



Characterization and use of a digital light projector for vision research

Orin Packer ^{a,*}, Lisa C. Diller ^a, Jan Verweij ^a, Barry B. Lee ^b, Joel Pokorny ^c,
David R. Williams ^d, Dennis M. Dacey ^a, David H. Brainard ^e

^a Department of Biological Structure, Box 357420, University of Washington, Seattle, WA 98195, USA

^b Department of Neurobiology, Max Plank Institute of Biophysical Chemistry, Göttingen, 37077 Germany

^c Visual Sciences Center, University of Chicago, Chicago, IL 60637, USA

^d Center for Visual Science, University of Rochester, Rochester, NY 14627, USA

^e Department of Psychology, University of California, Santa Barbara, CA 93106, USA

Received 26 May 2000; received in revised form 10 October 2000

Abstract

For creating stimuli in the laboratory, digital light projection (DLP) technology has the potential to overcome the low output luminance, lack of pixel independence, and limited chromaticity gamut of the cathode ray tube (CRT). We built a DLP-based stimulator for projecting patterns on the in vitro primate retina. The DLP produces high light levels and has good contrast. Spatial performance was similar to that of a CRT. Temporal performance was limited by the refresh rate (63 Hz). The chromatic gamut was modestly larger than that of a CRT although the primary spectra varied to a small degree with light output and numerical aperture. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Digital light projection; Stimulus; Spatio-temporal

1. Introduction

Psychophysicists and physiologists image stimuli on the retina and measure the response of the visual system. Common stimuli include spots, annuli, and sine wave gratings that are temporally modulated or drifted across the retina. The ideal stimulator would be capable of producing patterns containing spatial and temporal frequencies up to the resolution limits of the visual system, and would be able to reproduce the full gamut of physically realizable chromaticities. It would also have high light output, high contrast, support precise control of chromaticity and luminance, and be stable over time.

In recent years, new technologies have provided alternatives to conventional optical systems based on

incandescent light. For example, a light emitting diode (LED)-based stimulator has many ideal characteristics. LEDs have high luminance, are available with a wide range of dominant wavelengths, and can be switched very rapidly. Techniques have also been developed to make light output a linear function of input level (Swanson, Ueno, Smith, & Pokorny, 1987; Watanabe, Mori, & Nakamura, 1992). The drawback of an LED-based stimulator is that it is difficult to create spatial patterns.

The standard device for creating spatial patterns is the cathode ray tube (CRT). Unfortunately, a CRT delivers an order of magnitude less light than an LED-based stimulator. For a range of chromaticities, CRT light levels are not high enough to suppress fully rod activity (Shapiro, Pokorny, & Smith, 1996). Furthermore, the chromatic gamut is limited because the light emitted by CRT phosphors is spectrally broadband. Another drawback is that control of monitor output is complicated by interactions between separate phosphor channels and interactions between what is dis-

* Corresponding author. Tel.: +1-206-5430224.

E-mail address: orin@u.washington.edu (O. Packer).

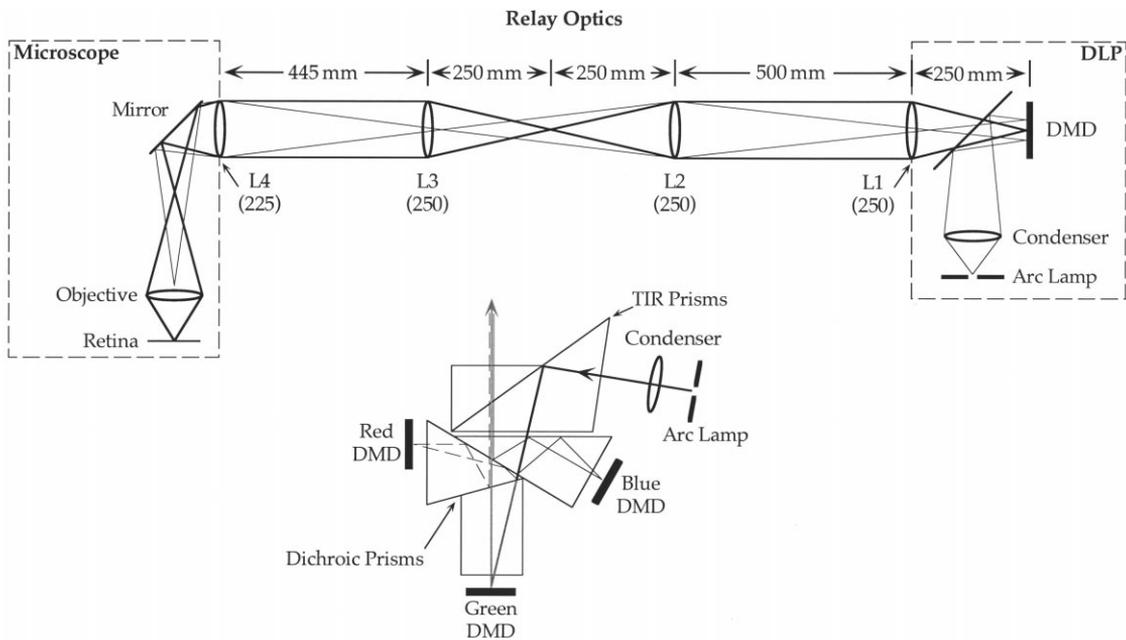


Fig. 1. A schematic diagram of the optics of the visual stimulator. The optics of one of the three chromatic channels of the digital light projector (dashed box on the right) are represented in highly simplified form. The internal optics are detailed in Hornbeck (1997). Light from the xenon arc lamp is relayed by lenses L1–L4 to a camera port in the microscope (dashed box on left). Images of the DMD (intersections of the heavy rays) are formed between lenses L2 and L3, in the back image plane of the objective, and on the retina. The arc is imaged (intersections of the light rays) near lens L3 and again at the entrance pupil of the objective ensuring that it will be out of focus on the retina. The inset (after Fig. 14 in Hornbeck, 1997) shows how the three primaries are created within the digital light projector. Total internal reflection (TIR) prisms steer the beam into three dichroic prisms which split it into long (r), middle (G), and short (B) wavelength bands. The R (thin dashed line), G (thick grey line), and B (thin solid line) beams each reflect off one of the three DMDs which further shapes the spectrum. The beams then reflect from internal surfaces of the dichroic prisms and become coaxial at the output. For clarity, only the axial rays are drawn.

played at nominally distinct screen locations (e.g. Brainard, 1989; Pelli, 1997). A CRT projector overcomes low light output but not the other drawbacks.

Recently, Texas Instruments has developed a digital display technology (Hornbeck, 1997) with the potential to combine the temporal and chromatic properties of an LED-based system with the spatial flexibility of the CRT display, but without the limitations of CRT technology. At the heart of the Digital Light Projector (DLP) is a Digital Micromirror Device (DMD), an array of tiny mirrors manufactured with integrated circuit technology. When coupled to a light source, an image is formed on the reflective surface of the DMD by modulating the amount of light reflected from each mirror. Color images are created by combining channels filtered for long, middle, and short wavelength bands of light. In theory, this technology can generate stimuli with high spatial and temporal resolution across a wide range of chromaticities.

We constructed a DLP-based visual stimulator for projecting spatial patterns on the *in vitro* primate retina during intracellular recording (Dacey & Lee, 1994; Dacey, Lee, Stafford, Pokorny, & Smith, 1996; Dacey et al., 2000a). Here we report its spatial, temporal, and chromatic properties and compare them to those of an ideal visual stimulator.

2. A DLP-based visual stimulator for retinal physiology

Although all visual stimulators image light on the retina and share many basic characteristics there are also many differences in design made necessary by experimental conditions. During intracellular recording from the *in vitro* primate retina, individual neurons are visually targeted for penetration by glass microelectrodes, a task requiring the high quality optics and mechanical stability of a microscope. Therefore, in our visual stimulator, the digital light projector is coupled by relay optics to a microscope.

2.1. Projector and microscope

The digital light projector is a VistaPro Plus manufactured in 1997 by Electrohome (now Christie Digital Systems). It is built around a Texas Instruments projection engine that contains three DMDs and their associated optics and electronics. Fig. 1 (inset, after Figure 14 in Hornbeck, 1997) shows how prisms split the light from a xenon arc with a luminous power of 1300 lm into long, middle, and short wavelength bands, each of which illuminates a separate DMD. Dichroic coatings applied to the mirrored surfaces of the DMDs further shape the spectrum of each band. The three channels

are recombined by the same prisms into a single coaxial output beam.

Each DMD is an 848×600 square array of mirrors, each of which is $16 \mu\text{m}$ on an edge with a $17\text{-}\mu\text{m}$ center-to-center spacing. A mirror can be independently tilted $\pm 10^\circ$ to an ‘on’ position in which the illuminating beam is reflected down the axis of the projection optics, or an ‘off’ position in which the beam is reflected into a light trap. The speed of mirror switching is critical since both intensity control and temporal resolution must fit within the temporal bandwidth of the mirrors.

Consider a stimulus that will be refreshed at the 63 Hz maximum refresh rate of the DLP (15.9 ms refresh period) and whose pixels have 10 bits of intensity resolution. Intensity is controlled by switching each mirror on for some proportion of each refresh period. To support 10 bits of intensity resolution, each mirror must be able to switch on and off in an interval that is $1/(2^{10})$ times the refresh period, in this case $15.5 \mu\text{s}$. This is just within the nominal $15 \mu\text{s}$ mechanical switching period specified for the mirrors. At the beginning of a refresh period, circuitry for each mirror reads a command telling it for what proportion of the refresh period to turn on. The circuitry then converts this proportion to a mirror switching sequence that has the desired on-proportion. The actual switching sequences are generated according to a proprietary algorithm designed to minimize visible flicker artifacts and are quite complex (Hornbeck, 1997).

Although the native intensity resolution of the DLP is 10 bits, only 8 bits of resolution are available to the user at any given time. Intensities are chosen from a lookup table whose entries have 10 bits of resolution. However, each lookup table is limited to 256 entries. Because Electrohome did not envision that scientists might want to rewrite these tables, they did not provide a rapid way to do so. Thus rewriting the tables takes 15–30 s and it is not generally feasible to modify lookup tables on the fly. The lookup table may, however, be adjusted between experiments to vary the palette of light levels available. For many applications, this is not a serious drawback.

The contrast ratio, the ratio of minimum and maximum light outputs is 180 when measured with an ANSI checkerboard pattern.

The microscope in our visual stimulator was an Olympus BH2 biological research microscope. The DLP image was introduced into the microscope optics through a camera port incorporated into the trinocular viewing head. Here, we consider only those aspects of the microscope design that are relevant to the visual stimulator.

2.2. Optical design

2.2.1. Imaging the DMD on the retina

The primary purpose of the optics is to image the DMD on the retina at the proper magnification as shown by a schematic of the optical design (Fig. 1). The DLP optics are within the Texas Instruments projection engine and are not accessible to the user. The right side of Fig. 1 is a highly simplified representation of the optics of one of the three chromatic channels showing that the optical distance between the DMD chip and the projection lens opening is 250 mm and that an image of the arc lamp is formed in the plane of the opening. The inset depicts how the three channels are formed and recombined. For clarity only the axial rays are shown. Although it is not immediately obvious, the optical distances from each DMD to the output plane are equal as are the optical distances from the arc lamp to the DMDs.

Intermediate images of the DMD are represented in Fig. 1 by the intersections of the heavier rays. Lenses L1 and L2 are a unity magnification relay, forming an image of the DMD midway between lenses L2 and L3. Their sole purpose is to increase optical path length. Note that the DMD and arc images formed between lenses L2 and L3 have the same geometry with respect to lens L3 that they have with respect to lens L1. Lenses L3 and L4 form a second intermediate image in the image plane of the objective. The focal lengths of lenses L3 and L4 were chosen so that the DMD image was minified by a factor of 0.9. Finally, the microscope objective forms an image of the stimulus on the retina. Because the illuminating beam is passing through the objective in the reverse direction, objective magnification is the inverse of the value printed on the barrel. Thus, the $4\times$ objective minifies the DMD image by a factor of 0.25 and the total minification of the microscope optics is 0.2 due to an additional lens in the vertical illuminator that minifies the image by a factor of 0.8. The final stimulus minification of 0.18 is the product of microscope minification and relay minification. The image of a single DMD mirror subtends $2.9 \mu\text{m}$ on the retina. This, in principle, allows the stimulation of single foveal cones. The image of the entire DMD, the product of the number of pixels in the DMD chip and the size of the image of an individual pixel, subtends 2.5 by 1.8 mm of retina (about 12 by 9° at the fovea). Because our measurements refer to visual stimuli used during physiology experiments, we will refer to stimulus dimensions by their angular extent on the monkey retina. The reader can convert to linear extent in micrometers using the conversion factor $200 \mu\text{m}/\text{degree}$.

2.2.2. Uniformity of illumination

A second goal of the optical design was to illuminate the retina uniformly by preventing spatial structure in the arc from appearing in the stimuli. The structure of the arc is partly scrambled by an integrating rod within the condenser assembly (not shown). In addition, the relay optics were designed to enhance uniformity. The condenser optics of the DLP form the first image of the arc in the plane of the projector's lens opening as shown by the intersections of the fine rays in Fig. 1. Lens L2 then relays the image ahead to just behind lens L3. Because lens L1 is so close to the arc image, it plays almost no role in arc imaging. The same is true for lens L3. The distance between the intermediate arc image and lens L4 was chosen so that the arc would be imaged in the entrance pupil of the objective and be out of focus on the retina.

2.3. Stimulus generation

A VSG3 stimulus generator from Cambridge Research Systems (UK) produced the video signal input to the projector. The VSG3 was designed for vision research and supports a wide variety of spatial, temporal, and chromatic stimuli. A custom computer program written in C for a Macintosh computer coordinated stimulus generation. The program read a file describing a series of stimuli and assembled the appropriate VSG3 commands. These high-level text commands were sent over a serial interface to the VSG3, which generated the video signal that drove the DLP.

When creating a visual stimulus, it is necessary to have precise control over the spatial and temporal properties that will produce the desired effect. Specifying the spatial and temporal properties of a stimulus created with a DLP is complicated by its ability to display video inputs with a wide range of frame rates and scanning frequencies. It does this by filtering the video signal in space and time and then resampling it to match the spatial resolution to the number of DMD mirrors and the temporal resolution to the DMD refresh rate. The filtering and resampling operations can distort the stimulus produced by the DLP. To minimize such distortion, we configured the VSG to produce a video signal matched to the native resolution of the DMDs. It is not clear from the documentation, however, whether this completely eliminates filtering and resampling of the video signal.

To produce a desired stimulus, the input–output relation between the digital level input to the VSG and the amount of light actually emitted by the DLP must be measured. The VSG3 stimulus generator is equipped with routines for automating these measurements. For our system, the input–output relation depends on the values programmed into the DLP's internal lookup tables. These are factory-programmed to produce a

nonlinear input–output relation that mimics the gamma function of a typical CRT. For our purposes, it was preferable to rewrite these lookup tables to produce a nominally linear input–output relation. We did this before we ran the VSG3 calibration routines.

3. Spatial, temporal, and chromatic properties

3.1. Spatial properties

The maximum spatial frequency that can be produced by our stimulator depends on sampling properties, optical quality, optical magnification, and the quality of electronic signal processing.

3.1.1. Sampling

Sampling theory (Shannon, 1949) says a stimulus can be reconstructed from a discrete set of samples only if every spatial frequency is sampled at least twice per cycle. Thus, the maximum spatial frequency that can be displayed by the DLP without spatial aliasing corresponds to a grating whose light and dark bars are formed by alternate rows of pixels. In the case of the $4 \times$ objective, a single cycle of this grating would be two pixels or $6 \mu\text{m}$ wide and have a spatial frequency of $0.167 \text{ cycles}/\mu\text{m}$. At the fovea of the macaque retina, this corresponds to $\sim 33 \text{ cycles/degree}$. It is important to note, however, that even below the Nyquist frequency sampling causes artifacts in the stimulus. A rather extreme example is shown in Fig. 5. Sampling artifacts become smaller relative to the waveform at lower spatial frequencies.

3.1.2. Spatial modulation transfer

In order to make use of the full range of spatial frequencies supported by the resolution of the device, the electronics and optics must transmit spatial frequencies below the sampling limit without too much attenuation. If the electronics and optics were perfect, spatial performance would be limited by diffraction at the smallest aperture in the stimulator. For a circular exit pupil, the diffraction-limited spatial modulation transfer function (sMTF) is

$$D(\underline{s}, \underline{s}_0) = 2/\pi [\text{Cos}^{-1}(\underline{s}/\underline{s}_0) - (\underline{s}/\underline{s}_0)\sqrt{(1 - (\underline{s}/\underline{s}_0)^2)}] \quad (1)$$

$f_N \underline{s} < \underline{s}_0$

where s is spatial frequency in cycles/degree and s_0 is the incoherent cut-off frequency (Goodman, 1996). In cycles/degree,

$$s_0 = \pi/180 \times a/\lambda \quad (2)$$

where a is exit pupil diameter and λ is wavelength. The exit pupil is the image of the limiting aperture viewed from the position of the detector used to make the measurements. When using the $4 \times$ objective, the limit-

ing aperture is the objective itself ($NA = 0.1$), which has a focal length of 30 mm and an exit pupil diameter of 6 mm.

The diffraction-limited sMTF of the visual stimulator is plotted in the upper curve of Fig. 2a. The optical cut-off is 189 cycles/degree and contrast at the 33 cycles/degree sampling limit is 78%. Thus, diffraction attenuates retinal image contrast only slightly at the spatial frequencies used to stimulate peripheral retinal neurons.

The sMTF of the visual stimulator was measured by drifting 100% contrast black and white sine wave gratings across a photodetector masked by a 6- μm aperture. The photodetector aperture was placed where the retina would normally lie. The open circles in Fig. 2a show modulation transfer plotted as a function of spatial frequency. This curve has been corrected for filtering by

the 6- μm aperture. The sMTF rolled off to a cut-off at 15 cycles/degree. In short, measured resolution was a log unit less than best possible resolution. This was surprising since the stimulator was designed with high quality optics. We confirmed the quality of the relay optics by imaging a 6- μm diameter point source on a CCD through the relay optics. The MTF, calculated from the Fourier transform of the point spread function and corrected for the finite size of the 6- μm aperture, was very close to the predicted diffraction-limited MTF of the microscope objective that was the limiting aperture in the relay optics. This rules out the relay optics themselves as the source of the spatial filtering.

Another potential limiting factor that we ruled out was the geometry of DMD mirror switching. The mirrors switch from on to off with a swing of $\pm 10^\circ$. In order for a mirror to turn completely off, the illuminating beam must swing from the optical axis past the entrance pupil of the projection optics. This requires limiting the angular extent of the light cone illuminating the DMD to less than the angular extent of mirror swing and is equivalent to limiting numerical aperture (NA) since, in air,

$$NA = \sin(\alpha) \quad (3)$$

where α is the angle between the optical axis and the edge of the cone (Smith & Atchison, 1997). For a cone with an angle α of $\pm 10^\circ$, NA is 0.17. This is higher than the NA of the visual stimulator and therefore does not limit resolution.

It thus appears that the contrast loss originates not within the optics, but in the signal that drives the DMD. We do not know whether this loss occurs during video signal generation in the VSG3 or during analog to digital conversion within the DLP.

Although the DLP cannot produce 100% contrast stimuli at high spatial frequencies, it is perfectly feasible to conduct experiments by using stimuli of lower contrast. The maximum usable spatial frequency is reached when the low sensitivity of the neuron and the low modulation transfer of the stimulator yield responses with an inadequate signal to noise ratio. In our hands, with the $4\times$ objective, this occurs at about 10 cycles/degree when recording from an H1 horizontal cell. If the contrast response of the cell being measured is known to be linear, contrast attenuation by the stimulator can be compensated by dividing the neuronal MTF by the stimulator MTF.

To generate higher spatial frequencies we reduce the stimulus magnification by use of a microscope objective with higher magnification. This increases the sampling limit and the maximum spatial frequency that can be imaged on the retina. For example, a $10\times$ objective reduces image magnification by a factor of 2.5 relative to the $4\times$ objective and increases the sampling limit to

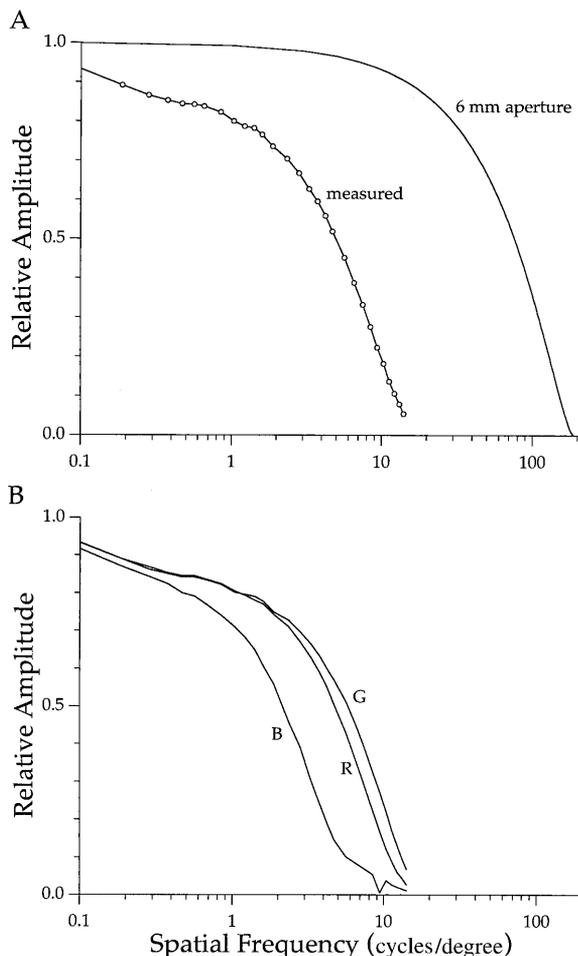


Fig. 2. The spatial modulation transfer function (sMTF) of the visual stimulator. (A) The diffraction-limited sMTF of an optical system with the same numerical aperture as the visual stimulator (solid line). It represents the best possible spatial performance. The lower curve (open circles) is the measured sMTF of the visual stimulator. (B) The sMTFs of the red (R), green (G), and blue (B) primaries of the stimulator measured after bringing a black and white sinusoidal grating to best focus.

83 cycles/degree. Measurements confirm that it also shifts the sMTF up the frequency axis by a factor of 2.5. Thus, it is possible to create high contrast luminance stimuli of high spatial frequency by selecting a microscope objective with the appropriate magnification. Note that the increase in spatial frequency range is purchased at a cost in available field size, since the reduction in magnification also reduces the size of the final projected image.

3.1.3. Axial chromatic aberration and the sMTF

Since each of the three DMDs is illuminated by different wavelengths, axial chromatic aberration could cause each DMD to be imaged in a different plane, differentially affecting the R, G, and B sMTFs. We tested for this by focussing a drifting black and white grating on a photodetector. The sMTFs of each primary were then measured in isolation and plotted in Fig. 2b. At spatial frequencies above 2 cycles/degree, the G primary had higher modulation transfer than the R primary while the B primary had the lowest modulation transfer at all spatial frequencies. At spatial frequencies above 1 cycle/degree these differences become significant. At 3 cycles/degree, the B primary has only half the contrast of the G primary. The sMTFs of the R and B primaries could be improved by refocusing until they were similar to the sMTF of the G primary at best focus, implicating axial chromatic aberration. Near the center of the DMD, the axial distances between the best focus of the R primary relative to the G and B primaries were 21 and $-17 \mu\text{m}$, respectively.

Differences in the sMTFs of the three primaries can significantly affect stimuli that require precise control of primary contrasts. Full field chromatic stimuli are immune to this problem, but a cone isolating grating of even relatively low spatial frequency created using the method of silent substitution (Rushton, Powell, & White, 1973; Estevez & Spekreijse, 1974) could have primary contrasts that are a factor of two different than their nominal values. Note that this effect will also occur with stimulators based on CRT technology.

3.1.4. Convergence

To provide high quality color images, all three DLP primaries must perfectly converge in the image plane. We measured convergence by measuring where images created by the three primaries fell. We imaged a $30\text{-}\mu\text{m}$ circular spot of light on a $30\text{-}\mu\text{m}$ aperture and measured the amount of transmitted light with a photodetector. The spot was presented sequentially at 25 equally spaced positions forming a square centered on the aperture. The horizontal and vertical distances between adjacent positions was $9 \mu\text{m}$. In the case of perfect convergence, the position at which the most light was transmitted would be the same regardless of which of

the three DLP primaries was used to create the spot. However, convergence was not perfect. In the center of the DMD, shifts as large as $2.6 \mu\text{m}$ were recorded. At the four corners the errors were larger, reaching a maximum of $10.7 \mu\text{m}$. Shifts were largest when the R primary was involved.

The relay optics of the stimulator were carefully aligned before testing to minimize their effects on mis-convergence. We also confirmed that software settings that control alignment of the three channels of the DLP were at their optimum values. The magnification difference due to chromatic aberration could be further minimized by using achromatizing optics that bring the primary wavelengths to a focus at the same location. Decreasing the magnification of the image on the retina also minimizes the impact of all of these effects.

Convergence error should be considered when producing stimuli with fine spatial detail. For example, the image of a single central pixel produced by the R channel would be offset from the image produced by the B channel by nearly the width of a pixel. For a pixel near the edge of the DMD, the offset would be ~ 4 pixel widths. In the center of the DMD, an R plus B grating whose bars were a single pixel wide and nominally in phase with each other would in fact be a grating whose R and B bars were 180° out of phase. Convergence error in the center of the DMD is similar to that for a high quality CRT monitor, while the error in the corners was somewhat worse. For example, Sony specifies that its CPD-GS series monitors have a 0.25-mm dot pitch and a convergence error of up to 0.3 mm in the center of the tube. In the corners, dot pitch increased to 0.27 mm while the convergence error remained at 0.3 mm .

3.1.5. Depth of focus

A final optical factor affecting the spatial performance of the stimulator is depth of focus, the amount of defocus that can be introduced before there is a substantial change in the sMTF. Fig. 3 shows luminance grating sMTFs for a series of focus settings ranging from $+100$ to $-100 \mu\text{m}$ in $50\text{-}\mu\text{m}$ increments made using the $4\times$ objective and with the optics at their maximum numerical aperture of 0.1. Defocus above (open symbols) and below (filled symbols) best focus (no symbols) produces similar contrast losses. Up to a spatial frequency of 1 cycle/degree, there is little change in modulation transfer even when the focus is off by $100 \mu\text{m}$. The reduction in amplitude reaches a factor of two at a spatial frequency of 10 cycles/degree when focus is off by $50 \mu\text{m}$ and 5 cycles/degree when the defocus is $100 \mu\text{m}$. Thus, although there is a progressive loss of contrast when focus is not optimum, we can focus accurately enough to preserve image contrast.

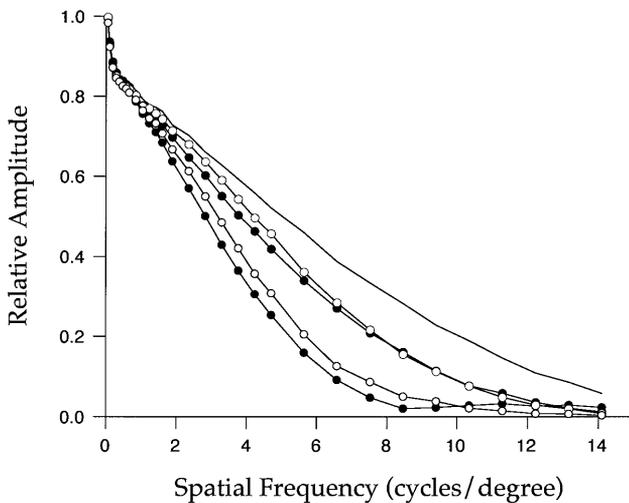


Fig. 3. The effect of defocus on spatial modulation transfer. The upper curve is the sMTF of the stimulator at best focus. The middle pair of curves are sMTF 50 μm above and below best focus. The lower pair of curves are sMTF 100 μm above and below best focus. In each pair, the open symbols represent defocus above best focus while the closed symbols represent defocus below best focus.

3.1.6. Pixel independence

In addition to contrast losses in the optics, there can also be losses caused by signal processing. For example, the contrast of a high spatial frequency grating displayed on a CRT is orientation-dependent (Pelli, 1997). When a grating whose bars are the width of a scan line is parallel with the scan lines, the electron beam is modulated once for each scan line. When the grating is perpendicular to the scan lines, the beam is modulated once for each pixel. If the temporal bandwidth of the electronics is too low, the contrast of the perpendicular grating will be lower than that of the parallel grating.

In theory the DLP need not be affected by this problem, but our measurements showed otherwise. Fig. 4a shows the total light output of the visual stimulator while displaying horizontal and vertical square wave gratings whose bars were 1 pixel wide. All of the bright pixels in each bar were set to some intensity and all of the dark pixels were turned off. The light output for a full field in which every pixel was turned on was also measured. This test was repeated over a range of bright pixel intensities from zero up to the maximum. For grating contrast at the two orientations to be identical, it is a necessary requirement that total light output be the same, since the number of bright pixels in both cases is half the total pixels. It is also required that the summed output of both gratings equal the light output of a full field in which every pixel is turned on. In fact, light output was less for one orientation than for the other. Furthermore, this effect was spatial frequency dependent. Fig. 4b shows that for a grating whose bars were 10 pixels wide, the effect was smaller. This result is consistent with a lack of temporal bandwidth in the

signal processing electronics since the rate of temporal modulation decreases when the bars are made wider. Thus, the VSG/DLP system exhibits a nonlinearity similar to a typical computer-controlled CRT monitor. To avoid this nonlinearity, we typically use gratings whose bars are oriented parallel to the rows of mirrors that are driven by the scan lines of the video signal. In particular, the sMTFs reported above were measured with such gratings. A bar 10 pixels wide displayed on the fovea through the $4\times$ objective corresponds to a spatial frequency of ~ 3 cycles/degree. For some applications, such as when the orientation tuning of a neuron is being measured, the effect of orientation on grating contrast above this spatial frequency can be considerable.

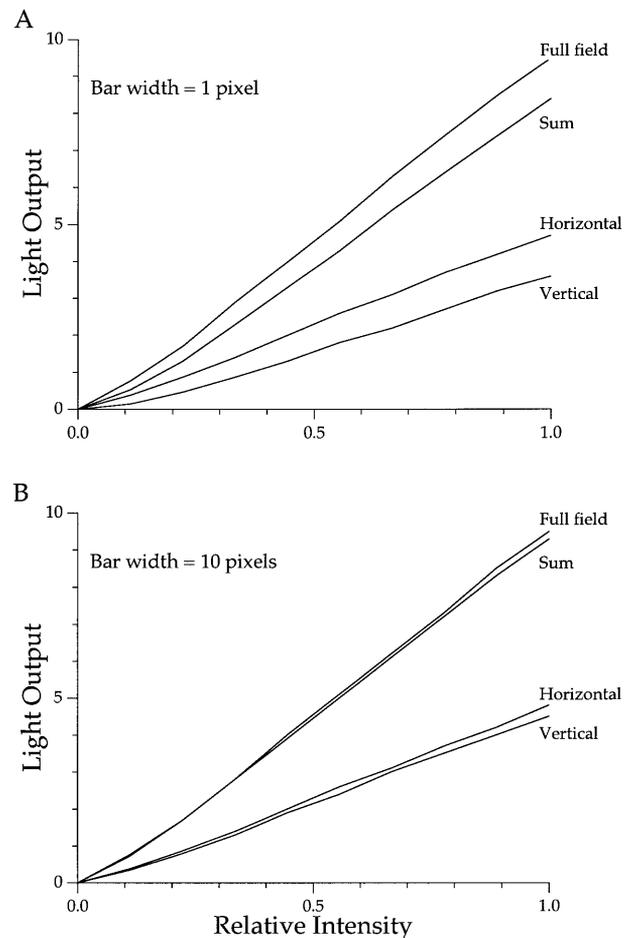


Fig. 4. Total DLP light output in arbitrary units measured by a photodetector for stimuli whose pixels were either 'off' or 'on'. Each curve represents a series of measurements in which the relative intensity of the on pixels was increased to maximum. (A) Full Field stimulus — all of the pixels were turned on. Horizontal stimulus — alternate rows of pixels were turned on. Vertical stimulus — identical to Horizontal except that the rows were rotated 90° on the DMD chip. Sum — the sum of Horizontal + Vertical. (B) Identical to (A) except that the width of the bars was 10 pixels rather than 1 pixel.

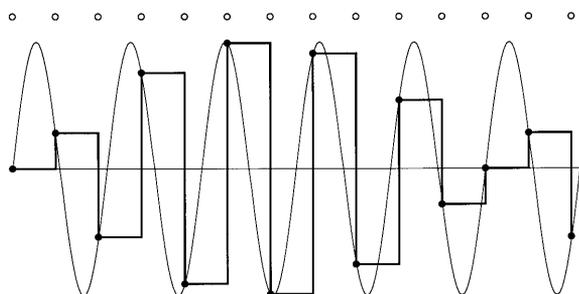


Fig. 5. A schematic diagram of a temporal beat. The x -axis is time and the y -axis is light output. The stepped waveform shows light output after sampling from an ideal sinusoidal waveform at the times indicated by the black dots. The beat is the waxing and waning contrast of this stepped waveform. The open circles show the evenly spaced sample times. The sinusoid is sampled 2.17 times per cycle and the beat frequency is 6 Hz.

3.2. Temporal properties

The maximum temporal frequency that can be produced by the visual stimulator depends on both the temporal sampling rate and temporal bandwidth. There have also been unpublished reports that the mirror switching strategy used by the DLP can produce temporal artifacts when the eye moves relative to the stimuli. We have not addressed this issue because in our experiments, the retina is always stationary relative to the DMDs.

3.2.1. Temporal sampling

The highest temporal frequency at which a stimulus can be extracted from the video signal, converted into digital form, and written into the memory associated with each DMD mirror is 63 Hz. Thus, to avoid aliasing, temporal frequency must be held to a maximum of 31.5 Hz.

Beating is a temporal phenomenon closely related to sampling that can produce an unexpected modulation of the stimulus. Beating is possible because the stimulus generated by the DLP is a series of discrete light output levels sampled from a continuous waveform. Beating happens when the samples are taken at different locations along the waveform from cycle to cycle. When viewed on an oscilloscope, the sampling locations appear to drift through the stimulus waveform. The result is a periodic change in stimulus contrast. Fig. 5 shows a sinusoidal stimulus whose amplitude is being updated at the times indicated by the evenly spaced open circles at the top. The stepped waveform represents the output levels that will be displayed by the DLP. In this example, output level is updated 2.17 times per cycle. Note that the early cycles of the stimulus are of low contrast. Later cycles increase in contrast and then start to decrease again. Stimulus contrast waxes and wanes with a temporal frequency, f , where

$$f = (s - n) \times tf \quad (4)$$

and s is refresh rate in samples/cycle, n is the nearest integral refresh rate and tf is the temporal frequency of the stimulus. In this example, the nearest integral sampling frequency is 2 (the Nyquist frequency) and the beat frequency is 1/6 of temporal frequency. Equation 4 shows that there is no beat if the temporal frequency evenly divides the refresh rate. For example, at the 63 Hz refresh rate, we use to minimize perceptible flicker, the higher such frequencies are 10.5, 12.6, 15.75, 21, and 31.5 Hz.

Perhaps more important than the simple presence of a beat is its amplitude. Amplitude is large whenever the refresh rate is low relative to the temporal frequency as in the above example. At high refresh rates and low temporal frequencies, the amplitude of the beat decreases. Thus, whenever high temporal frequencies are required, the simplest way to avoid beating is to use temporal frequencies that evenly divide the refresh rate.

3.2.2. Temporal modulation transfer

To achieve maximum stimulus contrast, the electronics that control stimulus intensity need to have sufficient temporal bandwidth. If the intensity change of a pixel that is alternating from maximum to minimum intensity is too slow, the pixel will neither decay to its minimum intensity before it is turned on again nor reach maximum intensity before it is turned off again. The resulting contrast loss gets worse with increasing temporal frequency.

The temporal response of the visual stimulator was measured by drifting a 2 cycle/degree sine wave grating across a photodetector masked with a 6- μm aperture placed in the plane of the retina. In Fig. 6, temporal modulation was plotted against drift rate using open symbols. As drift rate increased, the modulation trans-

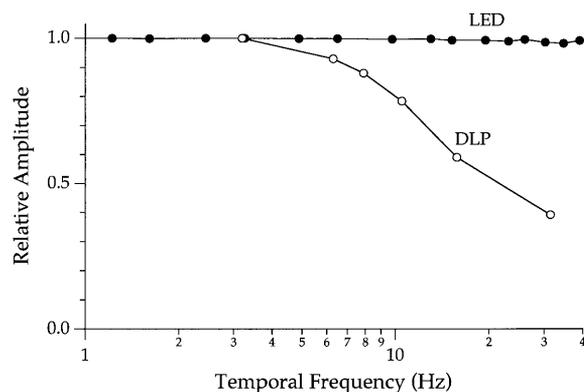


Fig. 6. The temporal modulation transfer function (tMTF) of two visual stimulators. Relative amplitude was plotted as a function of drift rate in Hertz for a 2 cycles/degree sinusoid drifted across a photodetector masked with a 6- μm aperture. The tMTF of the LED-based stimulator (filled circles) is flat out to 100 Hz. The tMTF of the DLP-based stimulator (open circles) rolls off to $\sim 40\%$ of maximum at the 31.5 Hz sampling limit.

fer of the visual stimulator decreased, falling to 40% of maximum at 31.5 Hz. As a comparison, the temporal modulation transfer of an LED-based stimulator in our lab is plotted with filled symbols. If a neuron has a linear response to temporal modulation and if the signal is sufficiently above the noise, the temporal properties of the stimulator can be removed from the temporal responses of a neuron by dividing the measured tMTF by the tMTF of the stimulator.

3.2.3. Frame delay

Finally, to accurately describe when a neuron responds to a stimulus we need to know the temporal relationship between the two. The VSG3 provides a synchronization pulse coinciding with stimulus onset that can be used to align response and stimulus. However, the video interface of the DLP takes time to transform the video signal into a form suitable for driving the DMDs, creating a delay between the generation of the pulse and the appearance of the stimulus on the DMD. We calculated this delay by plotting the temporal phase of a drifting grating as a function of drift rate. It had the signature of a pure delay, namely a straight line of negative slope passing through the origin. The delay calculated from the slope was 41.2 ms and was used to correct the temporal relationship between the stimulus and the response.

3.3. Chromatic properties

3.3.1. Primary spectra

Colors are created by mixing various proportions of the three DLP primaries. We measured the spectra of the primaries with a spectroradiometer and plotted their relative spectral energy distributions as a function of wavelength in Fig. 7a. The spectra overlap and have dominant wavelengths of 611, 549, and 470 nm. We refer to the three primaries as the R, G, and B primaries, respectively.

The chromaticity coordinates of these spectra are plotted on the CIE chromaticity diagram (Wyszecki & Stiles, 1982) of Fig. 7c. Equal energy white plots at (0.33, 0.33) and the spectrum locus (horseshoe-shaped contour) circumscribes the gamut of physically realizable chromaticities. Monochromatic lights at 380, 520, and 770 nm are labeled on the spectrum locus. These typically appear violet, green, and red, respectively. The chromaticity coordinates of the three DLP primaries are plotted as open circles. The triangular area that would be enclosed by drawing lines between the open circles represents the gamut of chromaticities that the DLP can create. Although the DLP can create chromaticities close to the long-wavelength limb of the spectrum locus, there is a large region of chromaticities that are outside the gamut of the device.

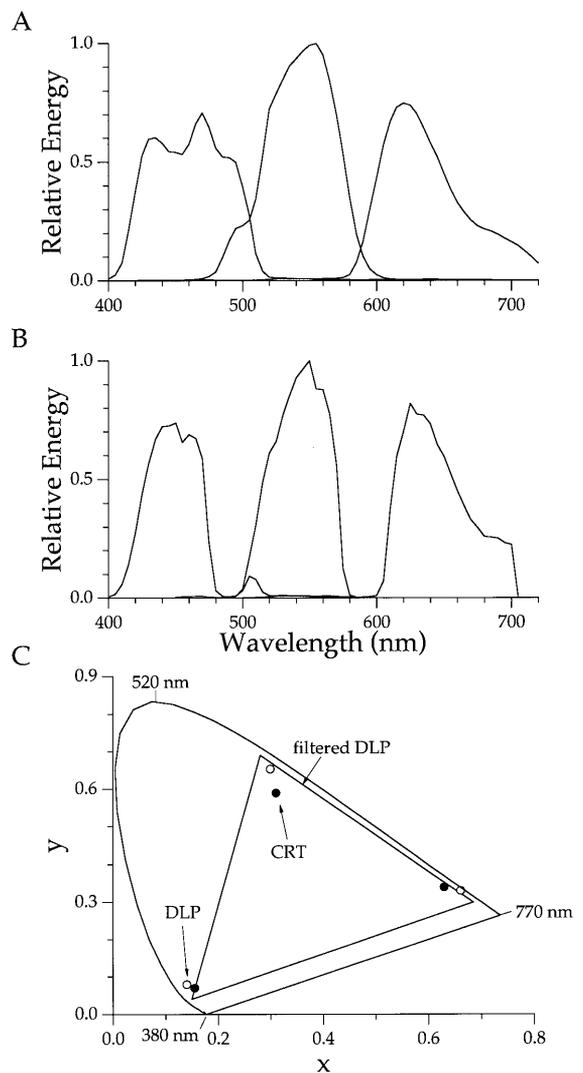


Fig. 7. (A) Relative native energy spectra of the short (R), middle (G), and long (B) wavelength primaries of the digital light projector (DLP). Dominant wavelengths are 611, 549, and 470 nm, respectively. (B) The DLP spectra after filtering by Intor 7 cavity notch filters. Dominant wavelengths shifted to 625, 547, and 460 nm, respectively. (C) A CIE chromaticity diagram showing the chromaticity coordinates of typical CRT phosphors (filled circles), the native spectra of the DLP (open circles), and the filtered spectra of the DLP (corners of triangle). The triangle connecting the coordinates of the filtered spectra represents the chromatic gamut that can be created by the DLP-based visual stimulator.

In order to maximize the gamut of chromaticities that could be created, we shifted the dominant wavelengths of the B and G primaries to shorter wavelengths and the dominant wavelength of the R primary to a longer wavelength. Because filtering in the DLP is done by coatings applied directly to optical surfaces, the spectra can be modified only by adding filters rather than by replacing existing filters. However, since the arc lamp is very intense, it is possible to shift the dominant wavelengths of the primaries modestly by shaving off parts of the existing spectra. On the other hand, it is

not feasible to shift the dominant wavelength of the G primary from 549 to 510 nm because the G primary has little energy at 510 nm.

We added custom-designed seven cavity spectral notch filters (Intor, # 592 Ref, # 494 Ref) to the collimated section of the beam between L1 and L2 (Fig. 1).

The refiltered primary spectra are shown in Fig. 7b. The overlap between the spectra has been eliminated and the dominant wavelengths of the reshaped R, G, and B spectra have shifted to 625, 547, and 460 nm, respectively. The chromaticity coordinates of the filtered primaries in Fig. 7c are the corners of the triangle shown explicitly in the figure. Because the filtered primaries are spaced slightly farther apart and are slightly closer to the spectrum locus, they enclose a modestly larger area of color space than either the unmodified DLP primaries or the phosphors of a typical CRT monitor (filled circles).

3.3.2. Chromatic uniformity

In order to accurately display colors, each DMD must be uniformly illuminated. Otherwise the relative light outputs of the three primaries vary from place to place on the DMD. Since the human eye is very sensitive to chromatic differences, we first inspected a nominally white field. Bluish and reddish splotches were apparent whenever the optics were stopped down to a numerical aperture of less than 0.1. We also measured the relative light output of each primary with a radiometer and calculated the number of photons that would be caught by each cone type from each primary. Measurements were made with a 400- μ m aperture centered on the DMD and at adjacent fields above, below, left, and right of the center field. Cone catch was calculated by taking the product of the quantal primary spectrum and the photopigment spectral sensitivity on a wavelength by wavelength basis and then summing over wavelength. When the numerical aperture was at a minimum, the differences in the cone catches for the three primaries were often larger than 50%. However, at the full numerical aperture of 0.1, the differences did not vary by more than 5% from location to location. Chromatic uniformity while using the DLP as a projector with a high numerical aperture projection lens was excellent, suggesting that nonuniformity would not be an issue when used in this way.

3.3.3. Spectral stability

Ideally, the shape of the DLP primary spectra would remain constant over the full range of stimulus parameters. In this case, the spectral output only needs to be measured for a single input level (typically the maximum) and the computation of device input values required to produce a desired chromaticity and luminance is straightforward (e.g. Brainard, 1995). On the other hand, if the shapes of the primary spectra change

with input level, this change must be characterized and a correction applied. The shapes of the DLP spectra depend somewhat on both input level and numerical aperture.

Fig. 8a shows the spectrum of a spot at two output levels. All three primaries were set to the same percentage of maximum and the numerical aperture of the optics was fixed at maximum (0.1). The two curves shown in Fig. 8a are the normalized spectra for the relative levels of 0.75 (thin line) and 1 (thick line). The shapes of the B and R primaries are very similar, but the G primary at maximum light output has 17% more energy at 525 nm relative to 570 nm than the G primary at relative level 0.75. The spectra for other levels fell between these extremes and no consistent shape change as a function of level was evident.

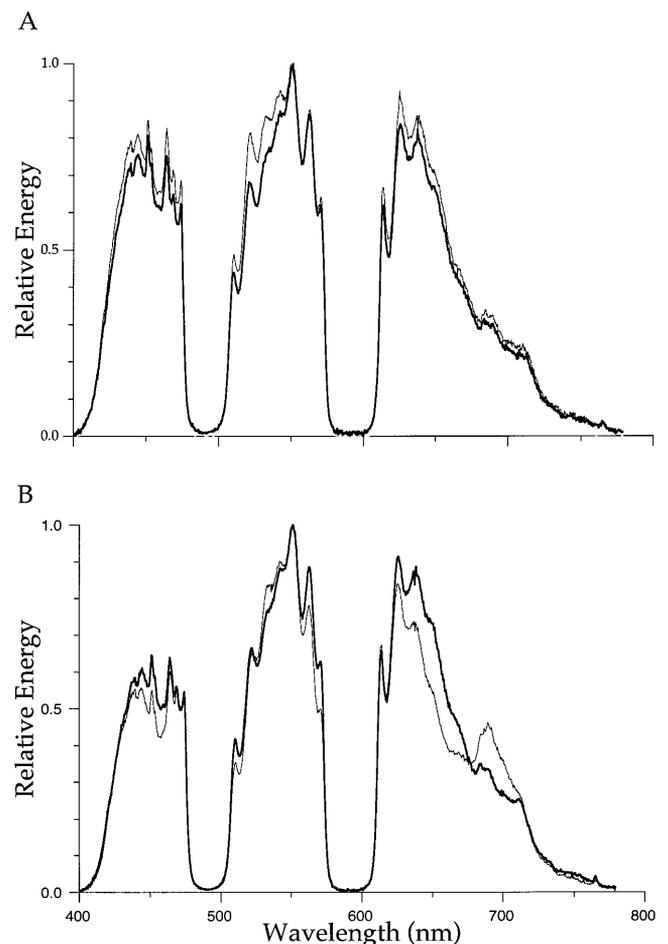


Fig. 8. The effects of output level and numerical aperture on DLP primary spectra. (A) The spectra of a spot at two output levels. All three primaries were turned on to either maximum output (thick line) or to 0.75 of maximum output (thin line). Numerical aperture was set to 0.1 in both cases. (B) The spectra of a spot as a function of numerical aperture. All three primaries were set to maximum output. Numerical aperture was set either to the maximum value of 0.1 (thick line) or to 0.01 (thin line). In both panels, the plotted spectra are normalized to have the same maxima.

Fig. 8b shows how numerical aperture affects the shape of the primary spectra. The thick line is the normalized spectrum measured through the full aperture of the stimulator (0.1) with all three primaries set to their maximum output level. The thin line represents the normalized spectrum measured through a numerical aperture of 0.01. Both the B and G primaries have generally similar shapes although there are local differences. On the other hand, the R spectra are quite different. The spectrum measured at full aperture has more energy between 630 and 675 nm, but less between 675 and 715 nm. The effect of numerical aperture may be handled by calibrating the stimulator using the same numerical aperture as will be used in the experiments.

We are unsure of exactly what causes the changes in spectral shape. However, the spectra of the dichroic filters in the DLP depend on the angle of incidence of the illuminating beam. Since the arc lamp is not a point source, there is a range of angles over which the illuminating light strikes the dichroic surfaces. This is true even in collimated sections of the beam. Reducing numerical aperture decreases the angular extent of the cone of light passing through the optics, eliminating light rays that may have been incident on dichroic mirrors at larger angles.

Similarly, each DMD mirror has a dichroic coating and intensity changes are produced by tilting each mirror relative to the illuminating beam. Averaged over a whole refresh period, a mirror that is tilting rapidly from on to off may reflect light with a slightly different spectrum than a mirror that switches less often.

The degree to which these spectral shifts matter depends on the experiment. For those experiments using the full spectral bandwidth of the stimuli, the shift in dominant wavelength caused by changes in output level or numerical aperture is only about 2 nm and is unlikely to be problematic for experiments that employ luminance modulations.

4. Summary and conclusions

In this paper, we measured the performance of a DLP-based visual stimulator. The promise of this design is that, in principle, it can provide high-resolution spatial and temporal control of a visual stimulus at very high light levels. Indeed, the xenon arc lamp used in the DLP is more intense than any other common light source, especially at short wavelengths, and the reflective design of the DMD mirrors minimizes light loss at the site of modulation. In practice, we found that we could indeed achieve high light levels and reasonable stimulus control. We have successfully used it to explore the spatial properties of the receptive fields of bipolar cells (Dacey et al., 2000a) and H1 horizontal cells (Packer & Dacey, 2000), as well as the spatio-chro-

matic properties of the receptive fields of the blue-ON small bistratified ganglion cell (Dacey, 2000) in macaque monkey retina. We believe performance similar to that of our stimulator can be expected from a DLP-based stimulator designed for psychophysics, although as noted above we have not studied stimulus artifacts that might arise when eye movements are possible. To close this paper, we summarize our main findings and then discuss the future prospects of the technology.

4.1. Spatial performance

The ideal visual stimulator should be able to create spatial frequencies at least as high as the 60 cycles/degree limit of foveal vision. By choosing the appropriate optical magnification, the DLP-based stimulator can achieve this goal. Because the spatial bandwidth of the stimulator is relatively low, no single magnification is ideal for recording from all retinal neurons. Nevertheless, adjusting image magnification can optimize spatial frequency range and stimulus size. With a $4\times$ objective, the highest spatial frequency that can be produced is ~ 10 cycles/degree. Because the stimulus is large, this configuration is ideal for stimulating peripheral neurons such as H1 horizontal cells that have large receptive fields but little response to high spatial frequencies. By using the lower image magnifications made possible with $10\times$ or even $40\times$ objectives, spatial frequencies up to the limits of normal vision can be created for recording from foveal neurons that respond well to high spatial frequencies but have small receptive fields.

The primary spatial advantage of the DLP for our application is the small size of the DMD chips, which makes it easier to design relay optics with the desired magnification and numerical aperture. Spatial nonlinearities that are not intrinsic to DMD technology were likely introduced by the electronics of the video interface.

4.2. Temporal performance

The temporal performance of the DLP-based stimulator did not reach the temporal limits of the visual system. Because of its limited refresh rate, the DLP can create stimuli with temporal frequencies no higher than 31.5 Hz. Fig. 6 shows that an LED-based stimulator has no loss of modulation up to at least 100 Hz because LEDs switch on and off very rapidly. Below its sampling limit, the tMTF of the DLP is less than that of LEDs.

4.3. Chromatic performance

One of the largest potential advantages of DLP technology was the ability to create a wider range of

chromaticities than is possible with a CRT. We use silent substitution to create stimuli that modulate single cone types (Dacey, Diller, Verweij, & Williams, 2000b). Creating high contrast stimuli that modulate L and M cones requires R and G primaries that are near the spectrum locus and as far apart in color space as possible. The CRT has a G phosphor of relatively long dominant wavelength that limits cone contrast. The G primary of the DLP is similar to that of the CRT leaving it only modestly more capable of producing the saturated blues and greens that generate the highest cone contrasts. The LED-based stimulator is more flexible than either CRT or DLP because LEDs can be chosen from a wide range of dominant wavelengths.

Although we went to some length to minimize chromatic nonuniformities, we were successful only at the maximum aperture (0.1) of the optics. At lesser apertures, the increased depth of focus allowed some structure in the illuminating optics to reach the stimulus causing visible chromatic splotches on a nominally white field. Note that for stimulators that use the DLP in its intended projection mode, spatial nonuniformity would not be a significant problem.

Other potential chromatic limitations such as convergence error and chromatic aberration can be minimized by sufficiently reducing the magnification of the DMD on the retina so that the spatial scale of the artifacts is too small to matter. These factors, along with shifts in the shapes of the primary spectra with output level and numerical aperture, need to be evaluated for each planned experiment, but for many uses will not be troublesome.

4.4. Improvements

Although the current generation of digital light projectors falls short of the full potential of the technology for generating visual stimuli, there is reason to expect improvements in future generations of machines. The digital television interface scheduled to be phased in over the next few years may eliminate spatial nonlinearities and provide more direct control over individual pixels. At least one company (Rochester Microsystems Inc., Rochester, NY) has developed a digital DMD interface that appears to provide direct control of individual pixels¹.

The temporal performance of the DLP is a function of refresh rate. Newer DLP engines are already faster (83 Hz) than the one we evaluated here (63 Hz). Given the speed with which integrated circuit technology is

advancing, it would not be surprising to see the DMD refresh rate equal that of the CRT.

Improvements in the chromatic performance of the DLP are more difficult to predict. Although the clever use of dichroic prisms to split and recombine the primaries, along with dichroic coatings to shape the spectra, makes possible a compact three channel projector, it is also the likely cause of spectral shifts with output level and numerical aperture.

Given today's technology, perhaps the best DLP-based stimulator for generating visual stimuli for in vitro physiology would use three single channel DLP engines that use filters rather than dichroic mirrors to shape the primary spectra and metal halide rather than xenon arc lamps. Their outputs would be combined using conventional beam splitters and their spectra shaped with interference or broadband filters. Lastly, each DMD would be driven by a custom digital interface. Such a stimulator would require a substantial amount of development work. However, it would maximize the potential of a visual stimulator based on DLP technology.

Acknowledgements

This work was supported by NIH grants EY10016 (D.H. Brainard), EY06678 (D.M. Dacey), EY00901 (J. Pokorny), EY01730 (Vision Research Core) and RR00166 (University of Washington Regional Primate Research Center). We thank Steve Buck for helpful discussions regarding calibration as well as the use of his radiometer.

References

- Brainard, D. H. (1989). Calibration of a computer controlled color monitor. *Color Research and Application*, 14, 23–34.
- Brainard, D. H. (1995). Colorimetry. In M. Bass, *Handbook of optics*, vol. 1. Fundamentals, techniques and optical design. New York: McGraw-Hill.
- Dacey, D. M. (2000). Parallel pathways for spectral coding in primate retina. *Annual Review of Neuroscience*, 23, 743–775.
- Dacey, D. M., & Lee, B. B. (1994). The 'blue-on' opponent pathway in primate retina originates from a distinct bistratified ganglion cell type. *Nature*, 367, 731–735.
- Dacey, D. M., Lee, B. B., Stafford, D. K., Pokorny, J., & Smith, V. C. (1996). Horizontal cells of the primate retina: cone specificity without spectral opponency. *Science*, 271, 656–659.
- Dacey, D. M., Packer, O., Diller, L. C., Brainard, D., Peterson, B., & Lee, B. B. (2000a). Center surround receptive field structure of cone bipolar cells in primate retina. *Vision Research*, 40, 1801–1811.
- Dacey, D. M., Diller, L. C., Verweij, J., & Williams, D. R. (2000b). Physiology of L and M cone inputs to H1 horizontal cells in primate retina. *Journal of the Optical Society of America A*, 17, 589–596.

¹ Other display technologies, particularly liquid crystal displays (LCDs), are also evolving rapidly and could be considered as alternatives to DLPs. Recent reports provide some information about the performance of LCD displays for vision research (Fairchild & Wyble 1998; Gibson & Fairchild, 2000).

- Estevez, O., & Spekreijse, H. (1974). A spectral compensation method for determining the flicker characteristics of the human colour systems. *Vision Research*, *14*, 823–830.
- Fairchild, M. D., & Wyble, D. R. (1998). Colorimetric characterization of the Apple Studio Display (Flat Panel LCD). Munsell Color Science Laboratory Report, Rochester Institute of Technology, Rochester, NY. Available online at <http://www.cis.rit.edu/mcsl/pubs/reports.shtml>.
- Gibson, J. E., & Fairchild, M. D. (2000). Colorimetric characterization of three computer displays (LCD and CRT). Munsell Color Science Laboratory Report, Rochester Institute of Technology, Rochester, NY. Available online at <http://www.cis.rit.edu/mcsl/pubs/reports.shtml>.
- Goodman, J. W. (1996). *Introduction to Fourier optics* (2nd edn). New York: McGraw-Hill.
- Hornbeck, L. J. (1997). Digital light processing for high-brightness high-resolution applications. Texas Instruments. Available online at <http://www.vxm.com/TIDL.html>.
- Packer, O., & Dacey, D. M. (2000). Receptive field structure of H1 horizontal cells in macaque monkey retina. *Investigative Ophthalmology and Visual Science*, *41*, S943.
- Pelli, D. G. (1997). Pixel independence: measuring spatial interactions on a CRT display. *Spatial Vision*, *10*, 443–446.
- Rushton, W. A., Powell, D. S., & White, K. D. (1973). The spectral sensitivity of 'red' and 'green' cones in the normal eye. *Vision Research*, *13*, 2003–2015.
- Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, *37*, 10–21.
- Shapiro, A. G., Pokorny, J., & Smith, V. C. (1996). Cone-rod receptor spaces with illustrations that use CRT phosphor and light-emitting-diode spectra. *Journal of the Optical Society of America A*, *13*, 2319–2328.
- Smith, G., & Atchison, D. A. (1997). *The eye and visual optical instruments*. New York: Cambridge University Press.
- Swanson, W. H., Ueno, T., Smith, V. C., & Pokorny, J. (1987). Temporal modulation sensitivity and pulse-detection thresholds for chromatic and luminance perturbations. *Journal of the Optical Society of America A*, *4*, 1992–2005.
- Watanabe, T., Mori, N., & Nakamura, F. (1992). A new superbright LED stimulator: photodiode-feedback design for linearizing and stabilizing emitted light. *Vision Research*, *32*, 953–961.
- Wyszecki, G., & Stiles, W. S. (1982). *Color science: concepts and methods, quantitative data and formulae* (2nd edn). New York: Wiley.