September 2006

A Silicon Retina that Reproduces Signals in the Optic Nerve

Kareem A. Zaghloul
University of Pennsylvania

Kwabena Boahen
University of Pennsylvania

A Silicon Retina that Reproduces Signals in the Optic Nerve

Abstract

Prosthetic devices may someday be used to treat lesions of the central nervous system. Similar to neural circuits, these prosthetic devices should adapt their properties over time, independent of external control. Here we describe an artificial retina, constructed in silicon using single-transistor synaptic primitives, with two forms of locally controlled adaptation: luminance adaptation and contrast gain control. Both forms of adaptation rely on local modulation of synaptic strength, thus meeting the criteria of internal control. Our device is the first to reproduce the responses of the four major ganglion cell types that drive visual cortex, producing 3600 spiking outputs in total. We demonstrate how the responses of our device's ganglion cells compare to those measured from the mammalian retina. Replicating the retina's synaptic organization in our chip made it possible to perform these computations using a hundred times less energy than a microprocessor—and to match the mammalian retina in size and weight. With this level of efficiency and autonomy, it is now possible to develop fully implantable intraocular prostheses.

A silicon retina that reproduces signals in the optic nerve

Kareem A Zaghloul[1] and Kwabena Boahen[2,3]

[1] Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104, USA
[2] Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA

E-mail: boahen@stanford.edu

Abstract
Prosthetic devices may someday be used to treat lesions of the central nervous system. Similar to neural circuits, these prosthetic devices should adapt their properties over time, independent of external control. Here we describe an artificial retina, constructed in silicon using single-transistor synaptic primitives, with two forms of locally controlled adaptation: luminance adaptation and contrast gain control. Both forms of adaptation rely on local modulation of synaptic strength, thus meeting the criteria of internal control. Our device is the first to reproduce the responses of the four major ganglion cell types that drive visual cortex, producing 3600 spiking outputs in total. We demonstrate how the responses of our device's ganglion cells compare to those measured from the mammalian retina. Replicating the retina's synaptic organization in our chip made it possible to perform these computations using a hundred times less energy than a microprocessor—and to match the mammalian retina in size and weight. With this level of efficiency and autonomy, it is now possible to develop fully implantable intraocular prostheses.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

One goal of understanding neural systems is to develop prosthetic devices that can someday be used to replace lesioned neural tissue. Designing a successful prosthesis that faithfully replicates the computations performed by a neural circuit requires a detailed understanding of that circuit's anatomic connections and functional computations. For such a prosthesis to be practical, the device must perform these computations as efficiently as, and at a physical scale comparable to, the lesioned network and must be independent of external control. Development of a retinal prosthesis is a logical first step in attaining this goal since the retina's circuits and computations are well understood.

Present attempts to engineer a viable retinal prosthesis have focused on the significant problem of efficient electrical stimulation of neurons along the visual pathway [1, 2].

Microelectrode arrays implanted epiretinally or subretinally evoke phosphenes in patients with visual loss (due to outer retinal degeneration) by relying on electrical stimulation of the remaining retinal cells to dictate firing patterns [3, 4]. Whereas the epiretinal approach relies on an external camera to capture visual information and on an external processor to recreate retinal computation, subretinal devices use photodiodes embedded in the electrode array to locally transduce light into stimulating current. Cortical visual prostheses address disease processes affecting structures post-synaptic to the outer retina [5, 6]. They are similar to epiretinal prostheses in that they also depend on external devices to capture and process visual information, but they must fully recreate thalamic function as well as retinal function.

While the emphasis on electrical stimulation technology is important in addressing the difficult problem of interfacing with the nervous system, a fully implantable retinal prosthesis would ideally capture all of the functions performed by the mammalian retina in one autonomous device. These neural computations can be performed at an energy efficiency and

physical scale comparable to biology by morphing neural circuits into electronic circuits [7]. Micron-sized transistors function as excitatory or inhibitory synapses or as gap junctions, thereby recreating the synaptic organization of the retina at a similar physical scale. The time scale and energy dissipation can be matched as well by operating these transistors in the subthreshold region, where they conduct nanoamperes or even picoamperes, just like small populations of ion channels do. As millions of transistors can be fabricated on a thumb-nail-sized piece of silicon using VLSI (very large scale integration) technology, this neuromorphic approach offers a fully implantable solution for neural prostheses.

The first effort to morph the retina into silicon, though widely acclaimed, suffered from several shortcomings. First, only outer retina circuitry was morphed: the cones, horizontal cells and bipolar cells [8]. Second, a logarithmic photoreceptor (cf cone) was used to capture a wide intensity range, but this degraded the signal-to-noise ratio by attenuating large amplitudes (i.e., signal) while leaving small amplitudes (i.e., noise) unchanged. Third, the local spatiotemporal average (cf horizontal cell) was subtracted to obtain contrast (cf bipolar signal), or, more precisely, the logarithm of contrast, but this made the signal-to-noise ratio even worse. Subsequent efforts [9] overcame these limitations by modulating synaptic strengths locally to control sensitivity and by including the cone-to-cone gap junctions to attenuate noise. But they still omitted the inner retina, which contains upwards of 44 cell types [10].

Here, we present a silicon retina modelled on neural circuitry in both the outer and the inner retina. It is constructed at a scale comparable to the human retina and uses under a tenth of a watt, thereby satisfying the requirements of a fully implantable prosthesis. By capturing both outer and inner retina circuitry using single-transistor synapses, the silicon retina we built passes only an intermediate range of frequencies. It attenuates redundant low spatiotemporal frequencies and rejects noisy high frequencies, much like the retina does. And by modulating the strengths of its single-transistor synapses locally, the device adapts to luminance and to contrast. It responds faster but more transiently as contrast increases, much like the retina does. This silicon retina outputs spike trains that capture the behaviour of ON- and OFF-centre [11] versions of wide-field transient and narrow-field sustained ganglion cells [12], which provide 90% of the primate retina's optic nerve fibres [13]. And, more significant for a prosthetic application, these are the four major types that project, via thalamus, to primary visual cortex.

In the remainder of this paper, we describe our model for the retinal circuitry, which includes a total of 13 neuronal types. We briefly outline our procedure for morphing neural circuits into electronic ones. And we compare the responses of our silicon retina's four output cell types to the mammalian retina.

2. Methodology

We implemented modulation by exploiting the exponential $I(V)$ relationship of the MOS (metal-oxide-semiconductor) transistor. In the subthreshold regime, the current from the drain terminal to the source terminal is the superposition of a forward component that decreases exponentially with the source voltage (V_s) and a reverse component that decreases similarly with the drain voltage (V_d); both components increase exponentially with the gate voltage (V_g). That is, $I_{ds} = I_0 e^{\kappa V_g}(e^{-V_s} - e^{-V_d})$ where $\kappa \approx 0.7$ is a nonideality factor; voltages are in units of $U_T = 25$ mV, at 25 °C (this equation describes the n-type device; voltage and current signs are reversed for a p-type [14]). Hence, the transistor converts voltage to current exponentially and converts current back to voltage logarithmically. Modulation occurs by changing the source voltage, which changes the transistor's transconductance. Current mirrors are added to reverse the direction of current when necessary.

Spike responses from silicon neurons on the chip are read out of the array by a digital arbiter [15], generating a sequence of ganglion cell addresses that are relayed off chip. The arbiter, which outputs the physical address and cell type of one ganglion cell at a time, is capable of relaying a new spiking address off chip every 35 ns. We display the colour-coded activity of the chip's four ganglion cell types on a monitor, which enables us to observe the entire array in real time. This video signal is generated by a chip with an array of leaky integrators, which convert spikes back into graded signals. We used a logic analyser (Tektronix TLA715, v 4.1.108, Beaverton, OR) to capture spike activity in response to visual stimuli projected onto a screen 136.5 cm in front of the chip (InFocus LP425, Wilsonville, OR). A CS mount video lens with a focal length of 8.5 mm (Edmund Scientific, Tonawanda, NY) was placed in front of the chip to focus images onto the chip surface. Spike activity from a particular column was analysed by restricting captured addresses to that column. Visual stimuli were programmed using Microsoft DirectDraw®. Periodic stimuli were defined by per cent Michelson contrast: $100 \times (I_{max} - I_{min})/(I_{max} + I_{min})$, where I_{max} and I_{min} are the peak and trough intensities (range = 0–100%).

Spatial frequency responses of ganglion cell activity (figure 4(b)) were fit with a balanced difference-of-Gaussian model. Briefly, we computed the frequency profile of a one-dimensional zero-mean inhibitory Gaussian with standard deviation, σ_{Inh}, and unit area subtracted from a one-dimensional zero-mean excitatory Gaussian with standard deviation, σ_{Exc}, and unit area. We compared this frequency response to our data and optimized σ_{Exc} and σ_{Inh} to give the best fit.

We reconstructed the visual image from the spikes (figure 6) by convolving ON and OFF sustained ganglion cell spike output with the same difference-of-Gaussian model, whose excitatory and inhibitory standard deviations were determined by the fit to the ganglion cell spatial frequency responses (σ_{Exc} and σ_{Inh}, figure 4(b)). We passed the outputs of this convolution through a temporal low-pass filter with a time constant of 22.7 ms, computing a new frame every 20 ms. We took the difference between images obtained from ON and OFF spikes and displayed it on a grey scale, with ON and OFF activity corresponding to bright and dark pixels,

respectively. Activity from transient ganglion cells did not enhance the resolution of the reconstructed image and was not included.

We derived how our silicon model's cone activity depends on its cone and horizontal space constants and on contrast [16]:

$$I_{CT} = \frac{2r}{r + 2 - 1/r + 2c}c,$$

where I_{CT} represents cone terminal activity and c represents stimulus contrast. r represents the ratio of the horizontal space constant to the cone space constant and is related to horizontal cell coupling, controlled globally in the circuit through V_{hh}. We fit the luminance adaptation curves (figure 7(c)) using this equation. We allowed r to increase with decreasing intensities since we obtained that data by exploiting the dependence of contrast sensitivity on horizontal cell coupling, thereby compensating for the effect of stray photocurrents. These photocurrents, which set an upper limit for the membrane time constants that can be realized, reduce the silicon retina's sensitivity by speeding up ganglion cell spike frequency adaptation and narrow-field amacrine cell pre-synaptic inhibition. We compensated for the latter effect by adjusting a second externally applied voltage that sets the narrow-field amacrine cell's (baseline) membrane leakage.

We derived how our silicon model's ganglion cell responses depend on input contrast for a flat spectrum (i.e., white noise) [17]. This stimulus was characterized by contrast per unit frequency d. The transient ganglion cell response, I_{Gt}, in spikes per second, is given by

$$I_{Gt} = S \left| \left(d\frac{j\tau_A\omega + \varepsilon(1-g)}{j\tau_A\omega + 1} \right) \left(\frac{1}{j\tau_0\omega + 1} \right)^2 \left(\frac{1}{j\tau_p\omega + 1} \right) \right|,$$

where $\tau_A \equiv \varepsilon\tau_{na}$, $\varepsilon \equiv 1/(1 + wg)$, τ_{na} is the narrow-field amacrine cell's time constant, g is the gain from the narrow-field amacrine cell to the transient ganglion cell and w is the wide-field amacrine cell modulated strength of narrow-field amacrine cell inhibition onto the bipolar terminal. The sustained ganglion cell response can be obtained by setting g to zero. We approximated the outer retina using a low-pass temporal filter with time constant τ_0. We also included the low-pass filtering behaviour of the chip's photoreceptors whose time constant is τ_p. We fit the four data sets (figure 8(c)) by allowing the system gain term, S, and the inhibition strength, w, to vary across different stimulus contrasts and fixed the remaining parameters. The best fits of this model to the four input contrast densities are shown as solid lines in figure 8(c). We found that the parameter values that fit these curves best were $\tau_p = 33$ ms, $\tau_0 = 77$ ms, $\tau_{na} = 1.0382$ s and $g = 1.07$. The increase in system gain, S, saturated over the four contrasts (352 to 1358 to 1711 to 1929), while the inhibition strength, w, increased exponentially (1.07 to 1.51 to 2.38 to 3.76).

3. Results

3.1. Modelling the mammalian retina

Our model for retinal circuitry (figure 1(a)) is based on identified synaptic interactions and local microcircuits


Figure 1. Modelling the retina. (a) Synaptic organization: cone outer segments (CO) supply photocurrent to cone terminals (CT), which excite horizontal cells (HC). Horizontal cells reciprocate with shunting inhibition. Both cones and horizontal cells are electrically coupled to their neighbours by gap junctions. Horizontal cells modulate cone to horizontal cell excitation and cone gap junctions (see the text). ON and OFF bipolar cells (BC) relay cone signals to ganglion cells (outputs) and excite narrow- and wide-field amacrine cells (NA, WA). They also excite amacrine cells that inhibit complementary bipolars and amacrines. Narrow-field amacrine cells inhibit bipolar terminals and wide-field amacrine cells; their inhibition onto wide-field amacrine cells is shunting. They also inhibit transient ganglion cells (OnT, OffT)—but not sustained ganglion cells (OnS, OffS). Wide-field amacrine cells modulate narrow-field amacrine cell pre-synaptic inhibition and spread their signals laterally through gap junctions. (b) Single-transistor synapses: electrical nodes represent neurons; we assume that they are electrotonically compact. *Inhibition* (bubble): increased voltage on the pre-synaptic node (purple) turns on the transistor and sinks more current (red) from the post-synaptic node (green), decreasing its voltage. The voltage applied to the third terminal (blue) modulates the transconductance (i.e., dI_{post}/dV_{pre}). A short line represents modulation. *Excitation* (arrowhead): current is sourced onto the post-synaptic node, increasing its voltage. In this case, the post-synaptic voltage modulates the conductance itself, shunting the current. We can change shunting excitation to shunting inhibition or modulated inhibition to modulated excitation by reversing the sign of either the pre- or post-synaptic voltage (using a p-type transistor instead of an n-type) or the sign of the current (using a current mirror). *Conduction* (bi-directional arrow): a bi-directional current flows between the two nodes (brown), whose voltages determine its forward and reverse components. Voltage on the third terminal modulates both (trans)conductances.

previously described in the literature, obtained using histological and physiological techniques. After constructing this model of the retina's synaptic connections, we use this model as a blueprint with which we assemble our silicon retina.

Our model for the outer retina is designed to realize luminance adaptation by adjusting synaptic strengths locally. Photocurrents from the cone outer segments in our model drive a network of cone terminals, which subsequently excite a network of horizontal cells. Because they are coupled through gap junctions, these horizontal cells compute the local average intensity in our model. We use this signal to modulate cone-to-cone coupling strength as well as the cone's membrane conductance (shunting inhibition). Using the local intensity signal to adjust these two synaptic strengths makes the cone terminal's sensitivity inversely proportional to luminance while preventing the changes in spatial frequency tuning that plagued previous attempts at light adaptation [9]. To compensate for the resulting signal attenuation at the cone

terminal, we also use our horizontal cell's local intensity signal to modulate the cone to horizontal cell synaptic strength. This auto-feedback mechanism, whereby our horizontal cell regulates its own input, is similar to that found in the retina [18, 19].

Cone terminals in our model drive two types of bipolar cells that rectify the cone signal into ON and OFF channels, thus reproducing the complementary signalling scheme found in the mammalian retina [11]. Bipolar cells subsequently excite amacrine cells with either narrow or wide fields. To ensure that the ON or OFF channels decorrelate their activity such that one channel is not fully active when the other channel is, bipolar and amacrine cells receive inhibition from the complementary channel in our model, similar to vertical inhibition found between the inner plexiform layer's ON or OFF laminae [20]. We make the signal at the bipolar cell's terminal more transient (high-pass filtered) than its cone input by applying sustained (low-pass filtered) inhibition from the narrow-field amacrine cell [21]. The bipolar cell terminal in our model also excites two types of ganglion cells, which we call transient and sustained. In transient ganglion cells, feed-forward inhibition from the narrow-field amacrine cells cancels residual sustained excitation from the bipolar terminal, similar to the synaptic complex found in mammalian retina [22].

Our model of the inner retina realizes contrast gain control, the control of sensitivity to temporal contrast, through modulatory effects of wide-field amacrine cell activity. These cells are excited by both ON and OFF bipolar cells and inhibited by both ON and OFF narrow-field amacrine cells in our model, similar to ON–OFF amacrine cells found in the retina [23]. By modulating their own inhibitory inputs as well as narrow-field amacrine cell inhibition at bipolar terminals, our wide-field amacrine cells compute temporal contrast. That is, their activity reflects the ratio between contrast fluctuations (high-pass signal) and average contrast (low-pass signal). As this temporal contrast increases, their modulatory activity increases, the net effect of which is to make our ganglion cells respond more quickly and more transiently. There is also an overall decrease in sensitivity due to the less sustained nature of the response. This adaptation captures properties of contrast gain control in the mammalian retina [17, 24].

3.2. Morphing the retina into silicon

We morphed our retinal model into a silicon chip by replacing each synapse or gap junction in our model with a transistor. One of its terminals is connected to the pre-synaptic node, another to the post-synaptic node and a third to the modulatory node. By permuting these assignments, we realize excitation, inhibition and conduction, all of which are under modulatory control (figure 1(b), see section 2). Their strengths are modulated locally, within the chip, except for a small number of biases globally controlled by the user (such as coupling between horizontal cells and between wide-field amacrine cells).

Morphing our model for the outer retina yielded the electronic circuit shown in figure 2(a). Omitting adaptation

in the phototransduction cascade, photocurrents linearly proportional to luminance discharge the cone terminal node, V_c, which we define as an increase in cone terminal activity. This drop in V_c produces a current that excites the horizontal cell network through an nMOS transistor followed by a pMOS current mirror. This excitatory current is modulated by horizontal cell activity, represented by V_h, and increases as V_h increases to realize auto-feedback. But this increased current also releases more charge onto V_c, thereby realizing horizontal cell inhibition of cone terminal activity. Thus, a single transistor implements two distinct synaptic interactions, one excitatory and the other inhibitory. Cone nodes are electrically coupled to their six nearest neighbours through nMOS transistors whose gates are controlled locally by horizontal cells, implementing our model of cone gap-junction modulation. Horizontal cells also communicate with one another, through pMOS transistors, but this coupling is controlled in this device by an externally applied voltage (V_{hh}).

Morphing our model for the inner retina yielded the electronic circuit shown in figure 2(b). Wide-field amacrine cell modulation of narrow-field amacrine cell inhibition is realized by applying voltages representing wide- and narrow-field amacrine activity to a transistor's source and gate terminals, respectively. This transistor drains current from the node that represents bipolar terminal activity, implementing pre-synaptic inhibition. It also sources current onto the node that represents wide-field amacrine activity, charging up that voltage, V_{wa}. This increase corresponds to inhibition of wide-field amacrine activity since, as V_{wa} increases, the strength of narrow-field amacrine inhibition, w, decreases. Conversely, as V_{wa} decreases, the strength of this inhibition increases. A p-type transistor and a current mirror realize excitation of the wide-field amacrine by the bipolar terminal. Additional details of this circuit, and the complete circuit schematics for our silicon retina, may be found elsewhere [25].

Morphing our model for the ganglion cell spiking neuron yielded the electronic circuit shown in figure 2(c). Details of this circuit are described elsewhere [26]. Briefly, input current charges up a ganglion cell membrane capacitor. As the membrane voltage approaches threshold, a positive feedback loop, modulated by V_{fb}, accelerates the voltage's rate of change. Once threshold is passed, the circuit generates a pulse (or spike) that is relayed to digital circuitry. The digital circuitry acknowledges receipt of the spike by sending a reset pulse which discharges the membrane. The reset pulse also dumps a quanta of charge onto a current mirror integrator through a pMOS transistor gated by V_w. Charge accumulating on the integrator models the build-up of Ca^{2+} within the cell after it spikes. This charge, which leaks away with a time constant determined by V_{tn}, draws current away from the membrane capacitor, modelling Ca^{2+}-mediated K^+ channels.

Our chip design was fabricated in a 0.35 μm minimum feature-size process, with its cell mosaics tiled at a scale similar to the mammalian retina (figure 2(d)). Phototransistors are tiled triangularly 40 μm apart; this spacing is only about two and a half times that of human cones at 5 mm nasal eccentricity [27]. The phototransistors are only 10 μm on a side, leaving ample space for post-synaptic circuitry,

Figure 2. Morphing the retina. (a) Outer retina circuitry: a phototransistor draws current through an nMOS transistor whose source is tied to V_c, which represents cone terminal (CT) activity, and whose gate is tied to V_h, which represents horizontal cell (HC) activity. This transistor passes a current proportional to the product of cone terminal and horizontal cell activities, thus modelling shunting inhibition from horizontal cells to cones. In addition, this current, mirrored through pMOS transistors, dumps charge on the horizontal cell node, V_h, modelling cone terminal excitation of the horizontal cells. V_L, a global bias set externally, sets the mean level of V_c. (b) Wide-field amacrine cell modulation: bipolar terminal (BT) activity (I_{bt}) excites a network of wide-field amacrine cells (WA) through a current mirror; it also excites the narrow-field amacrine cell (NA, excitation circuitry not shown). Wide-field amacrine cell activity modulates the strength of narrow-field amacrine cell feedback inhibition onto the bipolar terminal, subtracting a current wI_{na} from the bipolar cell's excitatory input I_{bc}. The same current is also subtracted from the wide-field amacrine cell's excitatory input, I_{bt}, thereby inhibiting it. V_{bq} controls the quiescent current supplied to the inner retina by the bipolar terminal and V_{aa} controls the extent of gap-junction coupling in the wide-field amacrine cell network. (c) Ganglion cell spiking circuitry: current I_{in} from the inner retina charges up a ganglion cell (GC) membrane capacitor. When the membrane voltage crosses threshold, the circuit produces a spike (Sp) that is relayed off chip by digital circuitry. This circuitry acknowledges receipt of the spike by sending a reset pulse (Rst) that discharges the membrane and dumps charge on a current mirror integrator that implements Ca^{2+}-dependent spike rate adaptation. (d) Chip design and human photoreceptor mosaic: each pixel, with 38 transistors on average, has a phototransistor (P), outer plexiform (synaptic) layer (OPL) circuitry, bipolar cells (BC) and inner plexiform layer (IPL) circuitry. Spike-generating ganglion cells (GC) are found in five out of eight pixels; the remaining three contain a narrow-field amacrine (NA) cell membrane capacitor. Inset: tangential view of human cone (large) and rod (small) mosaic at 5 mm eccentricity, plotted at the same scale (reproduced from [27]). (e) Functional architecture: signals from a central photoreceptor (not shown) and its six neighbours (CO) are pooled to provide synaptic input to each bipolar cell (BC). Each bipolar cell generates a rectified output, either ON or OFF, that drives a local IPL circuit. Sustained ganglion cells, which have a dendritic field diameter of 80 μm, receive input from a single local IPL circuit. Transient ganglion cells, however, receive signals from a central IPL circuit (not shown) and its six neighbours, and hence their dendritic field is 240 μm wide.

which is interspersed between them. Unlike neural tissue, silicon microfabrication technology cannot produce three-dimensional structures. Finally, to preserve the mammalian retina's functional architecture, the chip includes convergence from cones to bipolar cells and from bipolar cells to transient ganglion cells (figure 2(e)) [28]. The 3.5×3.3 mm^2 silicon die has 5760 phototransistors at a density of 722 per mm^2 and 3600 ganglion cells at a density of 461 per mm^2—tiled in $2 \times 48 \times 30$ and $2 \times 24 \times 15$ mosaics of sustained and transient ON and OFF ganglion cells.

3.3. Spatiotemporal filtering

Our silicon retina's ganglion cells respond to a restricted band of spatiotemporal frequencies, with transient cells displaying nonlinear spatial summation. In response to a drifting sinusoidal grating, spike trains from active ganglion cells of the same type differ significantly due to the cumulative effect of variability between transistors (CV = 20–25% for currents in identically sized and biased transistors) (figure 3(a)). In the entire array, 151 out of 360 ON-transient, 202/360 OFF-transient, 890/1440 ON-sustained and 792/1440 OFF-sustained ganglion cells exhibited no activity. Many of these ganglion cells were located near the edge of the chip, and as such, we did not investigate responses of ganglion cells there. Despite this heterogeneity, we were able to obtain results that match physiological data by averaging responses from all cells in a given column (figure 3(b)), much as physiologists average several trials from the same cell.
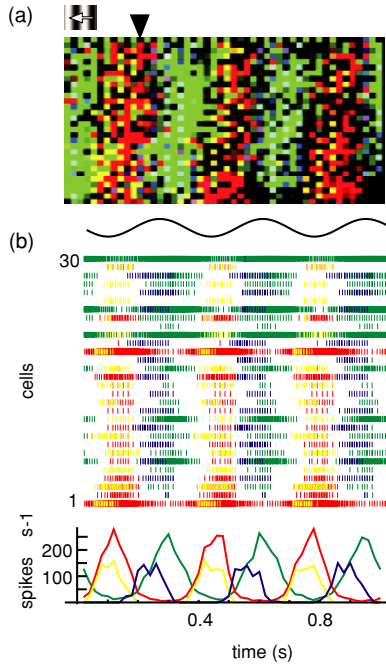
Figure 3. Quadrature representation. (a) A single false-colour frame rendered from spike activity in the entire 48×30 ganglion cell array, captured in response to a 3 Hz 50% contrast drifting sinusoidal grating (0.14 cpd) whose luminance varied horizontally across the screen and was constant in the vertical direction. We use a 50% contrast stimulus in all responses presented here unless otherwise noted. Ganglion cell outputs are colour coded as shown in figure 1(a). Where sustained and transient ganglion cells fire at the same spatial location, the colours overlap. Note that transient ganglion cells are tiled at half the resolution of sustained ganglion cells, and hence only appear in every other row and every other column, in a checkerboard-like pattern. (b) Spike raster (top) recorded from all cells in a single column (arrow in (a)) and histogram (bottom, bin width = 20 ms) of all ganglion cell responses for that column. Sustained cells (green and red) occur in every row (some are unresponsive) while transient cells (blue and yellow) occur in every other row, overlaid on the sustained ones. Spike-rate CVs (coefficient of variation) within this column, computed for all 30 sustained and 15 transient ganglion cells in this column regardless of activity, were 57% (OnT) and 162% (OnS). Complementary ON and OFF channels respond out of phase while transient cells lead sustained cells, exhibiting both earlier onset and shorter duration of firing. We computed the amplitude of the fundamental Fourier component of these histograms, which is plotted in all frequency responses presented, unless otherwise noted. The same applies to physiological data reproduced for comparison.

Our results reveal that both low and high spatial frequencies are attenuated. Sustained cells respond to a higher range of frequencies, as expected from their smaller receptive fields (figures 4(a) and (b)). When we varied the phase of a contrast-reversing sinusoidal grating, we observed frequency doubling in transient cells (figures 4(c) and (d)). This nonlinear summation is the fundamental distinction between narrow- and wide-field mammalian ganglion cells [29–31] and arises because the bipolar cell signals are rectified before they are summed [30].

Sustained cells in the silicon retina retain bandpass spatial filtering at all temporal frequencies (figures 5(a) and (b)). This pattern of spatiotemporal filtering matches the mammalian
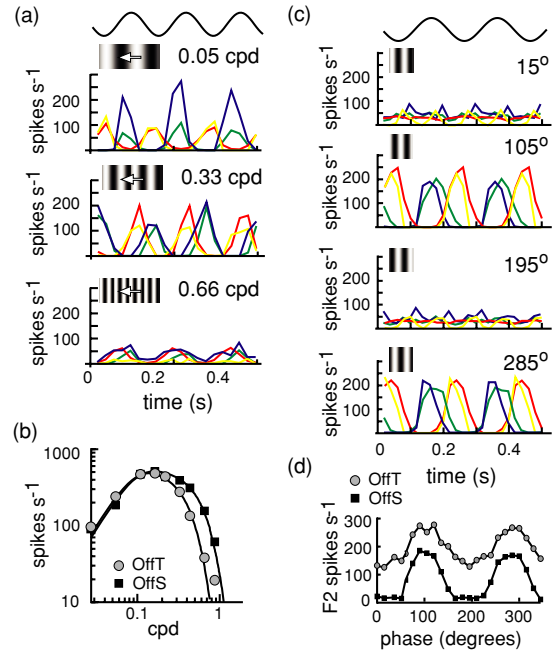

Figure 4. Spatial filtering and nonlinear summation. (a) Varying spatial frequency: responses to 7.5 Hz horizontally drifting sinusoids with three different spatial frequencies. The responses are strongest at an intermediate frequency, except for OnT cells, which showed an anomalous preference for low frequencies. (b) Spatial frequency tuning: OffT and OffS amplitudes are plotted for all spatial frequencies tested; they both peaked at 0.164 cpd, but OffS cells pass a higher range of frequencies. Solid lines are the best fit of a balanced difference-of-Gaussian model (OffT: $\sigma_{Exc}/\sigma_{Inh} = 0.20$; OffS: $\sigma_{Exc}/\sigma_{Inh} = 0.15$; see section 2) [48]. (c) Varying spatial phase: responses to a 5 Hz 0.33 cpd contrast-reversing grating at four different spatial phases. Transient cells show frequency-doubled responses at 15° and 195°. (d) Null test: amplitudes of the second Fourier component (F2) of OffT and OffS responses are plotted for all phases tested. The sustained cells' F2 response disappeared at certain phases, but it could not be nulled in the transient cells. Fluctuations in F2 amplitude arise from uneven spatial sampling in the silicon retina.

retina, except for a resonance found at very high temporal frequencies [32]. In the silicon retina, fast wide-field amacrine cell modulation augments slow horizontal cell inhibition to suppress low spatial frequencies, irrespective of whether they are presented at high or low temporal frequencies. And the optics and the cone–cone gap junctions blur high spatial frequencies, also irrespective of temporal frequency. As a result, the sustained cells pass a restricted band of spatial frequencies at all temporal frequencies.

On the other hand, the silicon retina's transient cells retain bandpass temporal filtering at all spatial frequencies (figures 5(a) and (c)). This pattern also matches the mammalian retina, except for the high-frequency resonance [32]. In the silicon retina, focused narrow-field amacrine cell inhibition augments diffuse horizontal cell inhibition to suppress low temporal frequencies, irrespective of whether they are presented at high or low spatial frequencies. And the cone membrane's capacitance smears high temporal frequencies, also irrespective of spatial frequency. As a
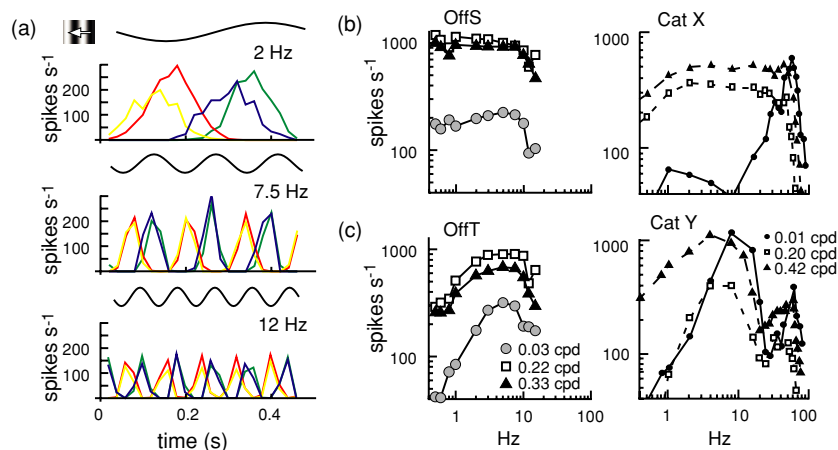
Figure 5. Temporal filtering. (a) Varying temporal frequency: responses to 0.22 cpd horizontally drifting sinusoids at three different temporal frequencies. The response is strongest at an intermediate frequency for transient cells, whereas sustained cell responses decline monotonically. (b) Sustained-cell temporal frequency tuning: responses of OffS and cat ON-centre X cells to low, medium and high spatial-frequency sinusoidal gratings drifting horizontally at different temporal frequencies (see the legend in (c)). Both pass all temporal frequencies below 10 Hz, except at low spatial frequencies. However, the cat data display a high-frequency resonance (cat data are reproduced from [32]). The ordinate here and in (c) represents responsivity, which is the amplitude of the fundamental Fourier component divided by the stimulus contrast. (c) Transient-cell temporal frequency tuning: same as in (b) but for OffT cells and cat ON-centre Y cells. Both pass a restricted band of temporal frequencies at all spatial frequencies below 0.33 cpd. Both (b) and (c) contain previously published data [25], reprinted here for comparison to mammalian data.

Figure 6. Response to a face. In the static image (top), only sustained ganglion cells respond. Reconstruction of the image from their activity (middle, see section 2) demonstrates fidelity of retinal encoding. In the moving image (bottom), transient ganglion cells respond as well, highlighting moving edges. The velocity of the image was approximately 26.96 deg s^{-1}. The mean spike rate was 19 spikes/cell/s. A similar version of this figure was published earlier [25], but did not include the response of the array to stimulus motion.

result, the transient cells pass a restricted band of temporal frequencies at all spatial frequencies

The overall effect of spatiotemporal filtering is best illustrated by natural stimuli (figure 6). Edges are enhanced by sustained ganglion cell activity in the static image. During rapid motion, transient ganglion cells capture this information with surprisingly little blurring. To confirm that the chip captures visual information, we reconstructed the natural stimulus from the sustained ganglion cell spike activity. Passing spike output through a spatiotemporal filter (see section 2) produces an image that is easily recognizable, even with only 30 × 48 pixels and just 0.4 spikes/cell/frame. This

result suggests that cortical structures receiving input from such a visual prostheses can extract useful visual information from the silicon retina's neural code through simple linear filtering.

3.4. Light and contrast adaptation

The silicon retina's ganglion cells adapt to mean luminance and encode stimulus contrast (figure 7). They maintain contrast sensitivity over at least one and a half decades of mean luminance. This intensity range was limited on the low end by leakage currents; these transistors pass a few picoamperes even when their gate voltage is zero. And it was limited on the high end by the projector in our experimental set-up (could not exceed 200 cd m^{-2}) and by stray photocurrents (light-induced leakage currents) in the silicon chip. To obtain the results presented here, we compensated for the effect of these photocurrents by changing two externally applied voltages that would otherwise require no adjustment (see section 2). Overall, the silicon retina's ganglion cell activity remained weakly correlated with absolute light intensity due to the residual effect of the stray photocurrents. Thus, as found in mammalian ganglion cell behaviour [33], responses at low contrasts are weaker at lower light intensities

The silicon retina's ganglion cells also adapt to temporal contrast. When presented with contrast-reversing gratings, the transient ganglion cells respond more quickly but more transiently with increasing contrast (figure 8(a)). And the peak firing rate tends to saturate at the highest contrast levels, as the responses became even more transient. This adaptation is similar to the contrast gain control observed in mammalian narrow-field sustained [34] and wide-field transient ganglion cells [49]. However, it was not as dramatic in the silicon retina's sustained cells, whose responses did not decay nor

Figure 7. Luminance adaptation. (a) Varying intensity: OnT and OnS responses to a sinusoidal grating (0.22 cpd) whose mean intensity was attenuated by amounts listed, using neutral density filters. Due to increases in sensitivity, the response amplitude hardly changes. The noisier responses at high intensity are due increased background activity, which tends to invoke synchronous firing due to cross-talk (i.e., ephatic interactions) in the silicon chip. (b) Cat ON-centre Y-cell intensity curves: the sinusoidal grating's (0.2 cpd) contrast varied from 1 to 50% and reversed at 2 Hz, for five mean luminances [33]. Here, and in (c), response versus contrast (small x-axis) curves are shifted to align the 50% contrast response with that particular mean luminance (large x-axis). Mean luminance is converted from trolands to cd m^{-2} based on a 5 mm diameter pupil (adapted from [33]). (c) OnT intensity curves: the sinusoidal grating's (0.22 cpd) contrast varied from 3.25 to 50% and reversed at 3 Hz, for four different mean luminances. Solid lines represent the best fit of an equation governing cone terminal activity (see section 2). As mean luminance decreased from 192 to 6 cd m^{-2}, the value assigned to the ratio of the horizontal space constant to the cone space constant, r, in this equation increased monotonically from 0.46 to 0.69, reflecting a reduction in peak spatial frequency response from 0.22 to 0.16 cpd (see section 2). This figure is similar to previously published data [25], but is replotted on a different scale for comparison to mammalian data.

saturate as much; they did, however, display a more rapid onset with increasing contrast. This difference between the silicon retina's sustained and transient ganglion cells suggests that narrow-field amacrine cell feed-forward inhibition enhances contrast gain control by making the response more transient (see section 4).

To better quantify the effect of contrast gain control [24], we measured the silicon retina's temporal frequency tuning at different contrasts. Our OFF-transient cells' peak response shifted to higher frequencies with increasing contrast, moving by an amount similar to that observed in the mammalian retina (figures 8(b) and (c)). But while this shift in tuning was accompanied by an overall strengthening of responses at all frequencies in our data, it was accompanied by preferential strengthening of high-frequency responses in the cat data.

4. Discussion

A fully implantable prosthesis requires a device that can independently extract the same visual information encoded by the mammalian retina at a similar physical scale and energy efficiency. Our silicon retina approximates the behaviour of the mammalian retina, in both linear response and nonlinear adaptations, validating this neuromorphic modelling approach for such applications. In addition, this real-time silicon model may be useful both in further testing specific hypotheses about the retina and in serving as a realistic retinal input to other downstream applications like cortical models, other artificial neural systems or robots. Yet although our silicon

retina qualitatively recreates the computations performed by the mammalian retina, there are some specific quantitative differences of note as well as some simplifications of retinal circuitry.

Our approach in constructing this silicon retina was to model synaptic connections found in the mammalian retina and to implement that model using transistor primitives. For example, reciprocal inhibition between bipolar and amacrine cells in complementary ON and OFF channels in our model mimics vertical inhibition between ON and OFF laminae [20] and serial inhibition found between amacrine cells [35] in the mammalian retina. Furthermore, our model extends that proposed by Victor and Shapley [24, 34] to include an anatomical substrate for computing the 'neural measure of contrast', suggesting that wide-field amacrine cells play this role in mammalian retina. However, there still remain some differences between our model and the functional architecture of the mammalian retina. Dopaminergic amacrine cells [36], and light sensing ganglion cells [37], are likely important in modulating mammalian retinal cone and horizontal gap-junction conductance, for example. These dopaminergic cells are not included in our model, and we instead rely on the local horizontal cell signal to modulate cone coupling.

We realized our goal of designing and fabricating a silicon retina that can operate and adapt independent of external control to a large extent, but there remains some degree of manual intervention necessary to make our chip work properly. To achieve autonomous operation in a final prosthetic application, all of the silicon retina's external biases were designed to be hard-wired to specific voltages.
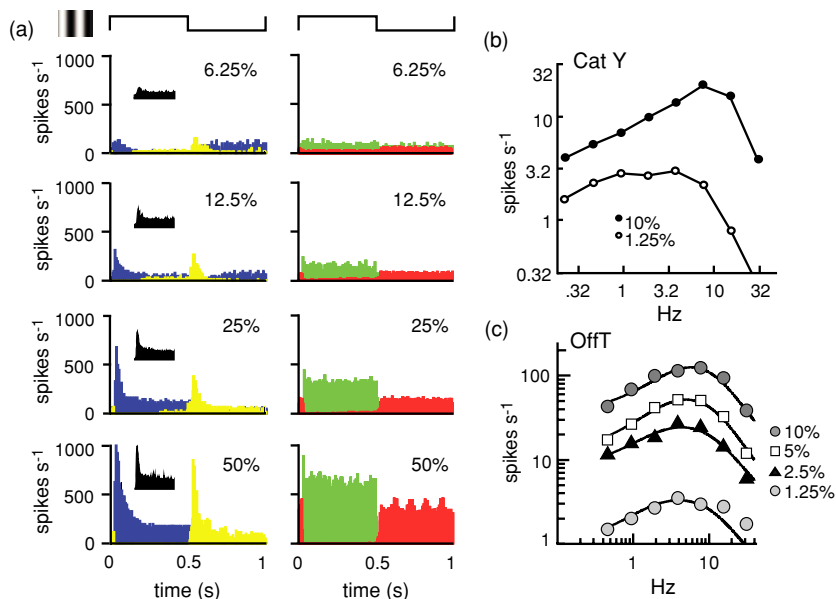
Figure 8. Contrast gain control. (a) Varying contrast: responses to a 1 Hz square-wave contrast reversal of a sinusoidal grating (0.22 cpd) at four different peak stimulus contrasts. Bin width is 4 ms. Responses increase sublinearly and change more rapidly with increasing contrast; these effects are more pronounced in the transient cells. Their responses decayed with a time constant that decreased from 28 to 22 ms as contrast increased from 6.25% to 50%. *Inset:* response of ON-centre cat X cell to half a cycle of the same stimulus [34]. (b) ON-centre cat Y-cell contrast-dependent temporal filtering: a stationary sinusoidal grating (0.25 cpd) whose contrast was determined by the sum of eight sinusoids was used. All eight sinusoids had the same amplitude, whose value relative to the background is stated. The amplitude of the fundamental Fourier component at these eight frequencies was plotted for two different amplitudes (reproduced from [24]). The peak sensitivity shifted from 3.9 Hz to 7.8 Hz as the contrast increased from 1.25% to 10%. (c) OffT contrast-dependent temporal filtering: the stimulus was the same except that we reduced the spatial frequency to 0.14 cpd and tested four different contrasts. The peak sensitivity shifted from 3.9 Hz to 7.8 Hz as the contrast increased from 1.25% to 10%. Solid lines are the best fit of an analytical model (see section 2), which indicated that an increase in the strength of narrow-field amacrine cell feedback inhibition from $w = 1$ to $w = 3.5$ could account for this change in temporal dynamics.

However, whereas the voltages applied to the biases that set mean cone terminal activity, mean bipolar terminal activity, mean ganglion cell activity and coupling strength in wide-field amacrine cells remained fixed, we had to manually adjust the voltages applied to the biases that set the coupling strength between horizontal cells and the bias that sets the narrow-field amacrine cell leakage current to compensate for light-dependent leakage currents, a shortcoming of silicon microtechnology at these small length scales. This fine tuning was only required for light adaptation. We did not adjust any bias voltages during any of the other experiments.

Leakage currents in the silicon substrate caused another discrepancy between our device and the mammalian retina. We had to set up the silicon retina's ganglion cells to have higher firing rates than their mammalian counterparts because a higher firing rate proves useful in limiting the effects of leakage currents in our circuit. Nevertheless, the baseline firing rate still tends to increase with light intensity due to stray photocurrents. The increase in background activity causes the silicon retina's temporal responses to become less sharp with increasing light intensity (figures 7(a) and (c)). Furthermore, cross-talk between the silicon neurons causes them to synchronize their firing (see figures 7(a) and 8(a)), especially in the presence of the high background activity induced by bright or large stimuli, distorting the response further. We expect to reduce these leakage currents in future designs so as to improve the silicon retina's operation.

Unlike mammalian retina, the silicon retina's ganglion cells fail to respond to temporal frequencies above 10 Hz (figure 5(a)). We had expected that because of our model's ability to temporally adapt, responses to higher temporal frequencies should be preserved. We find instead that high-frequency responses are eliminated in our device, which may explain why chip ganglion cells fail to exhibit the resonance seen at high frequencies in cat ganglion cells[4]. We suspect that this discrepancy may be a result of slow time constants in our phototransistors and outer retina circuit. When fitting our model to the data (figure 8(c), see section 2), we found that the time constants of the low-pass filters associated with the photoreceptors and outer retina were relatively long. This suggests that these sluggish responses in the feed-forward pathway impose a low-pass filter effect on the silicon retina's ganglion cell responses that is greater than in mammalian retina.

We also failed to observe contrast gain control in the silicon retina's sustained ganglion cells, unlike in the mammalian retina. This failure is likely related to the absence of any significant transient component in our sustained ganglion cells' responses. Mammalian retina's narrow-field sustained cells, on the other hand, do have a transient component, albeit less than that in their transient counterparts. In our model, we expected pre-synaptic inhibition to introduce

---

[4] But the lower temporal frequencies do not show this either. What is the biological data?

a transient component at the bipolar terminal by sharpening the sluggish outer retina response, as it is thought to do in the retina. However, the sharpening was apparently insufficient to compensate for the long time constants associated with the silicon retina's phototransistors and outer retina circuits. We believe that speeding up the outer retina will result in a transient component at the bipolar terminal, and hence produce contrast gain control in the sustained cells. Implementing feed-forward inhibition to a lesser degree in sustained cells in our silicon retina, like in the mammalian retina [22, 38], will also likely help in restoring the transient component and therefore contrast gain control in these sustained cells.

Transient ganglion cells overcome the sluggish outer retina response because feed-forward inhibition suppresses the residual sustained component in the bipolar terminal, leaving a mostly transient response. There remains a sustained component in these cells which is not fully compensated for by feed-forward inhibition, however, as demonstrated in their response to a drifting grating in figure 8(a). The overall decrease in transient behaviour in both sustained and transient cells due to the sluggish outer retina response, together with the ineffectiveness of spike frequency adaptation at the high firing rates we used [26, 39], makes our transient ganglion cells more similar to sustained mammalian ganglion cells than to their transient counterparts (see figure 8(a)).

We have used the same axon-hillock model to generate spikes in our transient and sustained ganglion cells (figure 2(c)). In mammalian retina, however, physiological studies have demonstrated different membrane properties in different types of ganglion cells [40, 41]. These differences in membrane properties may account for some of the differences seen in behaviour. Taking these differences in the spike generation into account may produce responses that better replicate those seen in mammalian ganglion cells.

We also find that, although transient ganglion cells in the silicon retina and in the mammalian retina both shift their peak responses to higher frequencies as we increase stimulus contrast (figure 8(c)), mammalian transient ganglion cells demonstrate a preferential strengthening of high-frequency responses as well (figure 8(b)). Analysis of our model indicates that modulation of pre-synaptic inhibition, as proposed by Victor and Shapley [24, 34], cannot alone account for such differential effects [17]. Thus, the discrepancy between our model and the mammalian retina indicates that there may be additional mechanisms for contrast gain control in the mammalian retina that account for the differential change in sensitivity at low and high temporal frequencies with increasing contrast (see [42–44]). Elucidating these mechanisms and implementing them in our model will help in generating ganglion cell responses that better match mammalian data.

Our artificial retina satisfies the requirements of a neural prosthesis by matching the biological retina in size and weight and using under a tenth of a watt. Rabbit retina uses 16.2 nW per ganglion cell (82 μmoles of ATP g^{-1} min^{-1} [45], or 88 mW g^{-1} min^{-1}, times 70 mg average weight, divided by 380 000 ganglion cells [10]). In contrast, our chip consumes 17 μW per ganglion cell (62.7 mW for the entire chip) at an average spike rate of 45 spikes s^{-1} per ganglion cell. Although this energy consumption is 1000 times less efficient than the mammalian retina, it still represents a 100-fold improvement over conventional microprocessors. A 1 GHz Pentium® processor operating at 10 W would dissipate 2.2 mW per ganglion cell to compute the response of a $13 \times 13 \times 13$ kernel ($X \times Y \times T$) updated at 100 times per second. With an upper limit on a proposed intraocular implant's power dissipation of 100 mW, a chip with the Pentium's® similar computing power could thus only compute the responses of under 40 ganglion cells, or a 6×6 array, which is too small for functional vision [3]. However, with the same 100 mW limit, our neuromorphic chip's energy efficiency allows it to compute the responses of 4000 ganglion cells, roughly a 60×60 array. We expect this energy efficiency to improve further, together with spatial resolution and dynamic range, as microfabrication technology advances.

5. Conclusion

Based on detailed knowledge of the retina's neuronal specializations, synaptic organization and functional architecture [28], we have constructed 13 neuronal types in silicon and linked them together in two synaptic layers on a physical scale comparable to the human retina. Furthermore, we have created a silicon retina that modulates its synaptic strengths locally. Our silicon retina realizes luminance adaptation, without using logarithmic compression, and contrast gain control independent of external control, thus capturing properties of retinal neural adaptation for the first time. Our success modelling neural adaptation using single-transistor primitives suggests that a similar approach could be used to morph other neural systems into silicon as well; this may eventually lead to fully implantable neural prostheses [46, 47] that do not require external interfaces.

Acknowledgments

References

[1] Rizzo J F et al 2001 Retinal prosthesis: an encouraging first decade with major challenges ahead Ophthalmology 108 13–4

[2] Margalit E et al 2002 Retinal prosthesis for the blind Surv. Ophthalmol. 47 335–56

[3] Humayun M S et al 1999 Pattern electrical stimulation of the human retina Vision Res. 39 2569–76

[4] Chow A Y et al 2001 Implantation of silicon chip microphotodiode arrays in the cat subretinal space IEEE Trans. Neural Syst. Rehabil. Eng. 9 86–95

[5] Normann R A, Maynard E M, Rousche P J and Warren D J 1999 A neural interface for a cortical vision prosthesis Vision Res. 39 2577–87

[6] Dobelle W H 1999 Artificial vision for the blind by connecting a television camera to the visual cortex *ASAIO J.* **46** 3–9

[7] Mead C A 1989 *Analog VLSI and Neural Systems* (Reading, MA: Addison-Wesley)

[8] Mahowald M and Mead C 1988 A silicon model of early visual processing *Neural Netw.* **1** 91–7

[9] Boahen K A and Andreou A G 1992 A contrast sensitive silicon retina with reciprocal synapses *Advances in Neural Information Processing Systems* ed J E Moody, S J Hanson and R P Lippmann (San Mateo, CA: Morgan Kaufmann Publishers) pp 764–72

[10] Masland R 2001 The fundamental plan of the retina *Nat. Neurosci.* **4** 877–86

[11] Kuffler S W 1953 Discharge patterns and functional organization of mammalian retina *J. Neurophysiol.* **16** 37–68

[12] Werblin F S and Dowling J E 1969 Organization of the retina of the mudpuppy *Necturus maculosus*: II. Intracellular recording *J. Neurophysiol.* **32** 339–55

[13] Rodieck R W 1988 The primate retina *Comp. Primate Biol.* **4** 203–78

[14] Tsividis Y P 1987 *Operation and Modeling of the MOS Transistor* (New York: McGraw-Hill)

[15] Boahen K A 1999 Point-to-point connectivity between neuromorphic chips using address events *IEEE Trans. Circuits Syst.* **47** 416–34

[16] Zaghloul K A and Boahen K 2005 An On–Off log domain circuit that recreates adaptive filtering in the retina *IEEE Trans. Circuits Syst.* **52** 99–107

[17] Zaghloul K A and Boahen K 2004 Optic nerve signals in a neuromorphic chip: I. Outer and inner retina models *IEEE Trans. Biomed. Eng.* **51** 657–66

[18] Kamermans M and Werblin F 1992 GABA-mediated positive autofeedback loop controls horizontal cell kinetics in tiger salamander retina *J. Neurosci.* **12** 2451–63

[19] Kamermans M, Fahrenfort I, Schultz K, Janssen-Bienhold U, Sjoerdsma T and Weiler R 2001 Hemichannel-mediated inhibition in the outer retina *Science* **292** 1178–80

[20] Roska B and Werblin F 2001 Vertical interactions across ten parallel, stacked representations in the mammalian retina *Nature* **410** 583–7

[21] Maguire G and Lukasiewicz P 1989 Amacrine cell interactions underlying the response to change in tiger salamander retina *J. Neurosci.* **9** 726–35

[22] Kolb H and Nelson R 1993 OFF-alpha and OFF-beta ganglion cells in cat retina: II. Neural circuitry as revealed by electron microscopy of HRP stains *J. Comp. Neurol.* **329** 85–110

[23] Freed M A, Pflug R, Kolb H and Nelson R 1996 On–Off amacrine cells in cat retina *J. Comp. Neurol.* **364** 556–66

[24] Shapley R and Victor J D 1979 The contrast gain control of the cat retina *Vision Res.* **19** 431–4

[25] Zaghloul K A and Boahen K 2004 Optic nerve signals in a neuromorphic chip: II. Testing and results *IEEE Trans. Biomed. Eng.* **51** 667–75

[26] Hynna K M and Boahen K 2001 Space rate coding in an adaptive silicon neuron *Neural Netw.* **14** 645–56

[27] Curcio C A, Sloan K R, Kalina R E and Hendrickson A E 1990 Human photoreceptor topography *J. Comp. Neurol.* **292** 497–523

[28] Freed M A and Sterling P 1988 The ON-alpha ganglion cell of the cat retina and its presynaptic cell types *J. Neurosci.* **8** 2303–20

[29] Enroth-Cugell C and Robson J G 1966 The contrast sensitivity of retinal ganglion cells of the cat *J. Physiol.* **187** 517–52

[30] Enroth-Cugell C and Freeman A W 1987 The receptive field spatial structure of cat retinal Y cells *J. Physiol.* **384** 49–79

[31] Demb J B, Zaghloul K A, Haarsma L and Sterling P 2001 Bipolar cells contribute to nonlinear spatial summation in the brisk-transient (Y) ganglion cell in mammalian retina *J. Neurosci.* **21** 7447–54

[32] Frishman L J, Freeman A W, Troy J B, Schweitzer-Tong D E and Enroth-Cugell C 1987 Spatiotemporal frequency responses of cat retinal ganglion cells *J. Gen. Physiol.* **89** 599–628

[33] Troy J B and Enroth-Cugell C 1993 X and Y ganglion cells inform the cat's brain about contrast in the retinal image *Exp. Brain Res.* **93** 383–90

[34] Victor J D 1987 The dynamics of cat retinal X cell centre *J. Physiol.* **386** 219–46

[35] Dowling J E and Boycott B B 1966 Organization of the primate retina: electron microscopy *Proc. R. Soc.* B **166** 80–111

[36] Jensen R J and Daw N W 1986 Effects of dopamine and its agonists and antagonists on the receptive field properties of ganglion cells in the rabbit retina *Neuroscience* **17** 837–55

[37] Berson D M, Dunn F A and Takao M 2002 Phototransduction by retinal ganglion cells that set the circadian clock *Science* **295** 1070–3

[38] Cohen E and Sterling P 1991 Microcircuitry related to the receptive field center of the on-beta ganglion cell *J. Neurophysiol.* **65** 352–9

[39] Boahen K A 1997 The retinomorphic approach: pixel-parallel adaptive amplification, filtering, and quantization *Analog Integr. Circuits Signal Process.* **13** 53–68

[40] Cohen E 2001 Synaptic mechanisms shaping the light-response in retinal ganglion cells *Prog. Brain Res.* **131** 215–28

[41] O'Brien B J, Isayama T, Richardson R and Berson D M 2002 Intrinsic physiological properties of cat retinal ganglion cells *J. Physiol.* **538** 787–802

[42] Rieke F 2001 Temporal contrast adaptation in salamander bipolar cells *J. Neurosci.* **21** 9445–54

[43] Kim K J and Rieke F 2001 Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells *J. Neurosci.* **21** 287–99

[44] Baccus S A and Meister M 1997 Fast and slow contrast adaptation in retinal circuitry *Neuron* **36** 909–19

[45] Ames A, Li Y Y, Heher E C and Kimble C R 1992 Energy metabolism of rabbit retina as related to function: high cost of Na^+ transport *J. Neurosci.* **12** 840–53

[46] Craelius W 2002 The bionic man: restoring mobility *Science* **295** 1018–21

[47] Zrenner E 2002 Will retinal implants restore vision? *Science* **295** 1022–5

[48] Rodieck R 1965 Quantitative analysis of cat retinal ganglion cell response to visual stimuli *Vision Res.* **5** 583–601

[49] Victor J D Private communication   **Q1**

Queries

(1) Author: Please provide the year in reference [49].

(2) Author: A single page can be printed in colour for 475 GBP. Each subsequent page costs an additional 95 GBP each. Please let us know if you are willing to cover the cost of printing any of your figures in colour, or if you would prefer colour to appear in the electronic version only.

Reference linking to the original articles

References with a volume and page number in blue have a clickable link to the original article created from data deposited by its publisher at CrossRef. Any anomalously unlinked references should be checked for accuracy. Pale purple is used for links to e-prints at arXiv.