Charles F. Neveu
neveu@artemis.arc.nasa.gov
Caelum Research Corporation
Intelligent Mechanisms Group
NASA Ames Research Center
Moffett Field, CA 94035

Lawrence W. Stark
Neurology and Telerobotics Lab
University of California, Berkeley

The Virtual Lens

Abstract

We describe a new type of feedback display based upon ocular accommodation, called the *virtual lens,* that maintains a focused projection of a CRT image on the retina independent of changes in accommodation, and that replaces the optical image-processing action of the crystalline lens with an arbitrary computable image transform. We describe some applications of the virtual lens in visual psychophysics and virtual environments.

1 Introduction

The cathode ray tube has become the retina of the mind's eye.

—Brian O'Blivion in *Videodrome*

Feedback displays are information sources that also receive feedback from the user and use it to modify the display. They are fundamental parts of virtual environments. The most familiar feedback display is the helmet-mounted display (Sutherland, 1968), which senses head position and uses this information to modify the image shown to the subject. Force-feedback joysticks (Hannaford, Wood, McAfee, & Zak, 1991; Szakaly & Bejczy, 1990) sense hand position through the joystick and offer resistance or force through the joystick back to the hand.

It is through feedback displays that virtual environment systems mimic reality—by sensing activity, analyzing it, and responding. Virtual environment systems can be programmed to mimic reality, or to warp things in ways that are not physically possible. Apart from the obvious entertainment value, creating non-physically realizable environments for neuromuscular systems allows one to perform experiments that were not possible before, precisely because they were, either practically or essentially, not physically realizable. Furthermore, competing models of how neurological control systems work will naturally tend toward agreement on how a neurological system will respond to a realistic environment; after all, the models were developed to explain these very responses. Such models must, however, predict different responses to unrealistic environments or they would be equivalent; thus, it is possible to distinguish them experimentally.

As organisms, we have a small number of senses: vision, hearing, smell, taste, touch, proprioception, and vestibular sensation. We likewise have a larger, but finite, number of responses that, in the real world, effect these senses. Vision is affected by head, eye, and limb position, hearing by head position and orientation, touch by physical contact, etc. The virtual lens is a fundamental type of feedback display, one that senses the accommodative state of the eye and uses that information to modify a visual display. The image processing effect of the

Presence, *Vol. 7, No. 4, August 1998,* 370–381

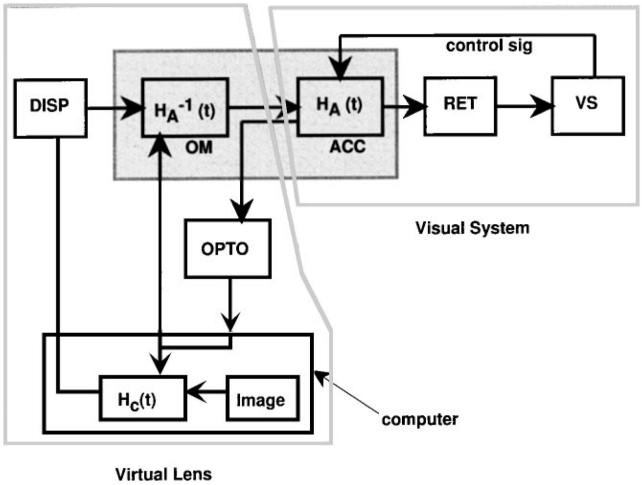

Figure 1. *Schematic diagram of virtual lens. Gray box shows optics H_A of crystalline lens of the accommodation system (ACC) and the optomechanical (OM) inverse optical system H_A^{-1} of virtual lens. DISP is the CRT display; RET is the retina of the user; VS is the visual system of the user which controls the user's accommodation system; OPTO is the optometer that measures accommodative state; Image is the image being viewed; H_C is the computed transfer function that replaces the optically cancelled H_A.*

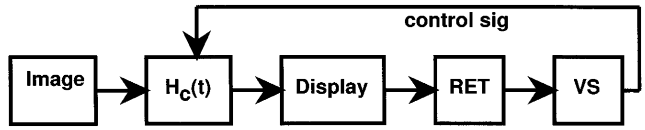

Figure 2. *Reduced schematic of virtual lens. Equivalent control loop with optics cancelled; computed transfer function $H_C(t)$ replaces transfer function of the accommodation system H_A.*

lens—*i.e.*, blurring—can be removed by placing a controllable, predistorting[1] optomechanical system (shown as $H_A{}^{-1}$ in Figure 1) between the subject and the display, and using accommodative state information to make its optical power equal and opposite to the accommodative response (H_A in Figure 1), thus optically removing the effect of accommodation. Simultaneously, the accommodative state information can be used to modify the CRT image *digitally*, blurring it, for example, to simulate the effect of the (now optically removed) lens—thus the term *virtual lens*. The image-processing action of the lens is optically removed and replaced by digital image-processing algorithms implemented in software (shown as H_C in Figure 2).

In this paper we briefly review the accommodation system, then explain the theory and implementation of the virtual lens. Finally, we describe some applications in the areas of neurological control theory, visual psychophysics, and virtual environments.

[1]In this paper we will use the word "distortion" in its more general signal processing sense to mean any kind of modification of the signal, *including blur*, rather than the more narrow optical sense: a class of aberrations that warp the image without blurring it.

2 Ocular Accommodation

2.1 Anatomy and Physiology

Ocular accommodation, in humans, refers to changing the power of the crystalline lens of the eye such that the object of regard, or *target*, is optically conjugate to the fovea—in other words, it's in focus. Helmholtz elucidated the theory of the mechanism of accommodation as generally accepted today. (See Figure 3.) The lens is suspended from its equator, not unlike the hub of a bicycle wheel suspended from the spokes, by a radial ring of fine fibers called the zonule of Zinn. The axial portion of the zonules attach to the lens capsule and radiate outward to the ciliary body, an annulus of tissue extending from the cornea to the ora serrata. The zonular fibers follow the ciliary body around the inside diameter of the globe, finally terminating at the ora serrata. These fibers are attached to the ciliary muscle, the active muscle of accommodation, at roughly their midpoint, effectively dividing the zonular fibers into two sections having different mechanical functions: an axial section from the lens to the ciliary muscle, and a peripheral section from the ciliary muscle to the ora serrata (Rohen & Rentsch, 1969; Rohen & Zimmerman, 1970; Rohen, 1979). The ciliary muscle is a unitary muscle and originates at the corneal-scleral spur, a ring of tough tissue extending inward from the limbus (the junction of the cornea and sclera). The vast majority (≈98%) of the muscle fibers proceed peripherally from their origin to

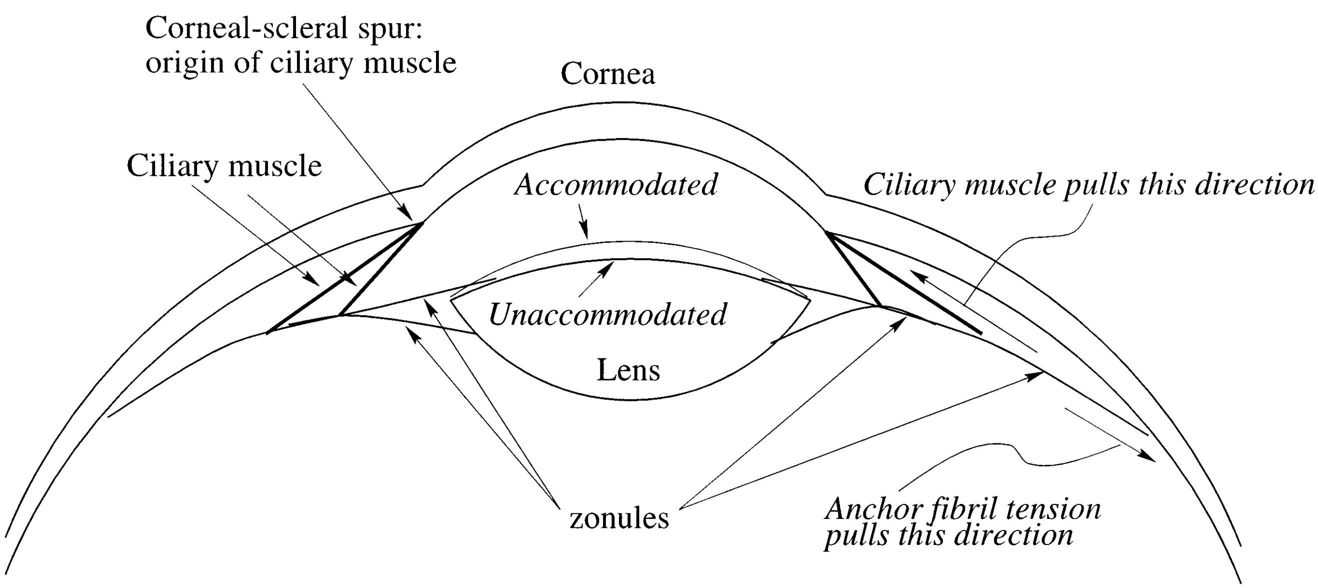

Figure 3. *Mechanism of accommodation. Unaccommodated lens for focusing on distant targets shown in left half of figure, accommodated lens for near targets shown in right half.*

their insertion onto the span fibrils of the zonule of Zinn at Bruch's membrane. Unlike most muscle systems, which have an active agonist and antagonist (e.g., flexor-extensor), the ciliary muscle has a passive antagonist, the peripheral fibers. In the unaccommodated state, the peripheral fibers exert tension on the ciliary muscle and the axial fibers. Thus, when the ciliary muscle contracts, it pulls *against* the peripheral zonules, increasing their tension, while pulling *with* the axial zonules, decreasing their tension. The decreased tension of axial zonules in turn decreases tension on the lens capsule, allowing the anterior surface of the lens to ''swell'' forward, increasing its curvature and hence its dioptric power.

2.2 The Neurological Control System of Accommodation

The focusing machinery of the eye is driven by neurological control systems. The binocular convergence system can drive the accommodation response, as can higher-level cognitive knowledge of—or misinformation about—target position. Accommodation can also be driven by blur. The blur accommodation system measures blur in the image and attempts to adjust the accommodation response so as to minimize blur, in effect much like an autofocus videocamera. Just how this ''autofocus'' mechanism detects blur is not well understood.

3 Virtual Lens

Phillips (1974; Phillips & Stark, 1977) developed an instrument that combined an electronic dynamic optometer (a device for measuring accommodation) with an electronically controllable image-projection system and an electronically controllable optomechanical system for modifying the optical distance from the subject to the target screen. (See Figure 4.) The dynamic optometer measured the subject's accommodative response and converted it into an electrical signal. The experimenter could control the target image quality (by way of blurring or focusing the image) and the target optical distance independently. The apparatus could be used to open the neurological control loop of blur accommodation by matching target distance to the clear vision distance of the subject, so that changes in the accommodation response had no effect on the degree or direction of

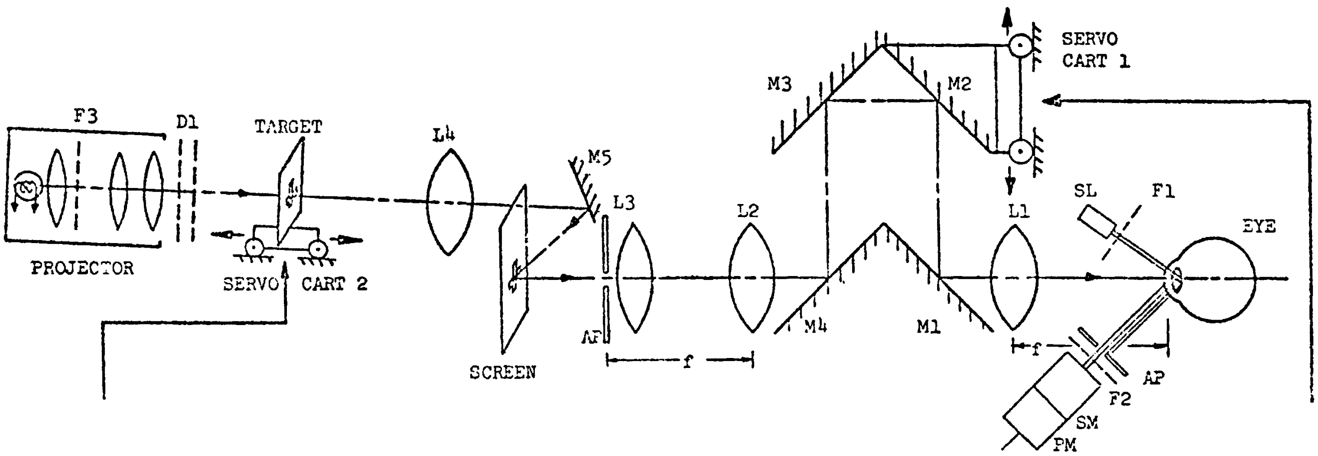

Figure 4. *Phillip's apparatus. Servo Cart1 controls target distance by moving mirrors M2 and M3 to change length of optical path; Servo Cart2 controls target blur by changing distance to target focusing lens L4. L1 is Badal lens; SL is slit lamp for illuminating lens; PM/SM block detects accommodative state.*

retinal defocus. Likewise, the apparatus could be set up to close the (normally open) loop of target blur, by using the accommodation signal to change the projector focus.

In principle, the *virtual lens*, shown schematically in Figure 1, is simply an extension of Phillips' apparatus, but with computer hardware and software replacing some of the apparatus's mechanical and optical components. The optics and servo cart for blurring the target image are replaced by digital image-processing software and the screen on which the target was projected is replaced by a high-resolution computer CRT. The blurred images can be precomputed at sufficiently small intervals (e.g., every 0.1 diopters) and then displayed on the CRT when appropriate, obviating the need for doing computationally expensive image-processing operations in real time.

It was noticed that the virtual lens offered another advantage over Phillips' apparatus: with standard optical methods, the point-spread function of a blurred image is approximately gaussian; by using digital-image convolution to produce the blurred images rather than optics, it is possible to use arbitrary point-spread functions for the convolution kernel. In fact, it is not necessary to limit the ''blur'' operations to linear convolutions at all; any arbitrary transform or distortion of the image can be used. Thus, the name ''virtual lens'': the combination of dynamic optometer measuring accommodation response, an optomechanical system to open-loop the accommodation system and keep the CRT image focused on the retina, and the computational modification of the image as a function of that response effectively replaces the optical filtering function of the crystalline lens with an arbitrary filter or other distortion, as shown in Figure 2. It is thus a type of ''mounted display,'' like an HMD, as shown in Figure 5, in which the image viewed is modified as a result of a measured motor action of the observer.

3.1 Design of virtual lens

3.1.1 Optics. The function of the optical subsystem of the virtual lens is to keep the image of the CRT display focused on the subject's retina, independent of changes in the subject's accommodation—in control theory terms, to open-loop the accommodation system. This is done by using real-time dynamic accommodation measurements to keep the target at the eye's clear-vision distance. This requires an accurate, stable dynamic optometer with a sufficiently high sampling rate, a computer to translate the accommodation signal into a signal capable of driving the virtual lens optical system, and an optical system that can function as the inverse of the accommodative response. This system is

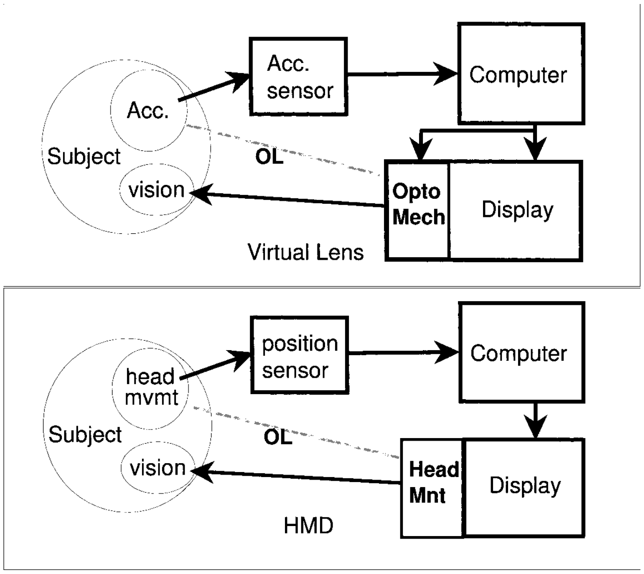

Figure 5. *Feedback displays: analogous nature of virtual lens and head-mounted display.* Top: *Virtual lens feedback display system.* Bottom: *Head-mounted feedback display system.* Acc. *is accommodation of subject;* Acc. sensor *is the optometer that measures accommodation in the subject; grey line labeled OL stands for open-loop and indicates that the neuromuscular control system on the left end of the gray line (accommodation or head movement) is open-looped by the mechanism on the right end of the line (optomechanical system or HMD)*

shown as H_A^{-1} in Figure 1. This is implemented by changing the optical path length of the target using a right-angle prism moved by a servo motor and the associated analog pre-amp and motor current amplifier, and the optical target setup. It is important to minimize the delay between change in accommodation and change in target position, and that the slew rate of the target be greater than maximum velocity of accommodation ($\approx$10 diopters / sec.) (Shirachi, Liu, Lee, Jang, Wong, & Stark, 1978). See Figure 6 for a block diagram of the virtual lens.

3.1.2 Dynamic Optometer. The optometer used was the Tomey QR-007, a commercial autorefractometer that uses the retinoscopic method to measure refraction. This QR-007 was specially modified by Tomey engineers to measure accommodation 16 times per second and output the raw phase data over a 9600-baud, built-in serial port. The well-known Nyquist theorem says that a 16 Hz sampling rate yields a maximum resolvable frequency of 8 Hz. Fortunately, the accommodation system is a very slow system: the power spectrum falls off as $1/s^2$, where s is the laplacian complex frequency, and most of the spectral energy lies below 2 Hz.

Delay: There are a number of components of delay in this version of the virtual lens. The largest component of delay in this system is currently the sampled data delay of 62.5 msec. (worst case) between a change in accommodation and its measurement by the optometer. Assuming for the moment a negligible delay between the analog sampling of phase information, its digitization and insertion into buffers for transmission through the serial port, the 9600-baud serial port imposes a limit of about 960 characters per second, or 96 accommodation measurements per second, or roughly an additional 10 msec. Timing loops in the display software give upper bounds of 8 msec. of additional delay, for a total of 80 msec.

3.1.3 Computer Controller. The computer controller was written in Objective-C using the NeXTStep Development environment. The architecture consisted of a number of connected objects and is shown schematically in Figure 7. The Tomey object is the interface to the Tomey QR-007 Refractometer. It translates raw phase data coming from the RS-232 serial port to sphere, cylinder, and axis refraction measurements, and sends this information on to the Target Distance and Target Display objects. The Target Distance object controls the optical distance of the target screen to the subject's eye. The Target Display object controls the image on the target screen. Each of these objects can operate in two modes, a *preprogrammed* mode and a *dynamic* mode. In the preprogrammed mode, the target distance or display is modified according to a stored data file (e.g., sine wave, square wave, ramp); in dynamic mode the real-time accommodation data from the Tomey is used to update the target, possibly in conjunction with preprogrammed data.

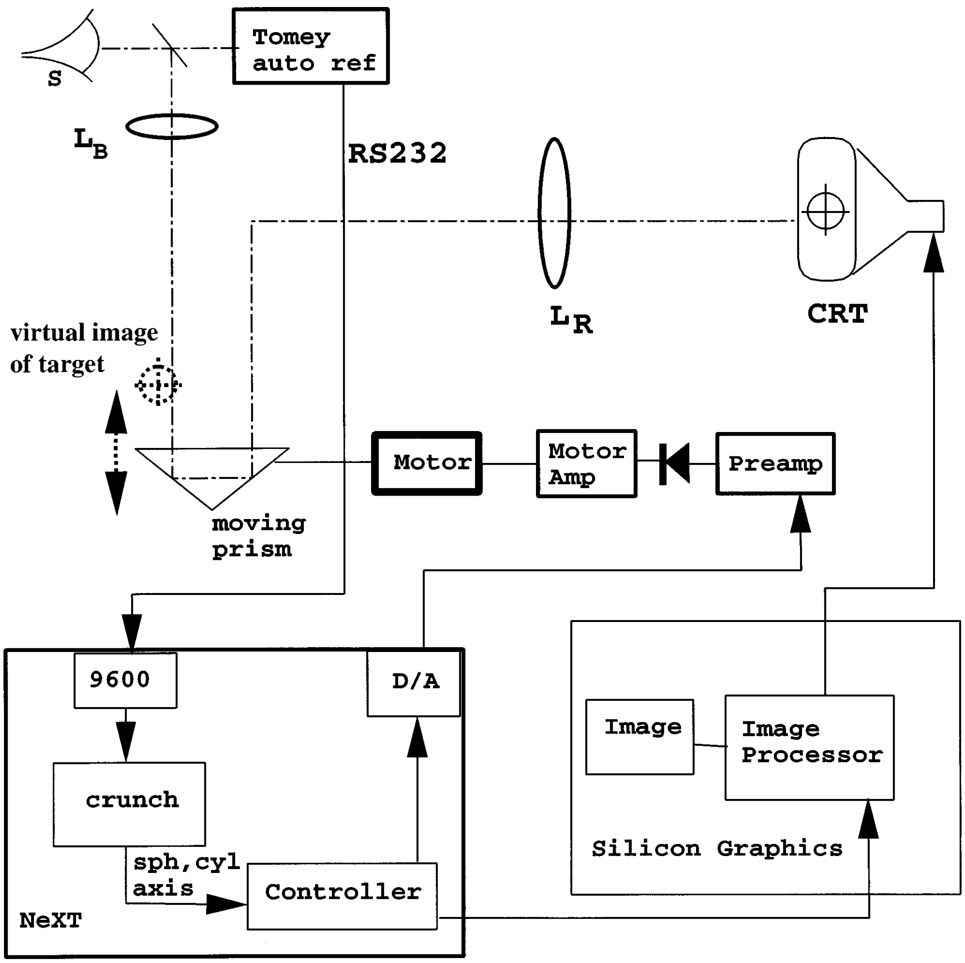

Figure 6. *Virtual lens. Block diagram of virtual lens. S is subject's eye; L_B is Badal lens; L_R is relay lens; I_{LR} is virtual image formed by relay lens.*

3.1.4 User Interface. The user interface is divided into three windows: the Target Distance Control window, the Display Defocus Control window, and a Virtual Lens Control panel for experiment-specific information, calibration parameters, and real-time display of accommodation response. (See Figure 8.)

3.1.5 Target Distance Control. The target distance control can run in two modes. In preprogrammed mode, target distance is a predetermined function of time, an independent variable. In dynamic tracking mode, the target distance is some function of the subject's accommodation response. In the current version of the software, the function is constrained to be linear, because it was easy to program, but of course there is in principle no reason the relationship can't be nonlinear.

Signal: The target distance signal output by the NeXT workstation is an analog sound signal from the sound-out jack. This signal consists of an audio-frequency carrier wave (440 Hz or 880 Hz are perfectly acceptable), amplitude-modulated by the target position signal. The voltage range of the output is $+/-0.5$ volts. This signal is then half-wave rectified by a single diode, creating a pulse-height encoded DC signal. This pulse-height encoded DC signal is converted into a current signal by the motor amp and sent to the servo motors; the result is a very linear system between about 10% and 100% modulation.

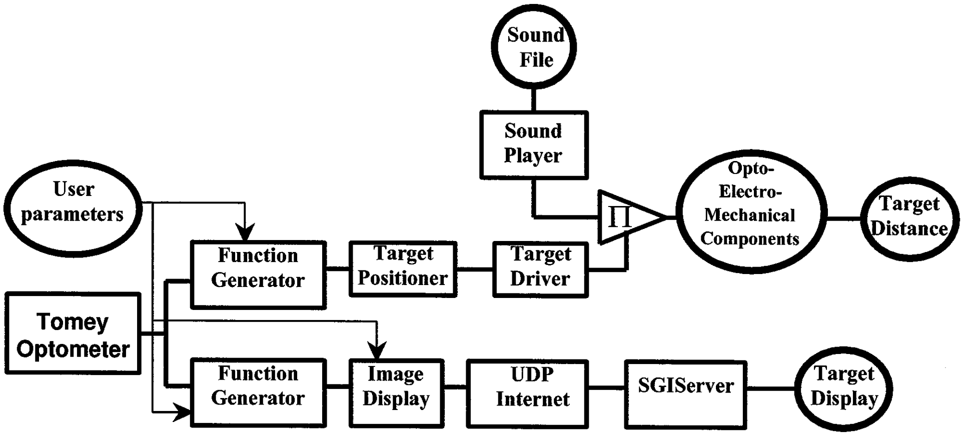

Figure 7. *Software architecture.*

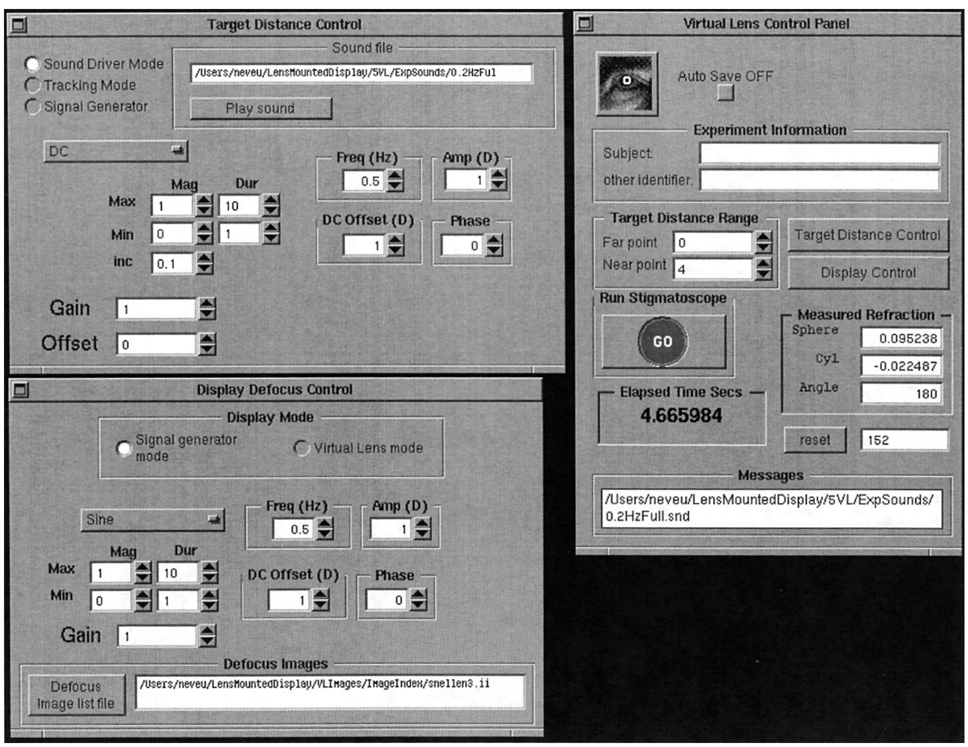

Figure 8. *User interface for virtual lens.* Right: *Main window;* Upper left: *Target Distance Control window;* Lower left: *Target Blur Control window*

Sound source: The source of the sound signal for controlling target distance can come from either a stored sound file (preprogrammed mode), or from a stored carrier wave, played in a loop and modulated in real time by sending setVolume messages to the SoundPlayer object (dynamic mode). This is done by the TargetDriver object. The TargetPositioner object accepts commands to change the target distance to a particular optical distance, specified in diopters. It uses calibration information supplied by the user to convert the command into a modulation in the range 0–100%.

3.1.6 Target Display Control. The target display control is analogous to the target distance control. It can run in a preprogrammed mode or a dynamic tracking mode. In preprogrammed mode, the displayed images follow a predetermined sequence; in the current version of the software, this takes the form of a signal generator, capable of generating sinewaves, pulses, or other time series, that are sent to the Display object to be translated into a particular image. In dynamic tracking mode, the signal from the signal generator represents a nominal target distance. Assume that the image-processing function used for blur is convolution with a gaussian, and our intent is to simulate defocus blur. When the nominal target distance is the same as the subject's clear-vision distance, the image displayed would be very sharp; when the subject's clear-vision distance differed greatly from the nominal target distance, the displayed image would be very blurred. Of course, any type of image distortion can be used, or indeed, any sequence of images.

The Display object accepts commands to change the target image distortion to one specified by B_T, where B_T is a focus error specified in diopters. For example, if the image distortion used is normal blurring, $B_T = -2.0$ diopters would specify an image that was low-pass filtered with a cutoff of approximately 15 cycles / degree, corresponding to 2.0 diopters myopic.

3.1.7 Optomechanical System. The lens L_R (f_L = 140 mm) is a relay lens that forms the virtual image I_{LR} of the CRT screen; this serves as the object for the lens L_B. The first nodal point of the subject's eye is made coincident with the second focal point of L_B (f_L = 100 mm), forming a Badal lens system (Badal, 1876). There are two main advantages to using a Badal lens system: first, there is a linear relationship between object distance from the Badal lens and accommodative stimulus in diopters; we use a +10 D (10 cm focal length) lens, enabling us to provide a maximum stimulus to accommodation of −10 D, equivalent to an object 10 cm from the eye. Since accommodation decreases with age roughly according to the formula $12 - age/4$, this range of accommodation generally cannot be produced by anyone past adolescence. The second advantage to a Badal lens system is that magnification of the target remains constant throughout its range of motion, eliminating size and intensity changes associated with changes in optical distance.

The position of the virtual image I_{LR} is moved with a trolley mechanism similar to that used by Phillips. The optical path makes a U-turn through the right-angle moving prism, a change of one unit in the prism position results in a change of two units of the optical path length. The prism has a full range of motion of 10.5 cm, so the range of target distance seen by the subject is −10.0 D to +11.0 diopters (0 diopters is optical infinity). The prism is moved by a +/− 22-volt servo motor and has a maximum slew rate of 70 cm / sec, giving a maximum rate of change of accommodative stimulus of 140 D / sec, considerably greater than the 10 D / sec maximum rate of the human accommodative system.

3.1.8 Display. A Silicon Graphics color RGB monitor was used as the CRT display. The display was positioned so that the image *f* of a 128 × 128 pixel square subtended a solid angle of 1° × 1°. This gives better than 60 cycles per degree of resolution while keeping the image a manageable size in pixels for image processing purposes.

4 Applications

4.1 Accommodative Control System

The virtual lens was originally developed to study accommodation in humans. Previous workers studying biological control systems have found it useful to "clamp" aspects of the system and observe the response, e.g., voltage clamping of neurons, patch clamping of cell membranes (Boulton, Baker, & Walz, 1995). Biological systems use feedback control for the same reasons that engineers use it in their designs: feedback can be used to stabilize and linearize otherwise unstable and nonlinear

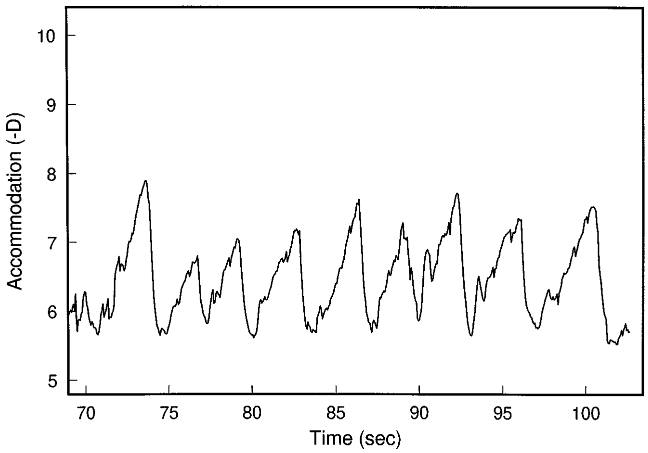

Figure 9. *Vergence clamping. Target vergence is maintained at 0.5 diopters nearer than actual accommodation measurement. Subject accommodates to near point, then relaxes, resulting in sawtooth-like signal.*

components. When the feedback loop is broken or modified, these characteristics become much more evident and can reveal a great deal about the internals of the control system.

The virtual lens can be used to ''error clamp'' the eye by tracking accommodation with target distance. In Figure 9 the target distance was kept slightly nearer than the subjects' accommodation, clamping the defocus error to ≈ -0.5 diopters; the resulting trace shows the accommodation ''chasing'' the target in a number of jumps, forming a ramp-like response, until the accommodative effort became too great, and the subject gave up and relaxed accommodation, only to have the target drop back in response, and the cycle begin again. This technique of breaking the feedback loop so that changes in response don't produce changes in the error signal is a well-known control system technique called *opening the loop* (Stark, 1968).

4.2 Visual Psychophysics: Blur Detection

The virtual lens can also be used to study the mechanism of blur detection in humans. A number of passive, image-based autofocus algorithms have been developed for computer vision and consumer electronics (primarily videocam) applications. These can be broadly categorized as variance-based and spatial frequency-based. A third possibility exists for human observers: signal-to-clutter maximization, where clutter is unwanted or distracting information. Since all of these mechanisms respond properly to defocus, it's necessary to use some other stimulus besides defocus to distinguish between these mechanisms. By replacing the distortion function, blur, with, say, a series of progressively more distorted images that have the characteristic that their variance is constant but their power spectrum is redistributed—a *variance-invariant transform*—we can distinguish between variance-based detection mechanisms and signal-to-clutter mechanisms. The statistical variance of an image I

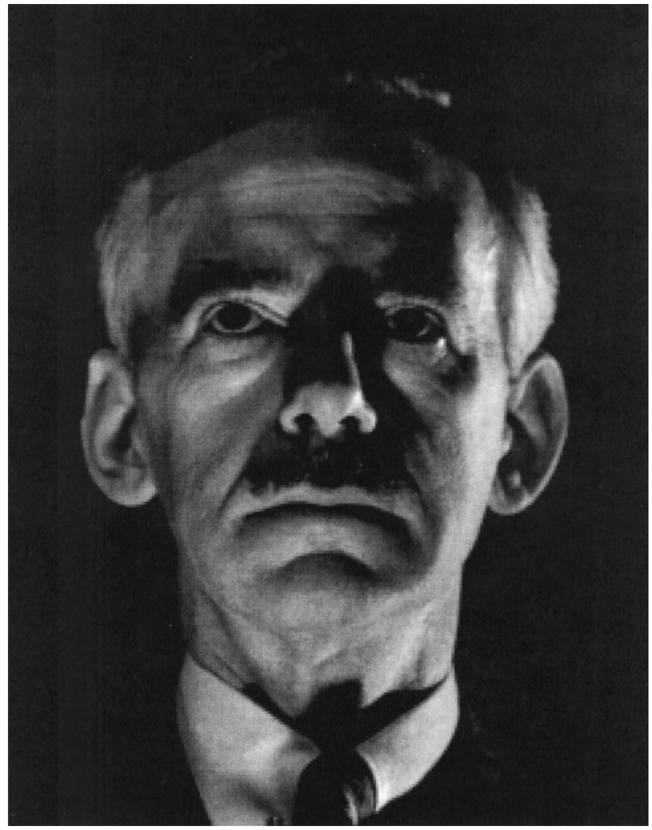

Figure 10. *Original oneill image.*

$$\sigma = \sum_{x \in I} \frac{(x - \bar{x})^2}{I_N - 1}, \tag{1}$$

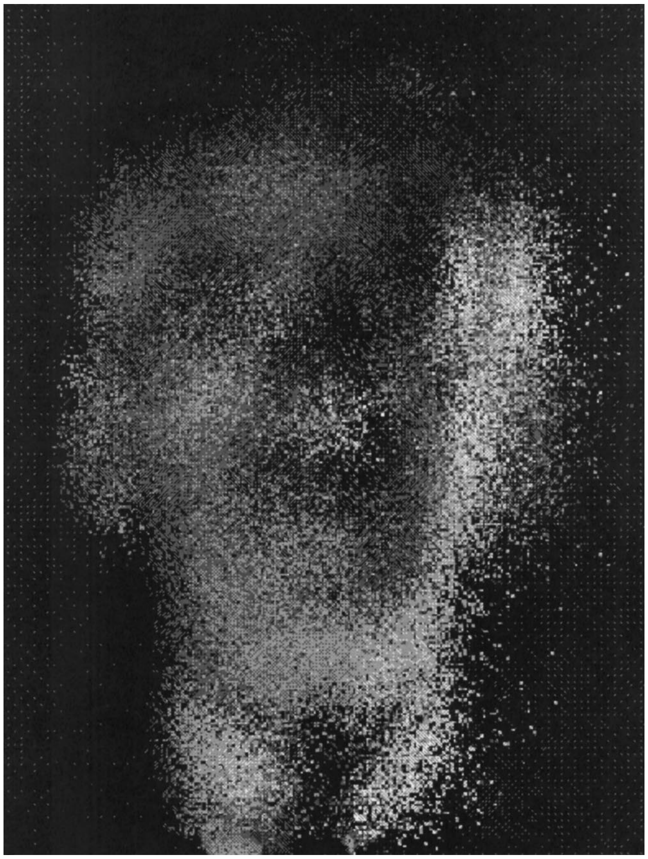

Figure 11. *Diffusion transform. A simple simulation of diffusion. Pixels are swapped with pixels some random distance and direction away; small mean-path length preserves large-scale structure and disrupts small-scale detail. The tendency of the pixels to drift to the upper-right is an artifact of the bottom-to-top, left-to-right iteration loop of the diffuse program.*

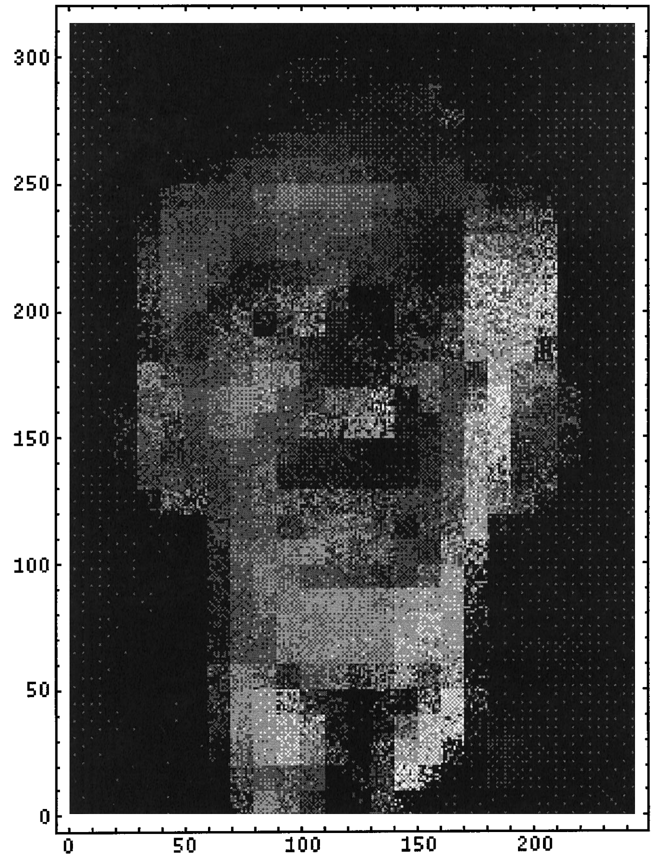

Figure 12. *Scramble transform. Image is divided into non-overlapping discrete tiles. Positions of pixels within a tile are randomized.*

where I_N is the number of pixels in the image, is independent of the relative positions of the pixels, and depends only upon their values. Any transform that changes the positions of pixels without changing their values will not change σ. The diffuse transform does a crude simulation of pixel diffusion: each pixel $x_{i,j}$ is swapped with another pixel $x_{i+p,j+q}$, where p, q, the *diffusion range* variables, are random variables uniformly distributed over some small range. Figure 10 shows the original oneill image before transformation. The results of the diffuse transform are shown in Figure 11 where p, $q \in [-5,5]$. A small diffusion range disrupts the local pixel neighborhoods while keeping the global, large-scale structure recognizable.

The scramble transform (Figure 12) divides the image into non-overlapping discrete tiles, and within these tiles randomizes the position of the pixels. This transform, like the diffuse transform, also maintains global structure while obscuring local details; in addition, the blockiness of the result is an artifact of the tiling operation, and adds distracting, irrelevant detail, as in Harmon's Lincoln (Harmon, 1973).

Parseval's theorem shows that the variance of an image is proportional to the total power in the power spectrum of the Fourier transform, and that if the total spectral power is kept constant, the variance remains constant. Total power is a function of the magnitude only and not the phase of the individual frequency com-

Figure 13. *Original Snellen image.*

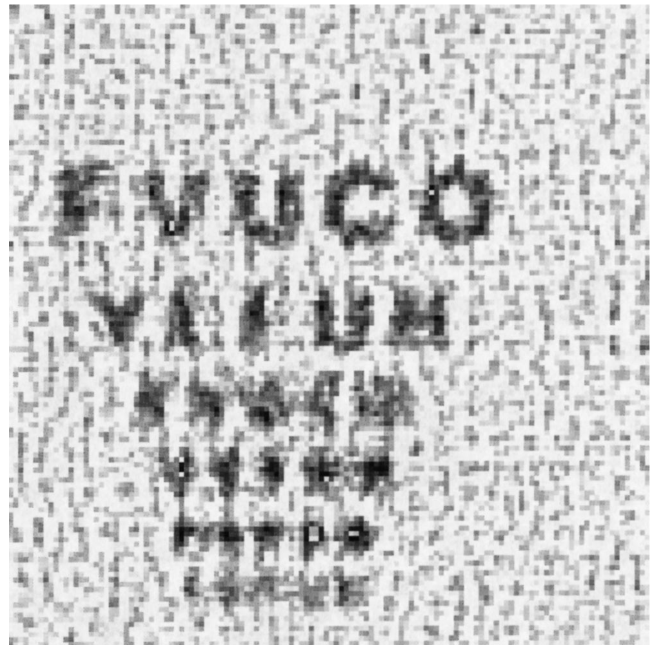

Figure 14. *Phaseblur transform. Noise is added to the phase component of frequency spectrum of the image. The power spectrum is unaltered.*

ponents, but the useful image information is contained primarily in the phase component. Thus, by randomizing the phase of frequency components (taking care, of course, to keep the resulting Fourier transform hermitian, ensuring a real-valued image upon inverse transformation), we can add noise to the image and distribute it in arbitrary spatial frequency patterns. From the original *Snellen* image (Figure 13), the phaseblur transform randomizes the phase of the high-frequency components of the image, so that the useful image information is limited to the low spatial frequencies, as shown in Figure 14.

4.3 Virtual Environment Applications

As was noted above, the virtual lens is a fundamental type of feedback display, in which an ocular-motor response is measured and used to modify a visual display. Our prototype was designed for visual psychophysics experimentation and would be impractical for virtual environment displays; in principle there is no reason the device could not be made small enough to be used with an HMD. Many HMDs have mechanical methods of adjusting the optical display distance; a motor-driven lens system, like those in high-quality, autofocus 35 mm SLR cameras, could function as the optomechanical system. The optometer requires a optical path through the pupil of the eye to the retina, and a means of illuminating the retina with an infrared reference beam, and a means of recording and measuring the beam's reflection.

4.3.1 Adaptive HMDs. The binocular convergence control system operates in synkinesis with the accommodation system, and can drive accommodation. Indeed, it is probably the primary driver of accommodation in normal, day-to-day binocular vision. In real-world vision, binocular convergence and accommodation are intimately linked: a near object has to be converged upon by both eyes, and must be accommodated to. This is shown diagrammatically in Figure 15. The diagonal line is called *Donder's line,* and represents the relationship between convergence demand and accommodative demand for real targets. All real targets fall on Donder's line; this relationship can be disturbed by placing lenses and / or prisms before the eye. Fortunately,

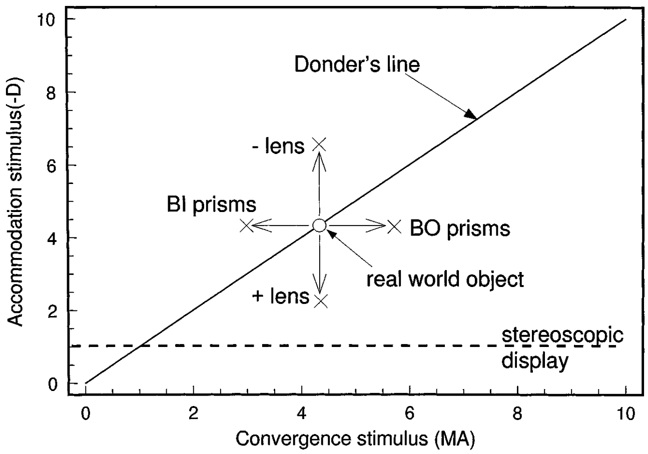

Figure 15. *Donder's line, real and virtual objects. This figure shows the relationship between accommodative demand and convergence demand. All real targets fall on the diagonal line labeled Donder's line. Viewing real targets through lenses or prisms changes the relationship, moving the targets off Donder's line; objects in stereoscopic displays fall on the dashed line. BO is base-out prisms; BI is base-in prisms.*

the human visual system is quite adaptable, and in practice there is a region around Donder's line, called the *zone of single, clear, binocular vision* within which a target can be converged and focused upon. Stereoscopic displays also create visual stimuli off Donder's line, as the stimulus to convergence—retinal disparity—is variable, and the stimulus to accommodation—the screen distance—is constant. A virtual lens could correct this phenomenon, by modifying the optical target to correspond to the convergence angle keeping the stimulus displayed on Donder's line.

4.3.2 Simulated depth of field. A virtual environment display could also incorporate accommodative state information to change depth cues; for example, a display could incorporate a virtual depth-of-field algorithm, which would blur those parts of the graphical display that are distant from the subject's accommodative plane.

Acknowledgements

We are grateful to Dr. Miyao and Dr. Ishihara for their assistance with the Tomey optometer, and also the engineers at Tomey: Eiichiro Nagasaka, Masao Yoshikawa, Susumu Nozawa, Toshiharu Morino.

References

Badal. (1876). *Ann oculist, 75,* 5–13. Jan.–Feb. 1876.

Boulton, A. A., Baker, G. A., & Walz, W. (1995). *Patch-Clamp Applications and Protocols (Neuromethods).* Humana Press.

Hannaford, B., Wood, L., McAffee, D., & Zak, H. (1991). Performance evaluation of a six-axis generalized force-reflecting teleoperator. *IEEE Transactions on Systems, Man and Cybernetics, 21*(3), 620–633.

Harmon, L. (1973). The recognition of faces. *Scientific American.* Nov. 1973.

Phillips, S., & Stark, L. (1977). Blur: A sufficient accommodative stimulus. *Documenta Ophthalmologica, 43*(1), 65–89.

Phillips, S. R. (1974). *Ocular Neurological Control Systems: Accommodation and the Near Response Triad.* (Doctoral dissertation, University of California).

Rohen, J. W. (1979). Scanning electron microscope studies of the zonular apparatus in human and monkey eyes. *Investigative Ophthalmology and Visual Science, 18,* 133–144.

Rohen, J. W., & Rentsch, F. J. (1969). The construction of the human zonular apparatus and its functional significance. *Albrecht von Graefes Arch Klin Ophthalmol, 178,* 1–19.

Rohen, J. W., & Zimmerman, A. (1970). Changes due to age of the human ciliary body. *Albrecht von Graefes Arch Klin Ophthalmol, 179,* 302–317.

Shirachi, D., Liu, J., Lee, M., Jang, J., Wong, J., & Stark, L. (1978). Accommodation dynamics: 1. Range nonlinearity. *American J. of Optometry and Physiological Optics, 55,* 631–641.

Stark, L. (1968). *Neurological Control Systems: Studies in Bioengineering.* Plenum Press, New York.

Sutherland, I. E. (1968). A head-mounted three-dimensional display. In *Proc. of the Fall Joint Computer Conference* (pp. 757–764).

Szakaly, Z., & Bejczy, A. K. (1990). Performance capabilities of a JPL dual-arm advanced teleoperations system. *Proc. of the 4th Ann. Workshop on Space Operations, Automation and Robotics* (pp. 30–41), Albuquerque, NM.