (Kuffler, 1953).

To determine whether our findings could be ascribed to one of these asymmetries, we repeated the decoding procedure using only populations of cells from a single cohort for which these physiological properties are matched. The strategy is illustrated in Figure 7A. The top row reproduces the distributions of each index's population means within the cohort, as shown in Figure 6C, and highlights the area in which the distributions overlap. For each index, we then chose 20-neuron populations randomly from this overlapping region, using a rejection sampling algorithm (Press et al., 2007). The bottom two rows of Figure 7A show the mean index values across ON and OFF populations drawn using the rejection sampling, confirming that this selection procedure yielded populations that were well matched, despite the mismatch in the original data set from which they were drawn.

Figure 7B shows how matching each physiological parameter affects decoding performance for the dark-object CM task. As can

be seen, matching populations for some properties (e.g., peak firing rate and nonlinearity index) reduces the decoding performance disparity between ON and OFF RGCs somewhat. However, in all cases, a substantial disparity remained—even when selecting matched populations required sampling from the extremes of the distribution. We also considered two other variants of the nonlinearity index for matching the populations, and these led to the same conclusion (data not shown): the nonlinearity index of Chichilnisky and Kalmar (2002), based on fitting exponential rather than spline models to cells, led to the same reduction in the decoding disparity as the nonlinearity index shown in Figure 7B, and an index based on the cumulative change in slope of the spline nonlinearity led to a slightly smaller reduction.

Since the response duration parameter (Fig. 7A, first column) can be used to subdivide ON cells into sustained and transient subgroups (Carciari et al., 2003), we performed an additional check for this parameter, separately analyzing the ON sustained (response duration >200 ms) and transient (response duration <200 ms) subgroups. For both subgroups, the asymmetry in decoding performance between ON and OFF populations remained.

Finally, we performed a parallel rejection-sampling analysis for the light-object CM task (for which there was no difference in decoding performance of the original ON and OFF cell populations). Not surprisingly, matching populations for each of the parameters did not change this finding (data not shown). Thus, no individual property studied here was responsible for the disparity in decoding performance between ON and OFF RGC populations.

Perhaps it is not so surprising that none of these individual properties account for the disparity, as these are only indirectly

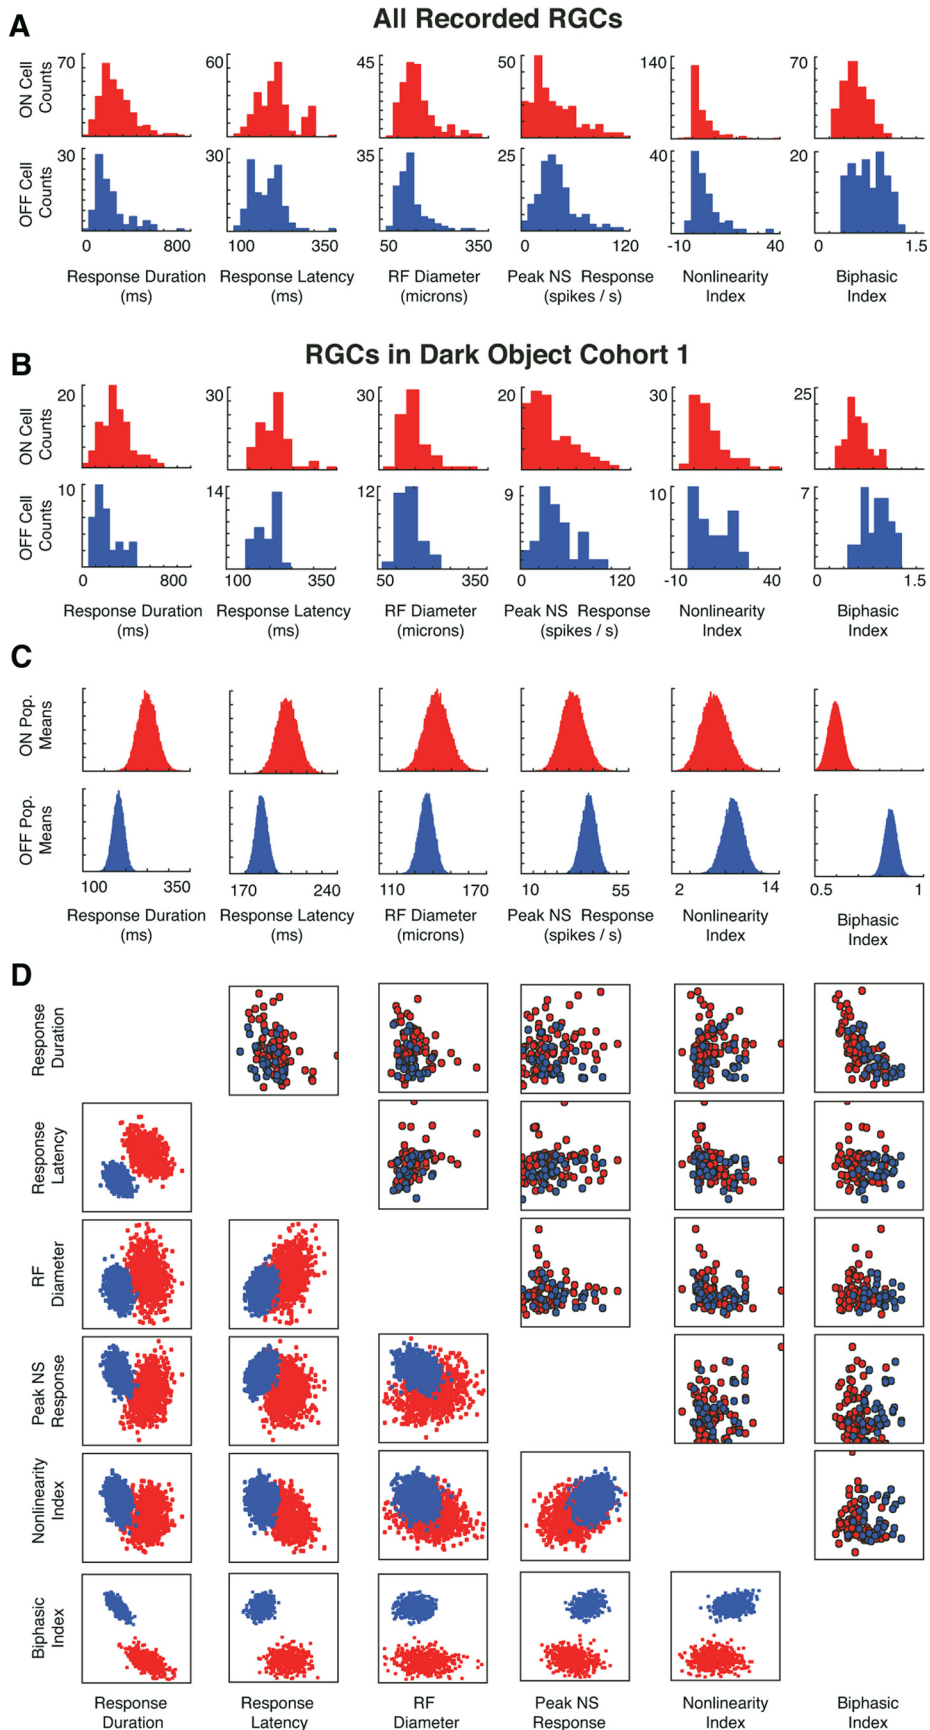


Figure 6. Asymmetries in physiological properties between ON and OFF RGCs. **A**, Distribution of six indices of physiological properties (response duration, response latency, receptive field diameter, peak firing rate, degree of nonlinearity, and biphasic index; see Materials and Methods for details) across all recorded ON ($N = 238$) and OFF ($N = 140$) RGCs. (*Figure legend continues.*)

Table 1. Summary of statistics for physiological characteristics shown in Figure 6A

	Response duration (ms)	Peak response latency (ms)	RF diameter (μm)	Peak firing rate (spikes/s)	Nonlinearity index	Biphasic index
ON ($N = 238$)						
Mean	256	208	148	36	4.5	0.61
Median	231	206	141	28	2.6	0.54
SD	126	34	41	27	5.6	0.41
Range	90–505	148–291	100–229	3–89	–0.2–15.5	0.32–0.91
OFF ($N = 121$)						
Mean	210	190	137	39	5.9	0.76
Median	185	157	134	35	4.5	0.73
SD	120	35	27	22	6.6	0.32
Range	85–472	143–235	101–196	11–84	–0.7–18.6	0.40–1.10
	***	***	**		*	***

For each parameter, the table contains the mean, median, SD, and 90th percentile range (5 to 95%) for ON (top) and OFF (middle) cells. For most parameters, there is extensive overlap in the distribution of values among individual ON and OFF cells, but the population mean values differ significantly for all parameters except peak firing rate (bottom).

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$ (unpaired t test, two tailed).

related to information transmission. The ability of a neuron to transmit information in a given context depends, fundamentally, on the range of firing rates that it can produce and how reliably they can be distinguished. This in turn depends on how the spatial properties, temporal properties, and nonlinear behavior of a neuron interact. Since the rejection sampling approach did not allow us to examine the interaction of parameters (because it would require selecting populations that were matched for two or more parameters simultaneously, and there was often very little overlap, as can be seen in Fig. 6D), we used a more integrated approach instead.

The approach we took was to use the LN model as a convenient way to dissect factors related to stimulus encoding and trace their interactions. We first examined how these interactions determine the range of firing rates produced by a cell, and next took into account how reliably the firing rates can be distinguished.

The factors contributing to the range of firing rates are shown for several example neurons in Figure 8. The linear filters for two ON cells (left) and two OFF cells (right) are shown in Figure 8A, and the combined effect of these filters and the ensuing static nonlinearity on signaling are shown in B. In each case, the first column shows the distribution of the signal output of the model linear filter for two stimuli (the light and dark CM stimuli). The second column shows the shape of the model spline nonlinearity. The abscissa corresponds to inputs to the nonlinearity (which is the signal that is the output of the linear filter). The entire axis covers the range of filtered signal values under NS conditions, since of the four stimuli used in this study (NS, WN, dark-object CM, and light-object CM), the range produced by NS stimuli was the greatest. The narrower range sampled by the CM stimuli is shown by the thickness of the line in the middle column: the thicker portion of the operating curve corresponds to the 5th to 95th percentile range of CM stimuli, and the thinner portion corresponds to the same percentile range for the NS stimulus.

Critically, the range of firing rates in response to a stimulus

←

(Figure legend continued.) **B**, Distribution of these indices for the cohort of cells tested with the dark-object CM stimuli at difficulties 2, 6, and 10; the ON and OFF RGCs display the same asymmetries as the entire population. **C**, Distribution of the mean values of the indices from populations of 20 cells drawn randomly from the cohort in B; as expected, the distributions show the same kinds of asymmetries as in A and B, but are more nearly Gaussian. **D**, Covariation of the indices for individual cells (top right) and for the means across random 20-cell populations (bottom left).

depends on how much of the operating curve of the nonlinearity is occupied by the linear filter signals, and where this range lies in relation to the shape of its operating curve. In some cases (the dark CM stimulus for Neuron 1 and both stimuli CM stimuli for Neuron 2), a substantial portion of the range sampled by the linear filter signals occupies a section of the operating curve that is flat. In other cases (both CM stimuli for Neurons 3 and 4), this range mostly occupies the rising portion of the operating curve. When a portion of the linear filter signal occupies a flat section of the operating curve, many input values are mapped to a single output value, and the range of output firing rates is reduced. In contrast, no such reduction occurs when the linear signal values are confined to the rising portion of the operating curve. In this way, the range of firing rates produced by a cell reflects an interaction between the linear filter and the shape and position of the nonlinearity's operating curve. The third column of the examples in Figure 8B shows the result of this interaction: the distribution of firing rates that emerge from the nonlinearity.

Because our focus is on differences in signaling motion for the light CM and dark CM conditions, the key point is whether these ranges differ for the two CM stimuli. Note that such a difference can only arise as a result of a nonlinearity: if the operating curve were a straight line, then light and dark CM stimuli would produce output ranges that have the same breadth (since the linear signal distributions shown in the first column are always mirror images of each other for the two stimuli of opposite contrast). However, the linear filter's characteristics are nevertheless crucial for determining the nature of this difference, because the linear filter determines the portion of the operating curve (e.g., flat portion vs rising portion) that is relevant for each stimulus type.

Across the cell populations, we find that this interaction between linear filter signals and the shape of the nonlinearity accounts for the overall difference in the way that ON and OFF cells carry information about light and dark CM stimuli (Fig. 8C,D). The first columns show that the range of the linear filter signals for light and dark CM stimuli are similar between the two cell classes. However (middle columns), a difference emerges when the effects of the nonlinearity are taken into account, because these ranges occupy different portions of the nonlinear operating curve. For dark stimuli, OFF cells have a greater firing rate range than ON cells (Fig. 8C), whereas for light stimuli (Fig. 8D), they are approximately similar.

This observed difference continues to hold when we take a noise estimate into account. Noise is given by the square root of the variance, which, under the Poisson assumption, is proportional to the mean firing rate. Thus, we use the range of firing rates divided by the square root of the mean firing rate as an estimate of firing rate signal-to-noise ratio (SNR). When we consider this measure of SNR, we find the same pattern as was seen for the firing rate range: OFF cells provide more signaling capacity than ON cells for the dark-object CM stimuli (Fig. 8C, last column; $p < 0.001$, unpaired t test, two tailed), whereas ON and OFF cells provide similar signaling capacity for the light-object CM stimuli (Fig. 8D, last column; $p \approx 0.3$). This overall pattern also holds if other percentile ranges are chosen as a measure of signal range (data not shown).

Finally, to confirm that the asymmetries in performance of ON and OFF populations are related to this measurement of the firing rate SNR, we applied the rejection-sampling analysis of Figure 7 to the SNR measures of Figure 8, C and D. Note that unlike the intrinsic neural properties included in Figures 6 and 7, the SNR measures are stimulus-specific and are calculated separately for the light- and dark-object CM stimuli. As mentioned

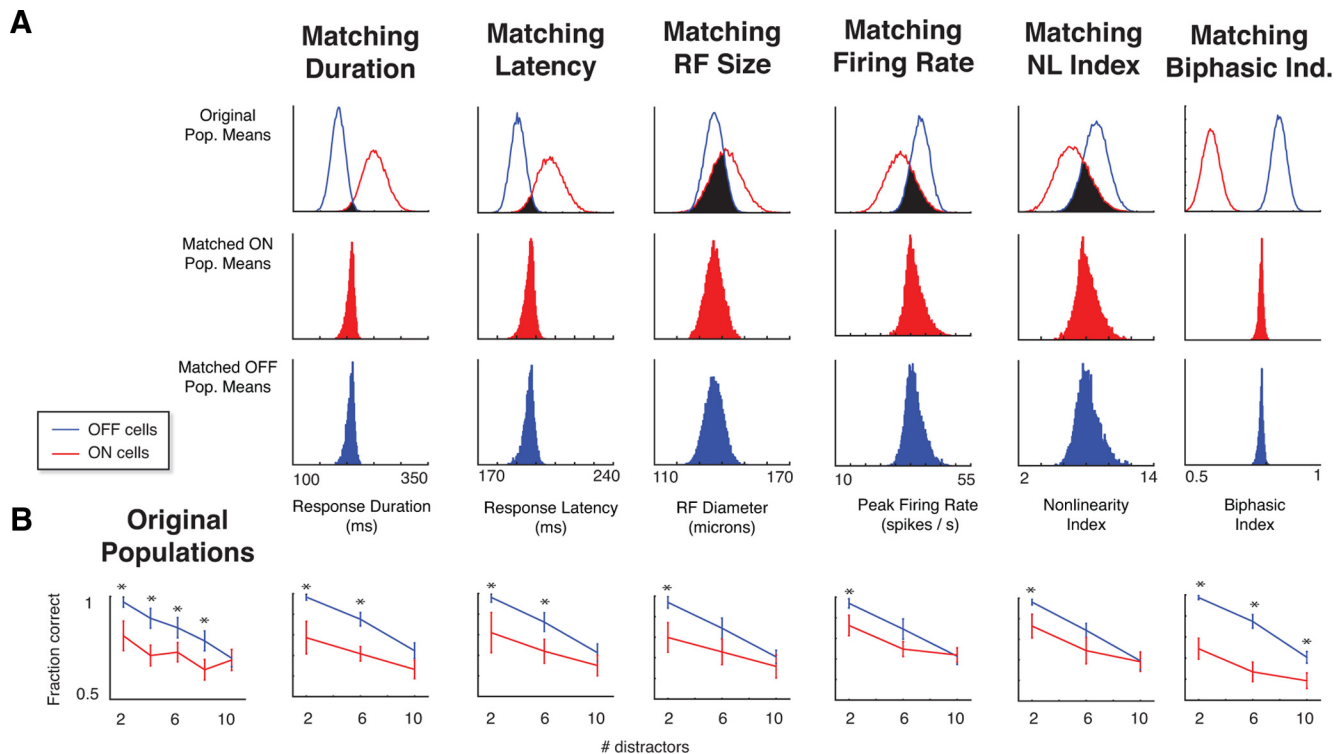


Figure 7. Differences in decoding performance between ON and OFF populations for the dark-object CM task persist when populations are equated for individual physiological parameters. **A**, The rejection-sampling procedure for selecting equated populations. For each index, the top row shows the distribution of the population of means of randomly selected 20-cell ON and OFF populations (as in Fig. 6C) in the cohort of cells tested with the dark-object CM stimuli at difficulty levels 2, 6, and 10. Populations that lie in the overlap (black region) are then selected for decoding analysis. The distributions of mean parameters across random populations selected from this region are shown in the second row (for the matched ON populations) and third row (for the matched OFF populations). **B**, Performance of the matched populations of **A** in the dark-object CM task decoding task. In each case, the ON cell populations continued to display deficits relative to OFF cell populations, meaning none of these parameters taken individually explain the performance gap. Asterisks (significance) and error bars were calculated as in Figure 3C.

above, unselected ON and OFF cells have a different distribution for the SNR index when calculated for dark-object CM stimuli (Fig. 8C). When populations of ON and OFF RGCs are selected so that they match according to this parameter (Fig. 9A), the performance difference between them is eliminated (Fig. 9B). For the light-object CM stimuli, the SNR distributions were similar for individual ON and OFF cells (Fig. 8D). Correspondingly, the population distributions overlapped heavily (Fig. 9C), and, as expected, choosing populations for which these indices were matched does not change performance (Fig. 9D). These results confirm that the above SNR measures, which reflect an interaction of linear and nonlinear response components with the stimulus, account for the essential difference between ON and OFF population that drives the performance asymmetries.

Discussion

Here, using the early visual system as a model of parallel processing, we identified a novel aspect of the role the different cell classes play in carrying information, and showed that it is explained by an interaction between the cells' linear and nonlinear properties. Previous results have focused on ways that linear (Sagdullaev and McCall, 2005; Balasubramanian and Sterling, 2009; Pandarinath et al., 2010) and nonlinear (Pitkow and Meister, 2012) properties in isolation play important roles in coding. Here we found a surprising asymmetry in population coding—that OFF cells can transmit motion information about moving light and dark objects, but that ON cells have a deficit at transmitting information about moving dark objects—and found that it reflects a difference in how the cells' linear and nonlinear properties work together, rather than how either work by themselves.

To obtain measures of these linear and nonlinear properties from physiological recordings of individual neurons, we parameterized their input–output relationships with an LNP model that faithfully captures retinal input–output relationships for a broad range of stimuli (for a general review of LNP models for retina and other systems, see Simoncelli et al., 2004; Paninski et al., 2007; Nirenberg and Pandarinath, 2012; Bomash et al., 2013). These models are phenomenological models, meaning that they serve to characterize ganglion cell response properties, even though their computational components do not specifically map onto physiological mechanisms. We can use them, as we do here, to understand how the operational characteristics of the circuitry (spatial and temporal sensitivity, nonlinearity, etc.) shape ganglion cells' specific visual processing properties, even without complete knowledge of the physiological mechanisms that underlie the operations (e.g., the specific synapses, ion channels, and wiring between cell types). Here, we showed that this approach works even when the kind of information that a cell carries cannot be ascribed to a single component of the phenomenological model, but is a result of how the components interact.

The role of models in this work

Aside from their essential role in this work (and in many other studies) of providing a parameterization of cell properties, models were used for two other crucial aspects of this study. The first of these is their role in the reconstruction experiments. The reconstructions were designed to assay the behavior of a retinal patch representing a significant fraction of the visual field, so that the number of ON and OFF RGCs required was quite large

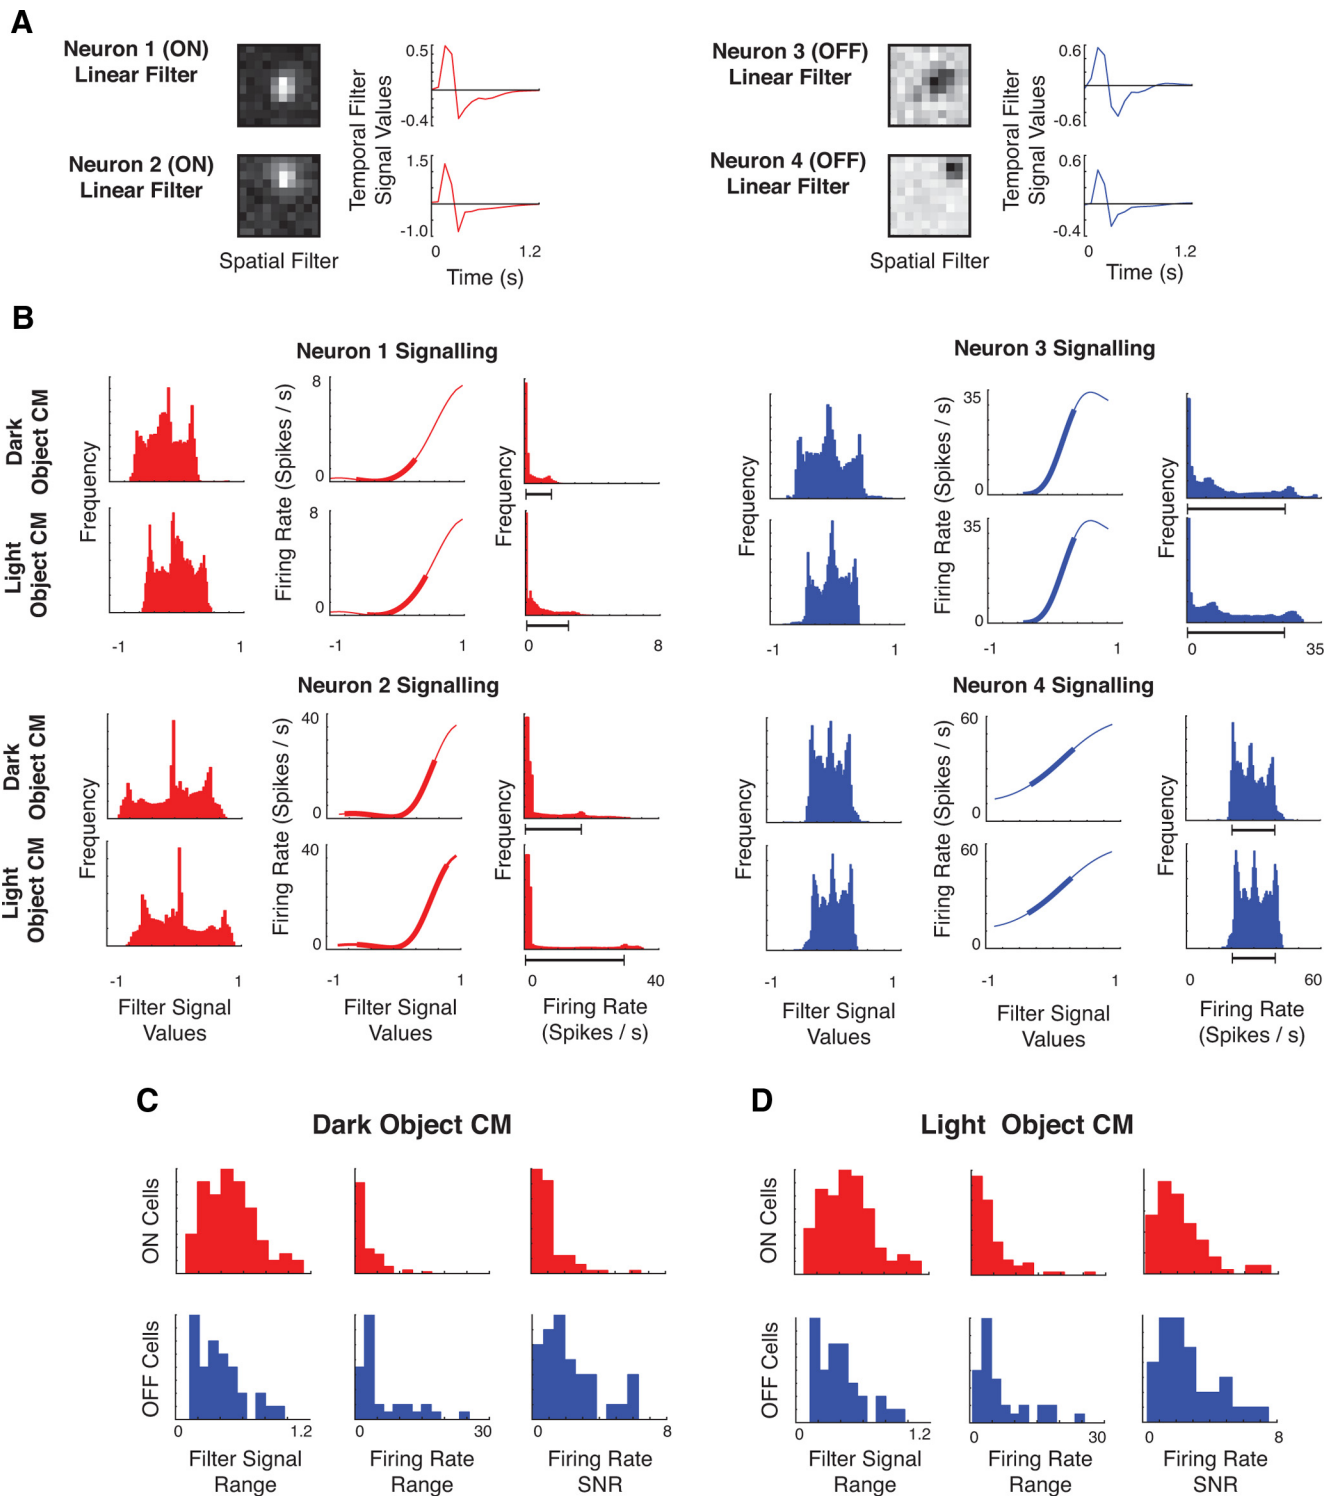


Figure 8. Differences in decoding performance between ON and OFF populations reflect an interaction between linear and nonlinear cell properties. **A** shows model linear filters (spatial and temporal components) for two ON cells (left) and two OFF cells (right). **B** shows how signaling is driven by these filters with the static nonlinearity for the four example neurons. For each of two ON neurons (left) and two OFF neurons (right), the first column shows the distribution of signals in the output of the linear filter L for the dark-object and light-object CM stimuli. The middle column shows each cell's nonlinearity N , with the heavy portion of the curve indicating the 90th percentile range of signals that are the outputs of the linear stage in the first column. The final column shows the distribution of firing rates, which correspond to the output of the nonlinearity. The firing rate range thus reflects both the shape of the nonlinearity and the range of filtered signal values that are incident on it from the linear filter. As is typical of the population, the ON cells have a narrower range of firing rates for the dark-object CM stimulus than for the light-object CM stimulus, while for the OFF cells, ranges are comparable (90th percentile ranges indicated by the dark gray bar). **C** (dark-object CM) and **D** (light-object CM) summarize this analysis across the population. For each kind of stimulus, ON and OFF cells have a similar range of signal values at the output of the linear filter (first column), but the range of firing rates (middle column) shows an asymmetry: for dark CM stimuli, it is narrower for ON cells than for OFF cells, whereas for light CM stimuli, it is broad for both. This difference (for dark objects but not for light objects) is maintained when the effects of Poisson-like variability are taken into account by a simple measure of SNR (last column).

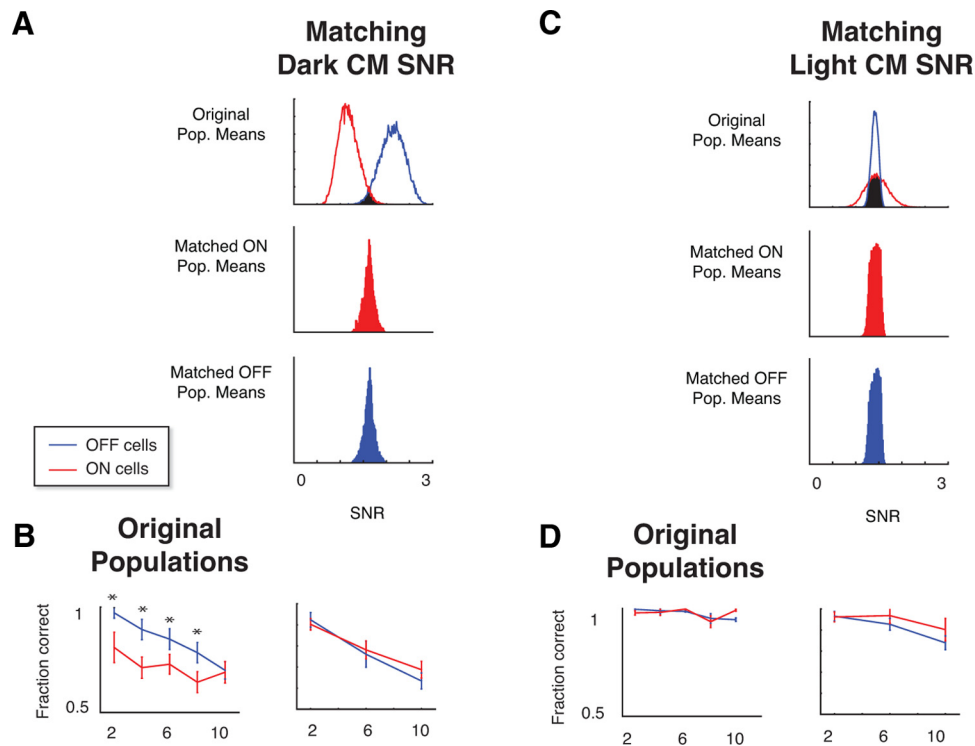


Figure 9. Differences in decoding performance between ON and OFF populations are eliminated when populations are equated for the firing-rate range SNR measure specific to the CM stimulus. **A**, The rejection-sampling procedure selects populations of ON and OFF cells for which the SNR measure on the dark-object CM stimulus is matched. **B**, Performance of the matched populations of **A** in the dark-object CM task decoding task. Unlike the measures tested in Figure 7, matching for the SNR measure eliminates decoding performance differences between ON and OFF populations. **C**, The SNR measure for the light-object CM stimulus exhibited less disparity between ON and OFF populations. **D**, Performance of the matched populations of **C** in the light-object CM task decoding task were similar to the randomly chosen populations, with no significant difference between ON and OFF cells. Asterisks (significance) and error bars were calculated as in Figure 3C.

(~20,000 in this case). Since this number of cells would be very difficult to obtain in experiments (and would involve the use of a large number of animals), we instead used models to simulate these responses. These reconstructions led us to the hypothesis that ON cells would display a selective deficit in signaling the motion of dark objects.

Additionally, models played a crucial role in designing experiments that would allow us to test the primary hypothesis about ON cells directly with recorded responses. By simulating experiments for different versions of a decoding task, we determined a design that would capture the phenomena under study (direction of motion) and that could be implemented using a quantity of cells, stimuli, and repetitions that is practical for a physiological experiment. This avoids experimental designs that would likely fail due to data limitation or lack of sensitivity.

However, we emphasize that although models played a key role in hypothesis generation and experimental design, the test of the hypothesis relied on empirical data: decoding the responses of RGCs obtained in physiological recordings.

Implications for psychophysics

Our main finding is that differences in the linear and nonlinear response components of RGC classes interact to yield asymmetries in their ability to carry information efficiently about light and dark objects. Because differences in the component properties of ON and OFF cells are common across vertebrates including primates, these differences might play a role in the ability of ON and OFF cells to enable perception of motion and other stimulus qualities in humans.

A direct psychophysical test of an RGC class's capabilities is difficult, as it would require that other cell classes be selectively

inactivated. However, a relevant natural experiment in humans does exist. A rare mutation in the GRM6 gene selectively inactivates ON bipolar cells. Individuals who carry this mutation lack the normal ON cell responses but have normal or near-normal visual acuity and, interestingly, no difference in perceiving light and dark words flashed under photopic conditions (Dryja et al., 2005). While perhaps puzzling at first, these results correlate with our findings here, as we find that OFF RGCs can reliably signal both light and dark objects thanks to interactions between the cell's linear and nonlinear characteristics. It stands to reason that (despite the overall differences we observe between mouse and primate in RGC temporal kinetics) a similar phenomenon may underlie the ability of the retinas of these individuals to signal the appearance of light and dark stimuli since stimulus appearance, like stimulus motion, leads to a transient signal that interacts with the linear and nonlinear characteristics of the cell. Note though, that our results bear specifically on the normal functioning of the retina, while the retinas of individuals studied by Dryja et al. (2005) were perturbed by mutation. Other factors related to this perturbation (the loss of the ON pathway) may therefore contribute: for example, the addition of transient ON responses (Rentería et al., 2006), changes in contrast sensitivity (Manookin et al., 2008), and changes in the baseline firing rate and degree of rectification (Zaghloul et al., 2003; Molnar et al., 2009) in the remaining OFF pathway.

Several psychophysical studies involving subjects with normal retinas have demonstrated asymmetries in light and dark perception: decrements are more easily detected than increments (Bowen et al., 1989; Kremers et al., 1993; Buchner and Baumgartner, 2007), and darks are processed more quickly than lights (Komban et al., 2011). It has been hypothesized that the physiological

basis of these observations lies in the differences between the nonlinearities in ON and OFF RGCs: that the more linear response of ON cells allows them to signal both increments and decrements, while the more rectified OFF cell responses only allow them to signal decrements (Chichilnisky and Kalmar, 2002). Based on these physiological findings, it stands to reason [as indicated in the study by Chichilnisky and Kalmar (2002)] that decrements have a perceptual advantage: decrements can be processed by either pathway, but increments are only processed by the ON pathway.

Our results indicate that the situation is more complex: a cell's nonlinearity cannot be viewed in isolation, as its effect on signaling depends on the range of input values delivered by the preceding processes; that is, the participation of a cell in a task can depend not only on its nonlinear function, but also on how its spatiotemporal sensitivity interacts with the characteristics of the stimulus. In this study, we find the same differences between ON and OFF RGC nonlinearities that have been reported previously (Chichilnisky and Kalmar, 2002; Zaghloul et al., 2003, Molnar et al., 2009), but for the moving object stimuli, the range of signals incident on them means that increments are processed by both cell classes, while decrements are processed primarily by one. Therefore, if the reported psychophysical asymmetries can be ascribed to whether one or both RGC classes are recruited in signaling, the perceptual advantage of darks over brights would reverse for stimuli of the appropriate spatiotemporal characteristics. Alternatively, if darks continue to have a perceptual advantage across the spatiotemporal gamut, one can infer that downstream factors (Jin et al., 2008, 2011a,b; Yeh et al., 2009; King et al., 2010) play a dominant role.

Relevance to other sensory systems

Our primary conclusion is that the role of a cell class in population coding is determined not by the cells' individual physiological properties, but rather by how their linear and nonlinear characteristics combine. This finding is likely to extend to other systems, as the basis of the finding was a simple interaction of the linear and nonlinear components of a phenomenological model with wide applicability. Particularly conspicuous candidates would include other systems where LNP cascade models have already been employed, such as the early auditory system (Calabrese et al., 2011), and systems that are known to have separate linear and nonlinear components relevant to stimulus encoding, such as in the electrosensory systems of some fish (Chacron, 2006).

References

- Applebury ML, Antoch MP, Baxter LC, Chun LL, Falk JD, Farhangfar F, Kage K, Krzystolik MG, Lyass LA, Robbins JT (2000) The murine cone photoreceptor: a single cone type expresses both S and M opsins with retinal spatial patterning. *Neuron* 27:513–523. [CrossRef Medline](#)
- Balasubramanian V, Sterling P (2009) Receptive fields and functional architecture in the retina. *J Physiol* 587:2753–2767. [CrossRef Medline](#)
- Bomash I, Roudi Y, Nirenberg S (2013) The virtual retina. *PLoS One* 8:e53363. [CrossRef Medline](#)
- Bowen RW, Pokorný J, Smith VC (1989) Sawtooth contrast sensitivity: decrements have the edge. *Vision Res* 29:1501–1509. [Medline](#)
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765. [Medline](#)
- Buchner A, Baumgartner N (2007) Text-background polarity affects performance irrespective of ambient illumination and colour contrast. *Ergonomics* 50:1036–1063. [CrossRef Medline](#)
- Burkhardt DA (2011) Contrast processing by ON and OFF bipolar cells. *Vis Neurosci* 28:69–75. [CrossRef Medline](#)
- Burkhardt DA, Fahey PK, Sikora M (1998) Responses of ganglion cells to contrast steps in the light-adapted retina of the tiger salamander. *Vis Neurosci* 15:219–229. [Medline](#)
- Cai D, DeAngelis GC, Freeman RD (1997) Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *J Neurophysiol* 78:1045–1061. [Medline](#)
- Calabrese A, Schumacher JW, Schneider DM, Paninski L, Woolley SM (2011) A generalized linear model for estimating spectrotemporal receptive fields from responses to natural sounds. *PLoS One* 6:e16104. [CrossRef Medline](#)
- Carcieri SM, Jacobs AL, Nirenberg S (2003) Classification of retinal ganglion cells: a statistical approach. *J Neurophysiol* 90:1704–1713. [CrossRef Medline](#)
- Chacron MJ (2006) Nonlinear information processing in a model sensory system. *J Neurophysiol* 95:2933–2946. [CrossRef Medline](#)
- Chichilnisky EJ, Kalmar RS (2002) Functional asymmetries in ON and OFF ganglion cells of primate retina. *J Neurosci* 22:2737–2747. [Medline](#)
- Dowling JE, Ripps H (1970) Visual adaptation in the retina of the skate. *J Gen Physiol* 56:491–520. [CrossRef Medline](#)
- Dryja TP, McGee TL, Berson EL, Fishman GA, Sandberg MA, Alexander KR, Derlacki DJ, Rajagopalan AS (2005) Night blindness and abnormal cone electroretinogram ON responses in patients with mutations in the GRM6 gene encoding mGluR6. *Proc Natl Acad Sci U S A* 102:4884–4889. [CrossRef Medline](#)
- Fee MS, Mitra PP, Kleinfeld D (1996) Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *J Neurosci Methods* 69:175–188. [CrossRef Medline](#)
- Gollisch T, Meister M (2008) Modeling convergent ON and OFF pathways in the early visual system. *Biol Cybern* 99:263–278. [CrossRef Medline](#)
- Grzywacz NM, Amthor FR (2007) Robust directional computation in on-off directionally selective ganglion cells of rabbit retina. *Vis Neurosci* 24:647–661. [Medline](#)
- Hartline HK (1938) The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *Am J Physiol* 121:400–415.
- Jacobs AL, Fridman G, Douglas RM, Alam NM, Latham PE, Prusky GT, Nirenberg S (2009) Ruling out and ruling in neural codes. *Proc Natl Acad Sci U S A* 106:5936–5941. [CrossRef Medline](#)
- Jeon CJ, Strettoi E, Masland RH (1998) The major populations of the mouse retina. *J Neurosci* 18:8936–8946. [Medline](#)
- Jin J, Wang Y, Swadlow HA, Alonso JM (2011a) Population receptive fields of ON and OFF thalamic inputs to an orientation column in visual cortex. *Nat Neurosci* 14:232–238. [CrossRef Medline](#)
- Jin J, Wang Y, Lashgari R, Swadlow HA, Alonso JM (2011b) Faster thalamocortical processing for dark than light visual targets. *J Neurosci* 31:17471–17479. [CrossRef Medline](#)
- Jin JZ, Wong C, Yeh CI, Gordon JA, Ruthazer ES, Stryker MP, Swadlow HA, Alonso JM (2008) On and off domains of geniculate afferents in cat primary visual cortex. *Nat Neurosci* 11:88–94. [CrossRef Medline](#)
- Komban SJ, Alonso JM, Zaidi Q (2011) Darks are processed faster than lights. *J Neurosci* 31:8654–8658. [CrossRef Medline](#)
- Kremers J, Lee BB, Pokorný J, Smith VC (1993) Responses of macaque ganglion cells and human observers to compound periodic waveforms. *Vision Res* 33:1997–2011. [CrossRef Medline](#)
- Kuffler SW (1953) Discharge patterns and functional organization in the mammalian retina. *J Neurophysiol* 16:37–68. [Medline](#)
- Manookin MB, Beaudoin DL, Ernst ZR, Flagel LJ, Demb JB (2008) Disinhibition combines with excitation to extend the operating range of the OFF visual pathway in daylight. *J Neurosci* 28:4136–4150. [CrossRef Medline](#)
- Masland RH (2001) The fundamental plan of the retina. *Nat Neurosci* 4:877–886. [CrossRef Medline](#)
- Meytlis M, Nichols Z, Nirenberg S (2012) Determining the role of correlated firing in large populations of neurons using white noise and natural scene stimuli. *Vision Res* 70:44–53. [CrossRef Medline](#)
- Molnar A, Hsueh HA, Roska B, Werblin FS (2009) Crossover inhibition in the retina: circuitry that compensates for nonlinear rectifying synaptic transmission. *J Comput Neurosci* 27:569–590. [CrossRef Medline](#)
- Moore DS, McCabe GP (2006) Introduction to the practice of statistics, Chapter 14. New York: W. H. Freeman.
- Nikonov SS, Kholodenko R, Lem J, Pugh EN Jr (2006) Physiological features of the S- and M-cone photoreceptors of wild-type mice from single cell recordings. *J Gen Physiol* 127:359–374. [CrossRef Medline](#)
- Nirenberg S, Pandarinath C (2012) Retinal prosthetic strategy with the capacity

- to restore normal vision. *Proc Natl Acad Sci U S A* 109:15012–15017. [CrossRef Medline](#)
- Nirenberg S, Bomash I, Pillow JW, Victor JD (2010) Heterogeneous response dynamics in retinal ganglion cells: the interplay of predictive coding and adaptation. *J Neurophysiol* 103:3184–3194. [CrossRef Medline](#)
- Oizumi M, Ishii T, Ishibashi K, Hosoya T, Okada M (2010) Mismatched decoding in the brain. *J Neurosci* 30:4815–4826. [CrossRef Medline](#)
- Pandarinath C, Victor JD, Nirenberg S (2010) Symmetry breakdown in the ON and OFF pathways of the retina at night: functional implications. *J Neurosci* 30:10006–10014. [CrossRef Medline](#)
- Paninski L, Pillow J, Lewi J (2007) Statistical models for neural encoding, decoding, and optimal stimulus design. *Prog Brain Res* 165:493–507. [CrossRef Medline](#)
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454:995–999. [CrossRef Medline](#)
- Pitkow X, Meister M (2012) Decorrelation and efficient coding by retinal ganglion cells. *Nat Neurosci* 15:628–635. [CrossRef Medline](#)
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP (2007) Random numbers. In: *Numerical recipes*, Ed 3, pp 397–401. Cambridge: MIT Press.
- Renteria RC, Tian N, Cang J, Nakanishi S, Stryker MP, Copenhagen DR (2006) Intrinsic ON responses of the retinal OFF pathway are suppressed by the ON pathway. *J Neurosci* 26:11857–11869. [CrossRef Medline](#)
- Sagdullaev BT, McCall MA (2005) Stimulus size and intensity alter fundamental receptive-field properties of mouse retinal ganglion cells *in vivo*. *Vis Neurosci* 22:649–659. [Medline](#)
- Schiller PH (2010) Parallel information processing channels created in the retina. *Proc Natl Acad Sci U S A* 107:17087–17094. [CrossRef Medline](#)
- Schwartz EA (1974) Responses of bipolar cells in the retina of the turtle. *J Physiol* 236:211–224. [Medline](#)
- Simoncelli E, Paninski L, Pillow J, Schwartz O (2004) Characterization of neural responses with stochastic stimuli. In: *Cognitive Neurosciences* (Gazzaniga MS, ed), pp 327–328. Cambridge, MA: MIT.
- Stone C, Pinto LH (1993) Response properties of ganglion cells in the isolated mouse retina. *Vis Neurosci* 10:31–39. [CrossRef Medline](#)
- Umino Y, Solessio E, Barlow RB (2008) Speed, spatial, and temporal tuning of rod and cone vision in the mouse. *J Neurosci* 28:189–198. [CrossRef Medline](#)
- Wang YV, Weick M, Demb JB (2011) Spectral and temporal sensitivity of cone-mediated responses in mouse retinal ganglion cells. *J Neurosci* 31:7670–7681. [CrossRef Medline](#)
- Wässle H (2004) Parallel processing in the mammalian retina. *Nature Rev Neurosci* 5:747–757. [CrossRef](#)
- Werblin F, Dowling J (1969) Organization of the retina of the mudpuppy, *Necurus maculosus*. II. Intracellular recording. *J Neurophysiol* 32:39–355.
- Xing D, Yeh CI, Shapley RM (2010) Generalization of black-dominant responses in V1 cortex. *J Neurosci* 30:13504–13512. [CrossRef Medline](#)
- Yeh C, Xing D, Shapley R (2009) “Black” responses dominate macaque primary visual cortex V1. *J Neurosci* 23:11753–11760. [Medline](#)
- Zaghloul KA, Boahen K, Demb JB (2003) Different circuits for ON and OFF retinal ganglion cells cause different contrast sensitivities. *J Neurosci* 23:2645–2654. [Medline](#)
- Zemon V, Gordon J, Welch J (1988) Asymmetries in ON and OFF visual pathways of humans revealed using contrast-evoked cortical potentials. *Visual Neurosci* 1:145–150. [CrossRef](#)