The Virtual Lens

Charles F. Neveu* † Lawrence W. Stark*

June 11, 1995

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 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**

Interest in virtual environments has spurred the development of new types of *feedback displays*. Feedback displays are information sources that also receive feedback from the user, and use it to modify the display. The most familiar feedback display is the helmet-mounted display [13], which senses head position and uses this information to modify the image shown to the subject. Force-feedback joysticks [1, 14] sense hand position through the joystick and offer resistance or force through the joystick back to the hand.

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 and state, 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 lens—*i.e.* blurring—can be removed by placing a controllable, predistorting opto-mechanical system (shown as H_a^{-1} in Figure 1) between the subject and the display, and using accommodative state information to make

---

*Neurology and Telerobotics Lab, University of California, Berkeley

†Author to whom correspondence should be addressed.

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).

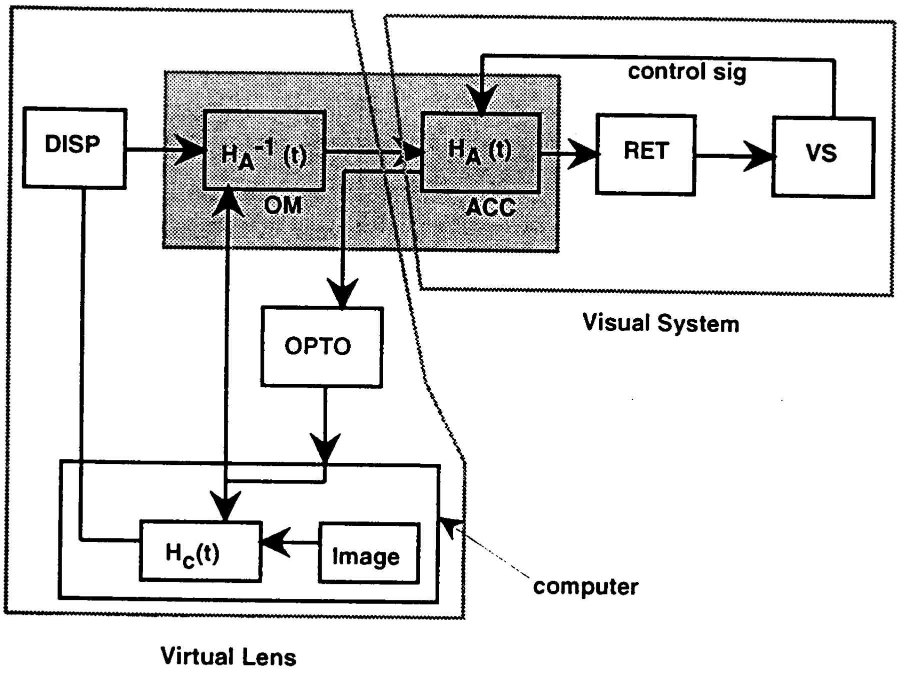

Figure 1: Schematic diagram of virtual lens
Gray box shows optics of crystalline lens and inverse optical system of virtual lens.

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.

2 Ocular accommodation

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

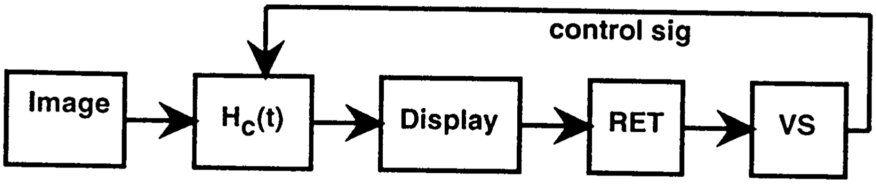

Figure 2: Reduced schematic of virtual lens.
Upper: Inverse optics of virtual lens cancel optics of accommodation; change in accommodation drives digital filter that modifies image.
Lower: Final equivalent control loop with optics cancelled; computed transfer function $H_c(t)$ replaces transfer function of the accommodation system H_A.

conjugate to the fovea, in other words, in focus. Helmholtz elucidated the theory of the mechanism of accommodation as generally accepted today, called the *dual, active, indirect* theory. The lens is suspended from its equator, not unlike the hub of a bicycle wheel suspended by 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, a 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 mid-point, 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 [8, 9, 7]. 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 ($\approx$ 98%) of the muscle fibers proceed peripherally from their origin to their insertion onto the span fibrils of the zonule of Zinn at Bruch's membrane (some spread laterally for a short distance in a load-sharing mechanism; these lateral components have been confused by many—from Müller on—to image a circumferential, sphincter-like muscle, although clearer thinkers from Helmholtz on have denied the existence of circumferential muscle). 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 extert 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. The theory is called *dual* because it involves both lenticular and extralenticular mechanical components, as opposed to theories that the lens was a kind of transparent muscle which changed its own shape [15]. It is *active* because the active contraction of the ciliary muscle causes an increase in the power of the lens, primarily because the front surface of the lens becomes more curved – the lens becomes rounder. it is *indirect* because the ciliary muscle does not directly compress the lens like a circumferential sphincter, but relieves passive tension originating primarily in the peripheral zonules (with possibly a small contribution from the elasticity of the choroid) that pulls on the lens capsule at the equator of the lens and flattens it.

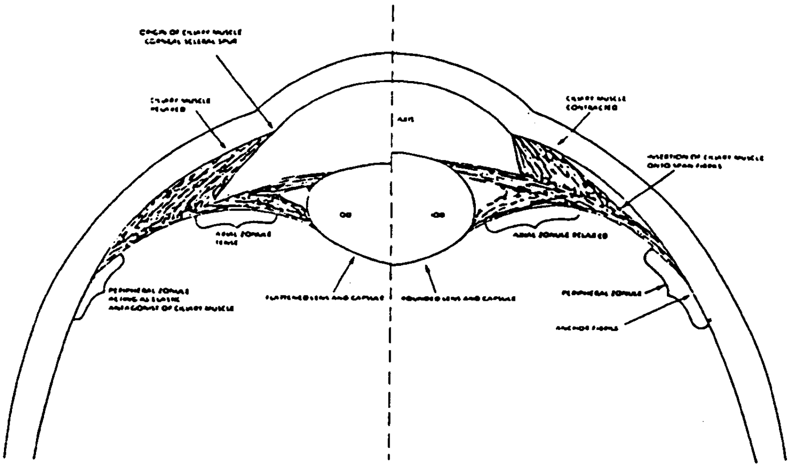

FIG. 4. The eye: sagittal section—showing important accommodative structures in anterior third of the eye (see text).

Figure 3: Mechanism of accommodation

3 Virtual lens

Phillips [6, 5] developed an instrument which 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. 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 control loop of retinal defocus by matching the accommodative state of the subject with target distance, so that changes in accommodation response had no effect on the degree or direction of 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.

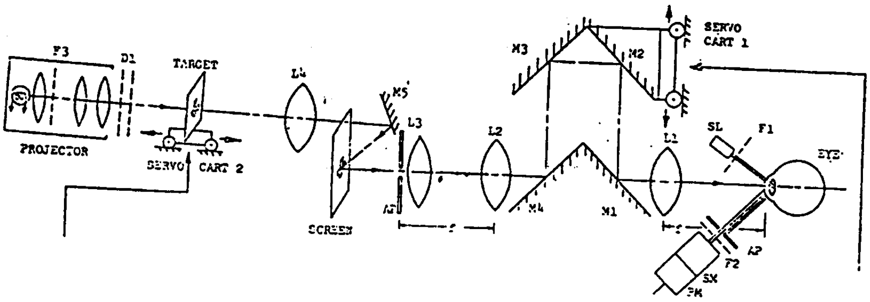

Figure 4: Phillip's apparatus

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 in advance 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 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 opto-mechanical 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 a head-mounted display, as shown in Figure 5, in which the image viewed is modified as a result of a measured motor action of the observer.

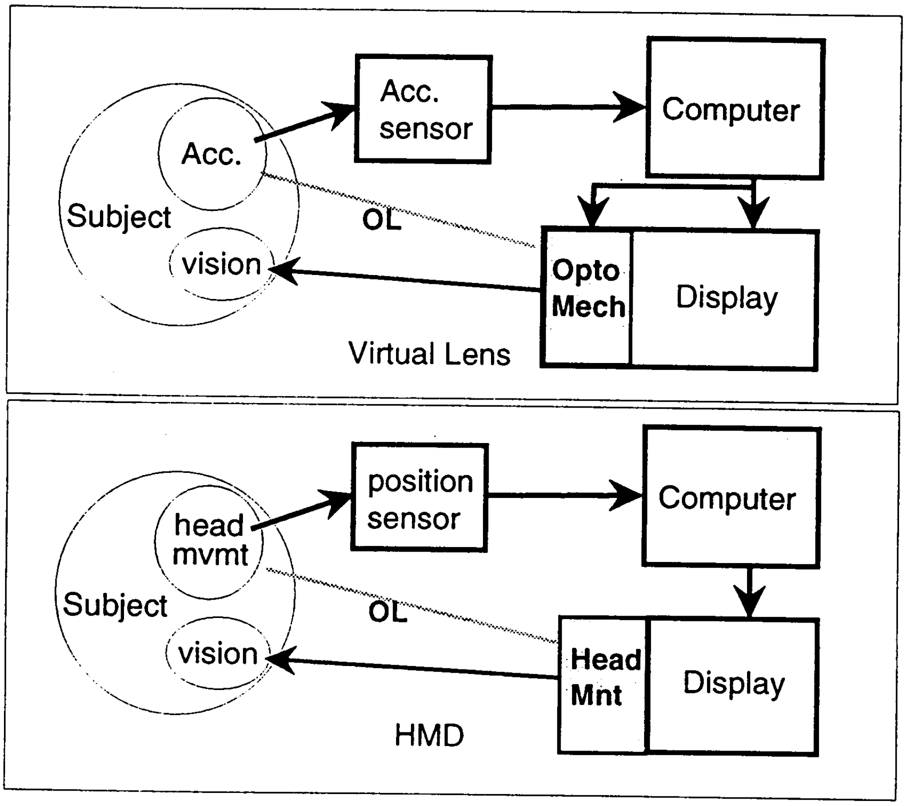

Figure 5: Feedback displays: Comparison of virtual lens and head-mounted display.

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 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. We explored two ways of doing this, with a pinhole pupil, and by active tracking and correction.

Pinhole pupil The pinhole pupil has a long history of use as a mechanism for opening the control loop of accommodation. It works by increasing the depth of field so that, for accommodation errors within the depth of focus, the image is essentially unchanged.

Unfortunately, artificial pinhole pupils have a number of drawbacks. The first problem is mechanical: the Tomey optometer needs a minimum 3mm pupil to be able to measure accommodation properly, so the pupil has to be constructed so that it is not blocking or interfering with the operation of the optometer. There are two solutions to this problem; the first, to use an artificial pupil made from infrared filter material. If the artificial pupil is centered on the optical axis of the Tomey optometer, the reflections from the edge of the pinhole give spurious reflections which cause the phase-sensitive optometer to give erroneous values It is possible to move the pupil off the optometer's optical axis and have the optometer shoot through the filter, just to one side of the pinhole, but then of course the optometer is well off the subject's visual axis, an arrangement for which it was not designed. A better solution to this problem was to use relay lenses to create a Maxwellian view, an optical system whose exit pupil is a pinhole whose image is placed coplanar with the entrance pupil of the eye, producing a virtual pupil at that plane and restricting the light from the stimulus entering the eye to pass through that virtual aperture. This solves the off-axis problem.

The second problem was more pedestrian: the image viewed by the subject was dim. "Conventional" targets — made of paper, plastic, etc. — illuminated with a white light source can overcome this difficulty by merely increasing the illumination, but our setup depended on the use of an off-the-shelf CRT (either the NeXT MegaPixel monitor or a Silicon Graphics Color Monitor, designed to be sufficiently bright to an individual seated less than a meter away with 3 mm - 5 mm pupils.

Finally, in order to have a wide enough depth of focus so that no change in the retinal image is seen with accommodation, it is necessary to use a ¡ 1.0 mm pupil [3]. With a pupil that size, however, the effect of diffraction becomes significant. The equation (small angle approximation) for the width on the retina of the Airy diffraction disc d_A created by a point source as a function of pupil diameter is

$$d_A(d_p) = 2f\frac{1.22\lambda}{N'd_p} \tag{1}$$

where f is the focal length of the eye, λ is the wavelength of light, N' is the index of refraction of the ocular medium and d_p is the diameter of the pupil. Using 20 mm for f, 550 nm for λ, 1.34 for N' and 0.6 mm for d_p, the diameter of the Airy disc

$$\begin{aligned} d_A(0.6) &= 2 \times 20 \times 10^{-3}\frac{1.22 \times 550 \times 10^{-9}}{1.34 \times 0.6 \times 10^{-3}} && (2)\\ &= 33.4\mu && (3)\\ &= 0.111^\circ && (4)\\ &= 6.68 \text{ arcmin} && (5) \end{aligned}$$

and functions as a low-pass filter with a cut-off of about 9 c/d. This is substantially blurred, as normal visual acuity is around 30 c/d. For comparison, we can estimate the size of the blur circle d_B in radians of a point source defocused 1D (ΔD); seen through a 3 mm pupil using the formula [11]

$$d_B = d_p|\Delta D| \tag{6}$$

which gives a blur disc of 0.003 radians or 10.4 arcmin.

Active Tracking The second way to open-loop the accommodation system is to use 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 which can function as the inverse of the accommodative response, this system is 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 that the delay between change in accommodation and change in target position be small, and that the slew rate of the target be greater than maximum velocity of accommodation ($\approx$ 10 diopters/sec)[10].

3.1.2 Dynamic Optometer

The optometer used was the Tomey QR-007, commercial autorefractometer that uses the retinoscopic method to measure refraction. In the retinoscopic method, a spot of light is shined through the pupil. and its reflection on the retina observed as the light is moved through a line perpendicular to the subject's line of sight. When the eye is over-accommodated relative to the distance of

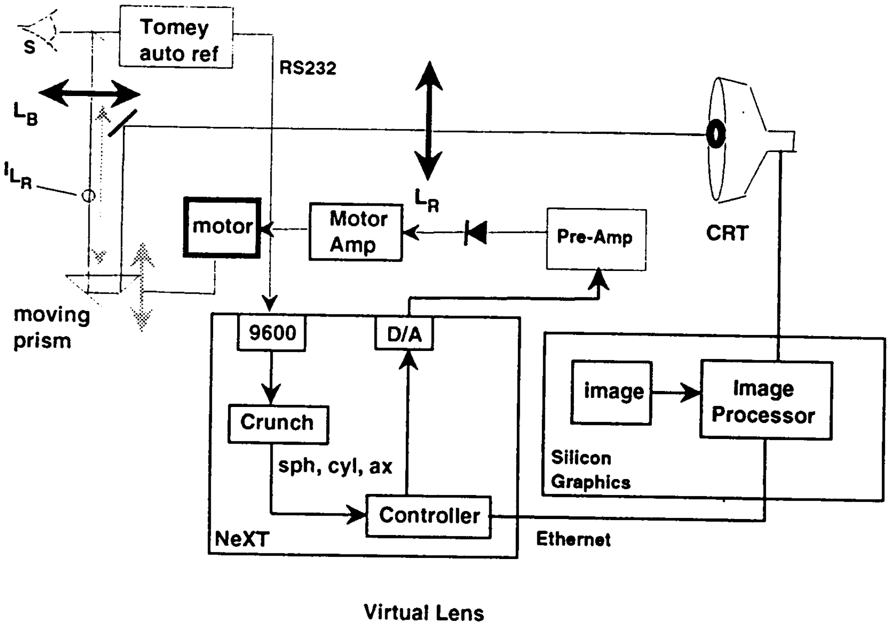

Figure 6: Virtual Lens

the light, the observed reflection off the back of the eye is seen to move in the opposite direction, when the eye is under-accommodated, the reflection moves in the same direction, and when the eye is perfectly accommodated to the distance of the light, the reflection doesn't move at all. When the light is moved through a known trajectory at a known distance from the subject, and the reflection from the retina imaged pair of photo-cells, speed of movement of the reflection can be measured and accommodative state determined. 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 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 $\frac{1}{s^2}$, and most of the spectral energy lies below 2Hz.

Delay There are a number of components of delay in this version of the virtual lens. The rate-limiting 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 ten msec. Timing loops in the display software give upper bounds of 8 msec. of additional delay, for a total of 80 msec., and possibly more. This delay is noticeable during large step changes in accommodation, and interferes with tracking ability; the rate limiting factor is currently the 16 Hz sampling rate of the optometer.

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 8. 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 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.

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.

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. The amplitude-modulation scheme is used because the sound hardware is not capable of producing a sustained DC level necessary for controlling the target position motor. This AM-modulated signal is amplified about 10 times by an operational amplifier, to about +/- 5.0 volts. This signal is then half-wave rectified by a single diode, creating an pulse-height encoded DC signal. The rectification has to be done after the signal is amplified, because a silicon diode is very non-linear in the region between 0 and 0.5 volts, and much more like an ideal diode and 0.5 volt voltage source in series in the region above 0.5 volts. This pulse-height encoded DC signal is converted into a current signal by the motor amp and sent to the servo motors; the carrier does not have to be electronically filtered out: inertia of the motor functions as the integrator! To borrow a metaphor from muscle physiology, the carrier is beyond the motor's "fusion frequency"; the result is a very linear system between about 10% and 100% modulation.

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%.

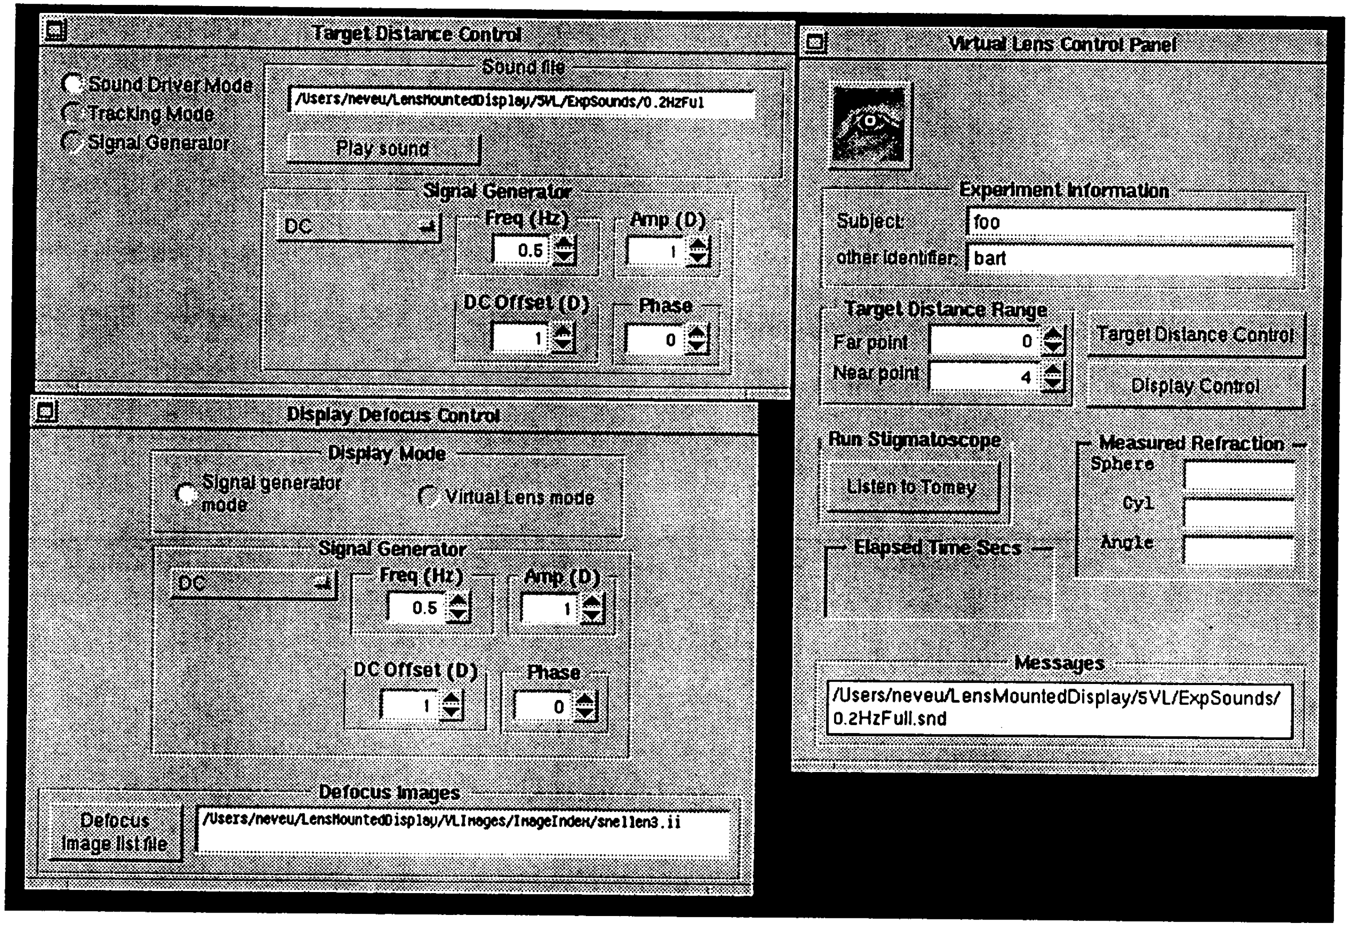

Target Distance Control

● Sound Driver Mode
○ Tracking Mode
○ Signal Generator

Sound file: /Users/neveu/LensMountedDisplay/5VL/ExpSounds/0.2HzFul

Signal Generator: DC; Freq (Hz): 0.5; Amp (D): 1; DC Offset (D): 1; Phase: 0

Display Defocus Control

Display Mode: ● Signal generator mode ○ Virtual Lens mode

Signal Generator: DC; Freq (Hz): 0.5; Amp (D): 1; DC Offset (D): 1; Phase: 0

Defocus Images: /Users/neveu/LensMountedDisplay/VLImages/ImageIndex/snellen3.ii

Virtual Lens Control Panel

Experiment Information
Subject: foo
other identifier: bart

Target Distance Range
Far point: 0
Near point: 4

Messages: /Users/neveu/LensMountedDisplay/5VL/ExpSounds/0.2HzFull.snd

Figure 7: User interface for virtual lens.

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, which 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 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 cut-off of about 15 cycles/degree, corresponding to 2.0 diopters myopic.

The Display object opens up a socket to a server called `SGIServer` running on a Silicon Graphics Personal Iris workstation. It sends commands over the socket using the UDP (Unreliable Datagram Protocol) protocol. These are fast, low-overhead messages; despite the name, they are practically quite reliable on a local area network; the name merely indicates that no flow control is implemented at this level of abstraction and reliability is the program's responsibility. `SGIServer` is a dumb image server; it only understands commands to load image files and to display an image by index. All semantic information about the correspondence between an error signal specified in diopters and a particular degree of image distortion is stored in the Display object.

3.1.7 Optomechanical System

The lens L_R (f_L = 140mm) is a relay lens that forms the virtual image I_{L_R} 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 = 100mm), forming a Badal lens system. 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 $12 - \frac{age}{4}$, this range of accommodation generally cannot be produced by anyone past adolescence. The second advantage to a Badal lens system is that magnification

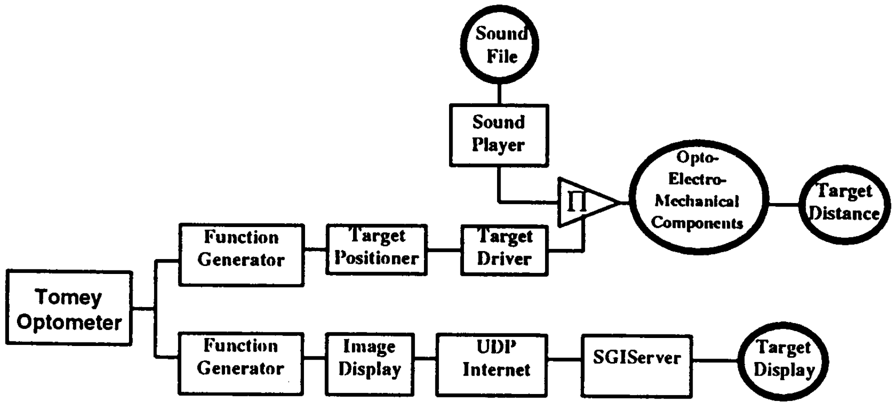

Figure 8: Software Architecture

of the target remains constant throughout its range of motion, eliminating size and intensity changes to associated with changes in optical distance.

The position of the virtual image I_{L_R} 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 2 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 +/- 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°.

4 Applications

4.1 Vision research

4.1.1 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, and glucose clamping in physiology are few prominent examples. 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 non-linear 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.

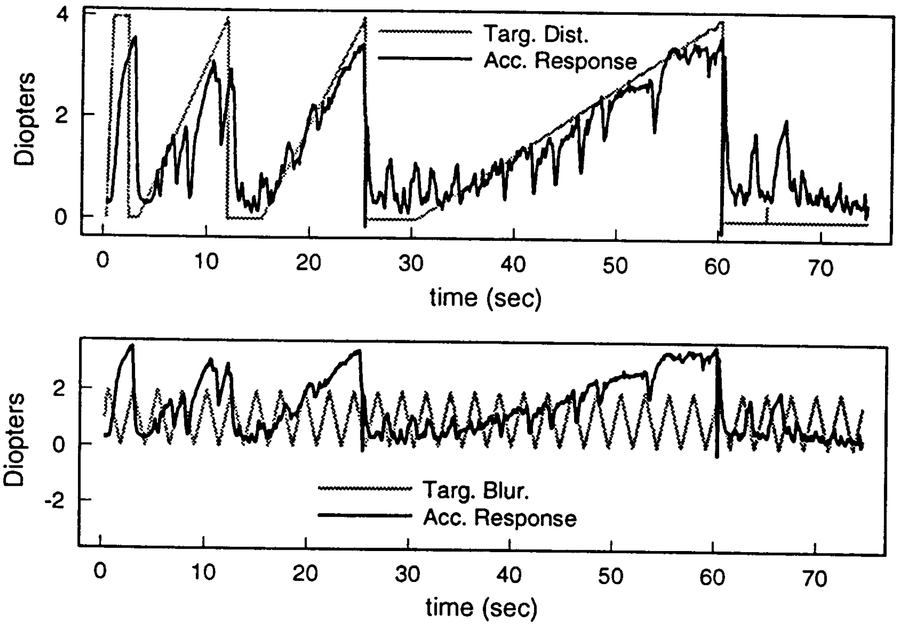

/Net/focus/Charles/neveu/LensMountedDisplay/Dissertation/SciPlotStuff/TargetBlur.SP

Figure 9: Dynamic target blur interferes with tracking of a dioptric stimulus. Subject tracked a dioptric ramp stimulus (upper plot, dotted line) of a target image which oscillated between sharp and blurred (lower plot, dotted triangle wave). Accommodative tracking was confused by the target blur; accommodative error (lower plot, solid line) is in sync with blur.

The virtual lens can be used to "error clamp" the eye by tracking accommodation with target distance. In figure 11 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

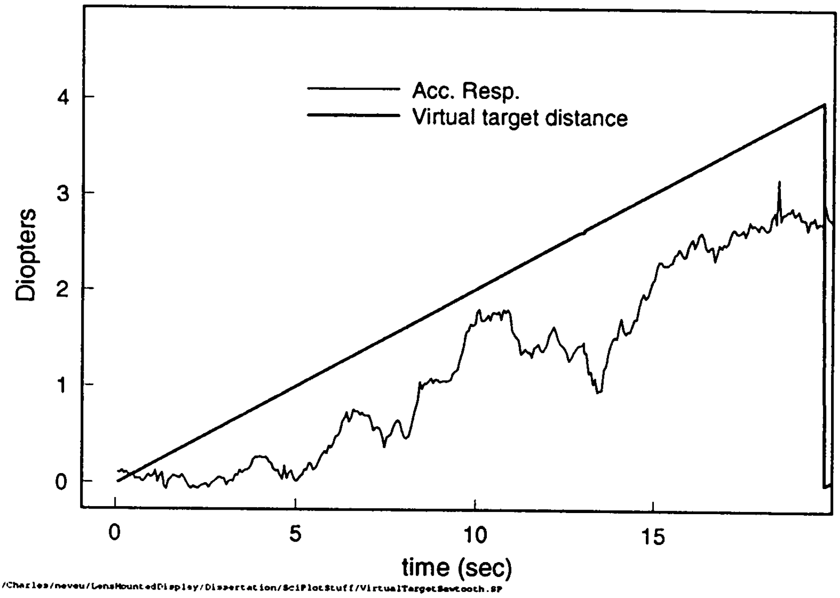

/Net/focus/Charles/neveu/LensMountedDisplay/Dissertation/SciPlotStuff/VirtualTargetSawtooth.SP

Figure 10: Tracking target blur in virtual lens mode.
Ramp stimuli are nominal target positions, solid line is accommodative response. Target blur seen by the subject was the error between nominal target position and actual accommodative response, i.e. when the subject matched the nominal target position, the image was sharp.

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* [12].

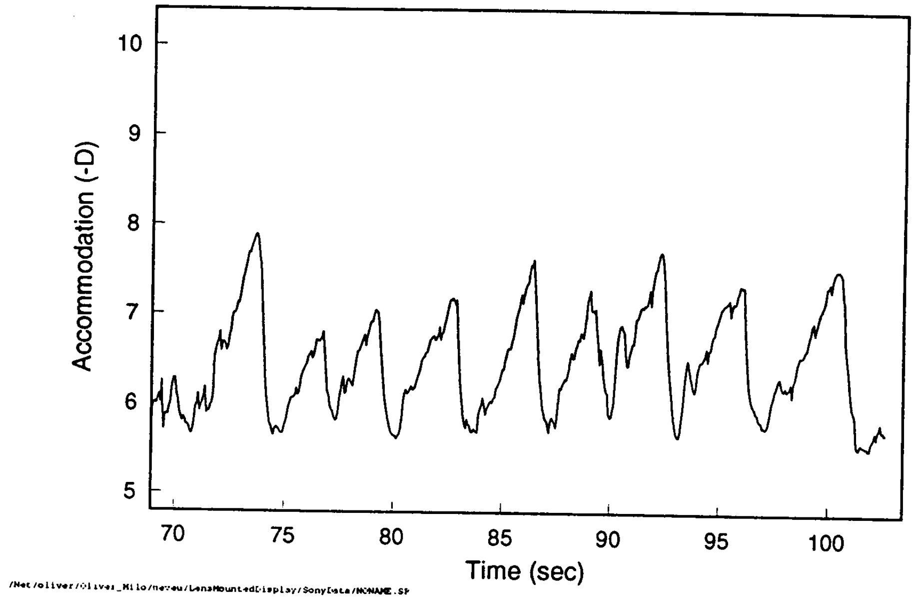

Figure 11: Vergence clamping with delay.

4.1.2 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 which have the characteristic that their variance is constant but their power spectrum is re-distributed—a *variance-invariant transform*—we can distinguish between variance-based detection mechanisms and signal-to-clutter mechanisms[4]. The statistical variance of an image I

$$\sigma = \sum_{x \in I} \frac{(x - \overline{x})^2}{I_N - 1}, \tag{7}$$

where I_N is the number of pixels in the image, is independent of the relative positions of the pixels, and depends only upon the 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. The Figure 12 shows the original oneill image before transformation. The results of the diffuse transform are shown in Figure 14 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 15) divides the image up into non-overlapping discrete tiles, and within these tiles, randomizes 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 [2].

Parseval's theorem shows that the variance 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 components, 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, insuring a real-valued image upon inverse transformation) we can add noise to the image and distribute it in arbitrary spatial frequency patterns. 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 16.

4.2 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 a head-mounted display. Many HMDs have mechanical methods of adjusting the optical display distance; a motor-driven lens system, perhaps like those in high-quality autofocus 35mm 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

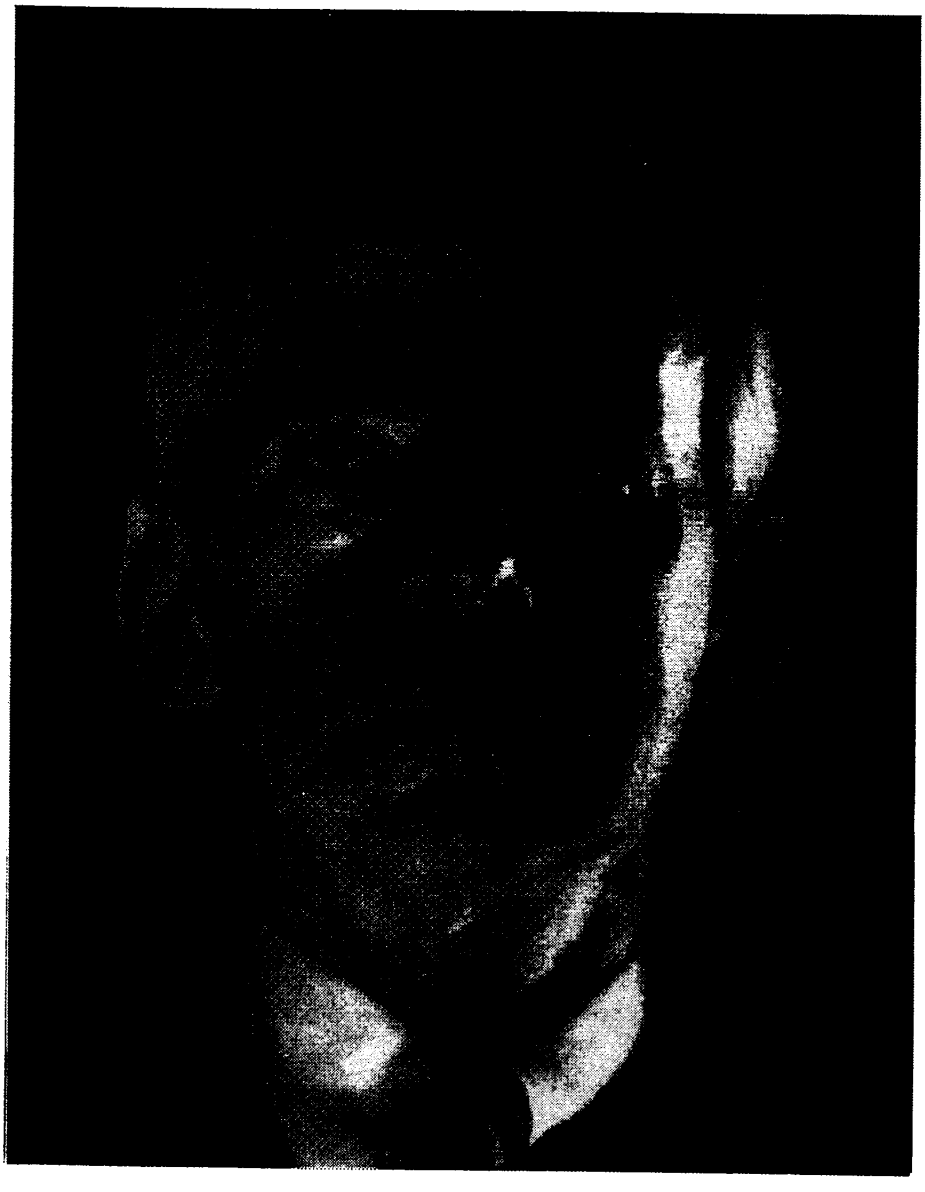

Figure 12: Original `oneill` image

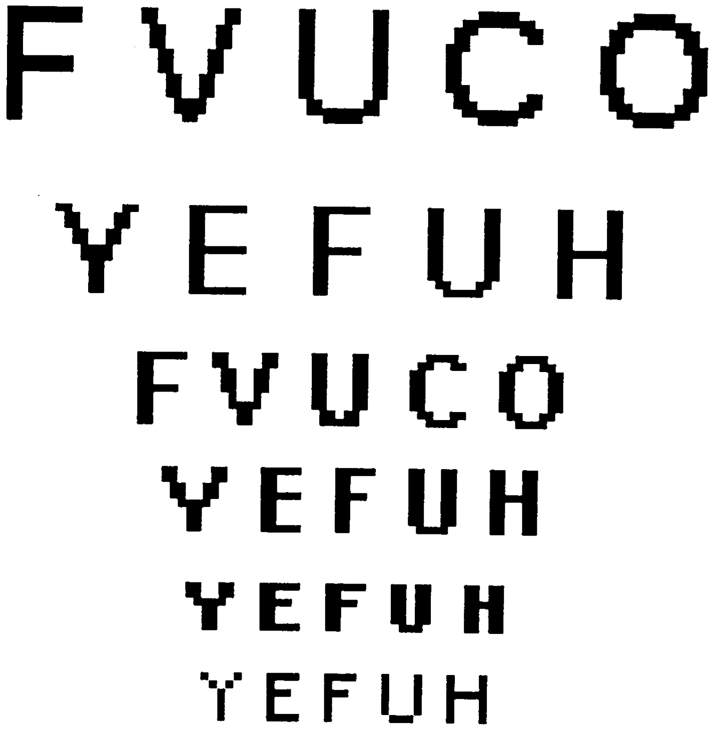

Figure 13: Original `Snellen` image

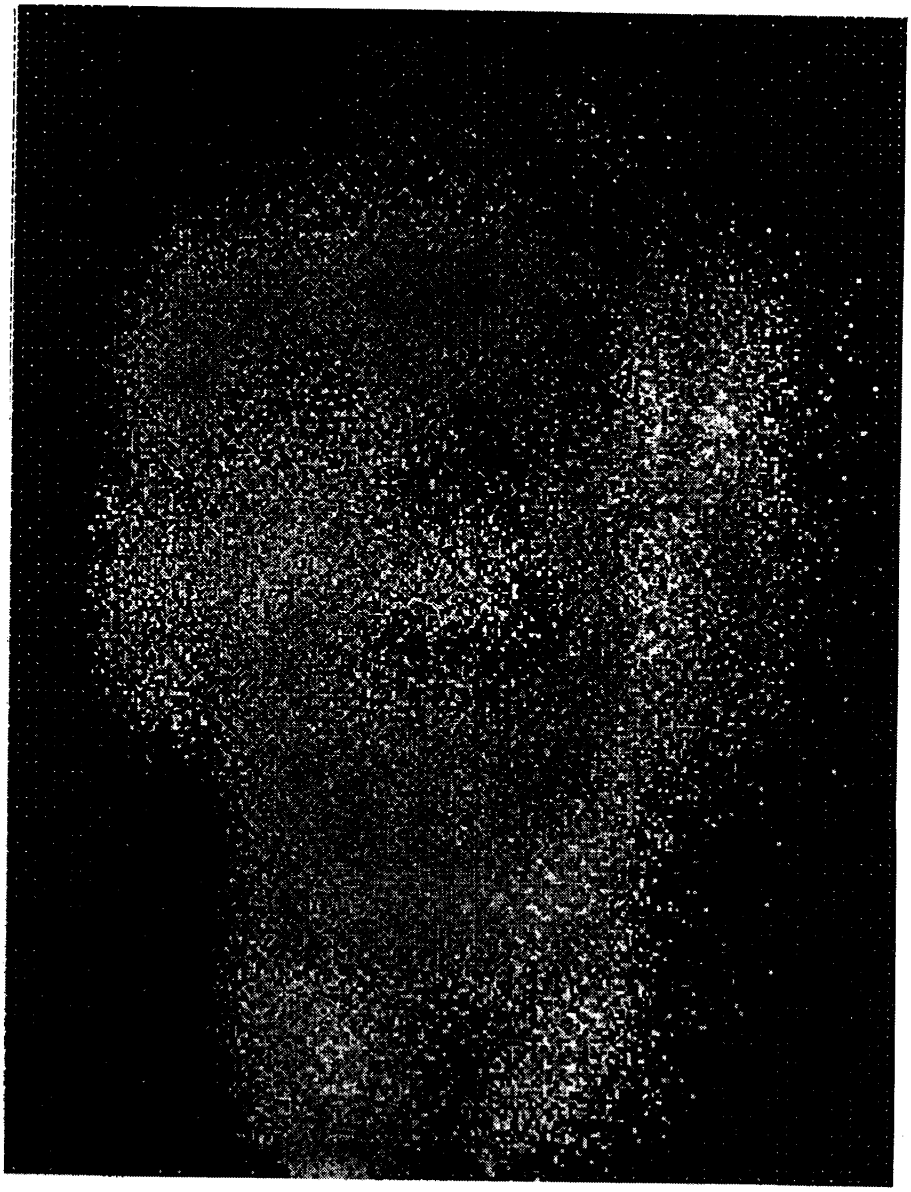

Figure 14: Diffusion transform. A crude simulation of diffusion. Pixels are swapped with pixels some random distance and direction away; small mean-path length preserves large-scale structure and disrupts small-scale detail. The tendency of the pixels to drift to the upper-right is an artifact of the bottom-to-top, left-to-right iteration loop of the `diffuse` program.

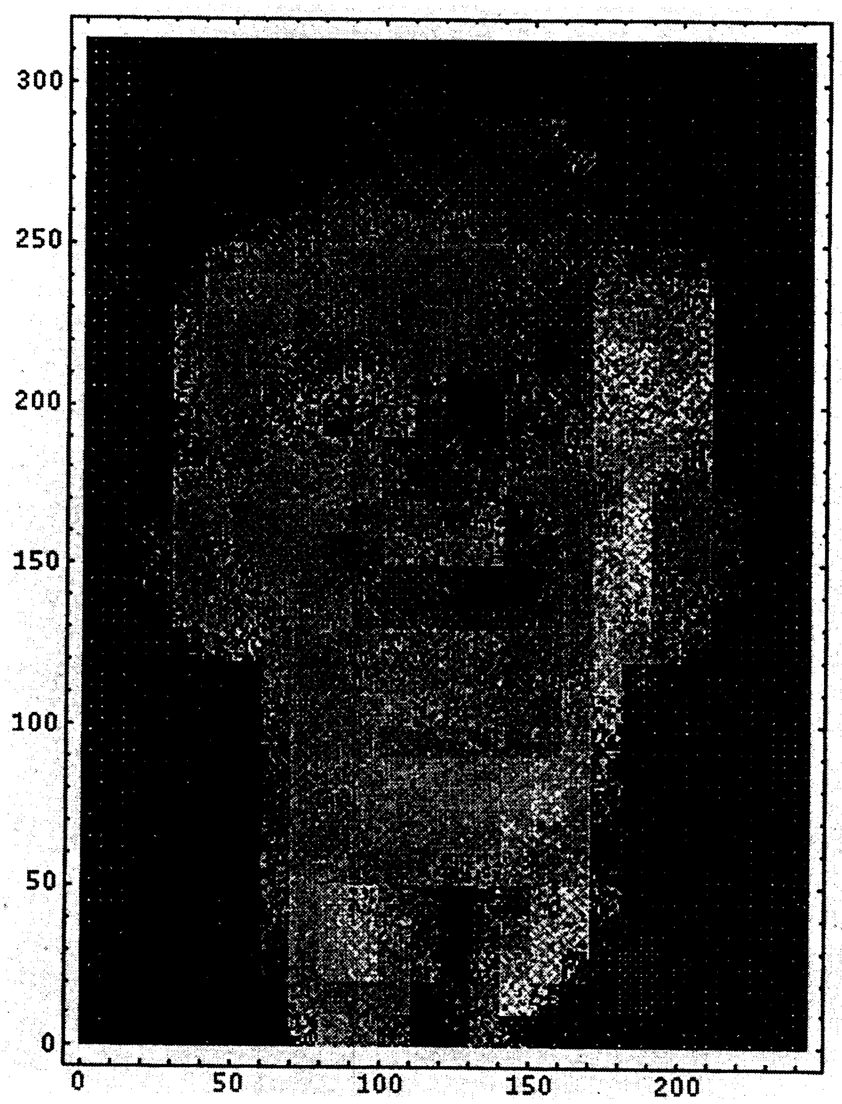

Figure 15: `Scramble` transform

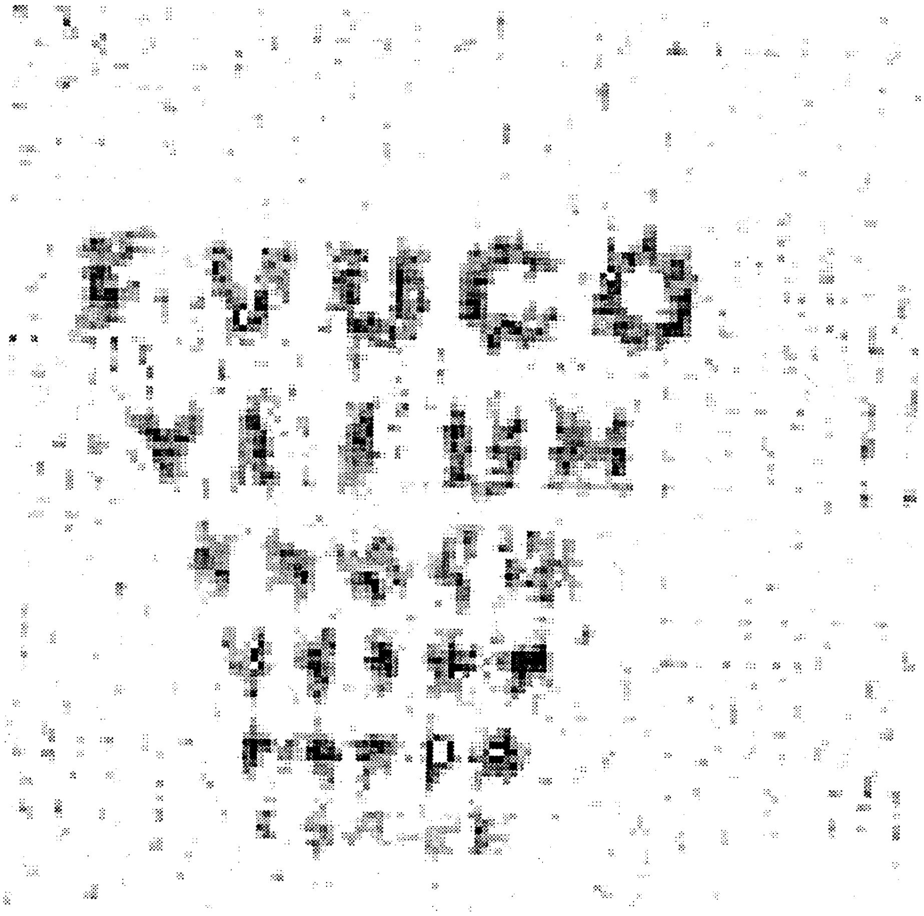

Figure 16: **Phaseblur** transform

beam, and a means of recording and measuring the beam's reflection.

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 17. 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, 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.

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

5 Acknowledgements

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

References

[1] B. H. et al. Performance evaluation of a six-axis generalized force-reflecting teleoperator. Technical Report 89-18, Jet Propulsion Laboratory, June 1990.

[2] L. Harmon. The recognition of faces. *Scientific American*, Nov. 1973.

[3] R. L. Jacobs, I. L. Bailey, and M. A. Bullimore. Artificial pupils and Maxwellian viewing systems. *Ophthalmic and Visual Optics Technical Digest*, 3:106–109, 1992.

[4] C. Neveu and L. Stark. Various variance-invariant transforms.

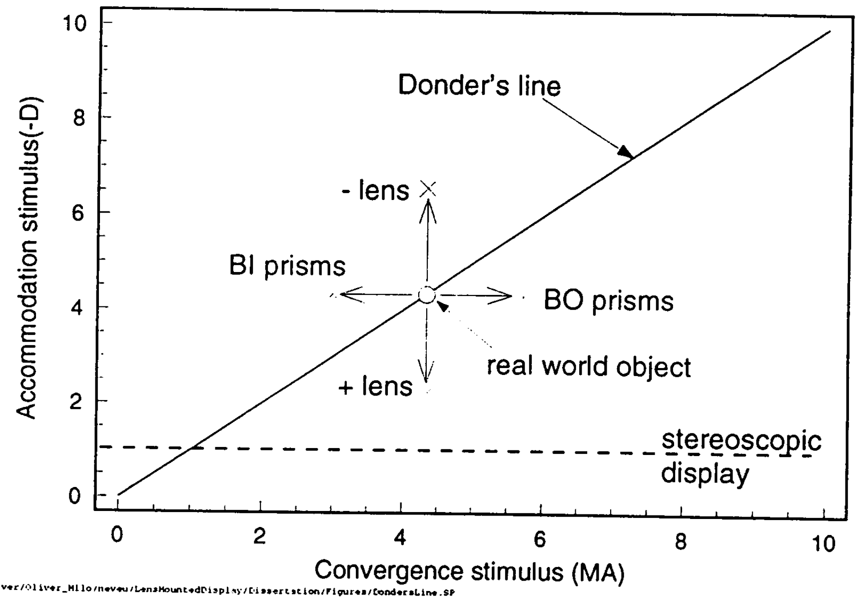

/Net/oliver/Oliver_Milo/neveu/LensMountedDisplay/Dissertation/Figures/DondersLine.SP

Figure 17: Donder's line, real and virtual objects

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

[6] S. R. Phillips. *Ocular Neurological Control Systems: Accommodation and the Near Response Triad.* PhD thesis, University of California, 1974.

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

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

[9] J. W. Rohen and A. Zimmerman. Changes due to age of the human ciliary body. *Albrecht von Graefes Arch Klin Ophthalmol*, 179:302–317, 1970.

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

[11] G. Smith. Angular diameter of defocus blur discs. *American Journal of Optometry and Physiological Optics*, 59:885–889, 1982.

[12] L. Stark. *Neurological Control Systems: Studies in Bioengineering.* Plenum Press, New York, 1968.

[13] I. E. Sutherland. A head-mounted three-dimensional display. In *Proc. of the Fall Joint Computer Conference*, pages 757–764, 1968.

[14] Z. Szakaly and A. K. Bejczy. Performance capabilities of a JPL dual-arm advanced teleoperations system. In *Proc. of the 4th Ann. Workshop on Space Operations, Automation and Robotics*, pages 30–41, Albuquerque, NM, June 1990.

[15] T. Young. On the mechanism of the eye. *Philos. Trans. R. Soc.*, 91:23–88, 1801.