Visual Neuroscience (1997), **14**, 1015–1027. Printed in the USA.

The use of m-sequences in the analysis of visual neurons: Linear receptive field properties

R.C. REID,[1,2,3] J.D. VICTOR,[1,2] AND R.M. SHAPLEY[1,3]

[1]The Rockefeller University, Laboratory of Biophysics, New York
[2]Cornell University Medical College, Department of Neurology and Neuroscience, New York
[3]New York University, Center for Neural Science, New York

(RECEIVED December 19, 1995; ACCEPTED March 5, 1997)

Abstract

We have used Sutter's (1987) spatiotemporal m-sequence method to map the receptive fields of neurons in the visual system of the cat. The stimulus consisted of a grid of 16 × 16 square regions, each of which was modulated in time by a pseudorandom binary signal, known as an m-sequence. Several strategies for displaying the m-sequence stimulus are presented. The results of the method are illustrated with two examples. For both geniculate neurons and cortical simple cells, the measurement of first-order response properties with the m-sequence method provided a detailed characterization of classical receptive-field structures. First, we measured a spatiotemporal map of both the center and surround of a Y-cell in the lateral geniculate nucleus (LGN). The time courses of the center responses was biphasic: OFF at short latencies, ON at longer latencies. The surround was also biphasic—ON then OFF—but somewhat slower. Second, we mapped the response properties of an area 17 directional simple cell. The response dynamics of the ON and OFF subregions varied considerably; the time to peak ranged over more than a factor of two. This *spatiotemporal inseparability* is related to the cell's directional selectivity (Reid et al., 1987, 1991; McLean & Palmer, 1989; McLean et al., 1994). The detail with which the time course of response can be measured at many different positions is one of the strengths of the m-sequence method.

Keywords: White noise, Reverse correlation, Cat, Lateral geniculate nucleus (LGN), Visual cortex, Simple cell

Introduction

A visual receptive field can be defined as the functional transformation from a visual stimulus, that varies in both space and time, to a time-varying neuronal output. All receptive-field studies are attempts to characterize this transformation between stimulus (S) and response (R):

$$S(x, y, t) \rightarrow R(t). \quad (1)$$

Any characterization of this transformation is potentially limited by the complexity and variety of the stimuli used to measure it. In this paper, we report on the use of a powerful technique for studying neural transformations using a rich spatiotemporal signal derived from an m-sequence, or maximal length shift-register sequence (Sutter, 1987).

A number of different methods have been used to study receptive fields with random or pseudorandom noise stimuli. These methods can be divided into two broad categories, which we will call *sparse-noise* or *dense-noise* methods according to the temporal signal used at each position. Sparse-noise stimuli are those which are zero most of the time, such as the random pulse sequence used by Jones and Palmer (1987). Stimuli consisting of single temporal sine waves may be considered as sparse stimuli in this context, because their Fourier transforms are zero except at a single frequency. Sparse stimuli, either in the time or the frequency domain, are most useful for characterizing quasilinear systems. Although they can be used to *detect* nonlinearities (Enroth-Cugell & Robson, 1966), they cannot probe the details of nonlinear interactions between stimulus components at two or more times or positions.

Dense-noise stimuli are those which are nonzero most of the time, and typically vary in a "complex" fashion. All of the dense-noise methods can therefore be used to characterize nonlinearities in receptive-field structure, since they probe the system with rich patterns of temporal (or spatiotemporal) inputs. Some dense-noise methods [e.g. the original white noise method of Wiener (1958)] are based on signals that are intended as laboratory approximations to a random process. Other dense-noise methods are based on particular sequences that, although they may resemble random processes, have a deterministic temporal structure which confers substantial practical advantages. These include time-domain methods, such as the Sutter's m-sequence approach (1987), and the frequency-domain methods, such as the sum-of-sinusoids method

Reprint requests to: R. Clay Reid, Harvard Medical School, Department of Neurobiology, 220 Longwood Avenue, Bldg B1-442, Boston, MA 02115, USA.

(Victor & Knight, 1979). This nonrandom category also includes the "sparse" m-sequence variant (Sutter, 1992*b*) and a recent hybrid method (Benardete & Victor, 1994). Dense-noise methods can be further classified on the basis of their power spectrum (white, that is, equal power at all frequencies *vs.* non-white) and on the basis of the probability distribution of the stimulus values (Gaussian *vs.* non-Gaussian). M-sequences are nearly white but they are binary (they take on only two different values), which is distinctly not Gaussian; sum-of-sinusoids are non-white but nearly Gaussian.

Dense- and sparse-noise stimuli have been used in other studies to measure the functional transformations in visual neurons. Gaussian-noise signals (or approximations to Gaussian noise) have been used to modulate the contrast of a single spatial stimulus, such as a spot or a sine-wave grating. In these studies (e.g. Victor et al., 1977; Shapley & Victor, 1981; Victor, 1987; Naka et al., 1985), the detailed *temporal* dynamics of single receptive-field structures, such as the center mechanism in a center/surround retinal ganglion cell, have been characterized precisely. In other studies, dense-noise stimuli have been used to map the *spatial* structure of receptive fields both in one dimension (Emerson et al., 1987) and in two dimensions (Citron et al., 1981; Jacobson et al., 1993). Sparse-noise stimuli have also been used to analyze receptive fields in two dimensions, but with these stimuli only one or two regions are modulated at any point in time (Jones & Palmer, 1987; DeAngelis et al., 1993*a*,*b*, 1995; McLean et al., 1994) so that nonlinear interactions could not be measured (but see Szulborski & Palmer, 1990; Ohzawa et al., 1990).

Linear analysis: Equivalence of sparse- and dense-noise methods

A linear receptive field is one that obeys the principle of superposition: the response to the sum of two stimuli is equal to the sum of the responses to the two stimuli in isolation. For neurons that sum their inputs in a roughly linear manner, the responses to dense- and to sparse-noise stimuli can therefore be analyzed and interpreted in a similar way. The receptive field is characterized as the stimulus that, on the average, tended to precede each action potential; it is the stimulus that tended to make the neuron fire. The most intuitive way of presenting this calculation is in the reverse correlation method of Jones and Palmer (1987; cf. Bussgang, 1952; Nuttal, 1957; Weiss, 1966; De Boer & Kuyper, 1968).

The reverse correlation algorithm works as follows. For each action potential, the spatial configurations of the preceding stimuli are added into separate running sums for each delay between stimulus and response. After a sufficient number of action potentials, these sums converge to spatial characterizations of the receptive field, parametric in time. This method thereby creates a standard receptive-field map for each delay. In an ON region, the average stimulus that preceded each action potential will have been lighter than the mean. In an OFF region, the average stimulus will have been darker. Since linear analysis cannot distinguish between excitation by one phase of a binary stimulus (again, one that takes on only two different values) from inhibition by the opposite phase, we will use the term ON to denote either ON-excitation or OFF-inhibition and OFF to denote OFF-excitation or ON-inhibition (see Emerson et al., 1987, for a discussion of the separate responses to the light and dark phases of a ternary stimulus).

For linear systems, the analysis of response properties with dense-noise stimuli is formally equivalent to the reverse correlation procedure for sparse-noise stimuli. For such systems, the response properties are completely described by the *first-order Wiener kernel*, $K_1(x,y,t)$, which is proportional to the cross-correlation of the visual stimulus and the neuronal response:

$$K_1(x,y,t) = 1/N \sum_{i=1}^{N} R(t_i)S(x,y,t_i - t) \qquad (2)$$

Informally, this cross-correlation can be considered an "optimal stimulus" of the same sort obtained with the reverse correlation procedure (Jones & Palmer, 1987). It is the average spatial stimulus preceding an action potential, as a function of delay, t. This cross-correlation is called the first-order kernel because it has an additional interpretation that extends its applicability. For nonlinear systems, it is the best-fitting linear approximation to the transformation from stimulus to response. Although this is often referred to as "the linear response," it must be emphasized that for nonlinear systems, this linear approximation may depend strongly on the stimulus set. Finally, it can also be seen as the spatiotemporal weighting function of the neuron's response to visual stimuli; it gives the strength of response to a stimulus presented at position (x,y) following a delay, t.

Nonlinear analysis of the responses to highly structured dense-noise stimuli

Nonlinear analysis based on dense-noise stimuli allows the study of more complex spatiotemporal interactions in the receptive field. It is important again to emphasize that for nonlinear systems, the sparse-noise and the dense-noise methods of analysis differ even when just the results of first-order reverse-correlation analysis are compared. Sparse-noise methods do not probe spatiotemporal interactions adequately, and thus the results of these interactions are intermixed with the quasilinear responses. Dense-noise methods are, within certain theoretical limitations (Victor 1992), effective probes of these interactions. They separate linear and nonlinear components of the response in a rigorously defined (but stimulus dependent) manner. For grossly nonlinear systems, such as those found in extrastriate cortex, it is unlikely that either dense or sparse noise inputs would give interpretable results, although modifications have been used with some success (Britten, 1995).

The detailed formalism of Wiener kernel analysis has been reviewed in a number of places (Marmarelis & Marmarelis, 1978; Sakai et al., 1990; Victor, 1992). Of the various temporal signals that have been used in linear and nonlinear analysis of the visual system, two are distinctly nonrandom: sum-of-sinusoids (Victor et al., 1977) and maximal length shift-register sequences or m-sequences (Sutter 1987). Kernels measured with these highly structured stimuli are good approximations to Wiener kernels measured with Gaussian white noise (Victor & Knight, 1979; Sutter, 1987; Victor, 1991) and are therefore useful for comparisons with kernels of model systems. The sum-of-sinusoids method has been used extensively in the visual system and results have been reviewed before (Victor, 1992). Here we describe the use of the m-sequences in the study of single visual neurons.

Maximal length shift-register sequences (m-sequences) are signals that are particularly suited to the study of multi-input systems for both practical and theoretical reasons. Practically, a two-dimensional m-sequence stimulus is easy to generate quickly on a computer. Just as importantly, analytical techniques developed by Sutter (1992*a*) allow for very fast calculation of first- and higher order kernels. Sutter and coworkers have used the m-sequence method to great advantage in the study of temporal and topo-

graphic relations in the VEP (Baseler et al., 1994) and electroretinogram in humans (Sutter & Vaegan, 1990; Sutter & Tran, 1992), and of cortical field potentials in the cat (Kitano et al., 1994). Our present study reports on the use of Sutter's method in the analysis of single neurons in the mammalian visual system.

Methods

Physiological preparation

The details of our physiological experiments have been described previously (Hochstein & Shapley, 1976; Reid et al., 1991). Adult male cats ranging in weight from 2.2 to 3 kg were initially anesthetized with an injection of ketamine (10 mg/kg, i.m.). Penicillin (500,000 units, i.m.) and dexamethasone (6 mg, i.v.) were given at the start of the experiment. During surgery, sodium thiamylal (Surital) was given i.v. as needed. Pupils were dilated with topical application of 1% atropine sulfate and the nictitating membranes were retracted with 10% phenylephrine. The eyes were protected with contact lenses with a 3-mm-diameter artificial pupil. Eye movements were minimized by attaching the eyes with cyanoacrylate glue to small metal posts connected to the stereotaxic frame.

Urethane (200–300 mg/kg, i.v., supplemented every 12 h) was delivered over the course of several hours as the maintenance anesthetic. The EKG was measured and, along with the blood pressure, was used to monitor the level of anesthesia. The animal was paralyzed with gallamine triethiodide (Flaxedil, 20–40 mg/h, i.v.) and artificially ventilated. Ventilation was adjusted so that end-expiratory CO_2 was near 3.5%. Core body temperature was monitored and body temperature maintained at 38° C. Recordings were made with plastic coated tungsten microelectrodes (Hubel, 1957).

Visual stimulation

The electronic visual stimulator employed in these experiments was designed in The Rockefeller University Laboratory of Biophysics (Milkman et al., 1980). The instrument produces a raster display of 256 picture elements (pixels) per line, 256 lines per frame, and 270 frames/s on a Tektronix 608 cathode ray tube monitor. For the Y-cell, the display was run at two frames per m-sequence step, for the simple cell, four frames per step. A look-up table within the display unit allowed us to modulate the contrast linearly. The 10 × 10 cm display was viewed, *via* a mirror, at a distance of 36 cm, resulting in stimulus subtending 16 deg × 16 deg of visual angle. The mean luminance was 100 cd/m^2.

M-sequences

Theoretical considerations

In this section, we consider a number of important mathematical properties of m-sequences; in the next section we will define them and outline how they are generated. A binary m-sequence is represented in a computer as a series of 0s and 1s of length $2^n - 1$, $(b_0, b_1, b_2, \ldots b_{2n-2})$. When employed as a visual stimulus, it is best considered as a list of 1s and −1s, or positive and negative contrasts around a mean. As noted below, 0 (binary) maps to 1 (contrast) and 1 (binary) maps to −1 (contrast). We consider this sequence to be the initial cycle of an infinitely long periodic stimulus, whose period is $2^n - 1$. The resulting signal, $(s_0, s_1, s_2, \ldots)$, has a number of desirable properties.

Property 1. The first-order statistics of an m-sequence are balanced, i.e. there is an almost equal number of 1s and −1s:

$$\sum_{i=0}^{2^n-2} s_i = -1.$$

Property 2. Its second-order statistics are also balanced, i.e. it is almost completely uncorrelated with itself for all time shifts:

$$\sum_{i=0}^{2^n-2} s_i s_{i+j} = -1, \qquad (j \neq 0)$$

The usual definition of white noise is that this autocorrelation function is equal to zero (at $j \neq 0$), which is equivalent to its power spectrum being equal, at all frequencies. The m-sequence is guaranteed to be white to within one part in $2^n - 1$, while a "random" sequence is typically white to within only one part in $\sqrt{2^n - 1}$. In a loose sense, therefore, an m-sequence can be considered a maximally white binary sequence of length $2^n - 1$.

Property 3. For time shifts less than n, an m-sequence has balanced statistics up to order n: every n-tuple of 1s and −1s (except for the string of n 1s, which is omitted) is represented exactly once.

Therefore, at least for short stretches of length n, a finite m-sequence has the same correlation properties as an infinite random stimulus.

Property 4. The product of an m-sequence with a time-shifted copy of itself yields the same m-sequence with a different delay:

$$s_i s_{i-j} = s_{i-k}$$

The relation between the js and ks is different for each m-sequence. It can be seen that Property 2 follows directly from Property 1 and Property 4. This final property of m-sequences will be useful when we consider the calculation of higher order kernels.

The nearly ideal second-order statistics (Property 2) of an m-sequence stimulus make it useful for the study of multi-input systems, such as a visual receptive field. If a sequence of sufficient length is used, the same sequence with different delays can be used to modulate different regions of the stimulus. In the present study, the stimulus was a 16 × 16 array of pixels, each modulated by the same m-sequence of order 16 (*pixel* refers to a uniform region of the stimulus; *machine pixel* will be used below to refer to the smallest point that the device displays). Successive pixels were presented with relative delays of 256 steps within the m-sequences. Property 2 guarantees that the stimulus is spatiotemporally white, rather than just temporally white. That is, for any pair of pixels and any fixed delay (other than 256 steps), the input signals will be the same (both light or both dark) as often as they are different (one light, one dark).

The computation of m-sequences

We have generated m-sequence stimuli by two different methods. Both methods are based on *shift registers*: n-bit integers that are updated by particularly simple rules. Successive steps of a shift register, R_i, are calculated using two pieces of information, the order, n, of the m-sequence and the feedback tap, T. The feedback tap is an n-bit number that generates the specific sequence. Each value of binary m-sequence (b_i) is equal to the lowest order bit of the associated term in the shift-register sequence (R_i).

Shift registers employ four bitwise calculations. A *bit shift* (>> or <<) moves all the bits of a binary number either to the left or the right, e.g.:

$$0111 >> 1 = 0011, \tag{3}$$

$$0111 << 1 = 1110. \tag{4}$$

The *bitwise exclusive or* (^) operation adds each binary digit of two numbers modulo 2:

$$0111 \wedge 1110 = 1001. \tag{5}$$

The *bitwise and* (&&) multiplies each binary digit:

$$1000 \,\&\&\, 1101 = 1000. \tag{6}$$

Finally, the *parity* operation $[P(..)]$ gives the sum of the digits in the binary number, modulo 2:

$$P(0111) = 1, \tag{7}$$

$$P(0110) = 0. \tag{8}$$

The more conventional shift-register calculation of an m-sequence (Golomb, 1982), which Sutter calls the *generating register* (1987) method, proceeds as follows:

$$R_{i+1} = R_i >> 1 + q\,P(R_i \,\&\&\, T), \tag{9}$$

where $q = 2^{n-1}$. In other words, the register is shifted one bit to the right and the nth bit is replaced with the *parity* of the *bitwise and* of the tap, T, and the original register, R_i.

Since the parity operator is unwieldy to implement in most computer languages, a second method, called the *tap register* method by Sutter (1987) is computationally more efficient. In this method, the register is first shifted to the *left*. If there was a 1 in the nth position of the original number, the *exclusive or* of T and the shifted register is taken:

if $(R_i \&\& q \neq q)$ (i.e. if bit $n-1$ is not set)

then $R_{i+1} = R_i << 1$,

if $(R_i \&\& q == q)$ (i.e. if bit $n-1$ is set)

then $R_{i+1} = (R_i << 1) \wedge T$.

For certain feedback taps, T, the sequence of register configurations repeats itself with period $2^n - 1$ (see Golomb, 1982). Since there are only $2^n - 1$ possible nonzero n-bit register configurations, these sequences are maximal in length.

The *m-sequence* itself can be defined as the *binary* sequence, b_i, derived from the above maximal length sequence by taking the lowest order bit of each successive n-bit integer:

$$b_i = R_i \,\&\&\, 1. \tag{10}$$

In other words, $b_i = 1$ if R_i is odd, $b_i = 0$ if R_i is even. The result of the tap register calculation is shown in Table 1 for an m-sequence of order 4 ($q = 1000$) and a feedback tap $T = 3 = 0011$.

Table 1. *Sixteen successive tap configurations for the tap-register method of m-sequence generation*[a]

$R_0 = 0001$	$b_0 = 1$
$R_1 = 0010$	$b_1 = 0$
$R_2 = 0100$	$b_2 = 0$
$R_3 = 1000$	$b_3 = 0$
$R_4 = 0011$	$b_4 = 1$
$R_5 = 0110$	$b_5 = 0$
$R_6 = 1100$	$b_6 = 0$
$R_7 = 1011$	$b_7 = 1$
$R_8 = 0101$	$b_8 = 1$
$R_9 = 1010$	$b_9 = 0$
$R_{10} = 0111$	$b_{10} = 1$
$R_{11} = 1110$	$b_{11} = 0$
$R_{12} = 1111$	$b_{12} = 1$
$R_{13} = 1101$	$b_{13} = 1$
$R_{14} = 1001$	$b_{14} = 1$
$R_{15} = R_0 = 0001$	$b_{15} = 1$

[a]The order $n = 4$ and the feedback tap $T = 3 = 0011$.

It turns out that any binary sequence $(b_0, b_1, b_2, \ldots, b_{2^n-2})$ constructed in this manner obeys a recursion relation $b_k \equiv b_{k-1}t_{n-1} + b_{k-2}t_{n-2} + \ldots + b_{k-n}t_0$, where $t_{n-1}t_{n-2} \ldots t_1t_0$ are the n digits in the binary representation of T, and $\equiv$ represents addition modulo 2. This constitutes yet a third way of generating the stimulus from the feedback tap, T, closely related to the first. This method is used in the third of the four approaches we have used to generate m-sequences, discussed below.

Although an m-sequence is generated in a computer by means of binary arithmetic, in much what follows it will be useful to make the substitution of $+1$ for 0 and -1 for 1:

$$s_i = -2b_i + 1. \tag{11}$$

Multiplication on the set $\{+1, -1\}$ is identical to addition modulo 2 on the set $\{0,1\}$, so this substitution is a natural one.

Image generation

While each step of the m-sequence is rapidly calculated, most computer driven visual display devices allow only a limited number of values to be updated each frame. Even updating a relatively modest 16×16 array of pixels can be quite difficult when these pixels must be modulated independently at a high rate. The simplicity of the m-sequence stimulus makes it easier to display, however, than would be a general two-dimensional random stimulus. The m-sequence stimulus, as used by Sutter (1987), employs a single m-sequence to modulate a large number of different pixels in a two-dimensional grid. The simplest case is for a rectangular array of pixels with r rows and c columns, where both values are powers of two. The m-sequence of order n is then evenly sampled in this two-dimensional array at intervals of p steps, where

$$p = 2^n/rc. \tag{12}$$

In this case, the spatiotemporal stimulus, $S(x, y, t_i)$, can be expressed in terms of the single-parameter m-sequence, s_i:

$$S(x, y, t_i) = s_{i+D(x,y)} \tag{13}$$

where the spatial offset term, $D(x,y)$, is given by

$$D(x,y) = px + pcy. \tag{14}$$

By far the simplest method to achieve this would be to have the number of pixels, rc, no larger than the size of the look-up table that translates the index of a point in the display (0 to 255 for an 8-bit display) into luminance. In this case, the look-up table could be changed each frame according to the successive values in the m-sequence. Since many display devices do not allow the entire lookup table to be modified each frame, this simplest strategy is not always possible. Because visual display devices have a variety of different structures, three more strategies will be presented.

A second strategy allows a virtually limitless number of pixels to be modulated simultaneously. It exploits the fact that many visual display devices (such as the a Number Nine board for the PC, Number Nine Computer Corporation, Cambridge, MA) have a large amount of video memory, of which a small portion is displayed each frame. The spatial origin can be freely chosen and each memory pixel can often be "zoomed" into many display pixels. The strategy requires that the m-sequence of order n must be evenly sampled in a two-dimensional array ($r \times c$) at intervals of p steps, as above.

The first frame of the stimulus is stored at the origin of the image memory in the first r rows and c columns (see Table 2). These rc pixels contain the positions $[0,p,2p,\ldots,(rc-1)p]$ in the m-sequence. If one continues in this manner, the first element of row $r+1$ has position 1 in the m-sequence, since the total length is $2^n - 1$, not $2^n = rcp$. Since p is relatively prime to $2^n - 1$ (it is a power of 2), the entire m-sequence will be sampled by continuing in this manner. Finally, if copies of the m-sequence are made $c-1$ columns to the right of the original and $r-1$ rows below, each frame of the entire stimulus ensemble is displayed simply by moving the origin of the display window relative to the bitmap origin (see Table 2). This method allows the entire stimulus ensemble to be stored in

$$(2c-1)[pr + (r-1)] \approx 2^{n+1} \tag{15}$$

memory locations as opposed to the $(rc)(2^n - 1)$ required by storing each frame independently (again, the zooming operation allows a single memory location to be shown on a larger screen location, for instance 16 × 16 machine pixels). More importantly, each frame is updated with one simple operation, the shift of origin, which is independent of the number of pixels. Rather than sending the luminance of each pixel to the display device every frame, only two numbers suffice, the x and y values of the origin. This strategy allows the display of up to 128 × 128 pixels every frame with a Number 9 board in a PC. This can be achieved by having a very long m-sequence (say $2^{19} - 1$, which still only requires on the order of 1 Mbyte of memory), which need not be run completely to record useful data. This is far more pixels than one needs to study a single receptive field but could be quite useful to study many receptive fields simultaneously.

A slight modification of this method is to keep the video display memory static, but to change the look-up table every frame, as in the first method. The entire m-sequence can be precomputed and stored in a single linear array of length $2^n + rc$, which can be filled in the manner illustrated in Table 2. A different portion of this array is copied directly into the look-up table each frame. This change of origin for the look-up table is exactly analogous to the change of origin for the spatial stimulus described in the preceding paragraphs.

A third more specialized method was used to create the 16 × 16 stimuli used in the present study. We used the Milkman visual stimulator described previously (Milkman et al., 1980) which, while quite versatile, does not allow stimuli to be moved freely in two dimensions. Instead, it has 16 one-dimensional spatial luminance profiles each of which is displayed through a separate window on the display screen. In a single profile, it was possible to display eight different square regions each frame. We used the 16 profiles and the windowing capability of the stimulator in the following way to create a 16 × 16 m-sequence display.

First, the screen was split in half vertically and each half divided into 16 horizontal strips, one for each spatial profile. For each frame of the stimulus, only half of the screen was displayed and the other half was dark. Thus 128 square pixels were visible at a time, 8 (horizontally) from each of the 16 profiles (stacked vertically). Since the frame rate of the display was quite high, 270 Hz, the alternation between 8 × 16 arrays was not perceptible; the array appeared to have 16 × 16 pixels. In the kernel calculations, this 4-ms time-shift was taken into account by interpolating between successive frames for pixels in the lagged half of the screen.

Finally, a fourth method for efficient generation of m-sequence displays, which has been used by Victor et al. (1994), exploits the fact that the binary sequence of values at each stimulus region satisfies the same recursion relationship as the m-sequence definition: $b_k \equiv b_{k-1}t_{n-1} + b_{k-2}t_{n-2} + \ldots + b_{k-n}t_0$ (again, $\equiv$ is interpreted as addition modulo 2). Because the identical relationship is obeyed by the values displayed at each region, this can be

Table 2. *Strategy for presenting an m-sequence stimulus on a display device with a two-dimensional memory*[a]

$c=4$				$c-1=3$			
0	4	8	12	16	20	24	$r=4$
16	20	24	28	32	36	40	
32	36	40	44	48	52	56	
48	52	56	60	1	5	9	
1	5	9	13	17	21	25	$pr=16$
17	21	25	29	33	37	41	
33	37	41	45	49	53	57	
49	53	57	61	2	6	10	
2	6	10	14	18	22	26	
18	22	26	30	34	38	42	
34	38	42	46	50	54	58	
50	54	58	62	3	7	11	
3	7	11	15	19	23	27	
19	23	27	31	35	39	43	
35	39	43	47	51	55	59	
51	55	59	0	4	8	12	
4	8	12	16	20	24	28	$r-1=3$
20	24	28	32	36	40	44	
36	40	44	48	52	56	60	

[a]In the terminology of the text: the order of the m-sequence is $n = 6$, four rows and four columns are displayed at any one time ($r = c = 4$), and the spacing between successive pixels is four steps ($p = 4$). The array shows the positions in the m-sequence (0–62, which correspond to b_0–b_{62}) as they are stored in a 19 × 7 block of memory. A different 4 × 4 region of the memory is displayed every frame. The dotted lines show the m-sequence positions displayed in the 5th and 43rd frames.

viewed as a recursion relationship for the display as a whole: $X_k \equiv X_{k-1}t_{n-1} + X_{k-2}t_{n-2} + \ldots + X_{k-n}t_0$, where X_k denotes the array of machine pixel values (0s or 1s) in the entire display at time step k. This relationship, which relates the state of the display at time k to its state at the preceding n times, can be iterated $k - n + 1$ times to provide a relationship between the display at time k and the display at the first n times: $X_k \equiv X_{n-1}u_{k,n-1} + X_{n-2}u_{k,n-2} + \ldots + X_0u_{k,0}$, where the coefficients $u_{k,m}$ depend in a complex but fixed manner on the tap set T.

To make use of this relationship, the display states at the first n times, $X_0, X_1, \ldots, X_n - 1$, are stored in n bitplanes of the display, and the coefficients $u_{k,m}$ are precomputed from T. Then, to produce the display at time k, the bitplanes m for which $u_{k,m} = 1$ are exclusively-*or*'d together. This can be done by loading the look-up table to implement the exclusive-*or* of these bitplanes, or, by starting with a look-up table which implements the exclusive-*or* of all bitplanes, and inactivating the bitplanes for which $u_{k,m} = 0$. In contrast to the previous methods, this method does not require that the regions be regularly spaced or that they sample the m-sequence at equal intervals, but does require that the display hardware can be configured as n bitplanes (where n is the order of the m-sequence). The number of bits of memory required is equal to the order of the m-sequence times the number of machine pixels required to resolve the shapes of the stimulus regions, independent of the number of stimulus regions to be displayed.

Choice of tap registers

For any given order, n, a large number of different n-bit feedback taps, T, generate m-sequences (Golomb, 1982). These are best found by brute force: by determining the length of the sequence that any given tap generates. By definition, if a sequence has length $2^n - 1$, it is an m-sequence.

Because of Property 4, however, taps must be chosen wisely. The potential problems arise because the product of the m-sequence stimulus at two nearby points at nearby times is equal to the m-sequence at another point with a different delay. That is, if one chooses any spatial offsets, δ_x and δ_y and a temporal offset, δ_t, then there is another set of offsets—ϵ_x, ϵ_y, and ϵ_t—such that

$$S(x,y,t)S(x+\delta_x, y+\delta_y, t+\delta_t) = S(x+\epsilon_x, y+\epsilon_y, t+\epsilon_t). \tag{16}$$

This relations follows directly from Property 4 and eqn. (13). Thus, if there were a nonlinear interaction between the response at a point, (x,y) and a nearby point, $(x + \delta_x, y + \delta_y)$, with a relative delay of δ_t, then the portion of the second-order kernel due to this interaction will be found within the *first-order* kernel at point $(x + \epsilon_x, y + \epsilon_y)$ with a relative delay of ϵ_t. Thus, if ϵ_t were small, then the first-order kernel measured at the point $(x + \epsilon_x, y + \epsilon_y)$ would be contaminated by that portion of the second-order kernel.

Because the m-sequence is sampled at regular intervals on a two-dimensional grid, ϵ_x, ϵ_y, and ϵ_t are independent of the point of origin (x,y), and are functions only of the particular m-sequence, the three variables that determine the spacing of the m-sequence display—r, c, and p—and of the three offsets: δ_x, δ_y, and δ_t. The portion of the second-order kernel due to interactions between points with an offset of δ_x, δ_y, and δ_t is "aliased" to an offset in the first-order kernel: ϵ_x, ϵ_y, and ϵ_t. For each triplet of relative interaction distances, δ_x, δ_y, and δ_t the second-order kernel with respect to positions (x,y) is aliased to a different triplet of offsets: ϵ_x, ϵ_y, and ϵ_t.

Therefore, an "interaction radius" can be chosen—say δ_x and $\delta_y \leq 4$ and $\delta_t \leq 5$—and then a number of different m-sequences can be tested in order to find one for which most ϵ_x, ϵ_y, and ϵ_t are acceptably large for each of the values of δ_x, δ_y, and δ_t within the interaction radius. Confusion between first- and second-order kernels can occur if ϵ_x and ϵ_y are small, but only if ϵ_t is small as well. Therefore, it is most important that all values of ϵ_t are significantly longer than the memory of the system. In the current study, the order of the m-sequence was 16, the display was 16×16 ($r = c = 16$), and successive pixels were 256 steps apart in the m-sequence ($p = 256$). For these values and for the interaction radius given above, the first three possible taps—$T = 45$, $T = 57$, and $T = 63$—were found to have few potential overlap of first- and second-order kernels. For order 15 m-sequences with a shorter relative delay between pixels ($p = 128$), overlap is more likely to be a problem. Of the first eight possible taps, the fifth ($T = 53$) was found to minimize potential overlaps.

Data analysis

The goal of the data analysis is to calculate the first-order kernel as a function of time and spatial position. At first glance this might appear to be computationally quite demanding. The output, r_i, must be correlated separately with the m-sequences from each of the pixels. If there are 256 pixels, 256 separate cross-correlations must be calculated. Since, however, each pixel of the display is modulated by the same signal with a variable delay [eqn. (13)], the correlation of the two-dimensional input with the output, r_i, can be calculated by means of a cross-correlation with this single input function, s_i:

$$k_\tau = 1/(2^n - 1) \sum_{i=0}^{2^n-2} s_{i-\tau} r_i, \tag{17}$$

The most natural units for both k_τ and r_i are spikes/s, or instantaneous rate. If there are n_i spikes in bin i, then $r_i = n_i/d$, where d is the duration of each stimulus frame.

This calculation is accelerated by means of the fast m-transform (Sutter, 1992*a*), which computes the correlation between any function and an m-sequence. Given this correlation between the m-sequence and the response, k_τ, the first-order kernel $K_1(x,y,t)$ [see eqn. (2)] can be looked up in a simple fashion:

$$K_1(x,y,t_i) = k_\tau \qquad \text{where } \tau = i - D(x,y), \tag{18}$$

where $D(x,y)$ is given by eqn. (14) and $t_i = id$. If the relative delay between pixels is long compared to the memory of the system (in the present study it was 256 steps/135 Hz = 1.9 s for the Y-cell, and 256 steps/67.5 Hz = 3.8 s for the simple cell), there will be no contamination between these pieces of the first-order kernel for different locations.

It is simplest to have the first bin of the neuronal output, r_0, be the firing rate during the first stimulus frame. Since there can be a response during this first frame, between time 0 and d, it is possible for the first bin of the response kernel, k_0, to be nonzero. This convention may therefore lead to the impression that latency of the responses is zero. It is nevertheless the simplest convention to choose, since it corresponds exactly to the reverse-correlation algorithm (Jones & Palmer, 1987): $K_1(x,y,t_i)$ is proportional to the average stimulus configuration at the time $t_i = id$ before each action potential.

The second-order kernel, $K_2(x_1, y_1, t_1, x_2, y_2, t_2)$, is proportional to the average value of the *product* of the stimulus at two positions, (x_1, y_1) and (x_2, y_2), and at two delays, t_1 and t_2. This average product of two stimuli represents the lowest order approximation of the system's deviation from linearity. It is similar in principle to the quadratic term in the polynomial expansion of a function of a single variable. Second-order responses are formally calculated as the correlation of the response, r_i, with the product of the stimulus, s_i, at two delays, τ_1 and τ_2,

$$k_{\tau_1,\tau_2} = 1/(2^n - 1) \sum_{i=0}^{2^n-2} s_{i-\tau_1} s_{i-\tau_2} r_i. \quad (19)$$

Here, the use of two subscripts indicates that k_{τ_1,τ_2} refers to a second-order kernel. Again, τ_1 and τ_2 are chosen so that $k_{\tau 1,\tau 2}$ represents the second-order response due to any pair of pixels with any two delays.

$$k_{\tau_1,\tau_2} = K_2(x_1, y_1, t_1, x_2, y_2, t_2). \quad (20)$$

Property 4 gives us the relation $s_{i-\tau_1} s_{i-\tau_2} = s_{i-\tau_3}$ for some τ_3. τ_3 is a function of the particular m-sequence and of the variables τ_1 and τ_2. Therefore, each value for the second-order kernel [eqn. (19)] can be found within the first-order kernel [eqn. (17)]:

$$k_{\tau_1,\tau_2} = k_{\tau_3}. \quad (21)$$

The second-order kernel, which would otherwise require an extremely lengthy calculation, is therefore calculated *via* the same fast m-transform that generated the first-order kernel (Sutter, 1987).

An additional strategy to minimize the contamination of first-order responses by second-order responses (see *Choice of tap registers*, above) is Sutter's "inverse repeat" method (1987, 1992*b*), which was not used in this study. If exactly the same m-sequence is used twice, but with opposite polarity (+1 for −1 and *vice-versa*), then all odd-order components of the first-order kernels will be the same, as calculated by eqn. (17), but the even-order will have the opposite sign. This is because, for an inverted m-sequence, the sign is inverted in Property 4 and therefore in eqn. (21). Therefore, contamination by second-order responses in particular will cancel out if the two first-order kernels are averaged.

Results

We will show detailed results from two cells in the cat visual system—a Y-cell in the lateral geniculate nucleus and an area 17 simple cell—picked in order to illustrate certain principles in the analysis of first-order response properties. These cells are fairly typical examples of both of their classes; we have studied many of each type as part of several different studies (Alonso & Reid, 1994; Reid & Alonso, 1995). The responses of both cells are dominated by the first-order kernels, which will be discussed in this paper. Although the cells both exhibited some significant nonlinearities, particularly truncation at zero spikes per second, second-order kernels will not be considered here.

Fig. 1 shows the "receptive field" of an OFF-center, ON-surround Y-cell (defined by its failure of the null test; Enroth-Cugell & Robson, 1966; Hochstein & Shapley, 1976) recorded in the lateral geniculate nucleus. More specifically, it represents first-order kernel, $K_1(x, y, t)$, for $t = 30$ ms. The light blue region, in the center, corresponds to the area where the neuron was excited by the dark phase of the stimulus or, equivalently, inhibited by the bright phase. The brighter the blue, the more strongly the neuron was excited by this dark phase. The diffuse dark red region, in the surround, corresponds to an area where it was weakly excited by the light phase of the stimulus (or inhibited by dark). Black corresponds to areas where the neuron was unaffected by visual stimuli. These data can be interpreted in several different ways. Most simply, Fig. 1 represents the average spatial configuration of the visual stimulus 30 ms before each action potential. In the center, the stimulus tended to be dark, in the surround, light. Conversely, each time one of the pixels in the center region was dark, the neuron was more likely (than average) to fire during the bin that began 30 ms later.

Clearly, the m-sequence method produces a useful picture of what would classically be called the receptive field (Fig. 1), and has the added benefit of producing different pictures for various delays between stimulus and response. Fig. 2 shows such receptive-field plots for the same Y-cell at 16 delays ranging from −7.4 to 104 ms in increments of 7.4 ms, the frame rate of the stimulus. The first frame, for negative times, is included to show the level of noise in the baseline. The values in this frame should be zero, since the firing of the neuron is uncorrelated with the input before it is presented. In the second frame, which corresponds to the response *during* the stimulus, a hint of an OFF response is visible. Dark stimuli in the receptive-field center produce a slight increase in the firing rate with a latency of less than 7.4 ms. The center response peaks at 22 ms, at which point the surround also becomes evident. The center responses then declines until changing sign at 37 ms. Such a biphasic response is not surprising for a "phasic" neuron, one that responds only briefly to a constant stimulus. The bandpass dynamics of a Y-cell center have been well described (Shapley & Victor, 1978), but they result in a seemingly paradoxical feature of the receptive field that should be emphasized: the *light* phase of the stimulus is strictly excitatory to the *receptive-field* center for delays longer than 39 ms. The surround shows the same biphasic response, but the zero crossing is later: around 52 ms.

It is evident from Fig. 2 that the time course of the response of this Y-cell is different for different regions in the receptive field. The center is fast and the surround significantly slower. This spatiotemporal inseparability (cf. Dawis et al., 1984 for the X-cell) is more salient when data are presented in the temporal rather than the spatial domain. Fig. 3 shows temporal cross sections through the first-order kernel, or the *impulse responses*, for the 32 pixels in Fig. 2 that gave the strongest responses. Each thin line in these plots represents the influence of one pixel on the firing rate of the neuron, parametric in time. Following the light phase of the 11 pixels in the center, the neuron was inhibited (negative) for the first 37 ms, followed by an excitatory rebound. The traces in Fig. 3(b) give the summed impulse responses for the 11 center pixels and also for all the remaining 245 pixels in the display, which were used to estimate the total surround response of the neuron. For the pixels in the surround, the light phase excited this neuron for 52 ms, followed by an inhibitory rebound. The surround clearly has a longer latency for this neuron. Since the center and surround are not strictly segregated spatially, the "surround" in Fig. 3(b) shows an initial weak OFF phase that should properly belong to the center mechanism.

Similar experiments were performed also to study visual cortical receptive fields. Fig. 4 shows the receptive field for a simple cell in area 17 at a delay of 44 ms. Four vertically elongated subregions are apparent, two ON (red) and two OFF (blue, the OFF

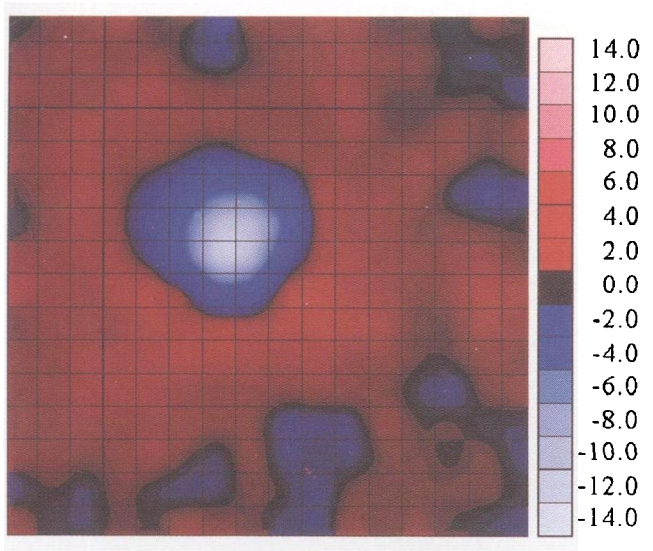

Fig. 1. Receptive-field representation (first-order kernels at a fixed time delay) for an OFF center Y-cell recorded in the LGN. Blue codes for OFF regions, red for ON. The brighter the blue or red, the stronger the response (see Fig. 3 for quantitative response values). The delay between stimulus and the response shown was 30 ms (or four frames at 135 Hz). Data collected for a complete cycle of two different m-sequences were averaged together (taps $T = 45$ and 57; 8 min each). In all figures, the receptive fields are smoothed by 1/2 pixel. The 16 × 16 entire stimulus subtended 16 deg. Scale is in impulses/s (see left).

Fig. 2. Receptive-field movie (spatiotemporal first-order kernel) for the same Y-cell shown in Fig. 1. Stimulus response delays range from −7.4 to 103 ms (from left to right, top to bottom). Same scale as in Fig. 1 (see below).

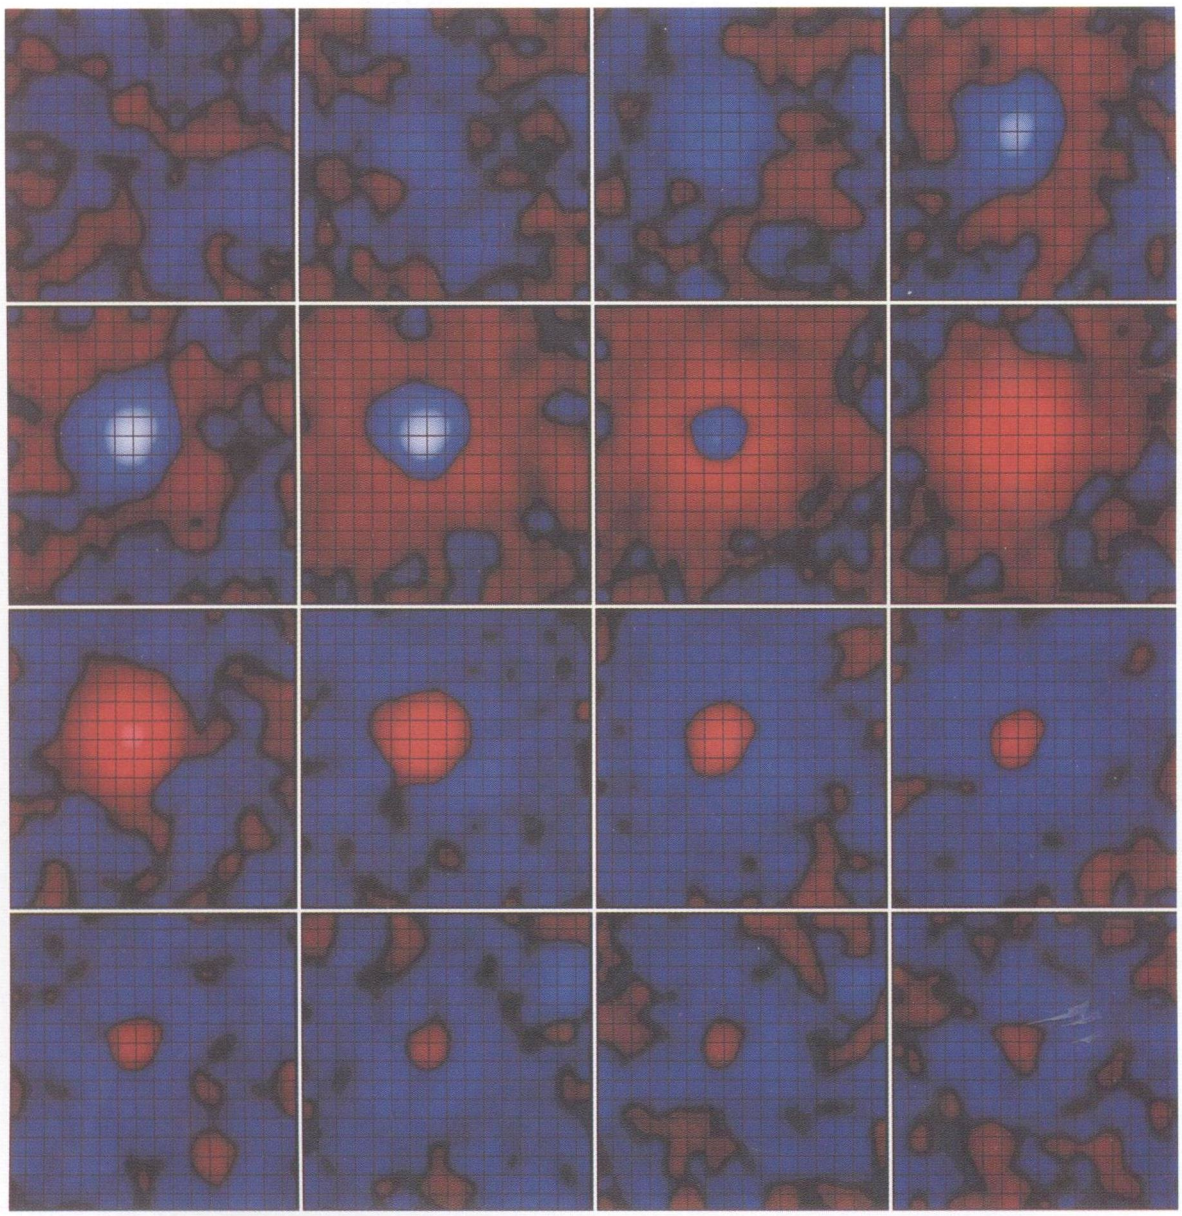

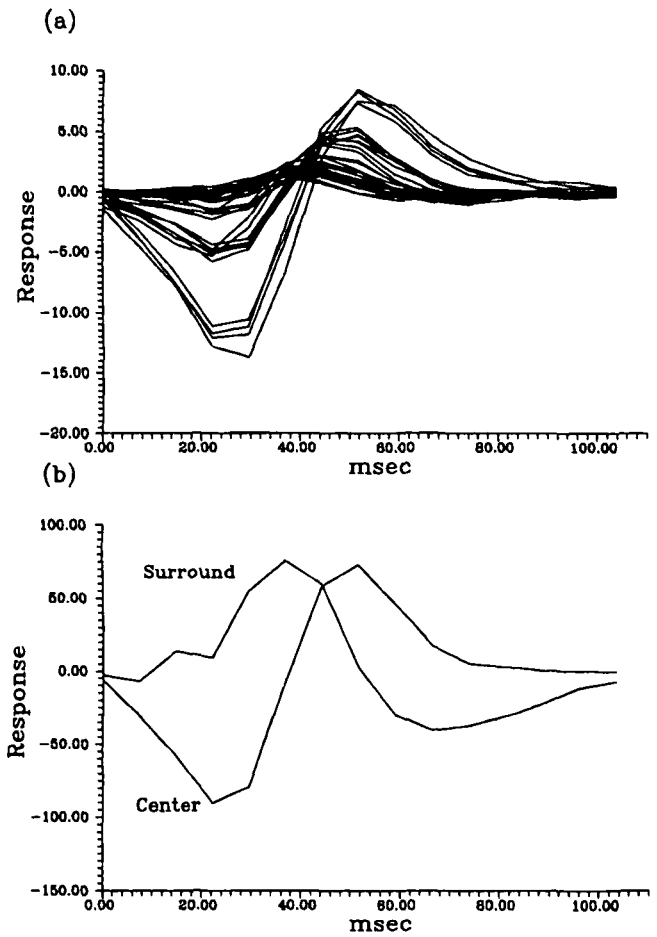

Fig. 3. (a) Temporal impulse responses for the same Y-cell for the 32 pixels that showed the strongest responses. Responses are measured in terms of the average increase in the firing rate, in impulses/s, following the light phase of the stimulus. (b) Summed impulse responses for the 11 pixels in the center and the 245 pixels in the surround.

region at the far left is very weak). Again, the first-order kernel is plotted so that red indicates ON response, blue OFF. The brighter (or whiter) the red or blue, the stronger the response. Loosely, this diagram represents the optimal stimulus that made the neuron fire. Again, this static response profile does not tell the whole story, since there is a different spatial receptive field for each delay between stimulus and response.

The receptive-field "movie" in Fig. 5 illustrates the spatiotemporal integration of this neuron. The response latency—the first delay for which there is a faint correlation between stimulus and response—is between 14.8 and 29.6 ms (the 14.8-ms bin), which is much longer than the 0.0 to 7.4 ms latency for the Y-cell. The duration of the response, at least 200 ms, is also significantly longer. Note that the frame duration of the stimulus is twice that used for the Y-cell: 14.8 *versus* 7.4 ms. This simple cell responded at the faster stimulus rate, but the receptive field—spread out over twice as many temporal bins—was significantly noisier.

The simple cell receptive field in Fig. 5 is clearly spatiotemporally inseparable: different response regions show different time courses. The most obvious example of this is the weak subregion at the far right which appears to have a longer latency than the other subregions. Again, spatiotemporal inseparability is better appreciated when the data are presented as temporal impulse response functions. A plot of the 32 strongest responses from individual pixels (Fig. 6a) shows the range of dynamics within the receptive field. Some pixels show fast responses that are clearly biphasic. Others are slower and nearly monophasic. This is more apparent when the responses from the pixels are summed separately along the eight columns that gave the strongest responses. The four largest sums (labeled 1, 3, 5, and 7 in Fig. 6b) show alternate ON and OFF responses with similar dynamics. The weak OFF region (number 1, on the far left in Fig. 5) is slightly slower than the others. More striking are the alternate columns (2, 4, 6, and 8 in Fig. 6c) that give weaker responses with a range of temporal waveforms, all of which are slower than those in Fig. 6b.

Fig. 7 shows the spatiotemporal receptive field summed along the vertical axis, with space along the one axis and time along the other. Cross sections along the time axis are equal to impulse responses shown in Fig. 6(b and c). As in the spatial receptive-field plots, ON responses are in red, OFF in blue. This representation emphasizes the clear spatial progression of the receptive field over time: it is oriented in space-time, a feature that implies a linear contribution to direction selectivity (Adelson & Bergen, 1985; Watson & Ahumada, 1985; Reid et al., 1987, 1991; McLean & Palmer, 1989; Albrecht & Geisler, 1991; Jagadeesh et al., 1993; DeAngelis et al., 1993*a*,*b*). As expected, this neuron was directionally selective in the direction predicted by the first-order kernel (from right to left).

Discussion

We have used the spatiotemporal m-sequence method of Sutter (1987) to map the first-order receptive fields of a Y-cell and a directional simple cell in the cat. While the m-sequence method is only one of several dense-noise stimuli that have been used to study the spatiotemporal structure of receptive fields (Citron et al., 1981; Emerson et al., 1987; Jacobson et al., 1993), it is probably the most efficient in terms of stimulus generation and data analysis.

The m-sequence approach provides a picture of receptive-field function that appears complete but, in fact, alterations of stimulus parameters determine which receptive-field structures are accentuated. The main variable is the "grain" with which the receptive field is probed, both in space and in time. The more pixels that are used to study a single receptive field and the faster the stimulus is modulated, the more detail is available in the linear receptive-field map. The tradeoff is that there is less effective stimulus energy. In either limit, a very fine or a very fast stimulus appears indistinguishable from a uniform gray. Practically, this means that the linear receptive field measured is lower in amplitude and closer to the noise.

In the spatial domain, increasing spatial resolution (per unit area) by a factor N, and shrinking the entire stimulus, will typically decrease signal size by a factor of N (assuming that the check size was small in comparison to the receptive-field center or surround—if it is larger than center or surround, then decreasing the check size might increase signal size). Assuming that the same m-sequence is used and that the neuron is linear, the signal-to-noise will also decrease by a factor of N, since the average number of spikes per bin will be unchanged.

In the temporal domain, in the limit of an interframe interval which is comparable to the integration time of the neuron, increasing temporal resolution by a factor N will decrease signal size by a factor of N (per unit time). Again assuming linearity, signal-to-noise will decrease by a factor of $N^{3/2}$. (There is an additional loss of a factor of $\sqrt{N}$ since the average

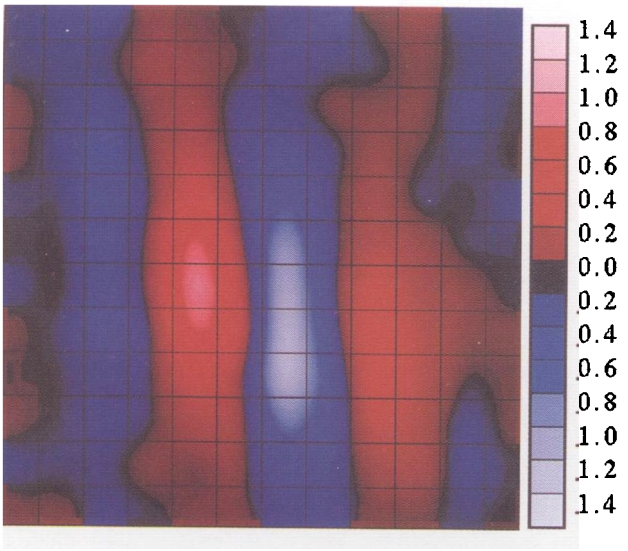

Fig. 4. Receptive-field representation (first-order kernels at a fixed time delay) for a layer 4 simple cell recorded in area 17. The delay between stimulus and response was 44 ms (or three frames at 67.5 Hz). Data collected for a complete cycle of two different m-sequences were averaged together (taps $T = 45$ and 57; 16 min each). The 16 × 16 entire stimulus subtended 16 deg, but only 12 × 12 pixels are shown. Scale is in impulses/s (see left).

Fig. 5. Receptive-field movie for the same simple cell shown in Fig. 4. Stimulus response delays range from −14.8 to 207 ms (from left to right, top to bottom). Same scale as in Fig. 4 (see below).

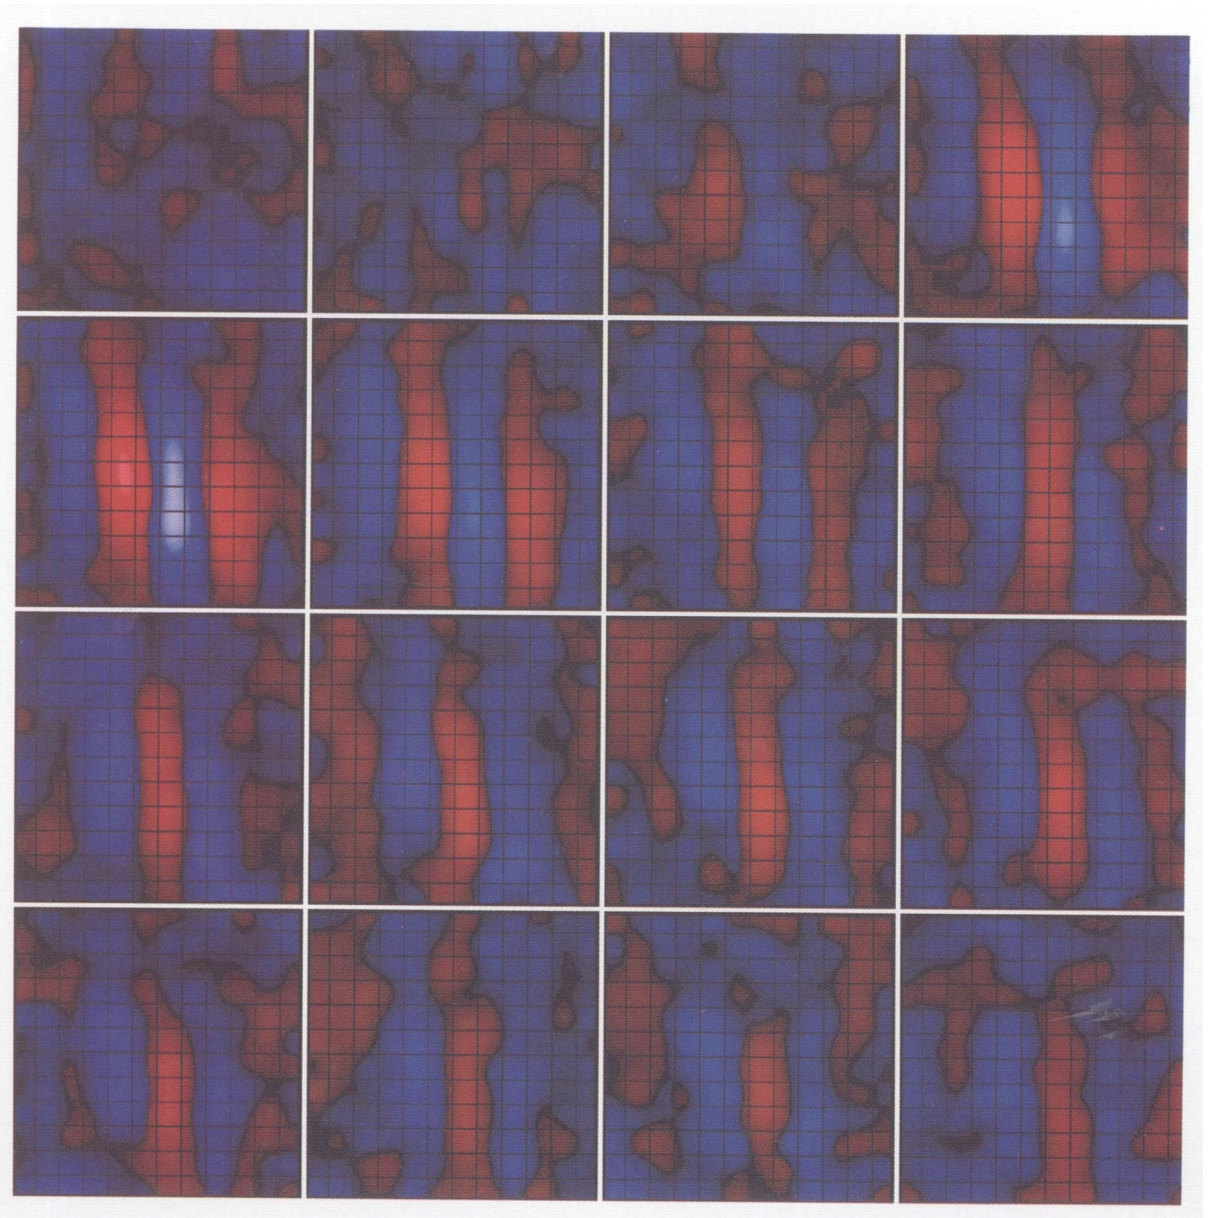

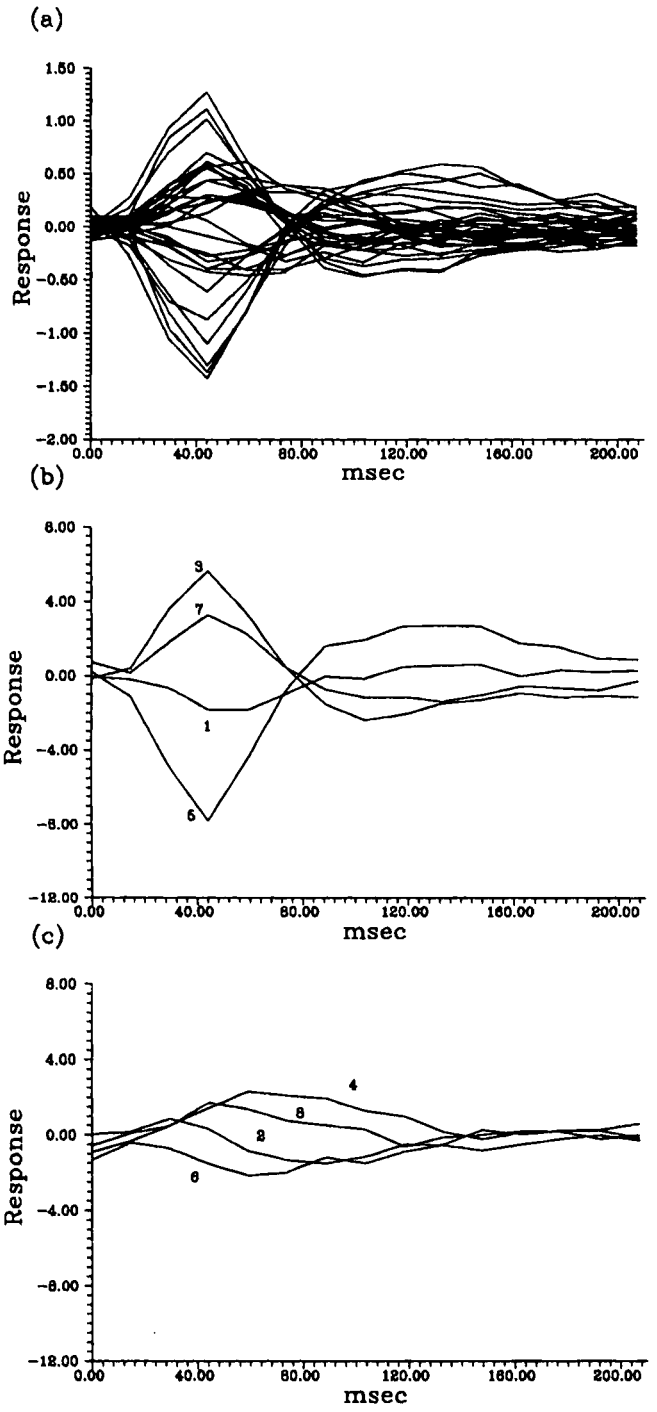

Fig. 6. (a) Impulse responses for the same simple cell for the 32 pixels that showed the strongest responses. (b) Summed impulse responses along the four vertical columns that gave the strongest responses. (c) Summed impulse responses along the interposed four vertical rows that gave somewhat weaker responses. Numbers refer to the columns, so that the positions can be determined (see Fig. 7).

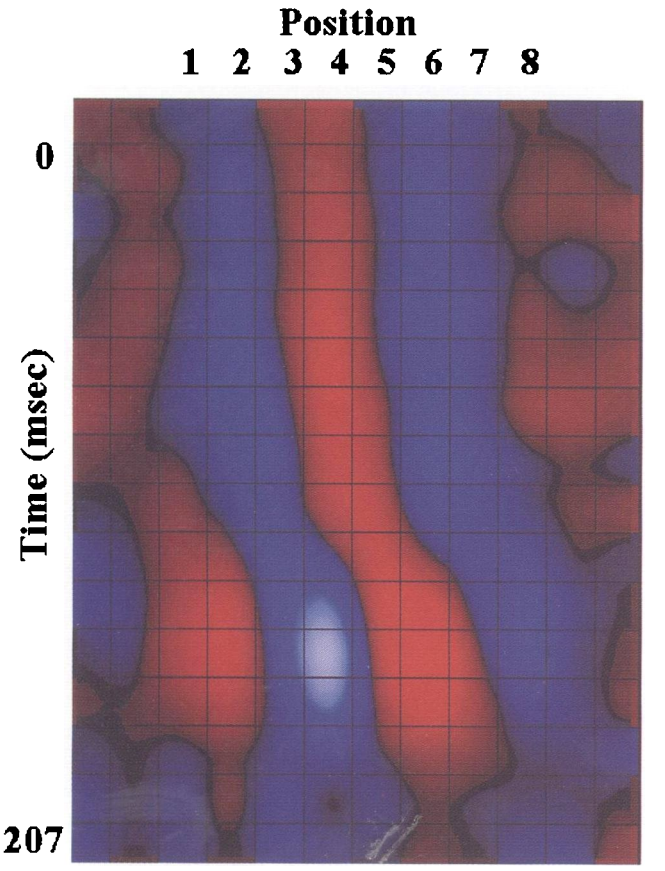

Fig. 7. Spatiotemporal plot of the simple receptive field, summed along the vertical axis, as in Fig. 6 (b and c). The spatial axis includes data from 12 columns of the stimulus (3/4 the screen), the temporal axis ranges from −14.8 to 207 ms. As for the spatial receptive fields, ON regions are shown in red, OFF in blue. The numbers correspond to the impulse responses shown in Fig. 6. Scale is five times greater than for Fig. 4 and 5 (−7 to 7 impulses/s).

number of spikes per bin will decrease by a factor of N.) However, since the m-sequence loop now runs N times faster, for the same amount of experiment time a factor of $\sqrt{N}$ will be gained back and the signal-to-noise will decrease by only a factor of N. All of this analysis assumes, of course, that the responses are stationary over time.

These analyses assume that the higher spatial or temporal resolution is achieved without changing the number of bins in the m-sequence. However, one may choose to increase the number of bins in the m-sequence in proportion to the resolution, in order to maintain the time between the temporal offsets $D(x, y)$ for each region, and thus maintain the same maximum lag time in each region's first-order response. This will change the minimal experiment duration and the signal-to-noise per m-sequence loop, but will not change the signal-to-noise per experiment time.

Finally, both the spatial analysis and the temporal analysis refer to estimates of single points on kernels. When kernels are used to predict model responses, they are typically summed. When spatial or temporal resolution is increased by a factor of N, a synthesis will use N-fold as many points. This partially mitigates (by a factor of $\sqrt{N}$) the decreases in signal-to-noise discussed above.

Choosing stimulus parameters always involves a tradeoff. For instance, the Y-cell data shown here were collected at a sufficiently fast frame rate (135 Hz) that the measured receptive field shows the detailed difference between the center and surround dynamics. The spatial scale, however, is such that the fine structure of the center mechanism (Levick & Thibos, 1980; Vidyasagar & Heide, 1984; Soodak et al., 1991) is not evident. If the pixels were any smaller, the weaker surround responses would be brought close to the noise.

The tradeoffs for more weakly responding cortical cells are always greater. For the simple cell shown, the spatial and temporal grain are nearly as coarse as possible: two pixels across each subregion, presented at 67 Hz. Receptive fields were also mapped at twice this frame rate, but the weakest subregions were barely distinguishable from the noise. Second-order responses decline more steeply with a decrease in stimulus energy, so the choice of parameters is even more critical if nonlinear interactions are of interest.

Our results on the linear receptive-field properties of the Y and the simple cell are not novel in themselves. Spatiotemporal inseparability of geniculate and cortical cells has been demonstrated in both the spatial domain (Citron et al., 1981 for *retinal* X-cells studied with white noise; McLean & Palmer, 1989 for simple cells studied with sparse-noise stimuli) and the spatial-frequency domain (Dawis et al., 1984, for geniculate X-cells; Movshon et al., 1978, Reid et al., 1987, 1991 for simple cells). Spatiotemporal dense-noise stimuli, however, have certain advantages over other mapping techniques. First, it is important to emphasize that the first-order response measured with dense-noise stimuli is an approximation to the first-order Wiener kernel (Marmarelis & Marmarelis, 1978; Victor & Knight, 1979). Dense-noise kernels may therefore be compared with the predicted Wiener kernels for the sorts of nonlinear networks found in real nervous systems (e.g. Victor & Shapley, 1979; Shapley & Victor, 1981; Emerson et al., 1987; Jacobson et al., 1993). The m-sequence method is a very efficient way of measuring these kernels, in terms of both stimulus generation and data analysis (Sutter, 1987). By contrast, it is justifiable to compare "kernels" measured with sparse-noise stimuli to the responses of theoretical models only when the measured system is known to be a linear system.

Finally, we have shown that linear analysis of m-sequence data provides a quantitative measure of classically described structures in Y-cell and simple-cell receptive fields. In general, detailed maps are created faster and with lower noise with the m-sequence technique than with sparse noise (Alonso et al., 1995). For this reason, the technique has proven useful in studies that rely on the detailed mapping of individual receptive fields (Reid & Shapley, 1992) or on the rapid mapping of multiple receptive fields simultaneously (Alonso & Reid, 1994; Reid & Alonso, 1995).

Acknowledgments

This work was supported by National Institutes of Health Grants 2 T32 GM 07739 and EY 10115 (R.C.R.), EY 1472 (R.M.S. & R.C.R.), EY 6871 (J.D.V. & R.C.R.), EY9314 (J.D.V.), EY 1428 (to Bruce Knight) and the Klingenstein Fund (J.D.V. & R.C.R.).

References

ADELSON, E.H. & BERGEN, J.R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America* **2**, 285–299.

ALBRECHT, D.G. & GEISLER, W.S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Visual Neuroscience* **7**, 531–546.

ALONSO, J.M. & REID, R.C. (1994). Coupling between neighboring LGN cells: Possible implications for simple receptive fields. *Society for Neuroscience Abstracts* **20**, 1476.

ALONSO, J.M., ATICK, J.J. & REID, R.C. (1995). The temporal responses of LGN receptive fields studied with white noise. *Investigative Ophthalmology and Visual Science* (Suppl.) **36**, S. 689.

BASELER, H.A., SUTTER, E.E., KLEIN, S.A. & CARNEY, T. (1994). The topography of visual evoked response properties across the visual field. *Electroencephalography and Clinical Neurophysiology* **90**, 65–81.

BENARDETE, E.A. & VICTOR, J.D. (1994). An extension of the m-sequence technique for the analysis of multi-input nonlinear systems. In *Nonlinear Vision: Determination of Neural Receptive Fields, Function, and Networks*, ed. PINTER, R. & NABET, B., pp. 87–110. Cleveland, Ohio: CRC Press.

BRITTEN, K.H. (1995). Spatial interactions within monkey middle temporal (MT) receptive fields. *Society for Neuroscience Abstracts* **21**, 663.

BUSSGANG, J.J. (1952). Crosscorrelation functions of amplitude-distorted Gaussian signals. *Technical Report 216, MIT Research Laboratory of Electronics.*

CITRON, M.C., KROEKER, J.P. & MCCANN, G.D. (1981). Nonlinear interactions in ganglion cell receptive fields. *Journal of Neurophysiology* **46**, 1161–1176.

DAWIS, S., SHAPLEY, R., KAPLAN, E. & TRANCHINA, D. (1984). The receptive field organization of X-cells in the cat: Spatiotemporal coupling and asymmetry. *Vision Research* **24**, 549–561.

DEANGELIS, G.C., OHZAWA, I., FREEMAN, R.D. (1993*a*). Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *Journal of Neurophysiology* **69**, 1091–1117.

DEANGELIS, G.C., OHZAWA, I., FREEMAN, R.D. (1993*b*). Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. *Journal of Neurophysiology* **69**, 1118–1135.

DEANGELIS, G.C., OHZAWA, I., FREEMAN, R.D. (1995). Receptive-field dynamics in the central visual pathways. *Trends in Neuroscience* **618**, 451–458.

DE BOER, E. & KUYPER, P. (1968). Triggered correlation. *IEEE Transactions on Biomedical Engineering* **15**, 169–179.

EMERSON, R. C., CITRON, M.C., VAUGHN, W.J. & KLEIN, S.A. (1987). Nonlinear directionally selective subunits in complex cells of cat striate cortex. *Journal of Neurophysiology* **58**, 33–65.

ENROTH-CUGELL, C. & ROBSON, J.G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology* **187**, 517–552.

GOLOMB, S.W. (1982). *Shift Register Sequences*. Laguna Hills, California: Aegean Park Press.

HOCHSTEIN, S. & SHAPLEY, R.M. (1976). Quantitative analysis of retinal ganglion cell classifications. *Journal of Physiology* **262**, 237–264.

HUBEL, D.H. (1957). Tungsten microelectrode for recording from single units. *Science* **125**, 549–550.

JACOBSON, L.D., GASKA, J.P., CHEN, H.-W. & POLLEN, D.A. (1993). Structural testing of multi-input linear-nonlinear cascade models for cells in the macaque striate cortex. *Journal of Physiology* **160**, 106–154.

JAGADEESH, B., WHEAT, H.S. & FERSTER, D. (1993). Linearity of summation of synaptic potentials underlying direction selectivity in simple cells of the cat visual cortex. *Science* **262**, 1901–1904.

JONES, J.P. & PALMER, L.A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology* **58**, 1187–1211.

KITANO, M., NIIYAMA, K., KASAMATSU, T., SUTTER, E.E. & NORCIA, A.M. (1994). Retinotopic and nonretinotopic field potentials in cat visual cortex. *Visual Neuroscience* **11**, 953–977.

LEVICK, W.R. & THIBOS, L.H. (1980). Orientation bias of cat retinal ganglion cells. *Nature* **286**, 389–390.

MARMARELIS, P.Z. & MARMARELIS, V.Z. (1978). *Analysis of Physiological Systems*. New York: Plenum Press.

MCLEAN, J. & PALMER, L. (1989). Contributions of linear spatiotemporal receptive field structure to velocity selectivity of simple cells in area 17 of cat. *Vision Research* **29**, 675–679.

MCLEAN, J., RAAB, S. & PALMER, L.A. (1994). Contribution of linear mechanisms to the specification of local motion by simple cells in areas 17 and 18 of the cat. *Visual Neuroscience* **11**, 271–294.

MILKMAN, N., SCHICK, G., ROSSETTO, M., RATLIFF, F., SHAPLEY, R. & VICTOR, J. (1980). A two-dimensional computer-controlled visual stimulator. *Behavior Research Methods and Instrumentation* **12**, 283–292.

MOVSHON, J. A., THOMPSON, I. D. & TOLHURST, D. J. (1978). Spatial summation in the receptive field of simple cells in the cat's striate cortex. *Journal of Physiology* **283**, 53–77.

NAKA, K.I., SAKURANAGA, M. & ANDO, Y.I. (1985). White-noise analysis as a tool in vision physiology. In *Progress in Clinical and Biological Research. Vol. 176, Contemporary Sensory Neurobiology*, ed. CORREIA, M.J. & PERACHIO, A.A., pp. 307–322. New York: Alan R. Liss.

NUTTAL, A.H. (1957). Invariance of correlation functions under nonlinear transformations. *Technical report, MIT Laboratory of Electronics.*

OHZAWA, I. & FREEMAN, R.D. (1986). The binocular organization of simple cells in the cat's visual cortex. *Journal of Neurophysiology* **56**, 221–242.

REID, R.C. & ALONSO, J.M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* **378**, 281–284.

REID, R.C. & SHAPLEY, R.M. (1992). The spatial structure of L, M, and S cone inputs to receptive fields in primate lateral geniculate nucleus. *Nature* **356**, 716–718.

REID, R.C., SOODAK, R.E. & SHAPLEY, R.M. (1987). Linear mechanisms of directional selectivity in simple cells of cat striate cortex. *Proceedings of the National Academy of Sciences of the U.S.A.* **84**, 8740–8744.

REID, R.C., SOODAK, R.E. & SHAPLEY, R.M. (1991). Directional selectivity and spatiotemporal structure of receptive fields of simple cells in cat striate cortex. *Journal of Neurophysiology* **66**, 505–529.

SAKAI, H.M, NAKA, K.-I. & KORENBERG, M.J. (1990). White-noise analysis in visual neuroscience. *Visual Neuroscience* **1**, 287–296.

SHAPLEY, R.M. & VICTOR, J.D. (1981). How the contrast gain control modifies the frequency responses of cat retinal ganglion cells. *Journal of Physiology* **318**, 161–179.

SHAPLEY, R.M. & VICTOR, J.D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *Journal of Physiology* **285**, 275–298.

SOODAK, R.E., SHAPLEY, R.M. & KAPLAN, E. (1991). Fine structure of receptive-field centers of X and Y cells of the cat. *Visual Neuroscience* **6**, 621–628.

SUTTER, E.E. (1987). A practical non-stochastic approach to nonlinear time-domain analysis. In *Advanced Methods of Physiological Systems Modeling, Vol. 1,* ed. MARMARELIS, V. Los Angeles, California: University of Southern California.

SUTTER, E.E. (1992*a*). The fast m-transform: A fast computation of cross-correlations with binary m-sequences. *SIAM Journal on Computing* **20**, 686–694.

SUTTER, E.E. (1992*b*). A deterministic approach to nonlinear systems analysis. In *Nonlinear Vision: Determination of Neural Receptive Fields, Function, and Networks*, ed. PINTER, R. & NABET, B., pp. 171–220. Cleveland, Ohio: CRC Press.

SUTTER E.E. & TRAN, D. (1992). The field topography of ERG components in man—I. The photopic luminance response. *Vision Research* **32**, 433–446.

SUTTER E.E. & VAEGAN. (1990). Lateral interaction component and local luminance nonlinearities in the human pattern ERG. *Vision Research* **30**, 659–671.

SZULBORSKI, R.G. & PALMER, L.A. (1990). The two-dimensional spatial structure of nonlinear subunits in the receptive fields of complex cells. *Vision Research* **30**, 249–254.

VICTOR, J.D. & SHAPLEY, R.M. (1979). Receptive field mechanisms of cat X and Y retinal ganglion cells. *Journal of General Physiology* **74**, 275–298.

VICTOR, J.D., SHAPLEY, R.M. & KNIGHT, B.W. (1977). Nonlinear analysis of cat retinal ganglion cells in the frequency domain. *Proceedings of the National Academy of Sciences of the U.S.A.* **74**, 3068–3972.

VICTOR, J.D. (1987). The dynamics of the cat retinal X cell centre. *Journal of Physiology* **386**, 219–246.

VICTOR, J.D. (1991). Asymptotic approach of generalized orthogonal functional expansions to Wiener kernels. *Annals of Biomedical Engineering* **19**, 383–399.

VICTOR, J.D. (1992). Nonlinear systems analysis in vision: Overview of kernel methods. In *Nonlinear Vision: Determination of Neural Receptive Fields, Function, and Networks*, ed. PINTER, R. & NABET, B., pp. 1–37. Cleveland, Ohio: CRC Press.

VICTOR, J.D., PURPURA, K., KATZ, E. & MAO, B. (1994). Population encoding of spatial frequency, orientation, and color in macaque V1. *Journal of Neurophysiology* **72**, 2151–2166.

VICTOR, J.D. & KNIGHT, B.W. (1979). Nonlinear analysis with an arbitrary stimulus ensemble. *Quarterly of Applied Mathematics* **37**, 113–136.

VIDYASAGAR, T.R. & HEIDE, W. (1984). Geniculate orientation biases seen with moving sine-wave gratings: Implications for a model of simple cell afferent connectivity. *Experimental Brain Research* **57**, 196–200.

WATSON, A.B. & AHUMADA, A.J., JR. (1985). Models of human visual-motion sensing. *Journal of the Optical Society of America* **2**, 322–342.

WEISS, T. (1966). A model of the peripheral auditory system. *Kybernetik* **3**, 153–175.

WIENER, N. (1958). *Nonlinear Problems in Random Theory.* New York: Wiley.