High-fidelity restoration of visual signals by electrical stimulation in the central primate retina

Alex R. Gogliettino[1,2,4,*], Sasidhar S. Madugula[1,2,3], Lauren E. Grosberg[2,4], Ramandeep S. Vilkhu[2,5], Jeff Brown[2,4,5], Huy Nguyen[2], Alexandra Kling[2,4], Paweł Hottowy[6], Władysław Dąbrowski[6], Alexander Sher[7], Alan M. Litke[7], and E.J. Chichilnisky[2,4,5,8]

1. Neurosciences PhD Program, Stanford University, Stanford, CA, 94305
2. Hansen Experimental Physics Laboratory, Stanford University, Stanford, CA, 94305
3. School University School of Medicine, Stanford University, Stanford, CA, 94305
4. Department of Neurosurgery, Stanford University, Stanford, CA, 94305
5. Department of Electrical Engineering, Stanford University, Stanford, CA, 94305
6. Faculty of Physics and Applied Computer Science, AGH University of Science and Technology, Krakow, Poland, 30-059
7. Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, CA, 95064
8. Department of Ophthalmology, Stanford University, Stanford, CA, 94305

*alex.gogliettino@gmail.com

Acknowledgements: We thank The California National Primate Research Center, SRI International, S. Morairty, J. Carmena, J. Wallis, K. Bankiewicz, J. Horton, C. Darian-Smith, W. Newsome and T. Moore for access to primate retinas; N. Shah, E. Wu, S. Cooler, M. Zaidi, P. Vasireddy, A. Phillips, M. Breidenbach, S. Mitra and F. Rieke for useful discussions and comments on the manuscript; J. Desnoyer, R. Samarakoon, and S. Kachiguine for technical assistance; K. Mathieson for development of 30 µm pitch electrode arrays. This research was supported by NIH NEI F31EY033636 (ARG), NIH NEI F30-EY030776-03 (SSM), Stanford Bio-X Bowes Graduate Student Fellowship (RSV), Stanford University Vice Provost for Undergraduate Education Small Grant (HN), Polish National Science Centre grant DEC-2013/10/M/NZ4/00268 (PH), The Polish Ministry of Science and Higher Education and its grants for Scientific Research (WD), Pew Charitable Trust Scholarship in the Biomedical Sciences (AS), a donation from John Chen (AML), Stanford Medicine Discovery Innovation Award, Research to Prevent Blindness Stein Innovation Award, Wu Tsai Neurosciences Institute Big Ideas, NIH NEI R01-EY021271, NIH NEI R01-EY029247, and NIH NEI P30-EY019005 (EJC).

Abstract

Electrical stimulation of retinal ganglion cells (RGCs) with electronic implants provides rudimentary artificial vision to people blinded by retinal degeneration, but cannot reproduce the neural code of the retina. Recent work has demonstrated more precise activation of RGCs using focal electrical stimulation in the peripheral macaque retina, but high-resolution vision requires the central retina. This work probes the effectiveness of focal epiretinal stimulation in the central macaque retina, using large-scale recording and stimulation *ex vivo*. The functional organization of the major RGC types was similar between the peripheral and central retina, and these cell types could be distinguished by their electrical properties. Electrical stimulation revealed similar RGC activation thresholds and reduced axon bundle activation in the central retina, but diminished stimulation selectivity. However, a novel stimulation algorithm revealed higher expected acuity in the central retina. These results support the possibility of reproducing high-acuity vision in the central retina with an epiretinal implant.

Introduction

Electronic neural implants have the potential to translate our understanding of the nervous system into clinical restoration of crucial sensory and motor functions. An important application is epiretinal implants, which electrically activate retinal ganglion cells (RGCs), causing them to send artificial visual signals to the brain. Although this technology is already a primary treatment for blindness resulting from photoreceptor degeneration (Weiland et al., 2011, 2016), present-day devices provide limited restoration of vision. One likely reason for this is that all currently available implants activate RGCs of different types simultaneously and indiscriminately, creating artificial neural signals that do not mimic the temporally precise and cell-type specific neural code of the retina.

A potential solution would be to develop a high-resolution implant that can precisely control the activity of the diverse RGC types, which are interspersed in the retina. Previous *ex vivo* work in the peripheral primate retina has shown that epiretinal electrical stimulation with high-density multi-electrode arrays can activate individual RGCs of the major types with single-cell, single-spike resolution (Jepson et al., 2013, 2014a; Sekirnjak et al., 2008), and that the distinct cell types can be distinguished by their recorded electrical signatures (Richard et al., 2016), raising the possibility of reproducing the neural code. More recently, a stimulation algorithm was developed that in principle can optimize the quality of vision restoration in conditions when stimulation is focal and precisely calibrated (Shah et al., 2019). However, it remains unclear whether these approaches can be applied effectively to the central retina, the principal target for electronic implants, because the properties and functional organization of the diverse RGC types and the achievable specificity of electrical stimulation in areas of high RGC density are not well understood.

Here, we test how effectively electrical stimulation can reproduce RGC signals for vision restoration in the central retina, and perform a direct comparison to the peripheral retina. First, using large-scale multi-electrode recordings and visual stimulation, we characterize the functional organization and visual neural code of the five numerically-dominant RGC types identified by their light responses – ON parasol, OFF parasol, ON midget, OFF midget and small bistratified cells (Dacey, 2004) – in the raphe region of the central macaque monkey retina (6.5-20° eccentricity). Next, we stimulate and record ON parasol and OFF parasol cells in the raphe and demonstrate that these two cell types can be identified and distinguished from ON and OFF midget cells based solely on their recorded electrical signatures, and can be activated with high spatial and temporal precision. Third, we use responses to electrical stimulation in parasol cells to estimate the visual perception that can be expected with a central-targeting implant that algorithmically optimizes electrical stimulation (Shah et al., 2019). Although the selectivity of stimulation was lower in the central compared to the peripheral retina, the unwanted activation of axons was reduced, and the expected image quality was substantially higher. These results support the possibility of high-fidelity vision restoration in the central retina with an epiretinal implant.

Results

To understand the quality of vision restoration that can be achieved in the central retina, we performed electrical recording and stimulation *ex vivo* in three steps that parallel the function of

a future implanted device. We first probed the functional organization and neural code of the major RGC types and the ability to distinguish them purely based on recorded electrical features, using large-scale recordings. Next, we determined the sensitivity and selectivity of RGC responses to electrical stimulation, using large-scale stimulation and recording, and compared them to results from the peripheral retina. Finally, using an optimization algorithm for electrical stimulation combined with a simulation of visual perception, we inferred the quality of an image that can be reproduced in the central retina by electrical activation of RGCs.

Identification of the major RGC types in the central retina

To determine the functional organization of major RGC types in the central retina, light responses were obtained from regions of macaque retina 1-4.5 mm from the fovea along the temporal horizontal meridian, in the *raphe* region (Vrabec, 1966). Recordings from the more commonly studied peripheral retina (5-12 mm temporal equivalent eccentricity) were also obtained for comparison. The preparation was visually stimulated with white noise while recording spikes from hundreds of RGCs simultaneously, using a multi-electrode array with 512 electrodes and 30 μm or 60 μm pitch (Litke et al., 2003). The spike-triggered average (STA) stimulus was computed for each cell, summarizing the spatial, temporal and chromatic properties of the light response (Chichilnisky, 2001; Chichilnisky and Kalmar, 2002). Cell types were then classified by clustering based on the spatial receptive field obtained from the STA, the time course of the STA, and the spike train autocorrelation function (Field et al., 2007).

This approach revealed the five numerically dominant primate RGC types -- ON parasol, OFF parasol, ON midget, OFF midget, and small bistratified (Fig. 1A) – which are readily identified by their light response characteristics and densities relative to other types (Dacey, 2004; Rhoades et al., 2019). Recordings from ON and OFF parasol cells were consistently nearly complete, as evidenced by the mosaic organization of their receptive fields (Fig. 1A). Recordings of ON and OFF midget cells and small bistratified cells were less complete (Fig. 1A). As expected, the spatial densities of these five cell types were higher than those of the corresponding cell types in the peripheral retina (Fig. 1B).

Neural code of the major RGC types in the central retina

Restoring visual signals requires a quantitative description of the visual neural code of major RGC types, which is incompletely understood in the raphe. To summarize the neural code, the linear-nonlinear Poisson (LNP) cascade model was fitted to responses to white noise (Chichilnisky, 2001; Chichilnisky and Kalmar, 2002), and the spatial receptive field, response time course and contrast response relationship from the model were examined (Fig. 1C-E). The receptive fields of raphe parasol and midget cells were on average smaller than those of peripheral parasol cells (Fig. 1D), as expected from anatomical studies (Dacey and Petersen, 1992; Watanabe and Rodieck, 1989). The response time course was slower in central parasol cells (Sinha et al., 2017), with a more exaggerated ON-OFF asymmetry (Fig. 1D). Contrast responses were more linear in raphe parasol and midget cells, with a less exaggerated ON-OFF asymmetry (Fig. 1D). These findings reveal modest light response differences between raphe and peripheral RGCs, building on prior work in the central retina ((Lee et al., 1990; Passaglia et al., 2002; Sinha et al., 2017; Solomon et al., 2002), see Discussion). However, the overall visual neural code is largely preserved.

Identification of major RGC types using intrinsic electrical features

High-fidelity vision restoration requires targeting distinct cell types without using their light responses for cell type identification, because light responses are degraded or absent during degeneration. Previous work in the peripheral retina showed that the major RGC types can be distinguished using their intrinsic electrical properties (Li et al., 2015; Madugula et al., 2022; Richard et al., 2016). However, it is unknown whether this is feasible in the raphe.

Raphe RGCs exhibited differences in electrical properties that can be leveraged for distinguishing cell types. First, the axon spike conduction velocity of each cell was estimated from its *electrical image* (the average spatiotemporal footprint of its spike (Litke et al., 2003; Petrusca et al., 2007)) by calculating differences in the time of recorded spikes at multiple electrodes located along the axon (see Methods). Parasol cells exhibited significantly higher conduction velocities than midget cells and could be distinguished reliably on this basis (Fig. 1F). Next, within the parasol and midget cell classes, principal components analysis (PCA) on the spike train autocorrelation functions revealed two clearly-defined clusters, which reliably distinguished the ON and OFF types (Fig. 1F). The degree of cell type separability using these features was comparable to that in the peripheral retina (Fig. 1G).

In what follows, this study will focus on ON and OFF parasol cells because the nearly complete recordings obtained from them (Fig. 1A) permit a view of the full visual signal that these populations could transmit in response to electrical stimulation.

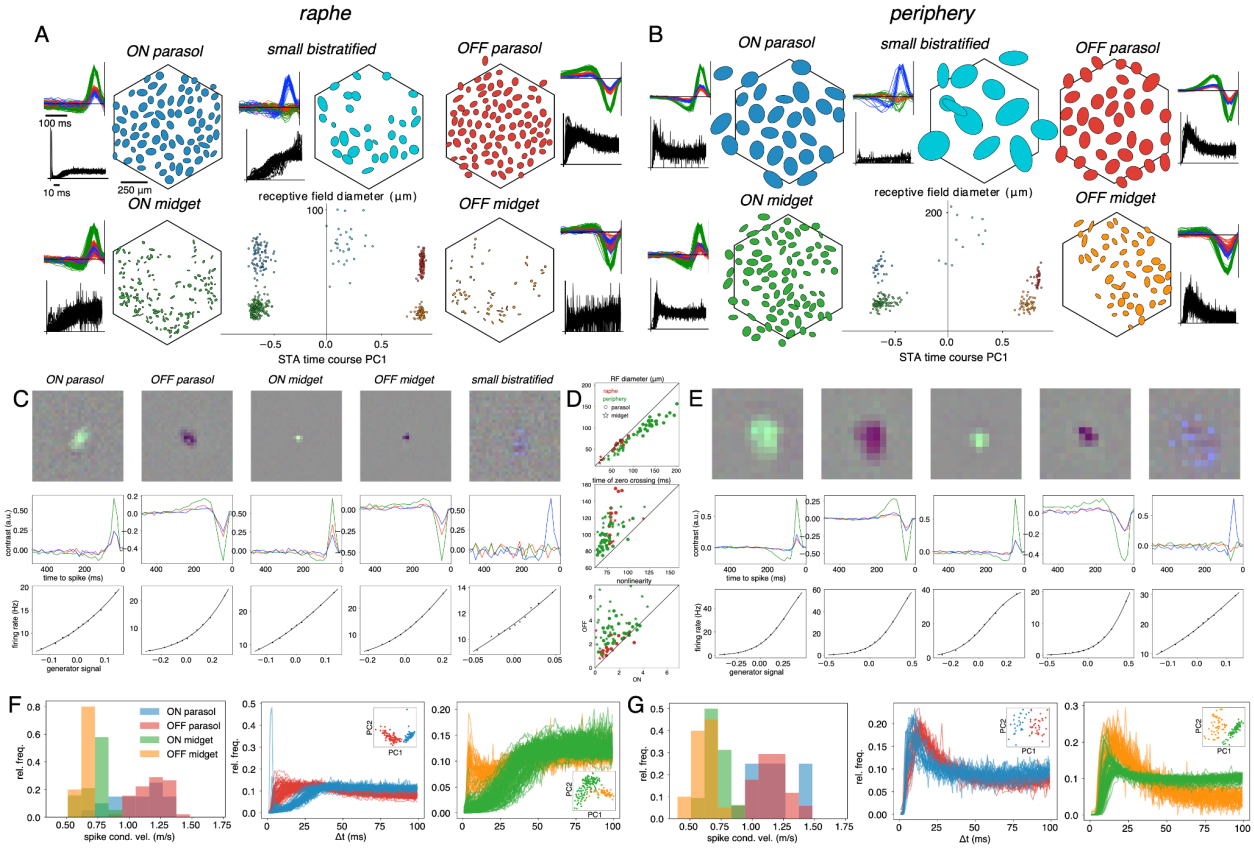

Figure 1. Identification of major retinal ganglion cell (RGC) types using light response and spiking properties in the raphe (left) and peripheral retina (right). A. Example recording from the raphe. Clusters of RF diameter and the first principal component (PC1) of the time course (and occasionally the PC1 of the spike train autocorrelation function, not shown) were used to separate and identify known functionally-distinct cell types. The receptive field (RF) mosaics, spike-triggered average (STA) time courses and spike train autocorrelation functions are shown for each cell type. The ellipses denote the 2σ boundary of a 2D Gaussian fit to the spatial component of the STA. Note: other cell types were also observed, forming different clusters, but for clarity, only data for the five numerically-dominant RGC types -- ON parasol, OFF parasol, ON midget, OFF midget, and small bistratified -- are shown. Black hexagonal outline denotes the approximate location of the electrode array. B. Same as A, but in the peripheral retina. C. Example spatial component of STA (top), STA time course (middle) and contrast-response relationship (bottom) for each of the five numerically-dominant cell types in a single raphe preparation. D. Comparisons of RF diameter (top), time of zero crossing of STA time course (middle) and nonlinearity index extracted from the contrast-response relation (see Methods, bottom) between raphe and peripheral ON and OFF parasol and midget cells. Each data point represents the median value from a single preparation. E. Same as C but for a peripheral preparation. F. Axon spike conduction velocities and spike train autocorrelation functions of ON parasol, OFF parasol, ON midget and OFF midget cells in the central retina (raphe region). These two features together can reliably distinguish these four cell types. One preparation was used to show cell type separability with axon spike conduction velocity whereas a separate preparation was used to show separability using the spike train autocorrelation function. G. Same as F but in the peripheral retina. A single peripheral preparation was used to show separability with axon spike conduction velocity and spike train autocorrelation function.

Selectivity of electrical stimulation of central parasol cells

The neural code of the retina is characterized by cell type-specific patterns of RGC activity (Roska and Meister, 2013), and the distinct RGC types are spatially intermixed in the inner retina. Therefore, reproducing the neural code accurately would typically require single-cell resolution control over RGC activity. To test whether this is possible in the raphe, spikes were identified in recorded traces immediately after application of a brief current pulse (triphasic, 0.15 ms) at each electrode, over a range of current levels (0.1-4.1 μA), each repeated 25 times (see Methods), using the 512-electrode array with 30 μm pitch. These small currents typically directly evoke a single spike in one or more RGCs near the electrode (Jepson et al., 2013; Sekirnjak et al., 2006, 2008). For each recorded cell, the evoked spike probability was calculated at each current level, and a sigmoid fitted to response probability as a function of applied current was used to estimate the activation threshold, i.e. the current that produced a spike with probability 0.5. Activation thresholds were similar between raphe and peripheral parasol cells (Fig. 2A). For each electrode, the axon bundle activation threshold was also calculated, using an algorithm that identified bi-directional propagation, a signature of axon activation (Grosberg et al., 2017) (see Methods). Although activation thresholds were similar between raphe and peripheral preparations, a larger fraction of raphe RGCs (78%) could be activated with a probability of 0.5 without activating axon bundles compared to the peripheral retina (50%, Fig. 2B), expanding on previous work (Grosberg et al., 2017) and presumably reflecting the lower density of passing peripheral RGC axons in the raphe region (Vrabec, 1966).

To summarize how precisely individual RGCs can be activated, a *selectivity index* was computed for each cell (see Methods). Across eight raphe preparations (Fig. 2C, top two rows), the average selectivity index of all ON and OFF parasol cells within the preparation ranged from 0.64-0.79 (Fig. 2D, top two panels). By comparison, the average selectivity in the peripheral retina (Fig. 2C, bottom two rows) was systematically higher (Fig. 2D, bottom two panels), ranging from 0.84-1.0 across preparations, as expected from the lower RGC density, and consistent with previous findings (Grosberg et al., 2017). In summary, single cell activation is often possible in the central retina, but is less frequently achieved than in the peripheral retina.

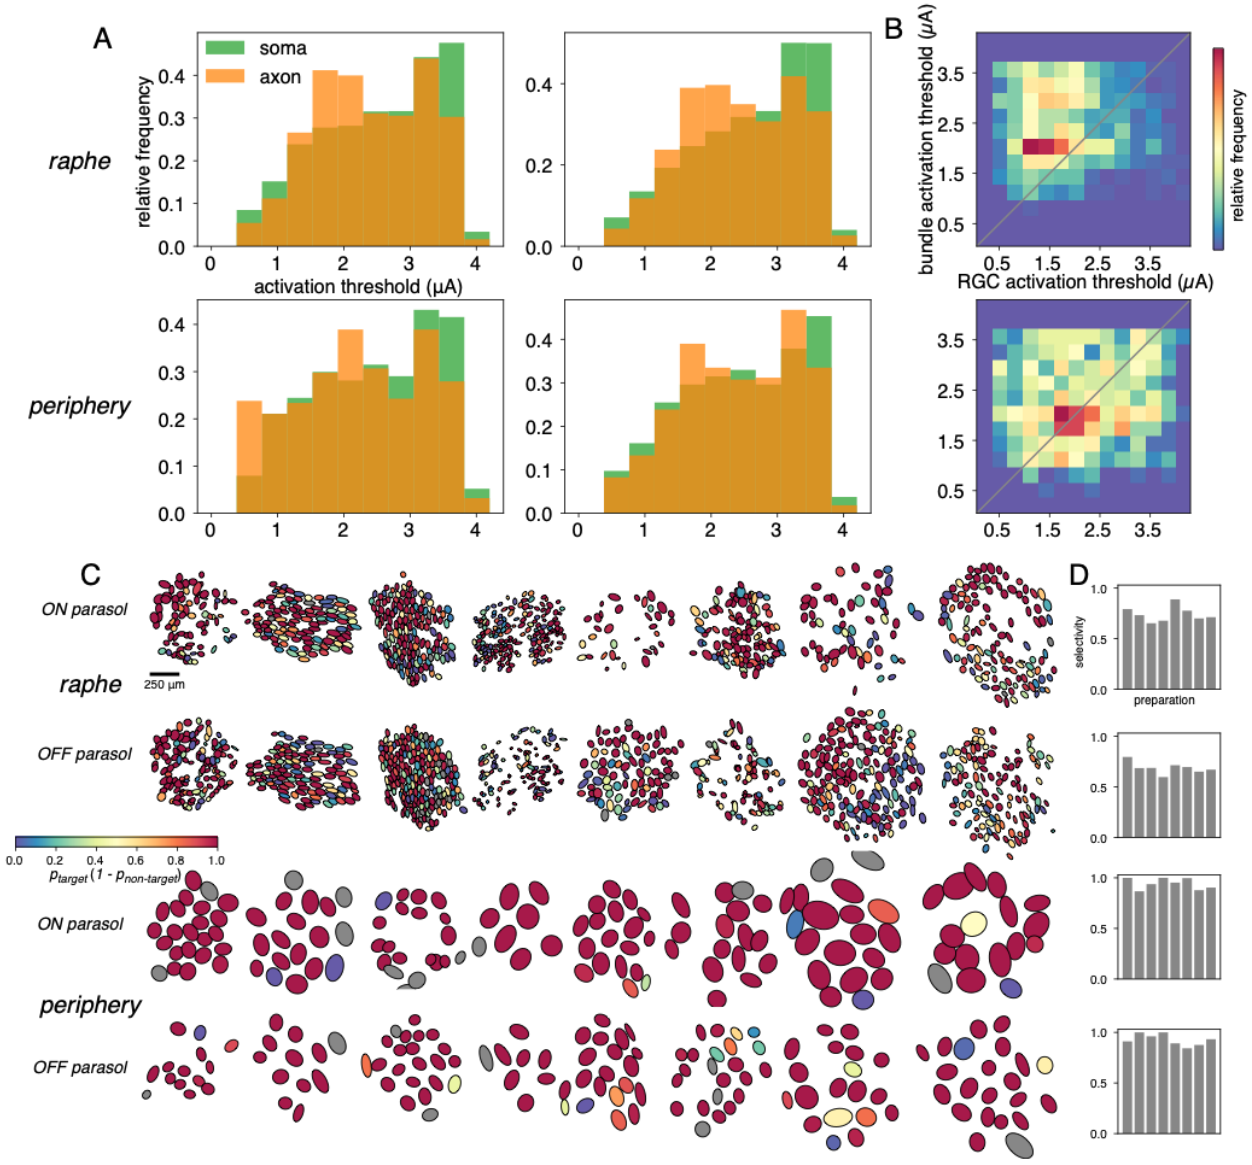

Figure 2. Electrical stimulation in the raphe and peripheral retina. A. Somatic and axonal activation thresholds in ON parasol cells and OFF parasol cells in the raphe (top) and peripheral retina (bottom). B. Two-dimensional histogram of parasol cell activation thresholds vs. axon bundle activation thresholds in the raphe (top) and the peripheral retina (bottom), pooled across preparations. Each data point accumulated in the histogram represents a comparison between the axon bundle activation threshold and the lowest RGC activation threshold observed on a given electrode, in a given preparation. C. ON parasol cell and OFF parasol cell receptive fields (RFs) for eight raphe (top) and eight peripheral retina (bottom) preparations. The RF of each cell is colored according to its selectivity index (see Results for details). RFs colored gray were excluded from electrical stimulation analysis because their evoked signals could not be distinguished from noise (see Methods). D. Average selectivity index values within each cell type and preparation.

Inference of high-fidelity vision restoration of in the central retina

To determine whether electrical stimulation in the raphe can restore high-resolution vision, an inference of visual perception was performed based on image reconstruction from measured

electrically-evoked responses, coupled with a recently-developed optimization algorithm for electrical stimulation (Shah et al., 2019). To obtain an estimate of optimal image reconstruction from RGC responses, LNP model responses to 40,000 flashed white noise images were generated (10,000 for four different pixel sizes). Linear reconstruction filters for each RGC were then computed using least squares regression between the images and model responses (Brackbill et al., 2020; Warland et al., 1997). Then, using the measured electrically-evoked responses of each cell at each stimulating electrode and current amplitude below axon bundle activation threshold (Fig. 2A,B), a spatiotemporal dithering strategy was used to select a sequence of rapidly-delivered electrical stimuli that would maximally decrease the expected mean-squared error (MSE) between the target image and the image reconstructed from the electrically-evoked responses (Shah et al., 2019).

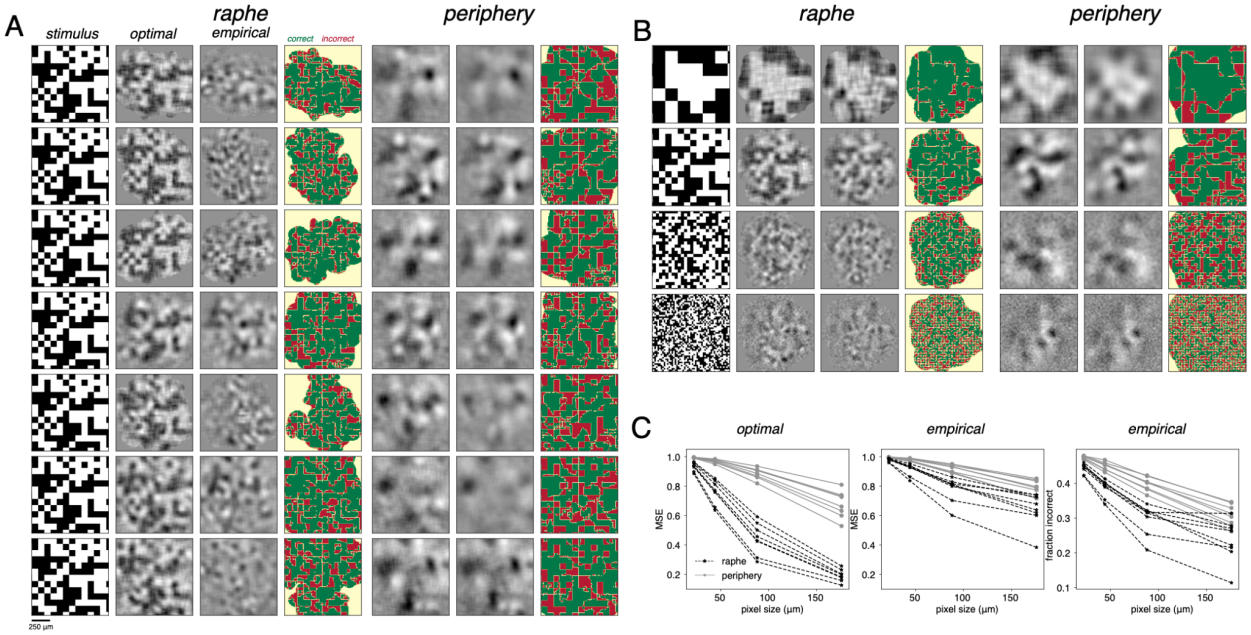

Figure 3. Inference of visual perception in the raphe and peripheral retina with white noise. A. Example linear reconstruction based on evoked RGC responses in individual raphe and peripheral preparations for a single white noise image (pixel size 88 µm). Each row is a distinct retinal preparation for each retinal location (two preparations per row). The first column shows the original image; the second column shows the optimal reconstruction (i.e. achievable with perfect control over the firing of each RGC); the third column shows the reconstruction that is achievable by optimized stimulation based on recorded evoked responses ((Shah et al., 2019), see Methods); the fourth column shows the pixels that were incorrectly reconstructed relative to the original image (red=incorrect, green=correct). Columns 5-7 show the same as columns 2-4 but in the peripheral retina. B. Reconstruction of white noise images with different pixel sizes (22, 44, 88, and 176 µm) in an example raphe (left) and peripheral (right) preparation. Each row is for a distinct white noise image. Columns show the same as in A. C. Normalized mean-squared error (MSE) in the optimal reconstructions (left) empirical reconstructions (middle) and fraction of incorrectly-reconstructed pixels in the empirical reconstructions (right). Each data point denotes the average of 15 images at each pixel size examined, for a single preparation.

Application of this method revealed substantial differences in reconstruction performance between the raphe and peripheral retina. Compared to the peripheral retina, the *optimal* reconstructions (i.e. the reconstruction that could in principle be achieved with perfect control over RGC activation) and the *empirical* reconstructions (i.e. reconstructions that could be achieved with measured activation probabilities) in the raphe were both systematically higher quality than those in the peripheral retina, with sharper edges and more distinct regions that

closely matched the structure of the stimulus (Fig. 3A). This was evident over a range of pixel sizes (Fig. 3B). Although image reconstruction was less faithful at smaller pixel sizes (Fig. 3B), image quality in the central retina was still higher at each pixel size examined.

To quantify reconstruction performance, two quantities were examined: the MSE between the reconstruction and original stimulus, and the fraction of pixels with incorrect polarities in the reconstruction. In the central retina, both MSE and the fraction of incorrect pixels were systematically lower compared to the peripheral retina, for each pixel size examined (Fig. 3C). This finding suggests that despite stimulation being less selective in the central retina (Fig. 2C,D), it is still sufficient to produce high quality visual perception. The difference in MSE and fraction of incorrect pixels between the optimal and empirical reconstructions was higher in the central retina (Fig. 3C), presumably due to the diminished selectivity relative to the peripheral retina (Fig. 2C D, see Discussion).

Reconstruction of behaviorally-relevant stimuli and natural images

Restoring vision that can support visually-guided behaviors, such as reading and object recognition, requires the perception of elementary visual features, such as text symbols and shapes. To test how well perception of such features can be restored in the central retina, reconstruction of alphanumeric characters and emoticons was performed using the approach described above (Fig. 4A). For a variety of characters, the original image was more visually recognizable in the raphe reconstructions compared to the peripheral retina reconstructions (Fig. 4A), suggesting that low-level visual features can be artificially restored with high fidelity in the raphe.

The natural visual world contains a vast range of objects and textures. To test how well the perception of naturalistic stimuli can be restored in the central retina using electrical stimulation along with a natural scenes prior, responses to 10,000 flashed natural images from the ImageNet database (Fei-Fei et al., 2009) were generated from the LNP model, and linear reconstruction filters were computed. Then, a previously described convolutional autoencoder was trained to enhance reconstructions obtained from the linear filters by learning a mapping from images linearly-decoded from RGC spike trains and natural images (Parthasarathy et al., 2017). The stimulation approach described above was applied to a set of natural images and the resulting linear reconstructions were enhanced using the autoencoder. Several aspects of the images (distinct shapes, textures and high spatial frequency content) were better captured by the raphe reconstructions compared to the peripheral retina reconstructions (Fig. 4B), suggesting that features of naturalistic images can also be restored in the raphe.

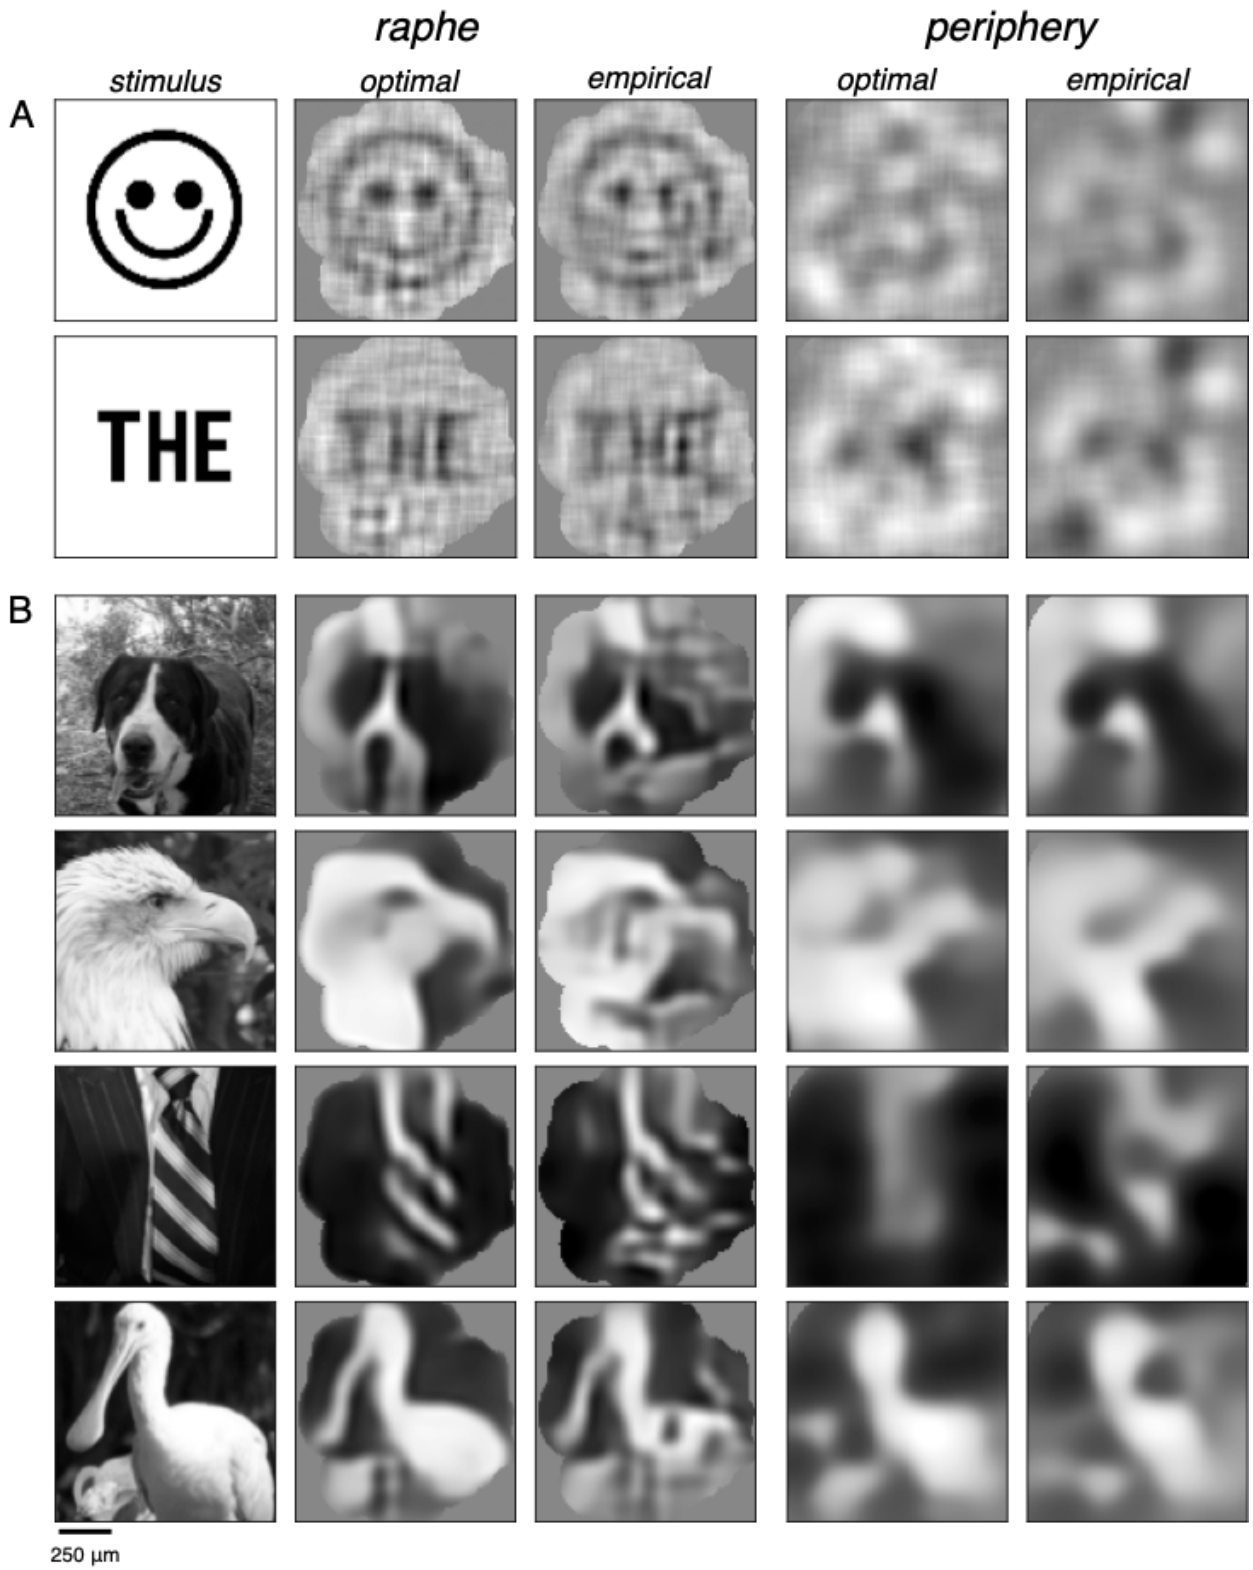

Figure 4. Inference of visual perception in the raphe and peripheral retina retina with text symbols and naturalistic images. A. Reconstruction of distinct text symbols in an example raphe (left) and peripheral (right) preparation. Each row corresponds to a distinct text symbol. The first column shows the original image; the second column shows the optimal reconstruction (i.e. achievable with perfect control over the firing of each RGC); the third column shows the reconstruction that is achievable by optimized stimulation ((Shah et al., 2019), see Methods). Columns 4 and 5 are

the same as 2 and 3 but for an example peripheral preparation. B. Same as C but for naturalistic images from the ImageNet database (Fei-Fei et al., 2009), after image enhancement with a convolutional autoencoder ((Parthasarathy et al., 2017), see Methods). Scale bar: 250 μm (1.25 degrees of visual angle).

Discussion

We examined the quality of vision restoration that is possible in the raphe region of the central macaque retina with epiretinal electrical stimulation. The functional organization and light response properties of the major RGC types in the raphe was similar to that of the peripheral retina (Fig. 1A-E), and the four numerically-dominant cell types could be distinguished using their recorded electrical features alone (Fig. 1F, G). Despite less precise control over electrical activation compared to the peripheral retina (Fig. 2C), the unwanted activation of axons was lower in the raphe (Fig. 2B), and the estimated quality of vision restoration was higher (Fig. 3), over a range of spatial scales (Fig. 3B, C), and with text symbols (Fig. 4A) and naturalistic images (Fig. 4B). These findings support the viability of targeting the central retina for vision restoration with a high-fidelity epiretinal implant.

Several approaches to vision restoration with electronic implants exist today (Weiland et al., 2011, 2016), including epiretinal (Zhou et al., 2013), subretinal (Palanker et al., 2020), and suprachoroidal (Bareket et al., 2017; Petoe et al., 2021) devices. Each approach has demonstrated the feasibility of partial restoration of visual signals, and in some cases, of useful visually guided behaviors. However, no device has been designed to restore the natural neural code of the retina at the level of individual RGCs and RGC types through calibrated electrical stimulation. A device that can do this must be able to record spontaneous activity and responses to electrical stimulation (Muratore and Chichilnisky, 2020). Recording spontaneous activity is required to determine the electrical image of each cell (Litke et al., 2003), which in turn reveals its geometry and can be used, along with firing statistics, for cell type identification (Fig. 1, (Richard et al., 2016)). Recording responses to stimulation is required to determine how current passed through each electrode influences the firing of each cell (Fig. 2A), which is critical in optimizing spatiotemporal patterns of stimulation (Jepson et al., 2014a; Shah et al., 2019). An implant of this kind is likely to be necessary for faithful reproduction of the neural code of the retina.

Relatively little is known regarding differences in light response properties between central and peripheral RGCs. Previous *in vivo* studies in the macaque retina have examined several chromatic, spatial and temporal response properties in central RGCs (Lee et al., 1990; Passaglia et al., 2002; Solomon et al., 2002). In particular, differences in responses to stimuli of varying temporal frequencies were reported between central and peripheral RGCs (Solomon et al., 2002). A more recent *ex vivo* study of macaque RGCs (Sinha et al., 2017) directly compared differences in response kinetics between central and peripheral RGCs and reported that central cells have slower light responses, likely as a result of differences in cone photoreceptor kinetics. While the overall functional organization and neural code appears to be similar in the central and peripheral retina (Fig. 1A-E), the differences observed here build on prior work and reveal ON-OFF asymmetries present in the central retina and how they compare to those in the peripheral retina (Chichilnisky and Kalmar, 2002).

The higher cell density in the raphe compared to the peripheral retina produced differences in visual reconstruction performance that were not easily predictable on first principles. In particular, in spite of reduced stimulation selectivity (Fig. 2C), the fidelity of visual reconstruction was higher in the raphe (Figs. 3,4), because of the greater reconstruction resolution afforded by more cells per unit area and the reduced activation of passing axons. However, the difference in performance between the optimal reconstructions (i.e. the inferred visual performance achievable with a device that has perfect control over the neural code) and the empirical reconstructions (i.e. the inferred visual performance possible with the observed stimulation selectivity), was systematically larger in the raphe (Fig. 3C). This finding is consistent with the lower observed selectivity in raphe preparations (Fig. 2C,D), which produces activation patterns that are less effective for visual reconstruction.

The selectivity of single-electrode stimulation used here is imperfect and thus cannot precisely reproduce the neural code, but could potentially be enhanced with multi-electrode stimulation strategies. In particular, previous work in the peripheral macaque retina demonstrated that selectivity can be enhanced at cellular resolution by applying customized bi-electrode and tri-electrode stimulation strategies as well as center-surround local returns (Fan et al., 2019; Jepson et al., 2014b). These findings build on a long history of multi-electrode stimulation approaches (Bonham and Litvak, 2008; Martens et al., 2011; Sweeney et al., 1990; Townshend and White, 1987), and demonstrate their applicability at the resolution of individual cells. The application of multi-electrode electrical stimulation to augment single-electrode stimulation in the raphe could potentially substantially improve reconstruction.

Avoidance of axon bundle activation could also potentially be enhanced. Although axon bundle activation was lower in the raphe relative to the peripheral retina, as expected based on the anatomical arrangement of axons (Vrabec, 1966) and on a previous finding (Grosberg et al., 2017), a substantial number of electrodes activated axons at lower current levels than RGCs, presumably because of their physical proximity. Thus, stimulation with these electrodes must be avoided: activation of axon bundles is likely to produce widespread, poorly controlled spiking in neurons located off the electrode array, with the potential for large, irregular phosphenes, as has been observed with much coarser stimulation approaches used in present-day epiretinal implants (Nanduri, 2011). Recent work, however, suggests an active axon avoidance strategy: in the peripheral retina, a bipolar, bi-electrode electrical stimulus often preserved somatic activation thresholds, while increasing axonal activation thresholds (Vilkhu et al., 2021), as predicted from biophysical models of cellular activation (Rattay, 1986, 1999; Rattay and Resatz, 2004). Such a stimulus could potentially be applied in a future device to minimize the activation of axon bundles without reducing the number of RGCs that can be targeted for stimulation.

Several limitations of this study could be addressed in future work.

(1) Analysis of parasol cells does not provide a complete picture of the neural code and image reconstruction that could be possible with a retinal implant. Parasol cells constitute ~15% of the overall RGC input to the brain (Dacey, 2004) with an even lower percentage in the central retina (Marshak, 2009). Midget cells (~50%) mediate high-acuity vision in primates, and color vision is thought to be mediated primarily by midget and small bistratified (~5%) cells (Dacey, 2004; Field et al., 2007, 2010). Although most midget and small bistratified cells were not identified in the present raphe recordings (Fig. 1A), they were likely stimulated – indeed, in the peripheral retina, ON

and OFF midget and small bistratified cells (which are collectively ~3-4x as numerous as parasol cells), are on average stimulated at the same current levels as parasol cells (Jepson et al., 2013). When all cells are included, the inferred reconstructions (Figs. 3,4) could potentially be more accurate (because additional cells confer additional information for image reconstruction) or less accurate (because stimulation would likely be less selective with additional cells included). The selection of parasol cells for analysis was motivated by the ability to record from complete populations and to reliably identify electrically their evoked responses because of their relatively large spikes. Improved spike sorting for reliably identifying spikes from all cell types is therefore an important direction for continued work.

(2) White noise visual stimulation for light response modeling is limited because its statistics differ greatly from the statistics of naturalistic stimuli, and thus could elicit RGC responses that are much different than would be observed in natural viewing conditions (Heitman et al., 2016). Thus, future work could focus on performing natural scenes stimulation for response modeling in the central retina (Heitman et al., 2016; McIntosh et al., 2016). With more complex stimuli, models that take into account additional properties such as cell-cell interactions (Mastronarde, 1983; Pillow et al., 2008) and spatial nonlinearities (Hochstein and Shapley, 1976; Shah et al., 2020) may become increasingly important. Machine learning approaches have also been shown to be valuable in capturing responses to natural scenes in the salamander retina more accurately than simple linear-nonlinear cascade models (McIntosh et al., 2016) and could potentially be applied to the primate retina.

(3) Pixel-wise mean-squared error (MSE), although a standard and intuitive starting point for image analysis, does not fully capture the quality of image reconstruction that is relevant for vision (Wandell, 1995). Future work could optimize the perceptual similarity of reconstructed images using metrics such as SSIM (Wang et al., 2004) during optimization and for the assessment of reconstruction quality, to more closely parallel the function of the visual system.

(4) Linear reconstruction (Brackbill et al., 2020; Warland et al., 1997), although simple and computationally tractable, is likely too simple an approximation for how the visual system extracts information from RGC spike trains to produce visual perception and visually-guided behaviors. Future work could apply more sophisticated nonlinear reconstruction methods for optimizing visual perception (Kim et al., 2021; Parthasarathy et al., 2017; Wu et al., 2022; Zhang et al., 2020), such as the convolutional autoencoder used to enhance the natural image reconstructions here ((Parthasarathy et al., 2017), Fig. 4B) or denoisers trained on naturalistic images (Wu et al., 2022), which may provide a more realistic implicit model for visual perception by incorporating natural scenes priors.

With these caveats in mind, the present results suggest, for the first time, the possibility of restoring high-resolution vision with an epiretinal implant that reproduces the neural code of the central retina. The approach can in principle also be applied to other circuits in the central nervous system, potentially improving the function of brain-machine interfaces of the future.

Methods

Experimental procedures

Eyes were obtained from terminally anesthetized macaque monkeys (*Macaca mulatta, Macaca fascicularis*) of either sex (32 males, 6 females) used by other researchers, in accordance with Institutional Animal Care and Use Committee guidelines. Immediately after enucleation, the eye was hemisected in room light and the anterior portion of the eye and the vitreous humor was removed. The posterior portion of the eye was stored in oxygenated, bicarbonate-buffered Ames' solution (Sigma) at 31-33°.

In infrared or dim lighting, small (~2x2 mm) segments of retina, along with the attached retinal pigment epithelium (RPE) and choroid, were separated from the sclera. In most preparations, the retina was then separated from the RPE; in one preparation (Fig. 1B), the RPE was left attached. The retina was placed retinal ganglion cell (RGC) side down on a custom high-density multi-electrode array and a transparent dialysis membrane was placed on the photoreceptor side of the retina. A total of 52 retinal preparations were used: 15 from the raphe region and 37 from the periphery. For raphe preparations, regions 1-4.5 mm (4.5-20°) eccentricity along the temporal horizontal meridian were obtained. For peripheral preparations, segments in both temporal and nasal regions of the retina were obtained, with eccentricities ranging 5-12 mm (22-56° temporal equivalent eccentricity). Once the retina was mounted on the array, it was superfused with oxygenated, bicarbonate-buffered Ames' solution at 33-35°C.

Visual stimulation and recording

The retina was typically stimulated with a white noise visual stimulus from a gamma-corrected cathode ray tube monitor refreshing at 120 Hz. The three monitor primaries were modulated independently for spectral variation, or in a coordinated fashion for a black-and-white stimulus. The stimulus refresh interval was either 8.37, 16.74, or 33.47 ms and the size of each stimulus pixel at the photoreceptor layer ranged from ~22-176 μm. For a single preparation (Fig. 1B), the retina was stimulated with a gamma-corrected organic light emitting diode monitor refreshing at 60 Hz. In this case, the stimulus refresh interval was 16.57 ms and the size of each stimulus pixel at the photoreceptor layer was ~36 μm. The recordings were obtained using a custom multi-electrode stimulation and recording system with 512 electrodes (5-15 μm diameter) with either 30 μm pitch, covering an area of 0.43 mm^2 or 60 μm pitch covering an area of 1.7 mm^2 (Hottowy et al., 2008, 2012; Litke et al., 2003). Raw signals from the channels were amplified, filtered (43–5,000 Hz), multiplexed, digitized at 20 kHz and stored for off-line analysis, as described previously (Litke et al., 2003). KiloSort (Pachitariu et al., 2016) was then applied to the raw recordings to separate and identify unique cells and their spike times.

Cell type classification

To classify distinct cell types based on their light response properties, the spike-triggered average (STA) was computed for each cell (Chichilnisky, 2001). The receptive field (RF) of each cell was estimated by determining the pixels within the STA with signal significantly above the noise level (Gauthier et al., 2009). The temporal component of the STA was computed by summing the primaries of pixels within each cell's RF. The spatial component of the STA was computed for each cell by taking the inner product between the time course and the full STA,

and the output was fitted with a 2D Gaussian. Clusters of light response properties (RF diameter, first principal component of time course) and spiking properties (first principal component of the spike train autocorrelation function, see below) were used to identify distinct cell types (Field et al., 2007). Recordings in most peripheral and some central retina preparations typically revealed the five numerically-dominant RGC types: ON parasol, OFF parasol, ON midget, OFF midget, and small bistratified. In most preparations, recordings from ON and OFF parasol cells were nearly complete, whereas recordings from ON and OFF midget and small bistratified cells were less complete.

Linear-nonlinear Poisson cascade model

To summarize the light response properties of the major RGC types, the linear-nonlinear Poisson (LNP) cascade model was fitted to responses to white noise (Chichilnisky, 2001; Chichilnisky and Kalmar, 2002). At the retinal eccentricities recorded from here, the RF diameters of raphe RGCs are approximately half the size of those in the periphery (Fig. 1A-E). Therefore, to ensure a fair comparison between model parameters in the raphe and periphery, comparisons were made using RGB stimuli (16.74 ms refresh interval) with a 22 μm pixel width in the raphe and a 44 or 55 μm pixel width in the periphery.

First, the RF of each cell was determined by fitting a 2D Gaussian to the spatial component of the STA (see above) and calculating the geometric mean of each standard deviation parameter multiplied by 2. To summarize the response time course of each cell, the STA time course of each cell was fitted with a difference of cascades of low pass filters (Chichilnisky and Kalmar, 2002) of the form $f(t) = p_1(t/\tau_1)^n \mathrm{e}^{-n(t/\tau_1-1)} - p_2(t/\tau_2)^n \mathrm{e}^{-n(t/\tau_2-1)}$ where $f(t)$ is the time course, t is the time index, and n, p_1, τ_1, p_2, and τ_2 are free parameters. Then, the time of zero crossing of the time course was estimated by numerically solving for the root of the fitted function using the bisection method. Last, the contrast response relationship for each cell was estimated from the STA, stimulus and spike train (Chichilnisky, 2001; Chichilnisky and Kalmar, 2002) and fitted with a smooth function of the form $f(x) = \alpha L(\beta x + \gamma)$, where x is the generator signal, L is the logistic cumulative distribution function, and α, β and γ are free parameters. The *nonlinearity index* for each cell was defined as the logarithm of the ratio of $f'(x_{+\sigma})$ and $f'(x_{-\sigma})$ where $f'(x)$ is the first derivative of the contrast response relation at x (computed numerically using the central difference method) and x_σ is the 1 standard deviation generator signal value estimated from the distribution of all generator signal values for that cell during the recording.

Identification of axonal and somatic electrodes

For each cell, the cellular compartments recorded by each electrode were identified from the average waveform of recorded spikes. First, the electrophysiological image (EI) was computed by averaging the spatiotemporal voltage signal on the array associated with each spike time, revealing the spatiotemporal “electrical footprint” of the average spike from the cell (Litke et al., 2003; Petrusca et al., 2007). Features of spike waveforms obtained from the EI were then used to classify them into distinct cellular compartments (Litke et al., 2003; Müller et al., 2015). In general, somatic and dendritic waveforms exhibit biphasic structure, whereas axonal waveforms exhibit triphasic structure, reflecting biophysical differences between compartments. Based on these known differences, a heuristic-based automated method was applied to label each

electrode as either “dendritic”, “somatic”, “axonal”, or “mixed” (either a superposition of somatic and dendritic or somatic and axonal waveforms). First, waveforms with a larger positive than negative phase were labeled as “dendritic”. Second, the waveforms at remaining electrodes were identified as “somatic”, “mixed”, or “axonal”, based on the ratio of the first positive peak to the second positive peak. Thresholds for each respective compartment were determined by two empirically observed inflection points (0.05 and 1.6 respectively) in the pooled distribution of ratios obtained from many cells and electrodes in peripheral recordings.

Analysis of electrical features for distinguishing cell types

Axon spike conduction velocity and the spike train autocorrelation function were used together to distinguish ON and OFF parasol and ON and OFF midget cells based on intrinsic electrical features alone (Richard et al., 2016). Each preparation in this analysis (Fig. 1F, G) was acquired using the electrode array with 30 μm inter-electrode spacing to ensure consistent estimates of spike conduction velocities.

Axon spike conduction velocity was estimated from the electrical images of ON and OFF parasol and midget cells, as described previously (Madugula et al., 2022). Briefly, axonal electrodes for each cell were identified (see above) and pairwise velocity estimates were computed by dividing the distance between each electrode by the time difference in the negative peak amplitude, for each pair of axonal electrodes. Then, a velocity estimate was computed by averaging the estimate from each pair of electrodes weighted by the product of the peak amplitudes.

The autocorrelation function of the spike trains from white noise visual stimulation were computed for each ON and OFF parasol cell and ON and OFF midget cell. Briefly, spike times were binned with 1 ms precision and a histogram of pairwise spike time differences (ranging from 1 to 100 ms) was calculated, and normalized by its L2 norm. To determine if ON and OFF types within the parasol and midget cell classes could be separated by features of the autocorrelation function, principal components analysis (PCA) was performed on a data matrix comprised of either ON and OFF parasol or ON and OFF midget cell autocorrelation functions, and projections onto the first two principal components were analyzed (Fig. 3F,G, insets).

Responses to Electrical stimulation

To identify the responses to electrical stimulation of each recorded RGC, a current pulse was delivered to each electrode on the array, in a random sequence, while recording responses on all electrodes simultaneously (Grosberg et al., 2017; Jepson et al., 2013). The pulse was triphasic, positive first, and charge-balanced (50 μs per phase; relative ratios 2:-3:1). The current amplitude range tested (second phase) was 0.1-4.1 μA (~40 amplitudes, logarithmic scale), with each amplitude repeated 25 times. In all electrical stimulation experiments, the electrode array with 30 μm inter-electrode spacing was used for all preparations.

An automated spike sorting method was used to identify RGC responses to electrical stimulation. First, for each stimulating electrode, RGCs were identified as candidates for electrically evoked activity if their EIs (obtained during visual stimulation) exhibited peak amplitudes larger 0.5 times the standard deviation of the recorded electrical noise on that electrode. Each RGC was then assigned a set of recording electrodes from its EI recording

signals with power greater than two times the power of electrode noise on each electrode to be used to determine responses. Cells that did not have a single electrode whose signal power was greater than 2 times the power of the electrical noise threshold were excluded from analysis because their signals could not be distinguished from noise. Then, for each current amplitude, across all RGCs and their respective recording electrodes, voltage traces recorded in the 3 ms period following stimulus application were grouped using unsupervised clustering. For each pair of clusters, signals were subtracted, aligned, and averaged to obtain differences, which reflected the cell(s) firing in one cluster but not the other. These residuals were iteratively compared to each of the recorded spike waveforms obtained from the EIs of all RGCs and their respective recording electrodes to identify the set of RGCs contributing to the responses in each trial for each amplitude.

This algorithm produces response probabilities as a function of current level, for each RGC and stimulating electrode, which is in general a monotonically increasing function (Hottowy et al., 2012; Jepson et al., 2013). The relationship between current amplitude and response probability for each RGC and stimulating electrode was then fitted by a sigmoidal curve of the form $p(a) = \frac{1}{\mathrm{e}^{-k(a-b)}}$, where $p(a)$ is the probability of activation, a is the current amplitude, and k and b are free parameters, using maximum likelihood estimation. The stimulation threshold, i.e. the current amplitude required to elicit a response probability of 0.5, was extracted from the fit (Jepson et al., 2013).

Identification of axon bundle activation thresholds

Axon bundle activation thresholds on each electrode were determined by an automated method based on a previously described algorithm (Tandon et al., 2021). The algorithm was modified to avoid bias resulting from differences in array geometries (this work used a smaller, hexagonal array whereas the algorithm described in (Tandon et al., 2021) was developed using a larger, rectangular array) and axon spike amplitude differences between central and peripheral RGCs (central RGCs in general have smaller spike amplitudes than peripheral RGCs). For each preparation, a threshold voltage was first determined to identify electrodes that recorded significant axonal signals in response to electrical stimulation, as follows. For each RGC recorded during white noise visual stimulation, the electrodes recording axonal signals were identified as described above and the average axonal spike amplitude was determined. The median axonal spike amplitude across all recorded RGCs was computed and was taken to be the threshold voltage. Next, to determine the axon bundle activation threshold, for each stimulus current applied, electrodes were first identified as either activated or inactivated, depending on whether the recorded signal was above the threshold voltage. Activity on the array was identified as an axon bundle activation event when the activated electrodes formed a contiguous path reaching at least two non-adjacent edges of the electrode array. The bundle activation threshold was then defined as the minimum current level at which an axon bundle activation event was evoked. Electrodes near the border of the array (outer two rings of electrodes) were excluded from analysis because their proximity to the edge precludes the ability to unambiguously distinguish RGC activity from axon bundle activity.

Analysis of selectivity

To summarize the selectivity of electrical stimulation, a *selectivity* metric was computed for each ON or OFF parasol cell. First, at each stimulating electrode, the response probability at each current level below axon bundle activation threshold was determined for the target cell (p_{target}) and every other non-target cell ($p_{non\text{-}target}$). Then, for each non-target cell, the maximum value of the quantity (p_{target})(1 - $p_{non\text{-}target}$) over current levels was determined. This was repeated for each non-target cell, and the minimum value of (p_{target})(1 - $p_{non\text{-}target}$) over non-target cells was determined at this stimulating electrode. This was repeated for each stimulating electrode, and the maximum value of (p_{target})(1 - $p_{non\text{-}target}$) over stimulating electrodes was defined as this cell's *selectivity index*.

Inference of visual perception

To estimate the quality of restored vision that could be achieved with the measured responses to electrical stimulation, a recently developed stimulation algorithm (Shah et al., 2019) was applied, in simulation. This method assumes that the brain reconstructs visual stimuli from RGC spike trains linearly and optimally, and that responses from individual RGCs that occur within a single "integration time window" of downstream visual processing are summed linearly to produce a percept.

In this framework, an optimal linear reconstruction filter is associated with each RGC that quantifies the contribution to the image reconstruction that occurs each time the cell fires a spike. To compute the optimal reconstruction filter for each cell, parameters from the LNP model were used to generate responses to 10,000 flashed white noise images lasting 100 ms. Specifically, the firing rate was simulated by taking an inner product between the STA and the stimulus, rectifying the result. The STAs were scaled such that the average parasol cell firing rate matched that observed experimentally (~3 spikes on average per 100 ms flash in the peripheral retina). The stimuli are denoted by the matrix $S \in \mathbb{R}^{\text{images x pixels}}$ and the responses in a matrix $R \in \mathbb{R}^{\text{images x cells}}$. A reconstruction matrix $W \in \mathbb{R}^{\text{cells x pixels}}$ was computed using linear least squares regression, via the normal equations: $W = (R^T R)^{-1} R^T S$. In this framework, the brain reconstructs an estimate of the true stimulus S from the RGC responses R by the linear operation $\hat{S} = RW$ (Warland et al., 1997). Different sets of filters were learned for white noise images with varying pixel sizes (22, 44, 88, and 176 μm pixels, 10,000 training images each) to test performance at a range of spatial frequencies. This procedure was repeated with grayscaled naturalistic images from the ImageNet database (Fei-Fei et al., 2009) (10,000 training images) to produce linear reconstruction filters for naturalistic stimuli.

To determine the reconstruction that could be achieved with the responses to electrical stimulation, RGC responses probabilities for each cell-electrode pair, at current levels below the threshold for axon bundle activation, were assembled in a matrix $D \in \mathbb{R}^{\text{cells x (electrodes x amplitudes)}}$, where the columns are distinct electrical stimuli and the rows are the responses of each cell. The electrical stimulus delivered at each time step was chosen such that the change in the image reconstruction due to the evoked responses of RGCs would maximally reduce the mean-squared error (MSE) between the reconstructed image and the target image. The algorithm sequentially chooses stimuli greedily at each time step to minimize the MSE until no further reduction in error occurs. No two stimuli resulting in a cellular response probability of more than 0.04 were delivered within 5 ms of each other to avoid the refractory

period, and each stimulation pulse was assumed to be 150 μs in duration, to be consistent with the experimental data and hardware configuration. This method was applied to a total of 60 white noise images (15 for each pixel size), 2 stimuli with text symbols/words, and 6 naturalistic stimuli from the ImageNet database (Fei-Fei et al., 2009), all separate from the image sets used for training.

To approximate the optimal reconstruction, or the reconstruction that could be achieved with perfect control over the activity of each RGC, a convex solver was used to solve for the optimal nonnegative responses under squared loss between the reconstruction and the true stimulus. This response vector was then discretized by rounding down each floating point response to an integer value. Then, for each stimulus, the resulting response vector was applied to the reconstruction matrix to produce the approximate optimal reconstruction.

Convolutional autoencoder

To enhance the linear reconstructions of naturalistic images, a previously described convolutional autoencoder (Parthasarathy et al., 2017) was trained on linear reconstructions and ground truth naturalistic image pairs. First, linear reconstruction filters were learned for each cell using least squares regression (see above) on a set of 10,000 training images from ImageNet. Then, linearly reconstructed images were generated for each image in the training set using the model generated responses and reconstruction filters. This linear reconstruction was then passed through an eight layer convolutional neural network with four downsampling layers and four upsampling layers. The Adam optimizer (Kingma and Ba, 2014) with MSE loss and a learning rate of 0.0004 was used for model training. A validation set of 1,000 images from ImageNet unseen by both the linear model and the autoencoder was used to evaluate model accuracy. Twenty training epochs with a batch size of 32 were used, after which validation error did not substantially decrease further.

Analysis of image reconstructions

To quantify reconstruction quality, two quantities were calculated. First, the set of all relevant pixels was determined for each preparation, indicating the regions covered by the RFs of all ON and OFF parasol cells. Then, the squared error between the reconstructed and original images was calculated over this set of relevant pixels and normalized by the squared L2-norm of the original image within this set of pixels. Next, the fraction of reconstructed pixels whose polarity was opposite that of the original stimulus was calculated. First, a subset of the aforementioned set of relevant pixels was determined, indicating regions covered by the RFs of ON and OFF parasol cells chosen to be stimulated by the algorithm. Then, in each reconstructed image, the polarity of each pixel in the reconstruction was compared to that in the original stimulus and the fraction of pixels with opposing polarities was calculated.

References

Bareket, L., Barriga-Rivera, A., Zapf, M.P., Lovell, N.H., and Suaning, G.J. (2017). Progress in artificial vision through suprachoroidal retinal implants. J. Neural Eng. *14*, 045002. .

Bonham, B.H., and Litvak, L.M. (2008). Current focusing and steering: modeling, physiology,

and psychophysics. Hear. Res. *242*, 141–153. .

Brackbill, N., Rhoades, C., Kling, A., Shah, N.P., Sher, A., Litke, A.M., and Chichilnisky, E.J. (2020). Reconstruction of natural images from responses of primate retinal ganglion cells. Elife *9*. https://doi.org/10.7554/eLife.58516.

Chichilnisky, E.J. (2001). A simple white noise analysis of neuronal light responses. Network *12*, 199–213. .

Chichilnisky, E.J., and Kalmar, R.S. (2002). Functional asymmetries in ON and OFF ganglion cells of primate retina. J. Neurosci. *22*, 2737–2747. .

Dacey, D.M. (2004). Origins of perception: retinal ganglion cell diversity and the creation of parallel visual pathways. In The Cognitive Neurosciences, M.S. Gazzaniga, ed. (Cambridge, MA: MIT Press), pp. 281–301.

Dacey, D.M., and Petersen, M.R. (1992). Dendritic field size and morphology of midget and parasol ganglion cells of the human retina. Proc. Natl. Acad. Sci. U. S. A. *89*, 9666–9670. .

Fan, V.H., Grosberg, L.E., Madugula, S.S., Hottowy, P., Dabrowski, W., Sher, A., Litke, A.M., and Chichilnisky, E.J. (2019). Epiretinal stimulation with local returns enhances selectivity at cellular resolution. J. Neural Eng. *16*, 025001. .

Fei-Fei, L., Deng, J., and Li, K. (2009). ImageNet: Constructing a large-scale image database. J. Vis. *9*, 1037–1037. .

Field, G.D., Sher, A., Gauthier, J.L., Greschner, M., Shlens, J., Litke, A.M., and Chichilnisky, E.J. (2007). Spatial properties and functional organization of small bistratified ganglion cells in primate retina. J. Neurosci. *27*, 13261–13272. .

Field, G.D., Gauthier, J.L., Sher, A., Greschner, M., Machado, T.A., Jepson, L.H., Shlens, J., Gunning, D.E., Mathieson, K., Dabrowski, W., et al. (2010). Functional connectivity in the retina at the resolution of photoreceptors. Nature *467*, 673–677. .

Gauthier, J.L., Field, G.D., Sher, A., Greschner, M., Shlens, J., Litke, A.M., and Chichilnisky, E.J. (2009). Receptive fields in primate retina are coordinated to sample visual space more uniformly. PLoS Biol. *7*, e1000063. .

Grosberg, L.E., Ganesan, K., Goetz, G.A., Madugula, S.S., Bhaskhar, N., Fan, V., Li, P., Hottowy, P., Dabrowski, W., Sher, A., et al. (2017). Activation of ganglion cells and axon bundles using epiretinal electrical stimulation. J. Neurophysiol. *118*, 1457–1471. .

Heitman, A., Brackbill, N., Greschner, M., Sher, A., Litke, A.M., and Chichilnisky, E.J. (2016). Testing pseudo-linear models of responses to natural scenes in primate retina. bioRxiv https://doi.org/10.1101/045336.

Hochstein, S., and Shapley, R.M. (1976). Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. The Journal of Physiology *262*, 265–284. https://doi.org/10.1113/jphysiol.1976.sp011595.

Hottowy, P., Dąbrowski, W., Skoczeń, A., and Wiącek, P. (2008). An integrated multichannel

waveform generator for large-scale spatio-temporal stimulation of neural tissue. Analog Integrated Circuits and Signal Processing *55*, 239–248. https://doi.org/10.1007/s10470-007-9125-x.

Hottowy, P., Skoczeń, A., Gunning, D.E., Kachiguine, S., Mathieson, K., Sher, A., Wiącek, P., Litke, A.M., and Dąbrowski, W. (2012). Properties and application of a multichannel integrated circuit for low-artifact, patterned electrical stimulation of neural tissue. J. Neural Eng. *9*, 066005. .

Jepson, L.H., Hottowy, P., Mathieson, K., Gunning, D.E., Dabrowski, W., Litke, A.M., and Chichilnisky, E.J. (2013). Focal Electrical Stimulation of Major Ganglion Cell Types in the Primate Retina for the Design of Visual Prostheses. Journal of Neuroscience *33*, 7194–7205. https://doi.org/10.1523/jneurosci.4967-12.2013.

Jepson, L.H., Hottowy, P., Weiner, G.A., Dabrowski, W., Litke, A.M., and Chichilnisky, E.J. (2014a). High-fidelity reproduction of spatiotemporal visual signals for retinal prosthesis. Neuron *83*, 87–92. .

Jepson, L.H., Hottowy, P., Mathieson, K., Gunning, D.E., Dąbrowski, W., Litke, A.M., and Chichilnisky, E.J. (2014b). Spatially patterned electrical stimulation to enhance resolution of retinal prostheses. J. Neurosci. *34*, 4871–4881. .

Kim, Y.J., Brackbill, N., Batty, E., Lee, J., Mitelut, C., Tong, W., Chichilnisky, E.J., and Paninski, L. (2021). Nonlinear Decoding of Natural Images From Large-Scale Primate Retinal Ganglion Recordings. Neural Comput. *33*, 1719–1750. .

Kingma, D.P., and Ba, J. (2014). Adam: A Method for Stochastic Optimization. arXiv.

Lee, B.B., Pokorny, J., Smith, V.C., Martin, P.R., and Valberg, A. (1990). Luminance and chromatic modulation sensitivity of macaque ganglion cells and human observers. J. Opt. Soc. Am. A *7*, 2223–2236. .

Li, P.H., Gauthier, J.L., Schiff, M., Sher, A., Ahn, D., Field, G.D., Greschner, M., Callaway, E.M., Litke, A.M., and Chichilnisky, E.J. (2015). Anatomical identification of extracellularly recorded cells in large-scale multielectrode recordings. J. Neurosci. *35*, 4663–4675. .

Litke, A.M., Bezayiff, N., Chichilnisky, E.J., Cunningham, W., Dabrowski, W., Grillo, A.A., Grivich, M., Grybos, P., Hottowy, P., Kachiguine, S., et al. (2003). What does the eye tell the brain?: Development of a system for the large scale recording of retinal output activity. 2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515) https://doi.org/10.1109/nssmic.2003.1351852.

Madugula, S.S., Gogliettino, A.R., Zaidi, M., Aggarwal, G., Kling, A., Shah, N.P., Vilkhu, R., Hays, M.R., Nguyen, H., Fan, V., et al. (2022). Focal Electrical Stimulation of Human Retinal Ganglion Cells for Vision Restoration. https://doi.org/10.1101/2020.08.23.263608.

Marshak, D.W. (2009). Retinal ganglion cells: Anatomy. In Encyclopedia of Neuroscience, (Elsevier), pp. 211–218.

Martens, H.C.F., Toader, E., Decré, M.M.J., Anderson, D.J., Vetter, R., Kipke, D.R., Baker, K.B.,

Johnson, M.D., and Vitek, J.L. (2011). Spatial steering of deep brain stimulation volumes using a novel lead design. Clin. Neurophysiol. *122*, 558–566. .

Mastronarde, D.N. (1983). Correlated firing of cat retinal ganglion cells. I. Spontaneously active inputs to X- and Y-cells. Journal of Neurophysiology *49*, 303–324. https://doi.org/10.1152/jn.1983.49.2.303.

McIntosh, L.T., Maheswaranathan, N., Nayebi, A., Ganguli, S., and Baccus, S.A. (2016). Deep Learning Models of the Retinal Response to Natural Scenes. Adv. Neural Inf. Process. Syst. *29*, 1369–1377. .

Müller, J., Ballini, M., Livi, P., Chen, Y., Radivojevic, M., Shadmani, A., Viswam, V., Jones, I.L., Fiscella, M., Diggelmann, R., et al. (2015). High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels. Lab Chip *15*, 2767–2780. .

Muratore, D.G., and Chichilnisky, E.J. (2020). Artificial Retina: A Future Cellular-Resolution Brain-Machine Interface. In NANO-CHIPS 2030: On-Chip AI for an Efficient Data-Driven World, B. Murmann, and B. Hoefflinger, eds. (Cham: Springer International Publishing), pp. 443–465.

Nanduri, D. (2011). Prosthetic vision in blind human patients: Predicting the percepts of epiretinal stimulation. University of Southern California.

Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and Harris, K. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. In NIPS Proceedings, (Neural Information Systems Foundation, Inc.),.

Palanker, D., Le Mer, Y., Mohand-Said, S., Muqit, M., and Sahel, J.A. (2020). Photovoltaic Restoration of Central Vision in Atrophic Age-Related Macular Degeneration. Ophthalmology *127*, 1097–1104. .

Parthasarathy, N., Batty, E., Falcon, W., Rutten, T., Rajpal, M., Chichilnisky, E.J., and Paninski, L. (2017). Neural Networks for Efficient Bayesian Decoding of Natural Images from Retinal Neurons. In Advances in Neural Information Processing Systems, I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds. (Curran Associates, Inc.),.

Passaglia, C.L., Troy, J.B., Rüttiger, L., and Lee, B.B. (2002). Orientation sensitivity of ganglion cells in primate retina. Vision Res. *42*, 683–694. .

Petoe, M.A., Titchener, S.A., Kolic, M., Kentler, W.G., Abbott, C.J., Nayagam, D.A.X., Baglin, E.K., Kvansakul, J., Barnes, N., Walker, J.G., et al. (2021). A Second-Generation (44-Channel) Suprachoroidal Retinal Prosthesis: Interim Clinical Trial Results. Transl. Vis. Sci. Technol. *10*, 12. .

Petrusca, D., Grivich, M.I., Sher, A., Field, G.D., Gauthier, J.L., Greschner, M., Shlens, J., Chichilnisky, E.J., and Litke, A.M. (2007). Identification and Characterization of a Y-Like Primate Retinal Ganglion Cell Type. Journal of Neuroscience *27*, 11019–11027. https://doi.org/10.1523/jneurosci.2836-07.2007.

Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P.

(2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995–999. https://doi.org/10.1038/nature07140.

Rattay, F. (1986). Analysis of models for external stimulation of axons. IEEE Trans. Biomed. Eng. *33*, 974–977. .

Rattay, F. (1999). The basic mechanism for the electrical stimulation of the nervous system. Neuroscience *89*, 335–346. .

Rattay, F., and Resatz, S. (2004). Effective electrode configuration for selective stimulation with inner eye prostheses. IEEE Trans. Biomed. Eng. *51*, 1659–1664. .

Rhoades, C.E., Shah, N.P., Manookin, M.B., Brackbill, N., Kling, A., Goetz, G., Sher, A., Litke, A.M., and Chichilnisky, E.J. (2019). Unusual Physiological Properties of Smooth Monostratified Ganglion Cell Types in Primate Retina. Neuron *103*, 658–672.e6. https://doi.org/10.1016/j.neuron.2019.05.036.

Richard, E., Goetz, G.A., and Chichilnisky, E.J. (2016). Recognizing retinal ganglion cells in the dark. Adv. Neural Inf. Process. Syst. *28*, 2467–2475. .

Roska, B., and Meister, M. (2013). 13 The Retina Dissects the Visual Scene. The New Visual Neurosciences 163. .

Sekirnjak, C., Hottowy, P., Sher, A., Dabrowski, W., Litke, A.M., and Chichilnisky, E.J. (2006). Electrical stimulation of mammalian retinal ganglion cells with multielectrode arrays. J. Neurophysiol. *95*, 3311–3327. .

Sekirnjak, C., Hottowy, P., Sher, A., Dabrowski, W., Litke, A.M., and Chichilnisky, E.J. (2008). High-resolution electrical stimulation of primate retina for epiretinal implant design. J. Neurosci. *28*, 4446–4456. .

Shah, N.P., Madugula, S., Grosberg, L., Mena, G., Tandon, P., Hottowy, P., Sher, A., Litke, A., Mitra, S., and Chichilnisky, E.J. (2019). Optimization of Electrical Stimulation for a High-Fidelity Artificial Retina. 2019 9th International IEEE/EMBS Conference on Neural Engineering (NER) https://doi.org/10.1109/ner.2019.8716987.

Shah, N.P., Brackbill, N., Rhoades, C., Kling, A., Goetz, G., Litke, A.M., Sher, A., Simoncelli, E.P., and Chichilnisky, E.J. (2020). Inference of nonlinear receptive field subunits with spike-triggered clustering. Elife *9*. https://doi.org/10.7554/eLife.45743.

Sinha, R., Hoon, M., Baudin, J., Okawa, H., Wong, R.O.L., and Rieke, F. (2017). Cellular and Circuit Mechanisms Shaping the Perceptual Properties of the Primate Fovea. Cell *168*, 413–426.e12. .

Solomon, S.G., Martin, P.R., White, A.J.R., Rüttiger, L., and Lee, B.B. (2002). Modulation sensitivity of ganglion cells in peripheral retina of macaque. Vision Res. *42*, 2893–2898. .

Sweeney, J.D., Ksienski, D.A., and Mortimer, J.T. (1990). A nerve cuff technique for selective excitation of peripheral nerve trunk regions. IEEE Transactions on Biomedical Engineering *37*, 706–715. https://doi.org/10.1109/10.55681.

Tandon, P., Bhaskhar, N., Shah, N., Madugula, S., Grosberg, L., Fan, V.H., Hottowy, P., Sher, A., Litke, A.M., Chichilnisky, E.J., et al. (2021). Automatic Identification of Axon Bundle Activation for Epiretinal Prosthesis. IEEE Trans. Neural Syst. Rehabil. Eng. *29*, 2496–2502. .

Townshend, B., and White, R.L. (1987). Reduction of Electrical Interaction in Auditory Prostheses. IEEE Transactions on Biomedical Engineering *BME-34*, 891–897. https://doi.org/10.1109/tbme.1987.326102.

Vilkhu, R.S., Madugula, S.S., Grosberg, L.E., Gogliettino, A.R., Hottowy, P., Dabrowski, W., Sher, A., Litke, A.M., Mitra, S., and Chichilnisky, E.J. (2021). Spatially patterned bi-electrode epiretinal stimulation for axon avoidance at cellular resolution. J. Neural Eng. *18*. https://doi.org/10.1088/1741-2552/ac3450.

Vrabec, F. (1966). The Temporal Raphe of the Human Retina. American Journal of Ophthalmology *62*, 926–938. https://doi.org/10.1016/0002-9394(66)91920-9.

Wandell, B.A. (1995). Foundations of vision: Behaviour, neuroscience, and computation (Sunderland, MA: Sinauer Associates).

Wang, Z., Bovik, A.C., Sheikh, H.R., and Simoncelli, E.P. (2004). Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. *13*, 600–612. .

Warland, D.K., Reinagel, P., and Meister, M. (1997). Decoding visual information from a population of retinal ganglion cells. J. Neurophysiol. *78*, 2336–2350. .

Watanabe, M., and Rodieck, R.W. (1989). Parasol and midget ganglion cells of the primate retina. J. Comp. Neurol. *289*, 434–454. .

Weiland, J.D., Cho, A.K., and Humayun, M.S. (2011). Retinal prostheses: current clinical results and future needs. Ophthalmology *118*, 2227–2237. .

Weiland, J.D., Walston, S.T., and Humayun, M.S. (2016). Electrical Stimulation of the Retina to Produce Artificial Vision. Annu Rev Vis Sci *2*, 273–294. .

Wu, E.G., Sher, A., Litke, A.M., Simoncelli, E.P., and Chichilnisky, E.J. (2022). Maximum a posteriori natural scene reconstruction from retinal ganglion cells with deep denoiser priors. bioRxiv https://doi.org/10.1101/2022.05.19.492737.

Zhang, Y., Jia, S., Zheng, Y., Yu, Z., Tian, Y., Ma, S., Huang, T., and Liu, J.K. (2020). Reconstruction of natural visual scenes from neural spikes with deep neural networks. Neural Netw. *125*, 19–30. .

Zhou, D.D., Dorn, J.D., and Greenberg, R.J. (2013). The Argus® II retinal prosthesis system: An overview. In 2013 IEEE International Conference on Multimedia and Expo Workshops (ICMEW), pp. 1–6.