A CHIP THAT FOCUSES AN IMAGE ON ITSELF

T. Delbrück
California Institute of Technology
Pasadena, California, 91125
e-mail: tobi@hobiecat.caltech.edu

In the modeling of neural systems, time is often treated as a *sequencer*, rather than as an *expresser* of information. We believe that this point of view is restricted, and that in biological neural systems, time is used throughout as one of the fundamental representational dimensions. We have developed this conviction partially because we model neural circuitry in analog VLSI, where time is a natural dimension to work with, and we believe there are deep similarities between the technology we use and the one nature has chosen for us.

Neurobiologists are beginning to explore neural control systems that self-generate sensory input. The focus chip we report here models the focusing system of our eye. The human focusing mechanism is a one-dimensional control system in which experimenters have access to both visual input and motor output signals. For our model, the primary hypothesis about this system is that control signals are *generated actively*, by the motor system in the course of control. We have built and partially characterized a model system, using analog VLSI circuit primitives already developed for other purposes, that incorporates this hypothesis. This chip focuses an image on itself, using time domain information about the quality of the optical image and the motion of the lens.

THE HUMAN ACCOMMODATION SYSTEM

The process by which the eye focuses an image onto the retina is called *accommodation*. The eye accommodates by distorting the curvature of the lens. When muscle fibers running radially outward from the lens contract, the increased tension on the lens flattens it, focusing farther away. When muscle fibers running circumferentially around the lens contract, the decreased tension on the lens allows it to bulge, focusing closer (Weale, 1960). In our model system the focus is changed by changing the distance between a rigid lens and the chip, much like focusing a camera.

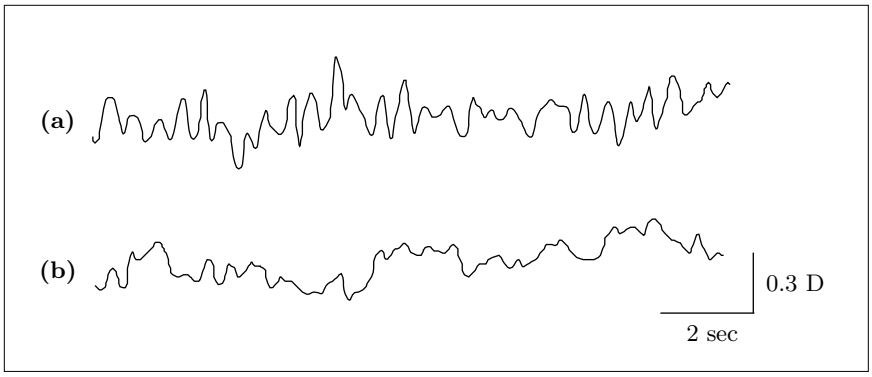

Figure 1. Recordings of fluctuations in human accommodation. These records were obtained using an infrared split-beam optometer. The optical distance to the target was 1 D. **(a)** Pupillary aperture was 7 mm. **(b)** Pupillary aperture was 1 mm. The retinal illumination was kept the same for each trial. (Adapted from Campbell *et al.* (1959) with permission.)

The stimulus for accommodation is not known, although there are results that positively indicate certain possibilities. The obvious possibility is the blur of the retinal image. The problem with imagining a control system for accommodation that uses retinal blur is that blur is an even-error signal: static blur does not say which way to change the accommodation to sharpen the retinal image. Other possibilities for the error signal that are odd-error have been proposed. In an elegant set of experiments, Campbell and Westheimer (1959) showed that chromatic and spherical aberration were sufficient odd-error signals in subjects with paralyzed accommodation. They did not show that these were necessary cues in subjects with normal accommodation reflexes. Given the possibility that image blur is one of the primary cues to accommodation, we might ask, exactly what functional of the image is used as the primary cue? The precise answer is unknown, though there are clues. For example, Fujii *et al.* (1970) showed that intensity gradient was more important than total contrast modulation, as a stimulus to accommodation.

Accommodation fluctuates even under steady-state conditions. The existence of these fluctuations has been known, or at least postulated, for a long time (see, for example, Helmholtz, 1924). Campbell *et al.* (1959) were among the first researchers to obtain recordings of these fluctuations. Figure 1 shows examples of these fluctuations; Figure 2 shows power spectra for the records in Figure 1. We can see that the amplitude of the fluctuations decreases with a smaller pupillary aperture, and their character changes. The fluctuations have a characteristic frequency that shows up as a pronounced bump in the power

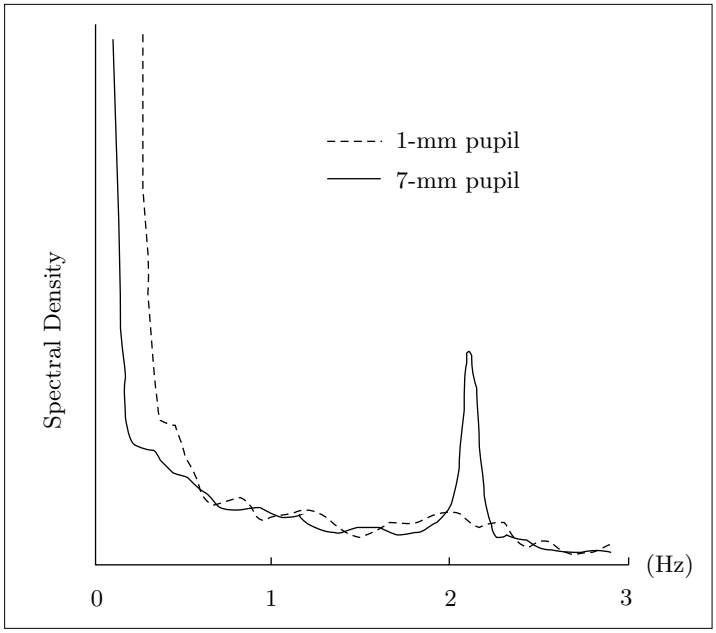

Figure 2. Power spectra for the records from Figure 1. Note that the peak in the power spectrum for the signal in Figure 1(a) disappears in the spectrum for Figure 1(b). (Adapted from Campbell *et al.* (1959) with permission.)

spectrum, usually near 2 Hz. The amplitude of the fluctuations increases as the optical distance to the target gets smaller (Denieul, 1982). For optical distances of 4 D (=25 cm viewing distance), the size of the fluctuations can grow larger than 0.1 D ($\approx \pm$ 0.5 cm fluctuation in focal distance). These oscillations are below perceptual thresholds under ordinary conditions, yet stimuli oscillating at below the perceptual threshold can drive the accommodation system (Kotulak and Schor, 1986b).

THE MODEL OF ACCOMMODATION USED BY THE CHIP

In our model, the measure of image quality is denoted by the term *sharpness*, or by the symbol s. The state of accommodation is denoted by the symbol l. In the current physical realization of our system, the lens is moved, rather than distorted, so this accommodative state is equivalent to the lens position, relative to the point at which the lens settles in the absence of any stimulus. The idea for our circuit came from a paper by Kotulak and Schor (1986). The

essence of the idea is simply stated: The sign of $\dot{s}$ † indicates whether accommodation is changing in the correct direction, and the sign of $\dot{l}$ indicates in which direction the accommodation is currently changing. The sign of the product $\dot{s}\dot{l}$ gives the *sign* of the error signal for $\dot{l}$.

Still open is the question of what to use as the error *magnitude*. Kotulak and Schor (1986) suggested that $\left|\dot{s}/\dot{l}\right|$ would be a reasonable choice; for the model reported here, $\tanh(\dot{s}\dot{l})$ is used as the error signal for $\dot{l}$. We integrate this error signal with respect to time, using the mass of the lens. The discussion section of this report notes some other possible uses of the error signal, besides integrating it with a mass.

We can now write a dynamical equation for the motion of the lens:

$$M\frac{d}{dt}(\dot{l}) = S\tanh\Big(\dot{s}\dot{l}\Big) - Kl - D\dot{l} + N\mathcal{N}(t) \tag{1}$$

An explanation of the terms in Equation (1) follows. First, the driving term is $S\tanh(\dot{s}\dot{l})$. Second, there are natural restoring spring (K) and damping (D) forces. Third, there is a noise term $N\mathcal{N}(t)$. The presence of this noise is essential to the operation of the system.

We imagine that the sharpness function $s(l)$ will peak at some l_0. The approximate form of this function, as computed by our chip, will be derived later; for now, we take this sharpness function to be a Gaussian of width σ, peaked around l_0:

$$s(l) = e^{-\Delta^2/2}, \qquad \Delta = \frac{l - l_0}{\sigma} \tag{2}$$

We see that Δ is the displacement from the correct focal point, in units of σ. The constant σ is the depth of field in this model system. Using this $s(l)$ in the dynamical equation (1) we obtain

$$M\ddot{l} = S\tanh\left(\frac{-\dot{l}^2}{\sigma}\Delta e^{-\Delta^2/2}\right) - Kl - D\dot{l} + N\mathcal{N}(t) \tag{3}$$

This equation represents a simple harmonic oscillator with the addition of noise and a novel forcing term. The noise term is essential, because in the absence of noise and in the presence of damping and restoring forces, the system will eventually settle down to $l = 0$, no matter what the form of the sharpness function $s(l)$.

The difference between our model and that of Kotulak and Schor (1986) is that these researchers use $\left|\dot{s}/\dot{l}\right|$ as the error magnitude. Their choice is sensible because it compensates, in a sense, for large $\dot{s}$ signals produced by rapid focus changes and not by positional focus errors. We use $\tanh(\dot{s}\dot{l})$ as our error signal because it is difficult to build a well-behaved four-quadrant analog divider. At

† A dot over a quantity indicates differentiation with respect to time.

the time this circuit was built, a product seemed more biologically plausible than did a quotient. Any error function $E(\dot{s}, \dot{l})$ that is positive in the first and third quadrants and negative in the second and fourth quadrants, will retain the correct sign-of-error properties.

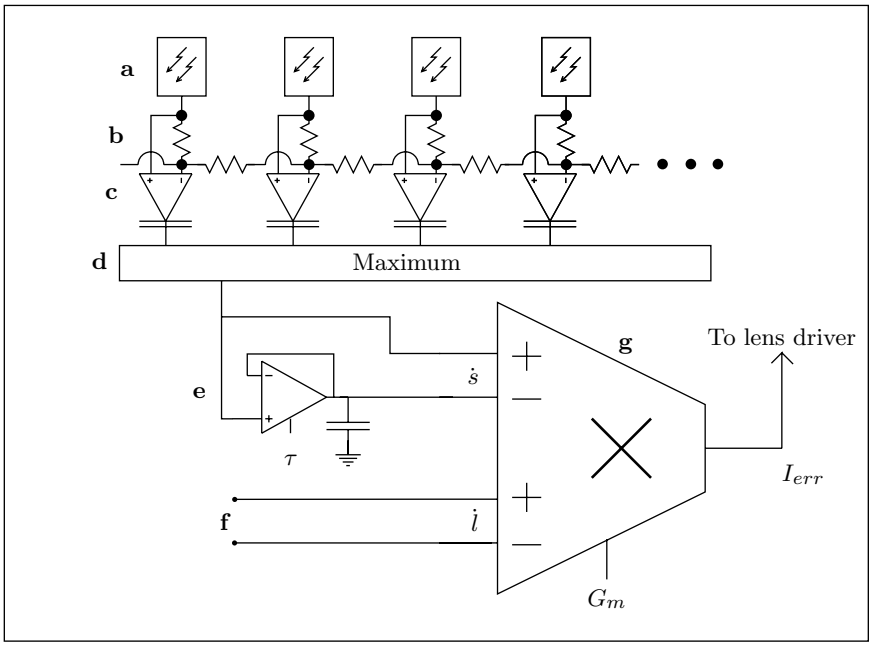

Figure 3. A schematic illustration of the circuitry of the chip.

THE CHIP CIRCUITRY

The chip consists of a one-dimensional silicon retina (Figure 3a,b), an image sharpness computation (3c,d), a differentiator (3e), and an analog multiplier (Figure 3g). All the individual circuits used in this chip have been described in detail elsewhere, so the description here will be confined to a brief functional discussion.

The Sharpness Computation

An image falls on the one-dimensional silicon retina (Figure 3**a,b**) (Mahowald and Mead, 1988). Each output of the retina is a differential voltage between the output of a logarithmic photoreceptor (**a**) and the spatial average computed with a resistive network (**b**). We call this differential voltage V_i for the i_{th} pixel. Each differential voltage is turned into a current by an absolute-value transconductance amplifier (**c**) (Mead, 1989). The i_{th} current is $I_i = I_b \tanh(|V_i/2|)$, where the units of voltage are $kT/q\kappa$. The body-effect factor κ is typically about 0.7. The bias current I_b, and hence the transconductance $G = \frac{I_b}{2kT/q\kappa}$, is set with an externally variable control (Mead, 1989). These I_i are fed into a circuit that computes a voltage s that is logarithmic in the maximum I_i (**d**). This circuit is an adaptation of the winner-take-all circuit (Lazzaro *et al.* 1989), in which the common inhibitory wire encodes the logarithm of the maximum input current. If several I_i are equal and are larger than all other I_j, s will be proportional to the logarithm of the sum of these I_i.

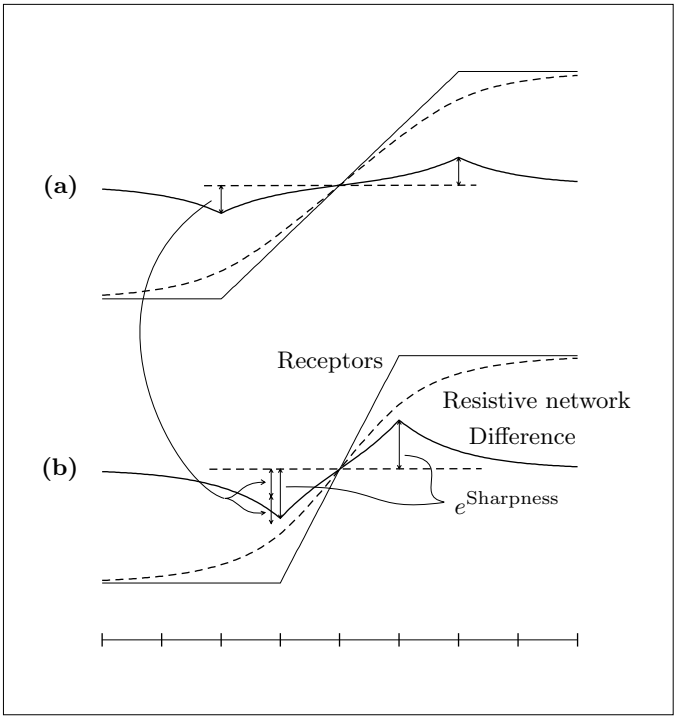

Figure 4. Graphical illustration of the sharpness computation. **(a)** Soft edge. **(b)** Sharp edge. The space constant λ of smoothing is the same for **(a)** and **(b)**. When the edge is twice as sharp the maximum difference between the receptors and the resistive net is nearly twice as large, the deficit being caused by the fact that the extent over which the edge is smeared approaches the scale of smoothing.

The sharpness computation is illustrated graphically in Figure 4. The resistive network computes a smoothed version of the log intensities. The space constant λ of smoothing in the resistive network is controllable. The sharpness s is the maximum difference between the local log intensity and the local spatial average computed by the resistive network. This maximum will occur at locations where the slope of the intensity profile changes. When the slope of the intensity profile changes, a distance of $O(\lambda)$ along the resistive network is required for the network to assume the new slope. The equation governing the behavior of a continuous one-dimensional resistive network is

$$\lambda^2 \frac{d}{dx}\left(\frac{dV}{dx}\right) = V(x) - I(x)$$

where $V(x)$ is the voltage on the network at location x, and $I(x)$ is the input voltage, in our case the log intensity (Mead, 1989). In order to change $\frac{dV}{dx}$ by some amount $\Delta(\text{slope})$, the difference $V - I$, integrated over a distance of $O(\lambda)$, must satisfy $\int_{O(\lambda)} (V - I)dx \sim \Delta(\text{slope})$. When the space constant is not too large compared with the extent of the blurred edge, the sharpness s will satisfy

$$s = \log\big(\max|V(x) - I(x)|\big) \sim \log\Big(\frac{\Delta(\text{slope})}{\lambda}\Big) \tag{4}$$

When λ is large compared with the extent of the edge, the reported sharpness will not depend on the edge sharpness. On the other hand, when λ is comparable to the receptor spacing, the differences $V(x) - I(x)$ will be small. Circuit offsets will more easily dominate the image induced signals, and will cause the sharpness output to assume a constant value close to the point of optimum focus. The optimum space constant was determined experimentally to be a few receptor spacings.

Since $\Delta(\text{slope})$ is inversely proportional to the distance of the lens from the focal plane (Figure 6), the slope of the sharpness function will be independent of the lens aperture or other geometrical parameters of the system. When the image becomes blurred to the point where image induced signals are comparable in size to circuit offsets, the sharpness signal will flatten out.

Future versions of this chip will probably compute the image sharpness measure by either simply summing the outputs of the absolute value amplifiers, or by computing the maximum first difference in the log intensities.

Time Domain Circuitry

A follower-integrator (Figure 3(**e**)), with externally controllable time constant τ, and transfer function $H(s) = 1/(\tau s + 1)$, computes a delayed version $\tilde{s}$ of the sharpness signal s; the difference $(s - \tilde{s})$ is a good approximation to the derivative $\dot{s}$ for frequencies below $1/\tau$ (Mead, 1989). An external sensor (**f**) gives the lens velocity as a differential voltage $\dot{l}$.

Error Signal Computation

The product of the differential voltage $\dot{s}$ and the velocity signal $\dot{l}$ is computed by a wide-range Gilbert multiplier (**g**) (Mead, 1989) to produce the error-signal current $I_{err} = S\tanh\big(\dot{l}\big)\tanh\big(\dot{s}\big)$. The multiplier bias current is once again externally controllable, and corresponds to the constant S in Equation (1). This function has characteristics very similar to the function $\tanh\big(\dot{l}\dot{s}\big)$ used in Equations (1) and (3). The primary difference is that $\tanh\big(\dot{l}\dot{s}\big)$ saturates more quickly as one moves away from the $\dot{l}$ and $\dot{s}$ axes, away from the origin. The current I_{err} is amplified externally, and is used to drive a solenoid attached to the lens. Finally, the mass of the optical arrangement acts to integrate this error signal with respect to time.

The current version of the chip consists of a 40-pixel array. It was fabricated through the MOSIS foundry in 2μ p-well technology. Each pixel is $165\mu m$ wide.

We tested the function of the sharpness sensor by focusing the image of an edge onto the chip, and varying the distance from the chip to the focal plane of the lens (Figure 5). The output peaked around the point of sharpest focus and fell off on either side, as expected. The width of the peak was consistent with geometrical calculations, as will be discussed later. The slope of the sharpness function was not affected by decreasing the aperture, since the slope of the logarithm of any linear function is identical.

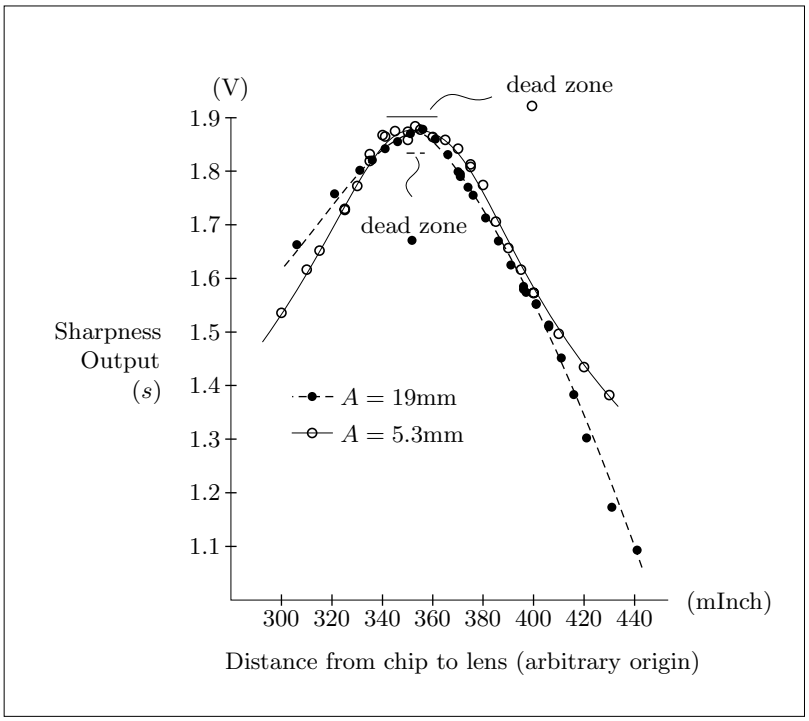

Figure 5. Output from sharpness sensor. The focusing target was a high contrast black and white edge. Using a smaller aperture resulted in a peak that was only slightly broadened, because the space constant λ was several times the receptor spacing d. The curves were hand fitted. The widths of the dead-zones were computed from the geometrical parameters $A = 19mm$ and $A = 5.3mm$, $f = 19mm$, and $d = 165\mu m$, shown in Figure 6. As the distance between the chip and the focal plane is increased to more than is shown in this figure, the sharpness output flattens out. This flattening is due to circuit offsets dominating image induced signals. For smaller apertures (larger depths of field), the flattening occurs farther from the focal plane.

Some theoretical characteristics of of the sharpness output s can be derived as follows. An image that is not in the focal plane of the lens can be represented as the original image convolved with a pill-box shaped kernal. The diameter of this pill-box is the diameter of the circle of confusion at the image plane (Horn, 1968, Horn and Sjoberg, 1981). This procedure uses only geometrical optics; on the scales with which we are concerned, diffractive effects are negligible. Figure 6(a) defines the geometrical parameters. A perfect step edge will be smeared out into an intensity distribution I (Figure 6(b)) given by

$$I = \frac{\pi}{2} - \sin^{-1}(u) - u\sqrt{1-u^2} \tag{5}$$

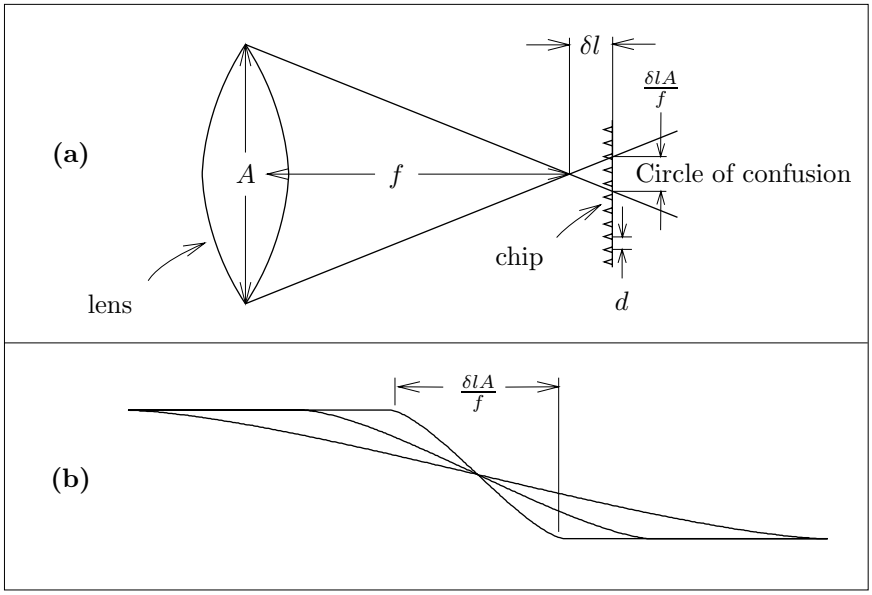

Figure 6. (a) Definitions of dimensions used in the text. The focal length f shown here is the effective focal length for the object being viewed; it is simply the distance from the lens at which the scene is in focus. d is the distance between receptors on the chip. δl is the distance between the chip and the focal plane. **(b)** Several blurred edges at various distances from the focal plane. The spatial profiles of the intensities are derived in the text.

When $\delta l \leq df/A$, s should take on a constant value, because the extent of the edge spans less than one receptor spacing. We compare this prediction with the measurements of the sharpness output shown in Figure 5.

Because measuring image sharpness is equivalent to some measure of the spectral power at the highest spatial frequencies, we can expect some effects of spatial aliasing. This aliasing will cause spurious changes in the reported sharpness, due only to lateral movement of the image, and not to changes in the focus. A scene that is rich in texture will not suffer these spurious changes in the reported sharpness; each edge in the image has a different offset relative to the receptor array, and the sharpness sensor chooses the edge with the maximum contrast, as seen by the array. The same principal will apply to a two-dimensional sensor array for a single edge, as long as the edge does not lie along one of the principal axes of the array. An array with randomly jittered pixel locations would be even better, since it has no preferred axes.

The most straightforward elimination of these spurious changes comes from filtering the image before it falls on the sharpness sensor, to eliminate spatial

frequency content above the Nyquist frequency for the receptor spacing. This filtering could occur in the human eye, where the optical cut-off frequency has been reported to be matched to the receptor spacing at the center of the fovea (Snyder and Miller, 1977).

Alternatively, we suggest that aliasing will only occur when the scene is in focus. Thus, alias-induced signals can be used as indicators of good focus. In general, the magnitudes of local time and space derivatives of the image, produced by lateral movement of the scene across the sensor, will serve as a good indicator of the focus. By integrating these signals over time and space we can obtain an extremely robust measure of the image quality. This is precisely the type of operation biological retinas can do very well. When the eye is fixating a scene, there are constant slow drift and rapid microsaccadic eye movements (Steinman *et al.* 1973). We suggest that these eye movements may generate signals that are used by the focusing mechanism of the eye.

A PHYSICAL REALIZATION OF THE SYSTEM

A schematic illustration of the system as it is now constructed is shown in Figure 7. The lens actuator is a solenoid with a ferromagnetic plug attached to the lens. The velocity of the lens is sensed with a linear variable transformer. The primary coil is excited with a DC current. The velocity is the differential voltage induced in the secondary coil.

Figure 9, shows records of the lens position obtained from this rather primitive setup using two different-sized apertures. Figure 10 shows the power spectra of these records.

DYNAMICAL PROPERTIES OF THE SYSTEM

Now that we have a dynamical system model we can easily test its explanatory power. Consider Figures 1 and 2; they show the effect a change in depth of field has on human accommodation fluctuations. Figure 8 shows what happens when the model system, represented by Equation (3), is subjected to a simulation of the same change in depth of field. The position l_0 of the focus target is shown by the thin solid line. Halfway through the simulation, the target jumps from one side of the zero point to the other. The zero on the vertical axis represents the equilibrium point in the absence of any focusing target. The only difference between the two simulations is in the width σ of the sharpness function, shown on the left. The results are tantalizingly similar to the records shown in Figure 1 for the human accommodation system. However, the behavior of the dynamical system represented by Equation (1) is dependent in a complex way on the values of the parameters. In the absence of the image-dependent term ($S = 0$), Equation (1) reduces to a simple harmonic oscillator driven by a stochastic noise process. Adding back in the sharpness

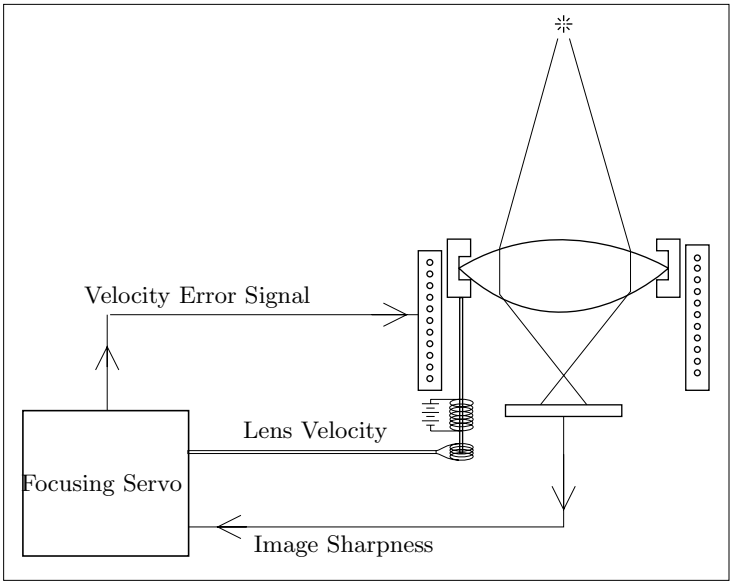

Figure 7. A schematic illustration of the physical interface with the optical system. The output of the chip drives a solenoid which is attached to the lens. The velocity of the lens is sensed with a variable transformer. The mass of the lens serves to integrate, with respect to time, the force signal produced by the chip.

term ($S \neq 0$) and selecting the correct set of parameters produces the behavior shown in Figure 8, but a different choice of parameters could have led to qualitatively different behavior.

We have distinguished three qualitatively different regimes of operation of our dynamical system. The first represents the behavior shown in Figures 1 and 2 for the human accommodation system and in Figure 8 for the simulations. In this regime, increasing the depth of field decreases the amplitude and coherence of the fluctuations in accommodation. In the second regime, not shown in this report, increasing the depth of field does not substantially change the character of the fluctuations. In the third regime, increasing the depth of field increases the amplitude and coherence of the oscillations. The third regime is shown in Figures 9 and 10 for data taken from our physical realization of the system.

To understand these effects, we make a simple analysis of the relative effects of the sharpness (S), damping (D), and depth of field (σ) parameters in the dynamical system represented by Equation (3), letting the spring constant (K) and noise (N) terms be small. The effect of the damping is to limit the saturated velocity of the lens. This velocity is attained when the nonlinear

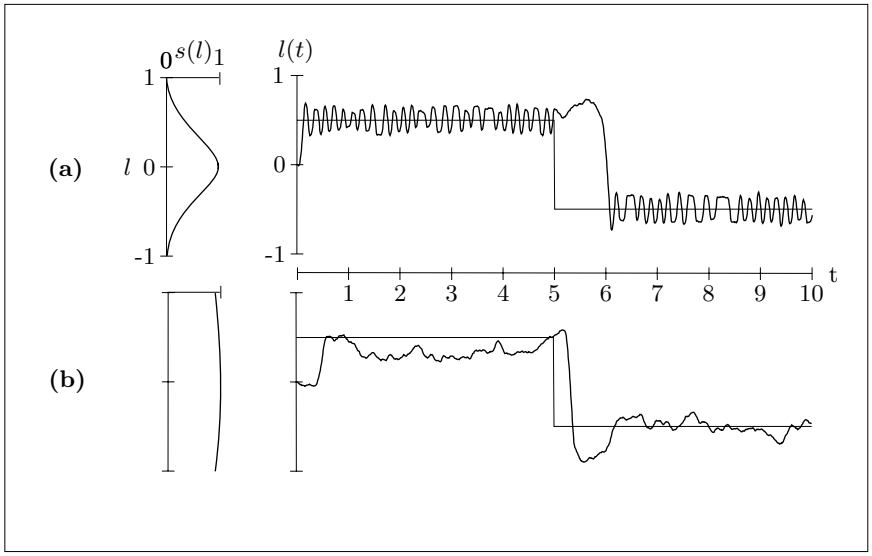

Figure 8. Results of simulations of the dynamical equation (3). The point of optimum focus l_0 switched from $+1/2$ to $-1/2$ halfway through the record. The noise source $\mathcal{N}(t)$ was a Gaussian process of variance 1. The values of the constants were $M = 2.5$, $S = 500$, $K = 1$, $D = 50$, $N = 100$. **(a)** $\sigma = 0.35$. **(b)** $\sigma = 2.5$. The form of the sharpness functions for parts **(a)** and **(b)** are shown on the left. The effect of the depth of field on the amplitude and coherence of the waveform is similar to that shown in Figures 1 and 2.

sharpness term, $S\tanh(\dot{s}\dot{l})$ is saturated and $\ddot{l} = 0$. In this state, $|\dot{l}| = S/D$. This condition will be self consistent when the sharpness term is saturated, which will, to first order, happen when

$$\dot{s}\dot{l} = \dot{l}^2 \frac{ds}{dl} \geq 1. \tag{7}$$

The sharpness function is steepest just outside the dead zone that is caused by finite receptor spacing and other optical blurring. There, $\delta l = df/A$, and from Equation (4), $\frac{ds}{d(\delta l)}\big|_{l=df/A} = \frac{A}{df}$. Using this value for $\frac{ds}{dl}$ and the saturated velocity S/D in Equation (7), we find that the sharpness term will be saturated when $\frac{S^2 A}{D^2 f d} \geq 1$.

As long as this inequality holds, the amplitude of the oscillations will not depend on the depth of field. Intuitively, the velocity will be saturated every time the system crosses the zero point, and an excursion will be halted by the

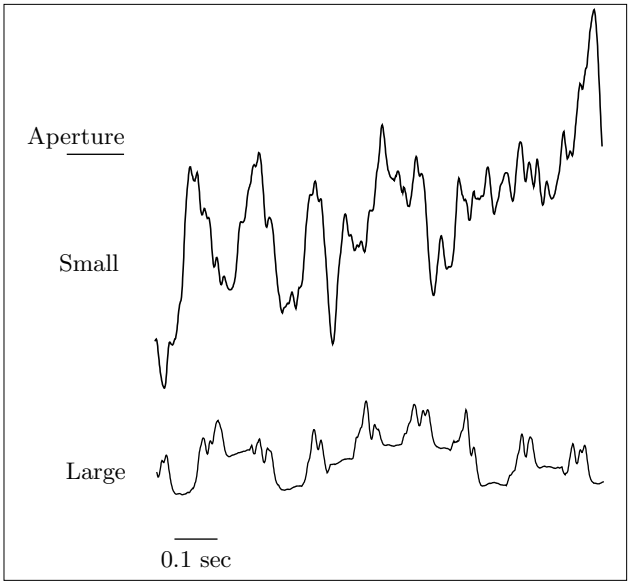

Figure 9. Records of lens motion obtained by integrating the velocity signal numerically.

effect of the saturated driving term. This situation corresponds with the second regime of operation mentioned previously.

When the inequality in Equation (7) is no longer satisfied (for example, when the aperture A becomes small enough), we obtain the first regime of operation, in which an increase in depth of field causes the fluctuations to lose their amplitude and coherence. Intuitively, the limiting velocity will no longer saturate the driving term. This leaves the system more susceptible to the built-in noise. This condition corresponds to the behavior shown by the human accommodation system in Figures 1 and 2, and by the simulation records in Figure 8.

The third regime of operation appears when a saturated excursion past the zero point ends because the sharpness derivative $\frac{ds}{dl}$ becomes small, and not because $\dot{l}^2$ gets small (Equation (7)). Figures 9 and 10 show records of the lens motion obtained from our physical realization of the system. We can see from these figures that decreasing the lens aperture increases the amplitude of the fluctuations, opposite to the behavior shown by human accommodation and to the simulation results. In our system, because the sharpness is encoded logarithmically, it is not the slope of the sharpness function that depends on the aperture, but rather the point of defocus where the sharpness function flattens out.

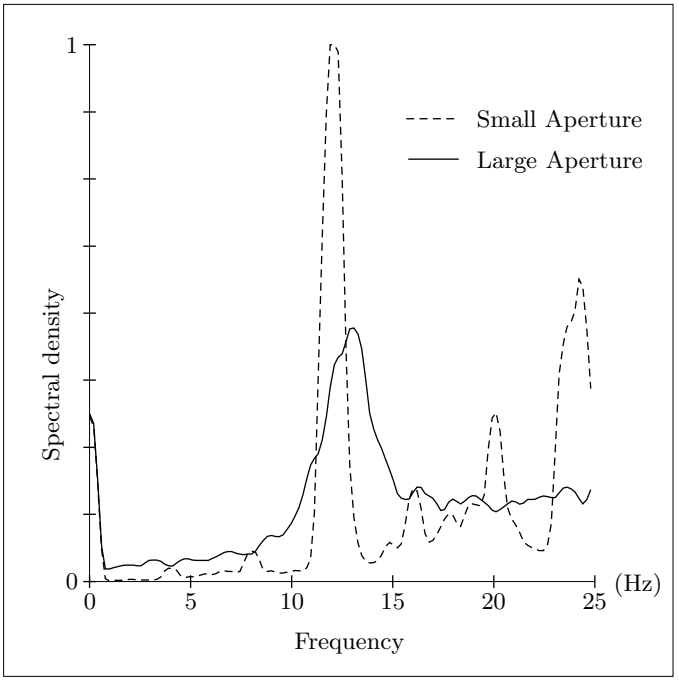

Figure 10. Power spectra for the records in Figure 9. The behavior is opposite to that shown by the human accommodation system in Figure 2. The reason for this difference is discussed in the text.

If the model is an accurate representation of the behavior of human accommodation, then we may conclude two principles of that system. First, the human accommodation system probably does not encode the sharpness logarithmically, as we do in our model system. The effect of a change of depth of field in the human system appears as a change in the slope $\frac{ds}{dl}$. Second, human accommodation is optimized so that the strength of the sharpness term is as small as it can be and still allow the system to operate under long depth of field. If the sharpness term was any larger, the fluctuations would be larger than they need to be under conditions of short depth of field.

DISCUSSION

The human accommodation system provides a simple example of a biological control system in which needed information may be generated actively by the motor system, in the course of control. The system described here represents a control system of a relatively unexplored variety. The system is unstable; small oscillations around the desired state are amplified until a nonlinearity becomes saturated. The system relies on noise for initiation of control. We might hope that these characteristics would circumvent many of the problems of gain control about which designers of control systems worry, but probably these problems are simply pushed into another arena.

We have omitted many features of the human accommodation system. The concept of volition is alien to our formulation; our system has no means of deciding that it would like to alter its focus in a particular direction. In the human accommodation system, volition certainly plays an important role in directing the apparatus toward the desired state. Also, our system has no concept of a linkage between vergence and accommodation. In the human accommodation system, there is strong coupling between vergence eye movements and accommodative response (Johnson *et al.* 1982). Our system's lens is mass dominated, and the damping forces are small relative to the restoring forces. The human lens system is probably spring dominated, with large damping forces (Ejiri *et al.* 1969). The role of the slow reaction time (1/3 sec.) in the human accommodation system has not been worked out, but could signify the presence of a neural integrator like that seen in the vestibulo-ocular reflex (Robinson, 1981). Alternatively, there might not be any neural or physical integration of the error signal; the error signal might affect the velocity of the lens directly. The use of a saturating nonlinearity is biologically plausible (Marg, 1955). The presence of noise is essential for operation of our system; human accommodation is quite noisy even in the absence of any stimulus (Johnson *et al.* 1984). The hypothesis that time domain information is being used is just that – a hypothesis, albeit an attractive one. A related hypothesis is that image sharpness is the primary image-quality cue that the human accommodation system uses. We have used sharpness in our model of accommodation, but this computational convenience does not indicate that other cues, such as chromatic or spherical aberration, might not be used as well.

Several features of the system we have built appear repeatedly in silicon models of the nervous system, and are worth pointing out. Quantities are scaled logarithmically, so that a large dynamic range is compressed into a workable operating range. Nonlinear aspects of operation can be advantageous. Time, as an intrinsic dynamical variable, appears naturally when we use analog computation. The active generation of time domain information may turn out to be useful in other contexts.

Acknowledgments

I thank Misha Mahowald, Carver Mead, John Harris, John Wyatt, and Berthold Horn for critical comments. I thank Hewlett-Packard for computing support, and DARPA and MOSIS for chip fabrication. This work was sponsored by the Office of Naval Research and the System Development Foundation.

References

Campbell, F.W. (1959). The accommodation response of the human eye. *Brit. J. of Physiological Optics.* **16:**188–203.

Campbell, F.W. and Westheimer, G. (1959). Factors influencing accommodation responses of the human eye. *J. Opt. Soc. Amer.* **49:**568–571.

Campbell, F.W., Robson, J.G., and Westheimer, G. (1959). Fluctuations of accommodation under steady viewing conditions. *J. Physiol.* **145:**579–594.

Denieul, P. (1982). Effects of stimulus vergence on mean accommodation response, microfluctuations of accommodation and optical quality of the human eye. *Vision Res.* **22:**561–569.

Ejiri, M., Thompson, H.E., and O'Niell, W.D. (1969). Dynamic viscoelastic properties of the lens. *Vision Res.* **9:**233–244.

Fujii, K., Kondo, K., and Kasai, T. (1970). An analysis of the human accommodation system. *Technology Reports of Osaka University.* **20:**221–236.

Helmholtz, H.v. (1924). *Treatise on Physiological Optics,* Vol. 1. Menasha: Optical Society of America, p. 191.

Horn, B. (1968). Project MAC: Focusing. *MIT Artificial Intelligence Memo.* No. 160.

Horn, B. and Sjoberg, R.W. (1981). The application of linear systems analysis to image processing. Some notes. *MIT Artificial Intelligence Memo.* No. 100.

Johnson, C.A., Post, R.B., and Tsuetaki, T.K. (1984). Short-term variability in the resting focus of accommodation. *Opthal. Physiol. Opt.* **4:**319–325.

Kotulak, J.C. and Schor, C.M. (1986). A computational model of the error detector of human visual accommodation. *Biol. Cybernetics.* **54:**189–194.

Kotulak, J.C. and Schor, C.M. (1986b). The accommodative response to subthreshold blur and to perceptual fading during the Troxler phenomenon. *Perception.* **15:**7–15.

Lazzaro, J., Ryckebusch S., Mahowald, M.A., and Mead, C.A. (1989). Winner-Take-All circuits of $O(n)$ complexity. In Touretsky, D.S. (ed), *Advances in Neural Information Processing Systems* 1. San Mateo, CA: Morgan Kaufman, pp. 703–711.

Mahowald, M. and Mead, C.A. (1988). A silicon model of early visual processing. *Neural Networks.* **1:**91–97.

Marg, E., Reeves, J.L. (1955). *J. Pot. Soc. Am.* **45:**926 (Fig. 1).

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

Robinson, D.A. (1981). The use of control systems analysis in the neurophysiology of eye movements. *Ann. Rev. Neurosci.* **4:**463–503.

Snyder, A.W. and Miller, W.H. (1977). Photoreceptor diameter and spacing for highest resolving power. *J. Opt. Soc. Am.* **67:**697–698.

Steinman, R.M., Haddad, G.M., Skavenski, A.A., and Wyman, D. (1973). Miniature eye movements. *Science.* **181:**810–819.

Weale, R.A. (1960). *The Eye and Its Function.* London: Hatton.