Supporting Information

Nirenberg and Pandarinath 10.1073/pnas.1207035109

SI Materials and Methods

Animals. The retinal degeneration (rd) mouse line was obtained from Charles River Laboratories (line C3H/HeN; strain code 025). The line is homozygous for the retinal degeneration allele Pdeb^{rd1} (the β-subunit of the cGMP phosphodiesterase) (1). The Channelrhodopsin-2 (ChR2)-expressing line was obtained from Jackson Labs (B6.Cg-Tg(Thy1-COP4/EYFP)9Gfng/J; stock no. 007615). The line is hemizygous for a construct consisting of ChR2 fused to Yellow Fluorescent Protein (ChR2-YFP) under the control of the Thy1 promoter (2).

To generate mice with both retinal degeneration and ChR2, we crossed the C3H/HeN line with the Thy1 ChR2-YFP line. The first-generation line was then backcrossed with the parent C3H/HeN line. The resulting line, identified using standard PCR genotyping, was then homozygous for the mutant Pdeb^{rd1} gene and hemizygous for ChR2-YFP; we refer to this line in this paper as the Thy1-ChR2 rd1/rd1 line. All procedures on experimental animals were carried out under the regulation of the Institutional Animal Care and Use Committee (IACUC) of the Weill Cornell Medical College of Cornell University and in accordance with NIH guidelines.

Prosthetic Device. The device consists of a camera, an encoder, and a stimulator. Images taken by the camera are converted into electrical pulses by the encoder (see next section), which runs on an OMAP 3530 processor (Texas Instruments). The pulse patterns are then converted to blue light pulses, which are delivered to the retina using a minidigital light projector (mini-DLP) (Pico projector, Texas Instruments), which we modified to deliver 0.7 mW/mm^2 of 475 nm light at the retina using a light-emitting diode (LED) (Xlamp XP-E blue LED; CRE, Inc.). The DLP contains a grid of mirrors (a digital micromirror device), whose positions can be switched with very high spatial and temporal resolution; each mirror is 7.56 × 7.56 μm, and monitor speed is 1,440 Hz. Signals from the encoder control the switching of the mirrors, allowing light from the LED to reach the retina. The high spatial and temporal resolution allows the delivery of pulses whose spatial extent can be smaller than the diameter of a cell and whose time course can be less than a millisecond. For the recording experiments, the prosthetic device was mounted above the recording chamber as shown in Fig. S5*A*.

For Figs. 2–4, the movies were taken in advance and presented from stored versions. They were passed through the encoder, as described above, and then presented to the retina via the modified mini-DLP. When presenting unencoded movies, as for the standard optogenetic method, the movies were presented with no conversion except from gray scale to blue scale (to excite the ChR2); i.e., the intensity values of the movies were fed directly to the modified mini-DLP's output.

Encoder. The encoder is a retinal input/output model generated from data taken from the normal retina. It is a linear-nonlinear (LN) cascade built to capture stimulus/response relations for a broad range of stimuli, followed by a Poisson spike generator (for general review of linear-nonlinear cascade models for retina and other systems, see refs. 3 and 4).

Firing rate at time t of the mth neuron in response to stimulus S is given by

$$\lambda_m(t) = N_m((S * L_m)(t)), \quad [1]$$

where * denotes spatiotemporal convolution, L_m is a linear filter corresponding to the mth neuron's spatiotemporal impulse response, and N_m is a function that describes the mth neuron's nonlinearity. L_m is parameterized as a product of a spatial function (weights at each pixel) and a temporal function [a sum of raised cosines (5, 6)]. The nonlinearities N_m are parameterized as cubic splines. We found the values of these parameters by maximizing the likelihood that the model produces the experimentally observed ganglion cell spike trains. To make the search more efficient, we take a two-step approach. We first assume that the nonlinearity is an exponential, as this takes us to the vicinity of a global maximum (7); we then use the parameter values obtained from this as a starting point to continue the search (for the linear filter and spline parameters). The parameters for each were found by coordinate ascent [i.e., by alternating stages of maximizing the log-likelihood with respect to (*i*) the filter parameters and (*ii*) the spline coefficients until a maximum was reached). For the data shown here, we treated the cells as independent. We also have a form of the model that includes a coupling term among the neurons, as in refs. 6, 8, and 9. Thus far, the results show that the form with no coupling gives reliable predictions (8, 9) and is easier to use for the prosthetic, but both are in hand (8, 9).

In brief, the model captures in a compact form the processing that the retina performs, including center-surround organization, temporal filtering, and spike generation. Briefly, images come in, they are convolved with a spatiotemporal filter (with positive and negative lobes), then are passed through a nonlinearity, and converted to spikes by a Poisson generator. The model captures the processing in a data-driven way in that the spatial and temporal coefficients of L_m and the shape of the nonlinearity, N_m, are determined directly from measured responses, giving them the accuracy to produce normal output.

Note that we refer to the "encoder" as the complete encoding device; it is, essentially, a model retina. Similar to the normal retina, the encoder consists of arrays of smaller encoders, one for each ganglion cell type. The different encoders all use the same model structure, the LN cascade given above; what makes them different from each other are the numerical values of the parameters. For example, there is a separate encoder for ON transient cells, for ON sustained cells, etc., or, in the monkey, for midget ON cells, for midget OFF cells, etc. As shown throughout the paper (Fig. 2 and Figs. S1 and S4), the input/output relations for these cell types are faithfully captured by their encoders. For further quantification on the effectiveness of the encoders, see ref. 10. In the building of the encoders, we measured performance using standard cross-validation, i.e., using out-of-sample stimuli drawn from the same statistical class as the training set. Here and also in ref. 10, we went further, testing performance using stimuli unrelated to the training set (a broad range of natural images, including faces, landscapes, children playing, etc.). These results demonstrate not only that the encoders perform well for out-of-sample stimuli with similar statistics, but also that they generalize to unrelated stimuli with different statistics.

To make the model effective on a broad range of stimuli, we used a nonstandard approach. Briefly, the classic, biological approach to improve retinal models has been to build adaptive mechanisms into them, for example, filters that allow the model to adapt to different image statistics (11, 12). These approaches include quasi-linear models that have components that explicitly adapt (e.g., parameters that depend directly on the statistics of the input; see ref. 11, where the time constant of a filter was made to depend on input contrast) or nonlinear models in which the adaptation is an emergent property of the nonlinear dynamics

(12). These strategies, however, are not easy to implement in a data-driven way: for the approach in ref. 11 and for similar approaches, the number of parameters is too large for the amount of data that can be obtained from retinal recordings, and, for the nonlinear model, the functional form to use for the dynamics (e.g., to cause it to accurately capture responses to a broad range of stimuli) is not clear. To address the problem, we developed, instead, a machine learning approach. With a minimal enlargement of the stimulus set (white noise and particular natural scene movies) we obtained a linear-nonlinear cascade model that generalized broadly (white noise, gratings, landscapes, people walking, cars driving, faces, animals, etc.). We reasoned that this might work because white noise and natural scenes are complementary: in both the temporal and spatial domains, natural scenes are much more heavily weighted toward low frequencies than white noise, and vice versa. The results showed that this strategy was very effective: the combined stimulus set allowed sampling of a diverse space of inputs that drove the optimization to a different location in parameter space than was found by either stimulus set alone (or by averaging the two).

With respect to light levels and contrast, these are addressed using a preprocessing step, whereby the light level and contrast range are rescaled to fall within the operating range of the encoders. The rescaling occurs through a front end preprocessor: briefly, the images pass through a camera, then a preprocessor—which does the rescaling to accommodate changes in light level and contrast in the environment—and then the encoder.

For the images/movies shown in these experiments, the encoders were built to operate in the following range on the basis of the images presented to the normal retina: mean light level was 1.7 $\mu W/cm^2$ (photopic range) and rms contrast was 0.27 $\mu W/cm^2$; intensity was measured using a New Focus 3803 Power Meter, and the spectrum of the monitor is given in ref. 13. (The light intensity coming out of the prosthetic is higher to drive the ChR2, as indicated above, under *Prosthetic Device*.)

Confusion Matrices. Confusion matrices were generated as in ref. 14. Briefly, for each presentation of a stimulus, s_i, a response, $\boldsymbol{r}$, was recorded. The response was decoded by choosing the stimulus most likely to have produced it, that is, the stimulus for which $p(s_i|\mathbf{r})$ was maximal. As is standard, $p(s_j|\mathbf{r})$ was calculated from $p(\boldsymbol{r}|s_i)$ using Bayes theorem, which states that $p(s_i|\boldsymbol{r}) = p(\boldsymbol{r}|s_i)p(s_i)/p(\boldsymbol{r})$. Because $p(s_i)$ was set equally for all stimuli in this experiment, $p(s_i|\boldsymbol{r})$ was maximized when $p(\boldsymbol{r}|s_i)$ was maximized. For each presentation of stimulus s_i that resulted in a response $\boldsymbol{r}$ that was decoded as the stimulus s_j, the entry at position (i,j) in the confusion matrix was incremented.

Each set of matrices was constructed from two sets of data: a training set and a testing set. The training set for all cases was the same, consisting of responses produced by the normal, wild-type retina. The test sets in Fig. 3 were out-of-sample responses from the normal retina (Fig. 3, *Top*), responses from blind retinas treated with the encoder-ChR2 prosthetic (Fig. 3, *Middle*), and responses from blind retinas treated with the standard optogenetic prosthetic (Fig. 3, *Bottom*).

The stimuli were movies of natural scenes (15 movies, each 667 ms in duration, each presented 60 times). For each stimulus, the response $\boldsymbol{r}$ was taken to be the spike train spanning 500 ms after stimulus onset. Spike generation was assumed to be inhomogeneous Poisson, allowing response probability for the 500-ms response to be calculated as the product of the probabilities for each bin in the response.

An array of bin sizes from 50 to 10 ms were used. The bin size shown in Fig. 3 was 50 ms; the smaller bin sizes are shown in Fig. S7, and the results, as shown, are very similar as those in Fig. 3. When decoding the responses for the populations of cells (the right-most matrices of Fig. 3 *A* and *B*), the response distribution for the population was computed by assuming that the cells were conditionally independent: this was done for all groups.

Performance was quantified using "fraction correct," which was the fraction of times the responses were correctly decoded, averaged over all stimuli. For each matrix, the fraction correct was the mean of the values on the diagonal.

Stimulus Reconstructions. As mentioned in the main text, for these experiments, we presented an image to the two prosthetic systems—our method and the standard optogenetic method—and recorded ganglion cell responses. To obtain a large enough dataset for the complete reconstruction, we moved the image systematically across the region of retina from which we were recording, so that responses to all parts of the image could be obtained with a single or small number of retinas. Approximately 9,800 ganglion cell responses were recorded for each image.

The reconstructions were carried out on a processing cluster in blocks of 10 × 10 or 7 × 7 pixels (checks). The decoding was performed using maximum likelihood; that is, for each block, we found the array of gray values that maximized the probability of the observed responses, $\boldsymbol{r}$ (following ref. 4 for high-dimensional searches). Briefly, we used gradient ascent techniques. We started at a random point in stimulus space, s_k. We evaluated the probability distribution $p(\boldsymbol{r}|s_k)$ for this stimulus and also calculated the slope of this probability distribution with respect to each dimension of the stimulus. We then created a new stimulus, s_{k+1}, by changing the stimulus s_k in the direction of increasing probability (as determined from the slope of the probability distribution). This process continued iteratively until we reached the peak of $p(\boldsymbol{r}|s)$. The reconstructions of each block were performed using multiple random starting points to confirm that they converged to the same peak.

The stimulus consisted of a uniform gray screen presented for 1 s, followed by the image of the baby's face for 1 s. The pixelation of the face, 35 × 32 pixels, was chosen so that the features of the face could be readily discerned. The choice of 35 × 32 pixels is consistent with the fact that facial recognition makes use of spatial frequencies at least as high as eight cycles per face, which requires at least 32 pixels in each dimension for adequate sampling (15). In Fig. 4, each pixel corresponded to 2.6 ° × 2.6 °of visual space, which corresponds to 12–20 ganglion cells in the mouse retina (16, 17) (center to midperiphery). Consistent with this, we reconstructed each image using 12–20 ganglion cell responses per pixel (per check), with the ratio of ON to OFF cells varying from location to location depending on the number of cells recorded from the array at a given location. See Fig. S3 for still images from a reconstructed movie using just OFF cells (as described in the main text *Discussion*).

For each stimulus, the response was taken to be the spike train spanning the 1-s stimulus and binned with 0.67-ms bins. Spike generation was assumed to be inhomogeneous Poisson, allowing response probability for the complete response to be calculated as the product of the probabilities for each bin.

Optomotor Tracking. The grating stimuli were passed through the encoder (specifically, an array of transient ON cell encoders; see main text for the reasoning) and then presented to the retina via a liquid crystal display (LCD), which was illuminated from behind with a bank of blue light-emitting diodes (LEDs) (to excite the ChR2), as shown schematically in Fig. S5*B*. For the standard optogenetic method, the grating stimuli were presented as is, except that the gray scale was converted to blue scale (to excite ChR2); i.e., the intensity values of the grating stimuli were fed directly to the LCD's output.

Eye tracking was measured following ref. 18. Briefly, the mouse was placed in an acrylic cylinder facing the LCD panel. Its head was positioned and held stable using a headpost. To collect eye-tracking data, the eye was imaged using infrared illumination and

an ISCAN infrared video camera. The ISCAN software finds the horizontal and vertical coordinates of the centers of the corneal reflection and the pupil (sampled at 60 Hz) and measures the difference between these two points along the horizontal and vertical axes. A schematic of the setup is shown in Fig. S5*B*.

Note that the baseline control for these experiments was a measure of eye movements from the same animal without the stimulus, as this captures the baseline drift that occurs with blind mice. As shown in ref. 18, normal tracking is highly variable among mouse strains, and it is not possible to obtain tracking behavior from the *rd* strain (because it is blind); for this reason, the eye movements that occur without the stimulus, and the eye movements that occur with the unencoded stimulus, served as the two controls.

Multielectrode Recording/Pharmacology. Extracellular recordings were made using a multielectrode array, as in refs. 13, 14, 19, and 20. Spike trains were recorded and then sorted into units (cells) using a Plexon Instruments Multichannel Neuronal Acquisition Processor. Neurotransmitter blockers were used to ensure the abolition of photoreceptor-driven responses, as in ref. 21. The following were used: 10 μM CPP [(±)-3-(2-carboxypiperazin-4-yl) propyl-1-phosphonic acid], 10 μM NBQX [1,2,3,4-tetrahydro-6-nitro-2,3-dioxo-benzo(f)quinoxaline-7-sulfonamide], which block NMDA and AMPA/kainate receptors, respectively, and 80 μM APB (2-amino-4-phosphonobutyrate), which blocks metabotropic glutamate receptors.

Preparing Retinal Sections and Whole Mounts. Mice were euthanized by CO_2, and the eyes were removed. Retinas were separated from the rest of the eye and fixed in 4% paraformaldehyde (wt/vol) in 0.1 M PBS, pH 7.4, for 15–30 min. Following three washes with PBS, the retinas were cryoprotected in graded sucrose solutions (15%, followed by 30%). Sections were then cut on a cryostat at 14 μm. For retinal whole mounts, the tissue was fixed in 4% paraformaldehyde/0.1 M PBS, washed in PBS, and mounted flat on a slide, lightly coverslipped.

SI Figures. As mentioned in the main text *Discussion*, the prosthetic method is effective because the encoders very reliably reproduce ganglion cell-firing patterns and because ChR2 is able to follow the patterns that the encoders produce. In the main text (Fig. 2), we demonstrated this by showing the firing patterns of the normal retina and the firing patterns of the degenerated retina when it was driven by the encoder-ChR2 prosthetic.

Here we break it down further and include the firing patterns of the encoders (Fig. S1). The top set of rasters shows the normal ganglion cell-firing patterns for several cells [using long (10-s) stretches]. The next set of rasters shows the firing patterns produced by the encoders for each of these cells (i.e., they show the output of the linear-nonlinear-Poisson model for each cell). Finally, the third set of rasters shows the firing patterns produced by ganglion cells from a degenerated retina when it is driven by the encoder-ChR2 prosthetic. In other words, the second set of rasters shows the extent to which the encoders follow the true cells, and the third set of rasters shows the extent to which the ChR2-expressing ganglion cells follow the encoders.

Fig. S2 shows the first 2 s in an expanded view, so the correspondence is visible at the level of individual spikes. We show this because it is interesting how well the ChR2-expressing ganglion cells follow the encoder's signals. Although the cells occasionally fail to fire, or fire doublets or triplets to a single light pulse, overall the correspondence is very tight. The implication of this is that the statistics of neuronal spike trains produced by natural stimuli are compatible with the properties of ChR2. We mention this because ChR2's usefulness for prosthetics has been a subject of discussion in the field because ChR2's properties are not as well-matched to visual stimuli as when they are presented directly, that is, unfiltered through neurons. For a review of the properties of ChR2, see ref. 22.

1. Bowes C, et al. (1993) Localization of a retroviral element within the rd gene coding for the beta subunit of cGMP phosphodiesterase. *Proc Natl Acad Sci USA* 90:2955–2959.
2. Arenkiel BR, et al. (2007) In vivo light-induced activation of neural circuitry in transgenic mice expressing channelrhodopsin-2. *Neuron* 54:205–218.
3. Simoncelli E, Paninski L, Pillow JW, Schwartz O (2004) Characterization of neural responses with stochastic stimuli. *The Cognitive Neurosciences*, ed Gazzaniga M (MIT Press, Cambridge, MA), pp 327–338.
4. Paninski L, Pillow J, Lewi J (2007) Statistical models for neural encoding, decoding, and optimal stimulus design. *Prog Brain Res* 165:493–507.
5. 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.
6. Pillow JW, et al. (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454:995–999.
7. Paninski L (2004) Maximum likelihood estimation of cascade point-process neural encoding models. *Network* 15:243–262.
8. Meytlis M, Bomash I, Pillow JW, Nirenberg S (2009) Assessing the importance of correlated firing using large populations of neurons. *Society for Neuroscience Annual Meeting*, Chicago. Prog. No. 165.3.
9. Nichols Z, Nirenberg S (2010) Correlations in complete retinal populations play a very small role in coding natural scene stimuli. *Society for Neuroscience Annual Meeting*, San Diego. Prog. No. 891.3.
10. Nirenberg S, Pandarinath C, Ohiorhenuan I (2011) Retina prosthesis. International patent WO2011106783. Available at http://worldwide.espacenet.com/publicationDetails/biblio?DB=EPODOC&ND=3&FT=D&date=20110901&CC=WO&NR=2011106783A2.
11. Victor JD (1987) The dynamics of the cat retinal X cell centre. *J Physiol* 386:219–246.
12. Famulare M, Fairhall A (2010) Feature selection in simple neurons: How coding depends on spiking dynamics. *Neural Comput* 22:581–598.
13. Nirenberg S, Carcieri SM, Jacobs AL, Latham PE (2001) Retinal ganglion cells act largely as independent encoders. *Nature* 411:698–701.
14. 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.
15. Rolls ET, Baylis GC, Leonard CM (1985) Role of low and high spatial frequencies in the face-selective responses of neurons in the cortex in the superior temporal sulcus in the monkey. *Vision Res* 25:1021–1035.
16. Jeon C-J, Strettoi E, Masland RH (1998) The major cell populations of the mouse retina. *J Neurosci* 18:8936–8946.
17. Williams RW, Strom RC, Rice DS, Goldowitz D (1996) Genetic and environmental control of variation in retinal ganglion cell number in mice. *J Neurosci* 16:7193–7205.
18. Cahill H, Nathans J (2008) The optokinetic reflex as a tool for quantitative analyses of nervous system function in mice: Application to genetic and drug-induced variation. *PLoS ONE* 3:e2055.
19. Jacobs AL, et al. (2009) Ruling out and ruling in neural codes. *Proc Natl Acad Sci USA* 106:5936–5941.
20. Nirenberg S, Meister M (1997) The light response of retinal ganglion cells is truncated by a displaced amacrine circuit. *Neuron* 18:637–650.
21. Lagali PS, et al. (2008) Light-activated channels targeted to ON bipolar cells restore visual function in retinal degeneration. *Nat Neurosci* 11:667–675.
22. Lin JY (2011) A user's guide to channelrhodopsin variants: features, limitations and future developments. *Exp Physiol* 96(1):19–25.

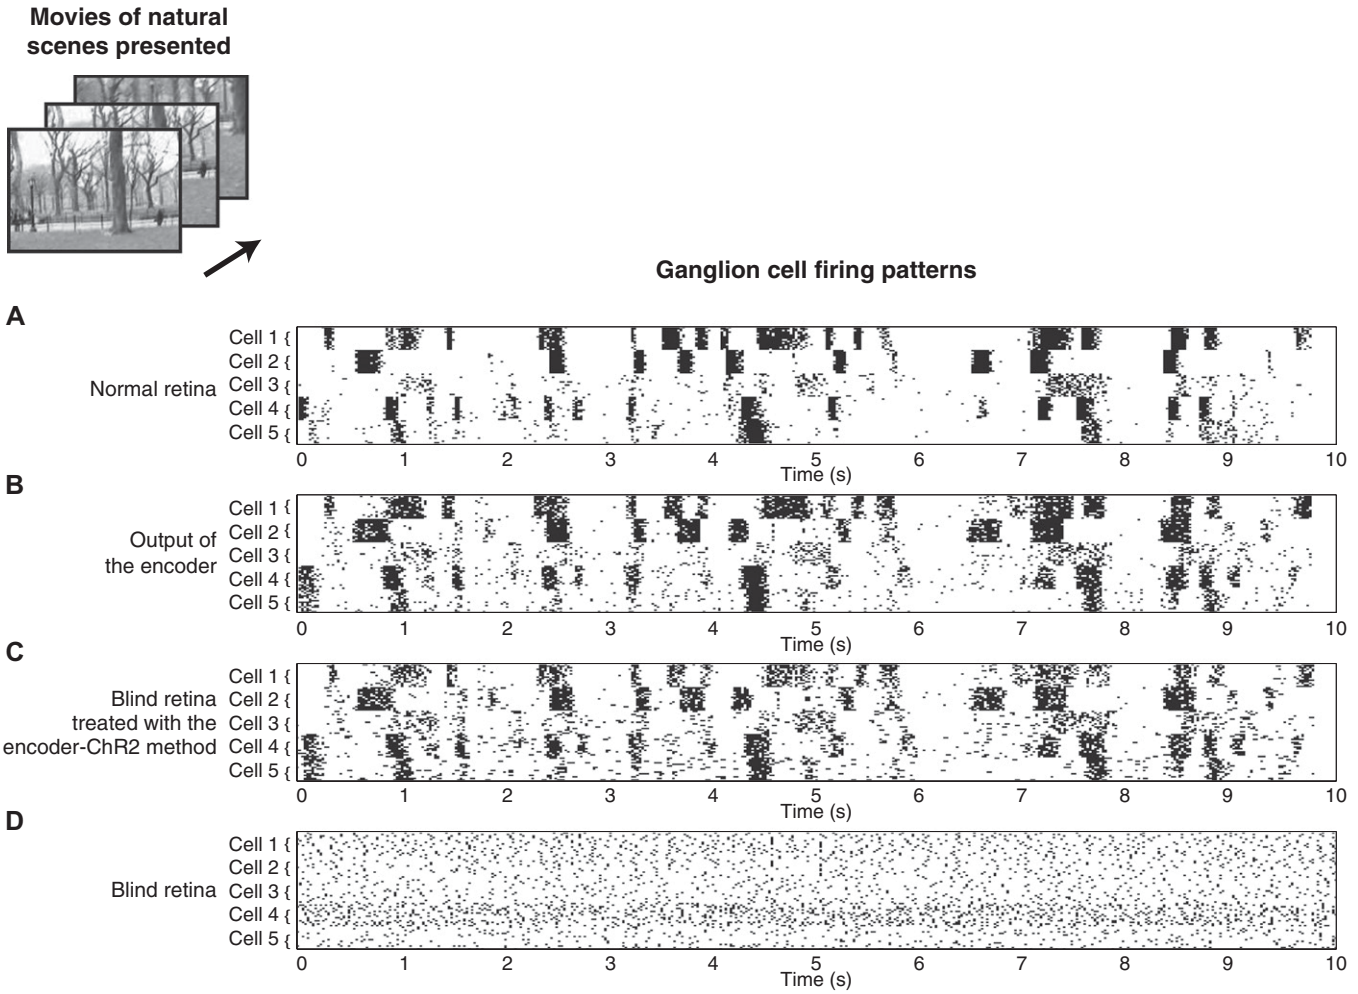

Fig. S1. Firing patterns from the normal retina, from the encoders, and from a blind retina driven by the encoder-ChR2 prosthetic. (*A*) Normal ganglion cell-firing patterns for several cells (in response to movies of natural scenes). This array of cells includes both ON transient and ON sustained cells (for other cell classes, see Fig. 2 of main text and Fig. S4). (*B*) The firing patterns produced by the encoders for each of these cells. (*C*) The firing patterns produced by ganglion cells from a blind retina when it was driven by the encoder-ChR2 prosthetic. As shown, the encoders faithfully reproduce normal ganglion cell responses, and the ChR2-expressing ganglion cells faithfully follow the signals produced by the encoders. (*D*) Recordings from a blind, degenerated retina that does not express ChR2 to reiterate that ganglion cells from retinas that do not express ChR2 do not respond to the stimulus.

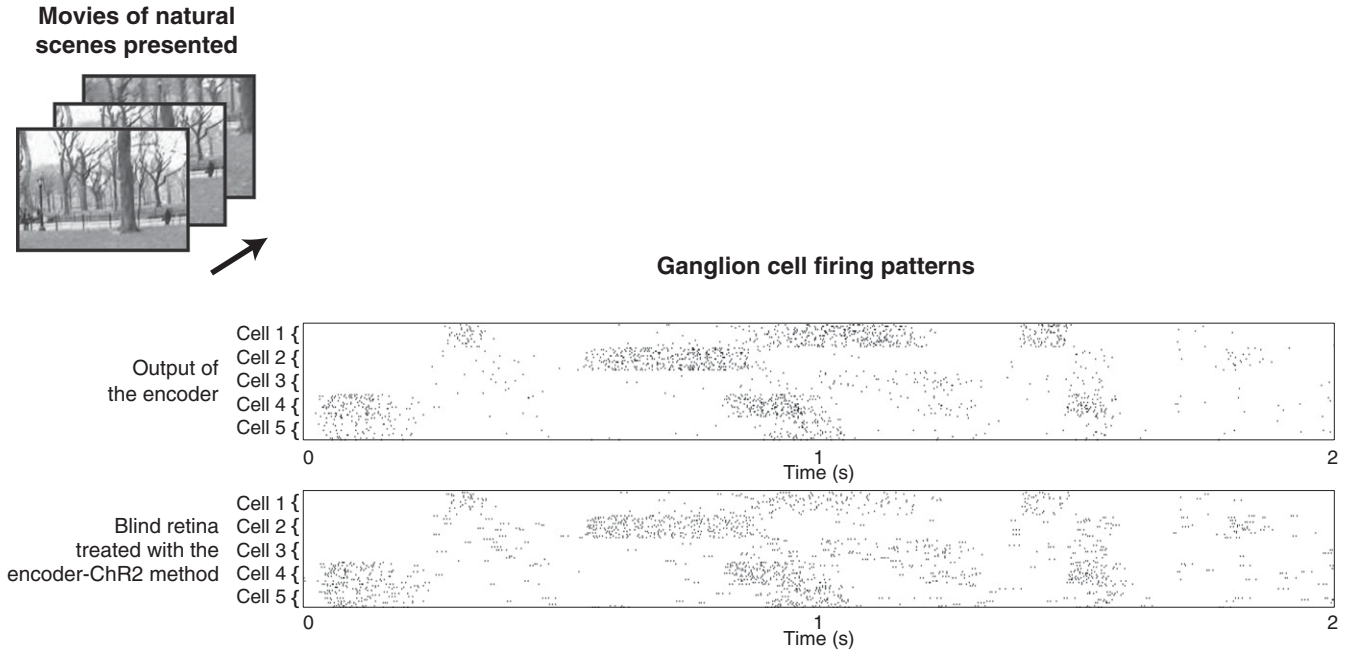

Fig. S2. Expanded view of the *B* and *C* in Fig. S1. The top set of rasters shows the firing patterns of the encoders (from Fig. S1*B*), and the bottom set of rasters shows the firing patterns of the ChR2-expressing cells (from Fig. S1*C*). Bin size is 2 ms; in Fig. S1, bin size is 25 ms. As mentioned, although cells occasionally fail to fire, or fire doublets or triplets to a single light pulse, overall the correspondence is very tight (*SI Materials and Methods*; see *SI Figures* for discussion.).

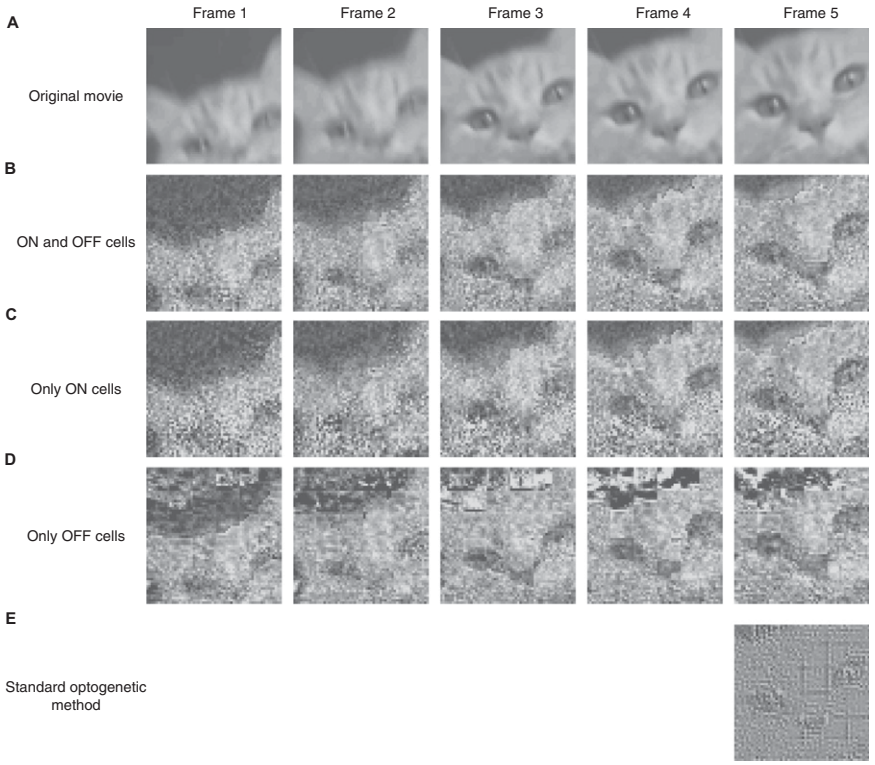

Fig. S3. Image reconstructions using only the responses of OFF ganglion cell encoders; this provides a measure of the quality of vision when only these cells are used. (*A*) Original images; these are frames from a short movie. (*B*) Images reconstructed using the responses of both ON and OFF ganglion cell encoders. (*C*) Images reconstructed using the responses of just ON ganglion cell encoders. (*D*) Images reconstructed using the responses of only OFF ganglion cell encoders. Each movie was reconstructed in blocks of 10 pixels by 10 pixels by 5 frames. As in Fig. 4 in the main text, the decoding was performed using maximum likelihood; that is, for each block, we found the array of gray values that maximized the probability of the observed responses (following ref. 4 for high dimensional searches). (*E*) Image reconstructed from the responses of real ganglion cells using the standard method (only ChR2, no encoder), using the same resolution as the images in *A*–*D*.

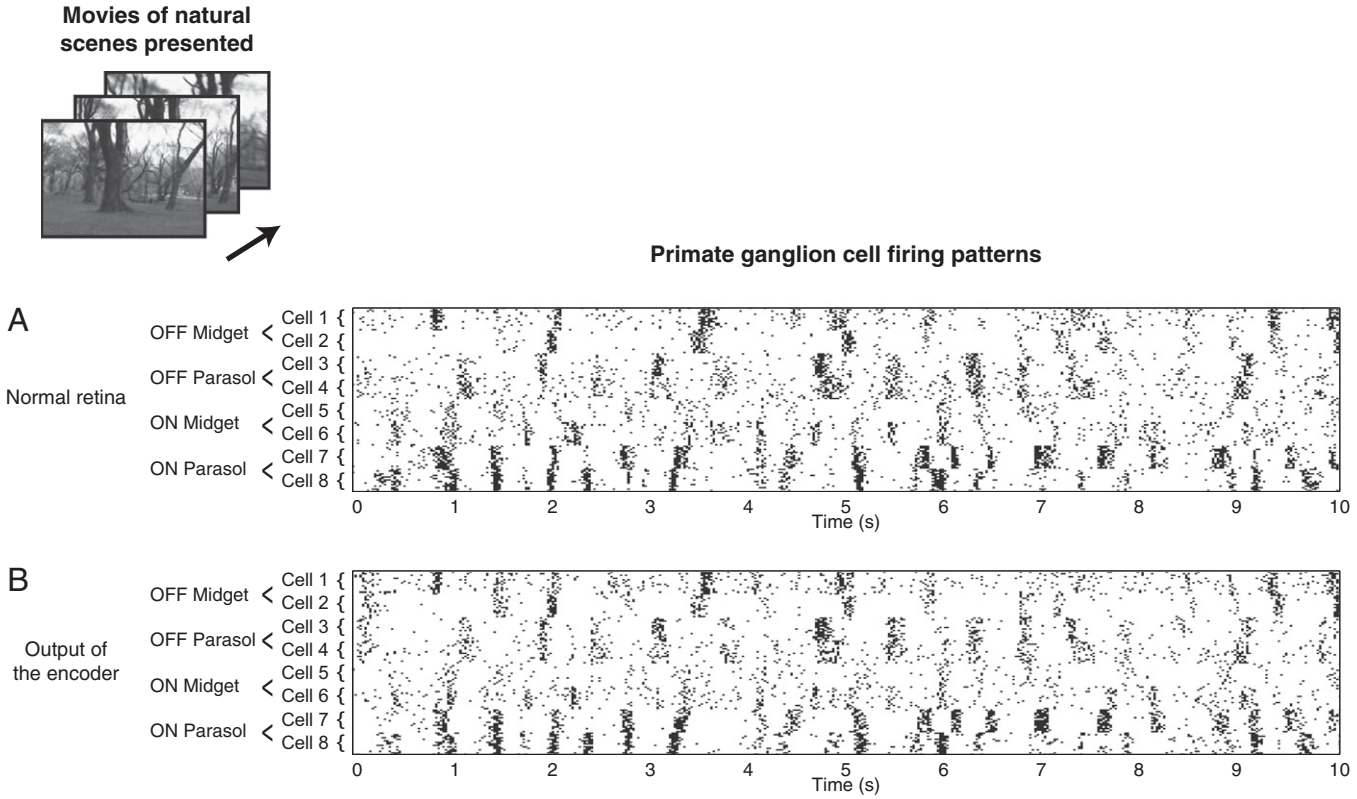

Fig. S4. Primate (macaque) retina: firing patterns from ganglion cells from normal retinas and from their encoders. (*A*) Normal ganglion cell-firing patterns in response to movies of natural scenes. ON and OFF midget and parasol cells are indicated; midget and parasol cells are classified on the basis of receptive field diameters, as in ref. 1. (*B*) The firing patterns produced by the encoders for each of the cells in *A*. Recordings from primate retina were carried out using a multielectrode array and the same conditions as described above for mouse retina, except that Ames' medium was used (1). The stimulus protocol was the same as for the mouse experiments, and the region of recording was midperipheral retina. For the models, we used the approach described above with the addition of a postspike feedback (2). Experiments were performed in accordance with the institutional guidelines for animal welfare.

1. Chichilnisky EJ, Kalmar RS (2002) Functional asymmetries in ON and OFF ganglion cells of primate retina. *J Neurosci* 22:2737–2747.
2. Pillow JW, et al. (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454:995–999.

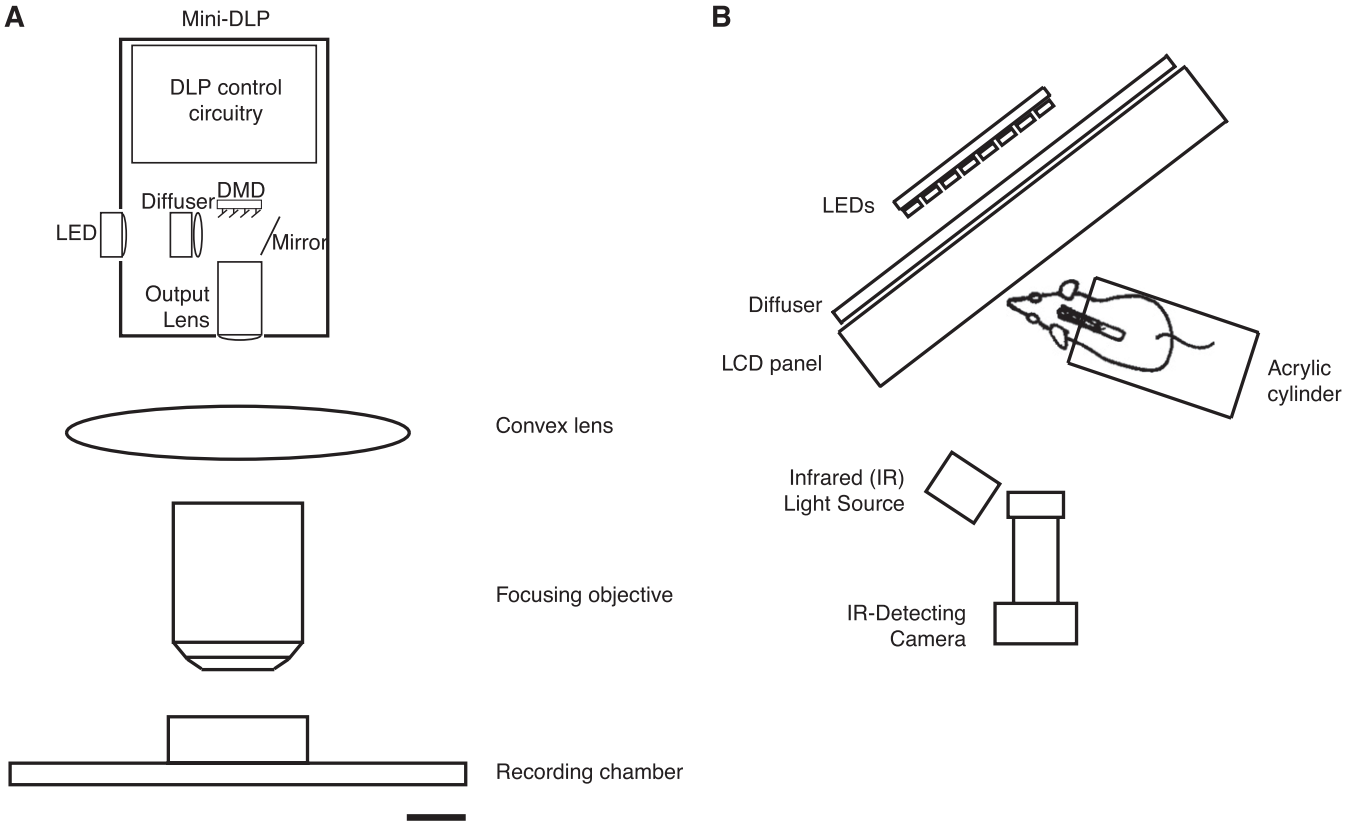

Fig. S5. Schematics of the stimulation setups. (*A*) For the multielectrode recording experiments, the mini-DLP was positioned above the recording chamber, as described in *SI Materials and Methods*. (Scale bar, 0.5 in.) (*B*) For the behavior experiments, the stimulus was delivered using an LCD panel, whose light source was a bank of LEDs, as described in *SI Materials and Methods*.

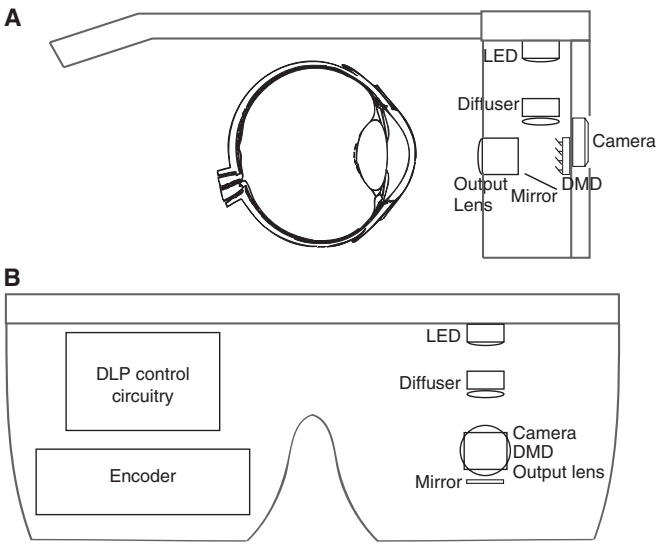

Fig. S6. Schematic of the prosthetic device incorporated into a pair of eyeglasses. All of the components, including the camera, the encoder, and the mini-DLP can fit into the eye pieces of a pair of eyeglasses. (*A*) Side view of the eyeglasses; the components are identical to those in Fig. S5, except the stimulator of the mini-DLP is rotated 90 ° (LED on top), and the lens is replaced by a lens appropriate for interaction with the eye's optics. The circuit board is also separated from the stimulator, which cannot be seen in the side view, but can be seen in *B*. (*B*) View from inside the eyeglasses, to show component layout.

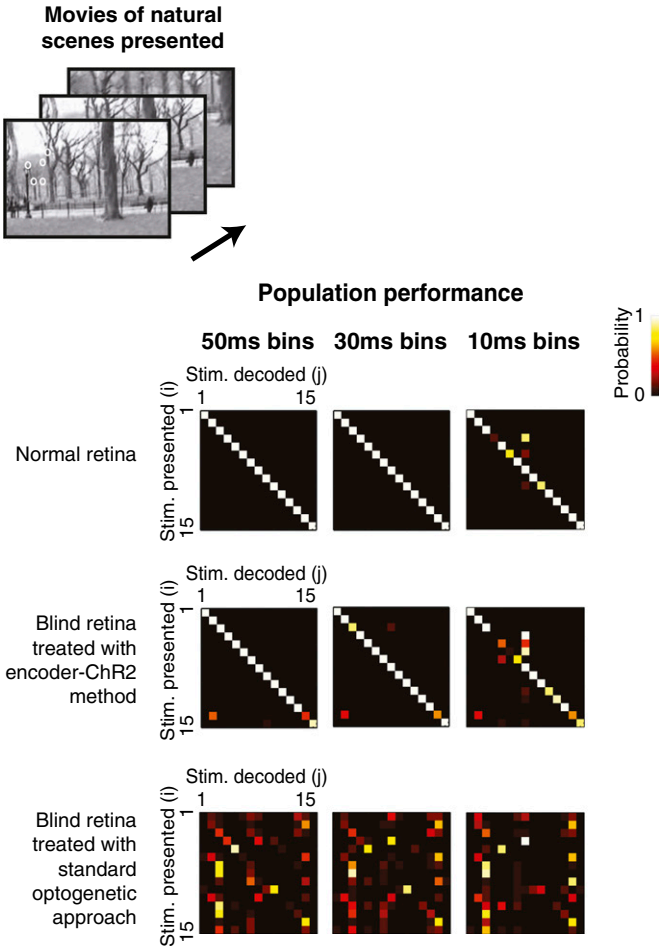

Fig. S7. Confusion matrices generated using an array of bin sizes. Ganglion cell responses were decoded using 50-, 30-, and 10-ms bins. As shown, the conclusions are the same as those presented in the main text, Fig. 3.

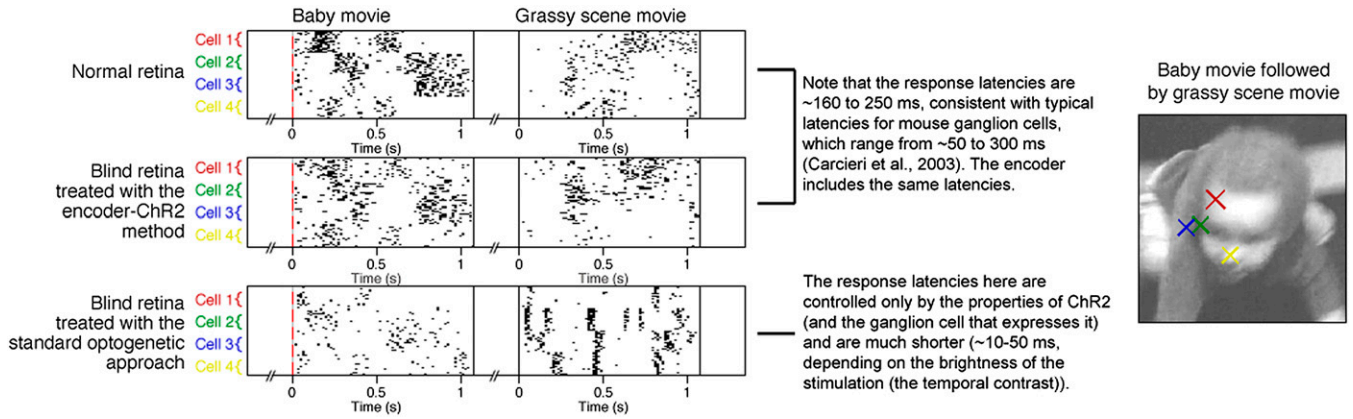

Movie S1. Animated stimulus/response relations. The movies show the complexity of the stimulus/response relations to natural images. To provide intuition, we used two movies: one of a baby playing and one of movement through a grassy scene. In the baby movie, the features of the images are relatively coarse (the shots are fairly close up so they cover large areas of the visual field), which makes the input/output relations relatively easy to follow (one can relatively easily discern the features in the movie that trigger the firing). In the grassy movies, the features are fine; this causes intuition to slip away because the full dimensionality of the model now matters. What is important is that the encoder performs well under both conditions, and we emphasize that these are just two short examples (see the many other figures in this paper, e.g., Fig. S4 above for primate responses). Note also that the latencies are different for the retinas in which encoding of the visual input is taking place (the normal retina and the blind retina driven by the encoder-ChR2 prosthetic) versus the retinas driven by the standard optogenetic method. The latencies for first two correspond to the normal latencies of ganglion cells from mouse retina, which range from 50 to 300 ms (1). In this particular movie, they range from ~160 to 250 ms. For the standard optogenetic method, the latencies are controlled only by the properties of ChR2 (and the ganglion cell that expresses it) and are much shorter [~10–50 ms, depending on the brightness of the stimulation (the temporal contrast)]. In each movie, four cells are shown; receptive field locations with approximate sizes are indicated by the colored Xs.

Movie S1

1. Carcieri SM, Jacobs AL, Nirenberg S (2003) Classification of retinal ganglion cells: a statistical approach. *J Neurophysiol* 90:1704–1713.