



Variance of high contrast textures is sensed using negative half-wave rectification

Charles Chubb^{a,*}, Jong-Ho Nam^b

^a *Department of Cognitive Sciences, Institute for Mathematical Behavioral Sciences, University of California at Irvine, Irvine, CA 92697-5100, USA*

^b *Department of Psychology, Sungshim Campus, Catholic University of Korea, Puchon, South Korea*

Received 6 October 1998; received in revised form 19 July 1999

Abstract

A rectifying transformation is required to sense variations in texture contrast. Various theoretical and practical considerations have inclined researchers to suppose that this rectification is full-wave, rather than half-wave. In the studies reported here, observers are asked to judge which of two texture patches has higher texture variance. Textures are composed of small squares, with each square being painted with one of nine different luminances. Different texture variances are achieved by manipulating the histograms of the texture patches to be compared. When the nine luminances range linearly from 0 to 160 cd/m², the transformation mediating judgments of texture variance takes the form of a negative half-wave rectifier: texture variance judgments are determined exclusively by the frequencies of luminances below mean luminance in the textures being compared. However, when the nine luminances range linearly from 60 to 100 cd/m², two of three observers use a full-wave rectifying transformation in making texture variance judgments; the third observer continued to use a negative half-wave rectifier. The unexpectedly asymmetric roles played by low versus high luminances in texture variance judgments suggest that the off-center system may play a dominant role in human perception of texture contrast. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Histogram contrast analysis; Texture; Second order processing

1. Introduction

One might guess that the color black would play a special role in visual processing. Although the level of light impinging on a given point of the retina can increase without bound, it cannot decrease below 0; in this sense, black is a visual absolute. That black is indeed treated specially by the visual system was shown by Whittle (1986) (see also Kingdom & Moulden, 1991). In this study, observers tried to detect which of two squares (0.93°), presented on a large background, was higher in luminance. The squares could be either brighter or darker than the background. As expected from previous results (Cornsweet & Teller, 1965; Leshowitz, Taub & Raab, 1968), for spot luminances near background luminance, sensitivity to the lumi-

nance difference conformed to a ‘dipper function’. That is, the threshold difference in luminance between the two squares at first decreased with deviation of square luminance from background, then began steadily to increase, following Weber’s law. For spots of luminance greater than the background, this pattern persisted across the range of spot luminances tested. However, spot luminances lower than the background yielded different results. Whittle discovered that, irrespective of background luminance, as spot luminances were decreased near to 0, observers became highly sensitive to small differences in luminance between the two spots. Evidently the visual system is equipped with special apparatus for processing luminances close to 0.

Despite these observations, models of texture processing have given no special priority to luminances near 0. Recent models (e.g. Bergen & Landy, 1991; Landy & Bergen, 1991) typically propose that texture-based judgments are mediated by a process in which (i) the visual system applies to the retinal input a battery

* Corresponding author. Tel.: +1-949-8241481; fax: +1-949-8242517.

E-mail address: cfchubb@uci.edu (C. Chubb)

of spatially local linear filters, variously tuned to spatial frequency and orientation, and (ii) rectifies the output from these filters. The rectified output of each filter thus provides a ‘neural image’ (Robson, 1980) in which high values signal high levels of a particular texture property. Researchers have tended to assume that the rectification used in these transformations is full-wave rather than half-wave. In particular, energy-based computations (e.g. Knutsson & Granlund, 1983; Bergen & Adelson, 1988) use a squaring (full-wave) transformation that discards information about white/black polarity in the input.

Here we document an unexpected luminance asymmetry in human texture perception. Specifically, in experiment 1 we show that when textures are composed of luminances spanning the range from 0 to 160 cd/m², perceived texture variance is highly sensitive to luminances between 0 and mid gray (or slightly lower than mid gray), and highly insensitive to luminances greater than mid gray. Two control experiments, 1.1 and (especially) 1.2, rule out the possibility that this performance results from the use of a nonlinear, luminance-sensing mechanism, combined with a variance-sensing mechanism, to perform the required judgments. On the other hand, as we show in experiment 2, when the luminances used to compose textures are restricted to a small range near middle gray (so that textures include no luminances near 0), then for two of our three observers, luminances above and below mean luminance play approximately symmetric roles in determining perceived texture variance.

2. Experiment 1

Patches of high variance texture are generally perceived as being higher in contrast than patches of low variance texture. This naturally leads to the presumption that perceived texture contrast is the subjective concomitant of some texture property akin to variance. This was our expectation prior to the experiment we now describe. The purpose of this experiment was to measure, for each luminance, v , the average impact $f(v)$ exerted on perceived texture variance by an occurrence in the texture of a texel (texture element) of luminance v . If the mechanism used by the observer to sense differences in physical texture variance were actually tuned to variance per se, then we should expect that occurrences in the texture of texels with extreme luminances (either white or black) should tend to increase perceived variance, whereas occurrences of mid-gray texels should tend to decrease perceived variance. In effect, the function f should be approximately parabolic, with its minimum at mid-gray.

2.1. Methods

2.1.1. Observers

The two authors (JN and CC) and one additional experienced psychophysical observer (EF) were used in this experiment. All observers had corrected-to-normal vision.

2.1.2. Apparatus

An IBM-compatible computer was used with an ATVista graphics system attached to an Ikegami DM516A monochrome monitor.

2.1.3. Linearization

Many monitors manifest spatial nonlinearities in the following sense: when two adjacent pixels are both assigned a given value v , the total amount of light emitted fails to be twice the amount of light emitted when a single one of those pixels is assigned v by itself. The Ikegami DM516A monitor is remarkably free from such spatial nonlinearities. Consequently, the following by-eye procedure suffices to achieve linearization. A fine grained checkerboard of luminances L_{hi} and L_{lo} was made to alternate in a coarse vertical square wave with bars of uniform luminance L_{mid} . The screen was then viewed from sufficiently far away that the fine granularity of the checkerboard was barely visible. At this distance, the square wave modulating between the two types of texture had a spatial frequency of approximately 4 c/deg. Since the texture itself could not be resolved, the square wave is visible only if the mean luminance of alternating texture bars is different. Thus, the luminance L_{mid} that makes the square wave vanish is equal to the average of the luminances L_{lo} and L_{hi} . We generated a lookup table by reiterating this procedure with different luminances L_{lo} and L_{hi} to determine, in each case, the L_{mid} midway between L_{lo} and L_{hi} .

2.1.4. Stimuli and task

All of the stimuli were composed of square texels, each subtending 6.68 min² at a viewing distance of 136 cm. Each texel was painted with one of the nine luminances in the set $\Omega = \{20i \text{ cd/m}^2 \mid i = 0, 1, \dots, 8\}$. A stimulus in a given trial comprised two abutting patches, each consisting of 60 rows by 30 columns of texels, yielding a total of 1800 texels per patch. Accordingly, we call any nonnegative, integer-valued function h of Ω a (patch) histogram if $\sum_{v \in \Omega} h(v) = 1800$. We write P_h for the texture patch with histogram h whose texel luminances are randomly permuted over the 60 row by 30 column array of texel locations. Thus, for any $v \in \Omega$, $h(v)$ gives the number of texels (with random locations) in P_h that are assigned luminance v .

The uniform histogram, assigning equal numbers of texels of each intensity to a given patch is

$$U(v) = 200 \quad \text{for all } v \in \Omega. \tag{1}$$

An integer-valued function ϕ of Ω is called a U -modulator if $U + \phi$ is a patch histogram. (Note that $\sum_{v \in \Omega} \phi(v) = 0$). In addition ϕ is called reversible if $U - \phi$ is also a patch histogram.

For observers CC and JN, conditions corresponded to the reversible U -modulators defined by the (first 9 columns of the) rows of Table 1. For EF, conditions corresponded to the U -modulators defined by the rows of Table 2. Each of CC and JN conducted 60 trials per

condition; EF conducted 200 trials per condition. On a trial using a given modulator, ϕ_i ($i = 1, 2, \dots, 15$), the observer attempted to judge which of patch $P_{U+\phi_i}$ versus $P_{U-\phi_i}$ had the greater sum of squared deviations from its mean luminance. Specifically, the observer received trial-by-trial feedback in attempting to judge which was greater:

$$SS_{U+\phi_i} = \sum_{v \in \Omega} (v - \mu_{U+\phi_i})^2 (U + \phi_i)(v)$$

vs.

Table 1
Data from Experiment 1 for observers CC and JN ^a

Modulator									CC		JN	
$\phi(0)$	$\phi(1)$	$\phi(2)$	$\phi(3)$	$\phi(4)$	$\phi(5)$	$\phi(6)$	$\phi(7)$	$\phi(8)$	C	I	C	I
30	7	-9	-18	-20	-18	-9	7	30	43	17	44	16
0	-15	-24	-26	-20	-11	6	30	60	24	36	12	48
0	22	19	1	-22	-37	-36	-7	60	45	15	33	27
50	-23	-24	-5	3	-5	-24	-22	50	49	11	45	15
-6	107	-45	-100	-21	64	28	-93	66	45	15	51	9
48	-70	91	-14	-110	-14	91	-70	48	49	11	44	16
23	50	-109	82	-21	-118	91	-35	37	42	18	39	21
31	-4	31	-98	80	-98	31	-4	31	54	6	51	9
60	30	6	-11	-20	-26	-24	-15	0	57	3	57	3
60	-7	-36	-37	-22	1	19	22	0	55	5	51	9
10	37	7	-31	-46	-31	7	37	10	48	12	46	14
66	-92	28	64	-23	-100	-45	108	-6	42	18	38	22
12	85	-109	-23	70	-23	-109	85	12	49	11	48	12
37	-35	91	-118	-21	82	-109	50	23	45	15	44	16
29	19	-49	62	-122	62	-49	19	29	47	13	48	12

^a The i th row defines and gives results for the i th experimental condition. The first nine columns of row i give the modulator ϕ_i used in the i th condition. On a trial from the i th condition, the observer attempted to judge which of $P_{U+\phi_i}$ vs. $P_{U-\phi_i}$ had greater texture variance. Column 10 (11) of row i gives the number of trials on which CC responded correctly (incorrectly) in condition i . Column 12 (13) gives the number of trials on which JN responded correctly (incorrectly).

Table 2
Data from Experiment 1 for observer EF ^a

Modulator									EF	
$\phi(0)$	$\phi(1)$	$\phi(2)$	$\phi(3)$	$\phi(4)$	$\phi(5)$	$\phi(6)$	$\phi(7)$	$\phi(8)$	C	I
22	6	-6	-13	-18	-13	-6	6	22	131	69
0	-11	-17	-19	-16	-8	5	22	44	55	145
0	17	14	1	-15	-28	-27	-6	44	122	78
37	-17	-18	-4	4	-4	-18	-17	37	126	74
14	28	-14	-31	-17	5	2	-17	30	125	75
26	-12	16	-12	-36	-12	16	-12	26	138	62
20	15	-28	9	-17	-35	16	-4	24	136	64
22	3	3	-31	7	-31	2	3	22	131	69
44	22	5	-8	-16	-19	-17	-11	0	186	14
44	-6	-27	-28	-15	1	14	17	0	139	61
7	28	5	-23	-34	-23	5	28	7	142	58
30	-17	2	5	-17	-31	-14	28	14	135	65
18	23	-28	-14	2	-14	-28	23	18	134	66
24	-4	16	-35	-17	9	-28	15	20	134	66
22	8	-15	4	-38	4	-15	8	22	135	65

^a The first nine columns of row i give the modulator ϕ_i used in the i th condition. Column 10 (11) of row i gives the number of trials on which EF responded correctly (incorrectly) in condition i .

$$SS_{U-\phi_i} = \sum_{v \in \Omega} (v - \mu_{U-\phi_i})^2 (U - \phi_i)(v), \quad (2)$$

where $\mu_{U+\phi_i}$ and $\mu_{U-\phi_i}$ are the mean luminances of the patches $P_{U+\phi_i}$ and $P_{U-\phi_i}$, respectively:

$$\begin{aligned} \mu_{U+\phi_i} &= \frac{1}{N} \sum_{v \in \Omega} (U + \phi_i)(v)v \quad \text{and} \\ \mu_{U-\phi_i} &= \frac{1}{N} \sum_{v \in \Omega} (U - \phi_i)(v)v, \end{aligned} \quad (3)$$

where $N = 1800$, the total number of texels in a patch.

On each trial, the observer fixated the center of a blank screen (80 cd/m²) and initiated the trial with a button press. The stimulus was then immediately displayed for 33 ms, after which the observer indicated with a button press which of the two patches, right versus left, appeared to have higher variance (in the sense of Eq. (2)). A feedback beep sounded after every incorrect response.

It was possible to give correctness feedback because each of the modulators ϕ_i , $i = 1, \dots, 15$, was constructed so that $SS_{U+\phi_i}$ was greater than $SS_{U-\phi_i}$. Indeed, the difference between $SS_{U+\phi_i}$ and $SS_{U-\phi_i}$ was selected so that, as assessed in pilot studies, it was a priori likely that performance would be in the threshold range (neither perfect, nor at chance) for that condition.

In addition, the modulators were chosen to span the space of all modulators, in order to be sure that the *perceived variance impact function* f could be estimated as a linear combination of ϕ_i , $i = 1, 2, \dots, 15$. For each $v \in \Omega$, $f(v)$ gives the mean impact exerted on the perceived variance of a texture patch by an occurrence in the patch of a texel of intensity v . It is this function f that our experiments aim to measure.

The experiment was conducted in blocks of 150 trials, in which each condition occurred ten times, with $P_{U+\phi_i}$ five times on the right, and five times on the left. Trials were randomly sequenced in the block. Observers CC and JN each performed six such blocks (60 trials per condition). EF performed 20 blocks.

2.2. Results and discussion

The results for CC and JN are shown in Table 1. For example, in Table 1, row i corresponds to the i th experimental condition: The first nine columns of row i define the modulator ϕ_i , with the j th column giving $\phi_i(20j \text{ cd/m}^2)$; the tenth column gives the number of correct responses by CC in condition i , and the 11th column gives the number of incorrect responses; the 12th and 13th columns give the correct and incorrect responses for JN. The results for EF are given in Table 2 using the same format.

2.2.1. Model

The obtained data were fit with a model under which the observer is assumed to

1. construct subjective estimates, $\overline{SS}_{U+\phi_i}$ and $\overline{SS}_{U-\phi_i}$ of $SS_{U+\phi_i}$ and $SS_{U-\phi_i}$ and

2. judge that $SS_{U+\phi_i} > SS_{U-\phi_i}$ just if $\overline{SS}_{U+\phi_i} - \overline{SS}_{U-\phi_i} + Y > 0$,

for Y a normal random variable with mean 0 and unspecified standard deviation.

The random variable Y is included to absorb trial-to-trial variability that does not depend on the modulator ϕ_i .

$\overline{SS}_{U+\phi_i}$ is assumed additively to combine random contributions from all the texels in patch $P_{U+\phi_i}$. For a given $v \in \Omega$, $P_{U+\phi_i}$ comprises exactly $(U + \phi_i)(v)$ randomly chosen texels with luminance v . It will be convenient to index these v -valued texels of $P_{U+\phi_i}$ (in arbitrary order) by $k = 1, 2, \dots, (U + \phi_i)(v)$, denoting the k th by $\tau_{v,k}$. Each of these v -valued texels is assumed to generate a random variable $X(\tau_{v,k})$. For a given v , the random variables $X(\tau_{v,k})$, $k = 1, 2, \dots, (U + \phi_i)(v)$, are assumed to be identically distributed with mean $f(v)$ and standard deviation $s(v)$. Moreover, all 1800 random variables generated by the texels of $P_{U+\phi_i}$ are assumed jointly independent. Finally, it is assumed that the subjective estimate $\overline{SS}_{U+\phi_i}$ of $SS_{U+\phi_i}$ is produced by summing the random variables generated by individual texels. However, we allow that texels in different locations of the stimulus may contribute with different weights (e.g. due to inhomogeneous distribution of attention). Thus for some nonnegative windowing function W mapping the set of all patch texels into \mathbb{R} , we assume

$$\overline{SS}_{U+\phi_i} = \sum_{v \in \Omega} \sum_{k=1}^{(U+\phi_i)(v)} W(\tau_{v,k}) X(\tau_{v,k}). \quad (5)$$

f is called the *perceived texture variance impact function* because it defines the mean impact exerted on perceived patch variance by an occurrence of a texel with a given luminance. The function s is called the *perceived variance noise injection function* because it defines the average amount of variability introduced into $\overline{SS}_{U+\phi_i}$ by an occurrence of a texel with a given luminance.

The plausibility of this model is discussed at length in Chubb (1999).

2.2.2. Interpreting the perceived texture variance impact function

What is the concrete meaning of the perceived texture variance impact function f ? It should first be noted that our methods allow us to determine f only up to an arbitrary positive scale factor and an arbitrary additive constant. That is, we can determine neither the mean value of f , nor the amplitude of f 's deviation from its

mean. All that we can determine are the relative deviations of f 's values from f 's mean value.

To understand the convention we use to scale f , it is useful to consider the following hypothetical texture variance judgment: imagine a texture patch Patch₁ with a uniform histogram, i.e. containing equal numbers of all nine luminances, randomly scrambled in a rectangular region. Now imagine producing another texture patch Patch₂, also with a uniform histogram, and then replacing a randomly chosen texel in Patch₂ with a texel of luminance v to produce a new patch Patch_{2,altered}. Then $f(v)$ reflects the expected difference in the perceived variance of Patch_{2,altered} compared to Patch₁.

The magnitude of $f(v)$ can be grasped in terms of the following hypothetical experiment. Imagine that the observer repeatedly judges which is higher in variance, Patch₁ versus Patch_{2,altered} (where patches are constructed independently on each trial). Suppose that $f(v) = 0.02$. This means that the alteration produces, on average, an increase in perceived patch variance equal to 0.02 standard deviations of the total noise degrading the observer's judgments. Thus, if $f(v) = 0.02$, then the observer will judge Patch_{2,altered} higher in variance than Patch₁ with probability slightly greater than 0.5. Specifically, the observer will judge Patch_{2,altered} higher in variance than Patch₁ with probability $\Phi(0.02) = 0.5080$, for Φ the standard normal cdf. On the other hand, if $f(v) = -0.013$, then the observer will judge Patch_{2,altered} higher in variance than Patch₁ with probability $\Phi(-0.013) = 0.4948$.

It is important to realize that the observer's goal is to synthesize a mechanism that is as effective as possible for making the required judgments of texture variance. This goal is not necessarily achieved by synthesizing a mechanism whose impact function is precisely parabolic in form (so as to be purely sensitive to texture variance). Rather, the goal is achieved by synthesizing a mechanism that is as sensitive as possible to texture variance. It may well be that the most sensitive available mechanism is also sensitive to other texture properties that are irrelevant to the variance judgment. In this case, the perceived texture variance impact function will deviate from the parabolic form to be expected from a mechanism tuned purely to texture variance.

2.2.3. Estimating the perceived variance impact function

For experiments using the design of the current studies, methods are provided in Chubb (1999) for obtaining a maximum likelihood estimate of f that is invariant with respect to all unmeasured model parameters. In particular, by requiring in each case the observer to compare $P_{U+\phi_i}$ versus $P_{U-\phi_i}$, one insures that the estimate of f is invariant with respect to the variance of Y (occurring in Eq. (4)) as well as the unmeasured standard deviations, $s(v)$, $v \in \Omega$.

For any guess f_{guess} at the impact function f , the likelihood $\mathcal{A}(f_{\text{guess}})$ gives the total probability of the obtained data under the assumption that $f_{\text{guess}} = f$. The process model sketched above ultimately predicts for each condition i , that, when presented with $P_{U+\phi_i}$ and $P_{U-\phi_i}$, the observer correctly judges $\text{SS}_{U+\phi_i} > \text{SS}_{U-\phi_i}$ with probability

$$\Phi(\phi_i \cdot f) = \Phi\left(\sum_{v \in \Omega} \phi_i(v) f(v)\right), \quad (6)$$

where Φ denotes the standard normal cdf.

Accordingly, \mathcal{A} is defined as follows. For each condition i , let k_i be the number of times the observer correctly judged $P_{U+\phi_i}$ greater in variance than $P_{U-\phi_i}$, and let n_i be the number of incorrect judgments in this same condition. Then under the model assumptions,

$$\mathcal{A}(f_{\text{guess}}) \approx \prod_i \Phi^{k_i}(f_{\text{guess}} \cdot \phi_i) (1 - \Phi(f_{\text{guess}} \cdot \phi_i))^{n_i}, \quad (7)$$

where the product ranges over all experimental conditions. The maximum likelihood estimate of f is the value of f_{guess} for which $\mathcal{A}(f_{\text{guess}})$ is maximized. The precise estimate depends upon the windowing function W (hence the approximate equality indicated in Eq. (7)); however, as discussed in Chubb (1999), this dependency is likely to be negligible in practice.

The reader is referred to Chubb (1999) for a more detailed discussion of the model, psychophysical methods and fitting procedures.

2.2.4. Testing the fits provided by the model

The model is characterized by the nine values of the impact function f . However, because the values of f are constrained to sum to 0, the model has only eight degrees of freedom.

For each observer, the fit provided by the model is assessed using a likelihood ratio test (e.g. Hoel, Port & Stone, 1971). A completely unconstrained model allocates a free parameter p_i to each experimental condition i , where p_i gives the probability that the observer judges correctly that $\text{SS}_{U+\phi_i} > \text{SS}_{U-\phi_i}$ when presented with patches $P_{U+\phi_i}$ and $P_{U-\phi_i}$. The maximum likelihood fit of this model to the data sets $p_i = k_i / (k_i + n_i)$ for each condition i , where k_i is the number of correct responses in condition i and n_i is the number of incorrect responses.

The constrained model defined by f is nested within the unconstrained model. That is, it is possible to express the unconstrained model as a function of a parameter set of which the parameters $f(0)$, $f(1)$, ..., $f(8)$ compose a subset. This is a precondition required for the likelihood ratio test. The likelihood ratio test compares the maximum likelihood \mathcal{A}_f of the nested model characterized by the impact function f to the maximum likelihood $\mathcal{A}_{\text{unconstrained}}$ of the unconstrained model. As shown by Wilks (1944) if the nested model

Table 3

Results of a likelihood ratio test of the model of Section 2.2.1, as applied to the data of Experiment 1^a

	Deg. of freedom	χ^2	P
CC	7	13.60	0.059
JN	7	9.31	0.231
EF	7	1.41	0.985

^a The model fits adequately for EF and JN, but the fit is questionable for CC.

captures the true state of the world, then the statistic $-2 \ln(A_f/A_{\text{unconstrained}})$ is asymptotically distributed as $\chi^2_{(v)}$ where the number of degrees of freedom v is equal to the number of free parameters in the unconstrained model minus the number in the nested model.

The results of this test are given in Table 3. The obtained P -values are all greater than 0.05. However, CC's P -value is only 0.059. This suggests that in making his judgments, CC may be sensitive to higher order interactions between texel luminances. To take a hypothetical example of such a higher order interaction: for luminances v and $v' \in \Omega$, it might be that the co-occurrence of high values of both $h(v)$ and $h(v')$ systematically elevates the perceived variance of patch P_h above

IMPACT OF TEXEL INTENSITY ON PERCEIVED TEXTURE VARIANCE




Fig. 1. Perceived variance impact functions (for three observers) obtained using a full range of texel luminances. The units of the vertical axis are standard deviations of the total noise compromising the observer's judgments. Thus, for example, the fact that for CC, $f(0) = 0.018$ indicates that, on average, an occurrence of a texel of luminance 0 increased the perceived variance of the patch in which it occurred by 0.018 standard deviations of the total noise. Note that impact functions take the form of negative half-wave rectifiers, rather than the full-wave rectifiers predicted by energy-based models of texture contrast perception.

what would be predicted merely from the impact function f . Such effects, if significant, require that interaction parameters be explicitly added to the model to adequately fit the data. We must therefore acknowledge the possibility that CC's impact function might change in form, were the model extended to include additional interaction parameters.

2.2.5. The perceived variance impact functions

The resulting perceived variance impact functions for the three observers are shown in Fig. 1. The units of the vertical axis are standard deviations of the total noise compromising the observer's judgments. Thus, for example, the fact that for CC, $f(0) = 0.018$ indicates that, on average, an occurrence of a texel of luminance 0 increased the perceived variance of the patch in which it occurred by 0.018 standard deviations of the total noise. The high values of $f(0)$ for all three observers indicate that, on average, a texel of luminance 0 tends to produce a maximal increase in perceived patch variance. For each of observers CC and JN, induced impact decreases linearly over the range of luminances from 0 up to 60 cd/m^2 (slightly below mid gray). For EF induced impact decreases linearly from 0 up to 80 cd/m^2 (mid gray). Most interestingly, however, for CC and JN (EF) all luminances greater than or equal to 60 cd/m^2 (80 cd/m^2) exert approximately equal impact on perceived patch variance.

2.2.6. Implications

In our introduction, we implied that if the observer were veridically sensing texture variance, then the perceived variance impact function f should be parabolic, with its minimum at mid-gray. To be precise, a veridical sensor of texture variance would use an impact function of the form $k_1\psi + k_2$, for constants k_1 and k_2 and

$$\psi(v) = v^2 - 8v + 9\frac{1}{3}, \quad v = 0, 1, \dots, 8. \quad (8)$$

(For $\mu_U = 4$, the mean of U , $\psi(v)$ is equal to $(v - \mu_U)^2 - 6\frac{2}{3}$, where, for convenience, we include the last term to make ψ sum to 0). In particular, as we prove in the appendix, for any U -modulator ϕ , $\text{SS}_{U+\phi} - \text{SS}_{U-\phi} = 2\psi \cdot \phi$, for

$$\psi \cdot \phi = \sum_{v \in \Omega} \phi(v) \psi(v). \quad (9)$$

The fact that the difference between $\text{SS}_{U+\phi}$ versus $\text{SS}_{U-\phi}$ does not depend on the difference in mean luminance between $P_{U+\phi}$ versus $P_{U-\phi}$ may be somewhat surprising. After all, $\text{SS}_{U+\phi}$ depends not only on $\phi \cdot \psi$, but also on the mean luminance $\mu_{U+\phi}$ of patch $P_{U+\phi}$. Similarly, $\text{SS}_{U-\phi}$ depends on both $\phi \cdot \psi$ as well as $\mu_{U-\phi}$. However, these dependencies of $\text{SS}_{U+\phi}$ and $\text{SS}_{U-\phi}$ on the respective patch mean luminances $\mu_{U+\phi}$ and $\mu_{U-\phi}$ conveniently cancel when one takes the difference $\text{SS}_{U+\phi} - \text{SS}_{U-\phi}$.

The perceived variance impact functions obtained in Fig. 1 are dramatically different in form from ψ . On the surface, the form of the perceived variance impact function suggests that the detection of texture variance is mediated not by a parabolic function akin to ψ , but rather by a negative half-wave rectifying function.

However, before proceeding to a more thorough discussion of these results, we must address a plausible objection. An examination of the modulators defining the conditions in Tables 1 and 2 reveals that all modulators other than those in rows 2 and 9 of Tables 1 and 2 are orthogonal to the linearly increasing modulator λ that assigns

$$\lambda(v) = v - 4, \quad v = 0, 1, \dots, 8. \quad (10)$$

However, it is obvious from the plots in Fig. 1 that the impact functions of all three observers have a strong negative correlation with the modulator λ . This is revealed by fact that all three observers do extremely poorly in the condition corresponding to row 2 of Table 1 (for CC and JN) or Table 2 (for EF), whereas all observers do extremely well in the condition corresponding to row 9. Indeed, these are the only conditions in which performance differs dramatically from the baseline condition shown in row 1. In this baseline condition, the histograms of the textures to be compared are modulated purely (up to rounding error) by $K\psi$ versus $-K\psi$, for histogram amplitude K selected

(on the basis of pilot experiments) to yield performance near threshold. In the condition of row 2, the distribution of the high variance texture is being modulated away from uniformity by $K(\psi + \lambda)$, whereas the low variance texture is being modulated away from uniformity by $-K(\psi + \lambda)$. In the condition of row 9, the distribution of the high variance texture is being modulated away from uniformity by $K(\psi - \lambda)$, whereas the low variance texture is being modulated away from uniformity by $-K(\psi - \lambda)$. The performance of all observers in these crucial conditions indicates that the impact function correlates very well with $\psi - \lambda$ and quite poorly with $\psi + \lambda$, suggesting that the impact functions of all observers are strongly negatively correlated with λ .

Note, however, that on 14 out of every 15 trials the observer was receiving feedback that would tend to confirm him in his use of an impact function with a strongly negative linear component. On only one out of every 15 trials (the trials in the condition corresponding to row 2) did the observer experience any penalty as a result of using an impact function with a large, negative linear component.

Thus, although it is true that observers used a negative half-wave rectifying impact function in experiment 1, the possibility remains that other options were available to them. In particular, as illustrated in Fig. 2, suppose observers possess two up-front transformations

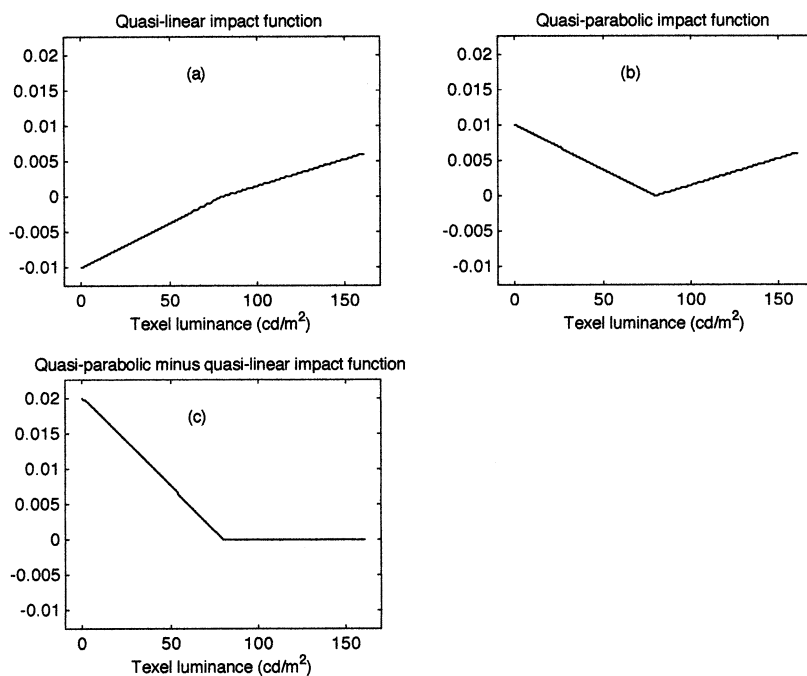


Fig. 2. How texture variance judgments might be made using a quasi-linear mechanism in combination with a quasi-parabolic (full-wave rectifying) mechanism. Suppose that observers possess two up-front transformations that are sensitive to texture variance: (a) a quasi-linear mechanism that responds slightly negatively to texture variance due to its (compressive) deviation from perfect linearity, and (b) a quasi-parabolic mechanism that is the primary source of sensitivity to texture variance. (c) To maximize their overall sensitivity to texture variance, observers might base their judgments on the difference between the responses of the quasi-parabolic and quasi-linear mechanisms. Such a strategy might well yield impact functions very similar to those obtained in experiment 1.

that are sensitive to texture variance: (a) a quasi-linear mechanism, cartooned in Fig. 2a, that responds slightly negatively to texture variance due to its (compressive) deviation from perfect linearity, and (b) a quasi-parabolic mechanism, cartooned in Fig. 2b, that is the primary source of sensitivity to texture variance (the hypothesized quasi-linear mechanism might result from taking the difference between on-center and off-center channel outputs, whereas the quasi-parabolic mechanism might result from taking the sum of on-center and off-center channel outputs).

To maximize their overall sensitivity to texture variance, observers might base their judgments on the difference, cartooned in Fig. 2c, between the responses of the quasi-parabolic and quasi-linear mechanisms. Such a strategy might well yield impact functions very similar to those obtained in experiment 1. Note that even though the impact function resulting from this strategy would deviate strongly from the ideal parabolic form, it might nonetheless optimize performance under the conditions of experiment 1 for the following reason: in experiment 1, observers are penalized for using the quasi-linear mechanism in only one condition out of 15; in all other conditions, taking the difference between the hypothesized quasi-parabolic and quasi-linear mechanisms would slightly improve performance above what could be achieved using the quasi-parabolic mechanism alone.

Perhaps, then, the wildly non-parabolic impact functions of Fig. 1 are an artifact of an ill-chosen set of experimental conditions. To investigate this possibility, we conducted two control experiments. In the first control study, we repeated the experiment for two of our three observers, using a new, randomly generated set of modulators, whose correlations with λ varied randomly. In the second (more stringent) control study, we tested two naïve observers, using a new set of modulators. The modulators in this second control study all contained large, carefully controlled quantities of λ ; however, the modulator set was balanced in that half the modulators correlated positively with λ , whereas half correlated negatively.

3. Experiment 1.1

3.1. Procedure

New modulators ϕ_i were generated as follows: seven random, mutually orthogonal modulators ξ_i ($i=1,2, \dots, 7$) were first generated, each of which was also orthogonal to ψ . Each ξ_i was then scaled to have a maximum absolute value equal to that of ψ . Then the 15 modulators used in the control conditions were produced as follows: ϕ_1 is set to $K\psi$, where the scale factor K is chosen so that performance at judging which

of $P_{U+K\psi}$ versus $P_{U-K\psi}$ has greater variance is approximately at threshold. For $i=2,3, \dots, 8$, ϕ_i was set to $K(\psi + \xi_i)$, and for $i=9,10, \dots, 15$, ϕ_i was set to $K(\psi - \xi_i)$.

The resulting modulators ϕ_i , $i=1, \dots, 15$, are given by columns 1–9 of the first 15 rows of Table 4. The second 15 rows of Table 4 are identical to the first 15 rows, except that all modulators have been scaled by a factor of 2/3 (ignoring rounding error) in order to increase the difficulty of the judgments. The 15 modulators of Table 5 are identical to the last 15 modulators of Table 4; only these modulators were used for observer EF. The method of construction insures that equally many of the modulators used in this study correlate positively with λ as correlate negatively.

For observer EF, stimuli were presented, as in experiment 1, in a totally mixed design. A single block contained 150 trials, ten of each condition. For CC five blocks were run in which the conditions given in the top 15 rows of Table 4 were mixed. Then five blocks were run in which the bottom 15 conditions were mixed.

3.2. Results and discussion

The results are shown in Tables 4 and 5. In Table 4 (Table 5), row i corresponds to the i th experimental condition for CC (EF); the tenth column gives the number of correct responses by CC (EF) in condition i , and the 11th column gives the number of incorrect responses.

The impact functions derived from these data are given in Fig. 3. As in the original experiment, the impact functions approximate negative, half-wave rectifiers. This approximation is classical in form for EF. His impact function decreases linearly over the luminances from 0 to 80 cd/m², then remains flat from 80 to 160 cd/m². For CC, the impact function deviates from the classical half-wave rectifier: his curve is essentially flat for all intensities except the two lowest. In essence, CC's judgments of texture variance depend only on the relative proportions in the textures being compared of the very darkest texture elements.

Although these findings support the results of experiment 1, they are far from conclusive. First, the individual modulators in experiment 1.1 tend to correlate much more weakly with λ than with ψ . Thus, as in experiment 1, there is hardly any penalty to using a variance-sensing mechanism that correlates fairly strongly with λ . Moreover, CC and EF had previously served as observers in experiment 1. They may well have carried over the decision procedure they had previously learned in experiment 1 to make the judgments required in experiment 1.1. To properly secure the result suggested by experiment 1, the following, more telling control study was conducted.

Table 4
Data from Experiment 1.1 for observer CC^a

Modulator									CC	
$\phi(0)$	$\phi(1)$	$\phi(2)$	$\phi(3)$	$\phi(4)$	$\phi(5)$	$\phi(6)$	$\phi(7)$	$\phi(8)$	C	I
30	7	-9	-18	-20	-18	-9	7	30	47	13
43	29	-14	-29	10	-29	-24	-3	17	53	7
4	16	1	-14	-12	-33	-10	-12	60	41	19
40	-23	7	-29	-3	-37	6	8	31	42	18
16	11	-24	-12	-13	-38	-6	37	29	48	12
41	-14	-20	12	-13	-16	-35	5	40	43	17
15	-2	5	-47	-11	6	-39	31	42	39	21
32	14	21	-7	-34	-33	-25	14	18	50	10
17	-14	-3	-8	-52	-8	7	18	43	39	21
56	-1	-19	-23	-30	-3	-7	27	0	54	6
20	38	-25	-7	-39	0	-23	7	29	50	10
44	4	6	-25	-28	2	-11	-23	31	55	5
19	29	3	-48	-31	-20	18	10	20	55	5
45	17	-22	11	-32	-42	21	-16	18	51	9
28	1	-39	-29	-8	-4	8	1	42	44	16
20	5	-6	-12	-14	-12	-6	5	20	46	14
29	19	-9	-19	6	-19	-16	-2	11	50	10
3	11	1	-9	-9	-22	-7	-8	40	38	22
27	-15	5	-19	-3	-25	4	5	21	42	18
11	7	-16	-8	-9	-25	-4	25	19	37	23
27	-9	-13	8	-9	-11	-23	3	27	45	15
10	-1	3	-31	-8	4	-26	21	28	41	19
21	9	14	-5	-21	-22	-17	9	12	38	22
11	-9	-2	-5	-36	-5	5	12	29	32	28
37	-1	-13	-15	-19	-2	-5	18	0	49	11
13	25	-17	-5	-25	0	-15	5	19	42	18
29	3	4	-17	-19	1	-7	-15	21	41	19
13	19	2	-32	-21	-13	12	7	13	45	15
30	11	-15	7	-20	-28	14	-11	12	49	11
19	1	-26	-19	-6	-3	5	1	28	41	19

^a The first nine columns of row i give the modulator ϕ_i used in the i th condition. Column 10 (11) of row i gives the number of trials on which CC responded correctly (incorrectly) in condition i .

Table 5
Data from Experiment 1.1 for observer EF^a

Modulator									EF	
$\phi(0)$	$\phi(1)$	$\phi(2)$	$\phi(3)$	$\phi(4)$	$\phi(5)$	$\phi(6)$	$\phi(7)$	$\phi(8)$	C	I
20	5	-6	-12	-14	-12	-6	5	20	63	37
29	19	-9	-19	6	-19	-16	-2	11	84	16
3	11	1	-9	-9	-22	-7	-8	40	59	41
27	-15	5	-19	-3	-25	4	5	21	61	39
11	7	-16	-8	-9	-25	-4	25	19	52	48
27	-9	-13	8	-9	-11	-23	3	27	64	36
10	-1	3	-31	-8	4	-26	21	28	53	47
21	9	14	-5	-21	-22	-17	9	12	87	13
11	-9	-2	-5	-36	-5	5	12	29	53	47
37	-1	-13	-15	-19	-2	-5	18	0	81	19
13	25	-17	-5	-25	0	-15	5	19	74	26
29	3	4	-17	-19	1	-7	-15	21	78	22
13	19	2	-32	-21	-13	12	7	13	68	32
30	11	-15	7	-20	-28	14	-11	12	76	24
19	1	-26	-19	-6	-3	5	1	28	49	51

^a The first nine columns of row i give the modulator ϕ_i used in the i th condition. Column 10 (11) of row i gives the number of trials on which EF responded correctly (incorrectly) in condition i .

INFLUENCE OF DIFFERENT LUMINANCES ON PERCEIVED TEXTURE VARIANCE

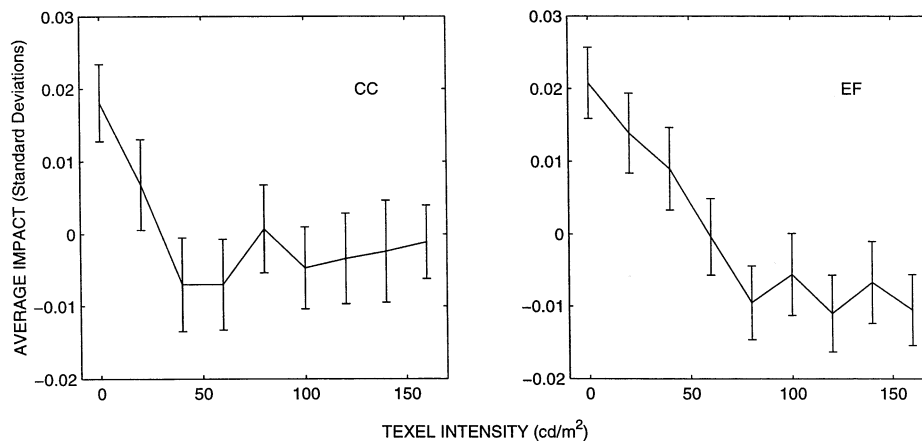


Fig. 3. Perceived variance impact functions (for two observers) obtained using a full range of texel luminances and a randomly constructed set of modulators. The units of the vertical axis are standard deviations of the total noise compromising the observer's judgments. Note that, as in experiment 1, impact functions take the form of negative half-wave rectifiers.

4. Experiment 1.2

4.1. Procedure

Two experienced psychophysical observers, both naïve to the purpose of the current experiment, were recruited. They were first given 480 trials of practice in judging which had greater variance, $P_{U+K\psi}$ versus $P_{U-K\psi}$, for eight values of K chosen so that some judgments were quite difficult, while others were relatively easy, with most values of K supporting performance in the neighborhood of threshold. The purpose of this training phase was to insure that observers had refined their method of judging unadulterated differences in variance before entering the phase of the experiment in which differences in texture variance were modified by other textural differences.

Following this training phase, the observer made 40 judgments in each of 27 conditions. The modulators ϕ_i , $i = 1, 2, \dots, 27$ were constructed as follows. A set of six random, mutually orthogonal modulators ξ_i ($i = 1, 2, \dots, 6$) was first generated, each of which was also orthogonal to each of λ and ψ . Each of ξ_i was then scaled to have a maximum absolute value equal to that of ψ . Then the 27 modulators used in the control conditions were produced as follows: ϕ_1 is set to $K\psi$, for K chosen so that the observer judged the variance of $P_{U+K\psi}$ greater than that of $P_{U-K\psi}$, with probability approximately 0.75. ϕ_2 was set to $K(\psi + \frac{1}{2}\lambda)$, and ϕ_3 was set to $K(\psi - \frac{1}{2}\lambda)$. Then for $i = 4, 5, \dots, 9$, $\phi_i = K(\psi + \frac{1}{2}\lambda + \xi_i)$, for $i = 10, 11, \dots, 15$, $\phi_i = K(\psi + \frac{1}{2}\lambda - \xi_i)$, for $i = 16, 17, \dots, 21$, $\phi_i = K(\psi - \frac{1}{2}\lambda + \xi_i)$, and for $i = 22, 23, \dots, 27$, $\phi_i = K(\psi - \frac{1}{2}\lambda - \xi_i)$. This construction yields a set of modulators ϕ_i such that

1. for all 27 conditions ($i = 1, 2, \dots, 27$), $\phi_i \cdot \psi = K$;

2. for the 13 conditions $i = 2, 4, 5, \dots, 15$, $\phi_i \cdot \lambda = K/2$;

3. for the 13 conditions, $i = 3, 16, 17, \dots, 27$,

$$\phi_i \cdot \lambda = -K/2.$$

The resulting modulators are defined by the first 9 columns of Table 6, with $\phi_i(j)$ given by the j th column of row i of Table 6. During this phase of the experiment, the observer continued to receive trial-by-trial feedback as to the correctness of texture variance judgments. Ten blocks, each comprising 108 trials (four from each condition), were conducted.

All pairs of textures except those from condition 1 differ decisively in luminance; however, these differences in luminance are uncorrelated with the correct response. Suppose that in experiment 1 observers were using the difference between a quasi-parabolic and a quasi-linear mechanism (as illustrated in Fig. 2) to generate their responses. In the current experiment performance will suffer under this strategy for the following reason: on each trial, the response of the hypothesized quasi-linear mechanism will be dominated by the relatively large λ component of the given modulator; however, the sign of this λ component is uncorrelated with the correct response. Thus, use of the quasi-linear mechanism will only introduce noise into observers' judgments. For this reason, the choice of modulators in the current experiment should force observers to eradicate dependence of their variance judgments on texture luminance (λ -content) if they possibly can.

4.2. Results and discussion

In the initial practice session, both observers attained a level of proficiency comparable to that attained by the observers in experiment 1. Results for the main part of

Table 6
Data from Experiment 1.2 for observers JH and DB ^a

Modulator									JH		DB	
$\phi(0)$	$\phi(1)$	$\phi(2)$	$\phi(3)$	$\phi(4)$	$\phi(5)$	$\phi(6)$	$\phi(7)$	$\phi(8)$	C	I	C	I
30	8	-9	-18	-21	-18	-9	7	30	30	10	21	19
15	-4	-16	-22	-22	-14	-1	19	45	23	17	16	24
45	19	-1	-14	-22	-22	-16	-4	15	30	10	22	18
-1	20	-22	-30	8	-30	-15	11	59	23	17	22	18
1	14	-4	-25	-35	-42	29	26	36	18	22	18	22
26	-34	0	-21	-3	-35	12	0	55	26	14	22	18
2	-2	-6	8	-39	-18	-13	5	63	27	13	10	30
6	0	14	-50	-29	4	-4	6	53	24	16	18	22
12	-14	6	-21	-18	-27	-25	49	38	21	19	17	23
31	-27	-11	-14	-52	2	13	27	31	23	17	19	21
29	-21	-28	-19	-8	13	-31	11	54	24	16	20	20
4	26	-33	-23	-39	6	-14	38	35	26	14	16	24
28	-6	-26	-52	-2	-11	10	32	27	23	17	18	22
24	-8	-46	6	-13	-33	2	31	37	25	15	25	15
18	6	-38	-23	-25	-2	23	-11	52	23	17	22	18
29	42	-7	-22	9	-38	-30	-12	29	31	9	22	18
31	36	11	-18	-35	-49	14	4	6	35	5	20	20
56	-11	15	-14	-2	-43	-3	-23	25	30	10	22	18
32	21	9	16	-40	-26	-28	-17	33	26	14	26	14
36	23	29	-42	-31	-3	-19	-16	23	27	13	23	17
42	9	21	-13	-18	-35	-40	26	8	31	9	26	14
61	-5	4	-7	-50	-6	-2	4	1	26	14	18	22
59	1	-13	-11	-8	5	-46	-11	24	33	7	24	16
34	49	-18	-15	-40	-1	-29	15	5	34	6	25	15
58	17	-11	-44	-4	-18	-5	10	-3	30	10	25	15
54	15	-31	13	-13	-41	-13	9	7	33	7	23	17
48	29	-23	-16	-25	-9	8	-34	22	32	8	25	15

^a The first nine columns of row i give the modulator ϕ_i used in the i th condition. Column 10 (11) of row i gives the number of trials on which JH responded correctly (incorrectly) in condition i . Column 12 (13) of row i gives the number of trials on which DB responded correctly (incorrectly) in condition i .

experiment 1.2 are given in columns 10–13 of Table 6. Column 10 of row i gives the number correct in the condition using modulator ϕ_i for observer JH; column 11 gives the number incorrect. Column 12 gives the number correct for observer DB, and column 13 gives the number incorrect.

Both observers spontaneously noted that the task became much harder when the actual experimental trials commenced, following the practice trials. Thus, the introduction of uninformative luminance differences between the textures to be compared made the task more difficult, suggesting that the mechanism observers adopted for judging texture variance during the initial practice session was partially sensitive to texture luminance.

Moreover, during the actual experiment, observers were unable to eradicate this sensitivity to luminance from the mechanisms they used to judge texture variance. This is shown first by the fact that observers showed no tendency toward improvement over the ten blocks of the experiment. Second, as in experiment 1, the forms of the perceived texture variance impact functions for observers JH and DB (Fig. 4) remain

elevated for the lowest luminances, and flatten out over all higher luminances (although this pattern is shown only weakly in the impact function for DB). It should also be noted that these impact functions are much lower in amplitude than those obtained in experiments 1 and 1.1. This indicates that the impacts exerted on perceived texture variance by texels of different luminances were generally reduced in effectiveness (relative to the noise compromising performance) compared to the prior experiments.

If our observers had access to a full-wave rectifying mechanism (such as that illustrated in Fig. 2b) for purposes of judging texture variance, then they should have used it in the current experiment. Such a full-wave rectifying mechanism would have been immune to the uninformative texture luminance differences introduced, and would thus have enabled significant improvement in performance. The data strongly suggest that observers did not have access to such a full-wave rectifying mechanism. On the contrary, neither observer was able to fashion a mechanism that was very effective at performing this task. Moreover, the impact functions of the mechanisms they ended up using continued to approximate negative half-wave rectifiers.

INFLUENCE OF DIFFERENT LUMINANCES ON PERCEIVED TEXTURE VARIANCE




Fig. 4. Perceived variance impact functions (for two naïve observers) obtained using a set of modulators systematically constructed to enable observers to exclude texture luminance as a useful cue for making texture variance judgments. The units of the vertical axis are standard deviations of the total noise compromising the observer's judgments. Note that even in this case, observers persist in using negative half-wave rectifying impact functions, although the response of a negative half-wave rectifying mechanism is strongly influenced by texture luminance. This suggests strongly that the most sensitive device observers can synthesize for purposes of judging texture variance uses a negative half-wave rectifying transformation.

We conclude that under the conditions of these experiments, texture variance judgments do not reflect texture variance per se; rather, such judgments use a system that applies a negative half-wave rectifier to the input. For most observers (CC, JN in experiment 1, CC in experiment 1.1, JH and DB in experiment 1.2) this negative half-wave rectifier has a non-zero threshold. However, for EF in experiments 1 and 1.1, the rectifier is classical in form, with a linear decrease (over the entire range of luminances below the mean) that falls to 0 precisely at mean luminance and remains flat over the range of luminances above the mean.

These results suggest the intriguing possibility that (under the conditions of the current studies) texture variance judgments are based exclusively on the response of the off-center system. However, several other hypotheses must be considered.

1. It may be that the form of the perceived texture variance impact function results from the operation of an instantaneous, compressive retinal nonlinearity.
2. Alternatively, the form of the perceived texture variance impact function may be a by-product of divisive gain control processes.

Under hypothesis 1, the observer is veridically judging texture variance; however, the effective luminances of texels have been altered at the retina. For simplicity, suppose this hypothetical nonlinearity increased linearly with slope 1 over the range of luminances less than mid-gray, but was nearly flat over the range of luminances greater than or equal to mid-gray. In this case, luminances less than mid-gray would be veridically represented following retinal processing; however, the

post-retinal activation levels produced by all luminances greater than mid-gray will be approximately equal. In this case, we might expect the left side of the perceived variance impact function to resemble the left side of a parabola (as it roughly does). However, because the effective luminances of all luminances greater than mid-gray are approximately equal, they must all exert approximately equal impact on perceived texture variance. This might account for the flattening of the perceived variance impact function over the range of luminances greater than or equal to 80 cd/m².

Note, however, that if this compressive nonlinearity is supposed to be general to texture processing (as one might expect of a retinal nonlinearity), then texels with luminances greater than or equal to 80 cd/m² should all play approximately equivalent roles in any texture judgment. In particular this should be the case when the observer is attempting to judge the mean luminance of a patch of texture. If, in fact, the luminance of a texel were transformed by a significant, compressive nonlinearity, then it should be possible to construct textures with identical mean physical luminances which appear dramatically different in luminance. Consider, for example, a square checkerboard of black and white checks viewed on a uniform, mean gray background. When viewed from far enough away that checks cannot be resolved, this patch will be invisible (since it has the same physical mean luminance as the background). However, according to the 'up-front, compressive nonlinearity' hypothesis, as one draws closer, and checks become resolvable, not only should the checkerboard become visible as texture, it should also appear darker

IMPACT OF TEXEL INTENSITY ON PERCEIVED TEXTURE LUMINANCE

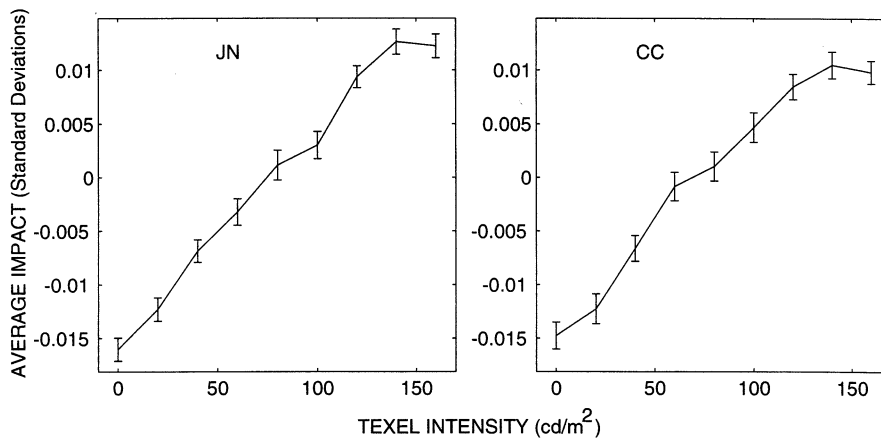


Fig. 5. Perceived luminance impact functions (for two observers) obtained using a full range of texel luminances. As in Fig. 1, the units of the vertical axis are standard deviations of the total noise compromising the observer's judgments. The average impact on perceived patch luminance exerted by a texel of luminance v is approximately a linear (increasing) function of v across the range of luminances from 0 up to 140 cd/m^2 , where the function saturates.

than the gray background, because the effective luminances of the white texture elements are decreased by the hypothetical, compressive nonlinearity. This does not occur.

To document this fact, we have performed an experiment (roughly analogous to the current, texture variance experiment, but somewhat different in methodology) to assess the relative impacts of different texel luminances on perceived texture luminance (Nam & Chubb, 2000). In that experiment, the observer's task was to judge which of the two texture patches had higher luminance (rather than variance). The results of this study are shown in Fig. 5. Plotted are the perceived luminance impact functions for the two authors. The average impact on perceived patch luminance exerted by a texel of luminance v is approximately a linear (increasing) function of v across the range of luminances from 0 up to 140 cd/m^2 , where the function saturates (i.e. the impact exerted by 160 cd/m^2 texels is no greater than that exerted by 140 cd/m^2 texels).

Thus, for purposes of judging texture luminance, all texel luminances except the highest are approximately veridically coded: the impact on perceived patch luminance exerted by a given texel is approximately proportional to the luminance of that texel.

This result supports the position of He and MacLeod (1998), who propose that the compressive retinal nonlinearity implicated by experiments using high frequency interference fringes (e.g. Sekiguchi, Williams & Packer, 1991; MacLeod, Williams, & Makous, 1992; Chen, Makous & Williams, 1993; MacLeod & He, 1993; He & MacLeod, 1996) is not an instantaneous nonlinearity (of the sort proposed contemporaneously by Fechner (1966) (originally published in 1860) and

also by Maxwell (1860)), but is rather due to the operation of a rapid gain-control mechanism.

This view suggests that the form of the texture variance impact function is probably not due to a general, instantaneous compressive distortion of effective texel luminances.

We must also consider, however, the possible effects of gain control processes on the perceived variance impact function. It is generally thought that a dynamic, retinal luminance pattern $S(x,y,t)$ yields a cortical input function of the following sort (e.g. Barlow, 1965; Barlow & Levick, 1969; Shapley & Enroth-Cugell, 1984):

$$I_S(x,y,t) = \frac{S(x,y,t) - \text{local_average}_S(x,y,t)}{\text{local_average}_S(x,y,t)} \quad (11)$$

where $\text{local_average}_S(x,y,t)$ is the average luminance of S across some small neighborhood around (x,y) shortly prior to t . (Makous' (1997) careful review of neurophysiological literature suggests that gain control mechanisms in the retina are probably only subtractive, implying that the divisive normalization accomplished by the denominator of Eq. (11) is probably performed postretinally.)

There has been some debate about the scope of the pooling performed by local_average_S . Several groups (Freeman & Badcock, 1996; He & MacLeod, 1998) have argued that $\text{local_average}_S(x,y,t)$ may well be restricted in its spatial pooling to individual cones. This view, however, is controversial. For example, the results of Chen et al. (1993) suggest that the retinal nonlinearity revealed by interference fringe experiments involves interactions between cones. If precortical gain control is purely temporal in its averaging, then any transient departures from a fixed adapting luminance (such as

are used in the current stimuli) should produce approximately proportional retinal activations.

On the other hand, a precortical gain control process that involved divisive normalization by a mean luminance, $\text{local_average}_S(x,y,t)$, computed across multiple cones over time might yield results qualitatively similar to those in Fig. 1. To see this, suppose $\text{local_average}_S(x,y,t)$ is computed over a nontrivial spatial region. In this case, as the mean luminance of S is decreased, the denominator on the right side of Eq. (9) will tend to decrease, effectively magnifying deviations of luminance from the mean. Conversely, as mean luminance increases, deviations of luminance from the mean are compressed.

Consider then a modulator ϕ such that $U + \phi$ has the same variance as $U - \phi$, but lower mean luminance. A gain control mechanism characterized by Eq. (9) could amplify the deviations of luminance from the mean in $P_{U+\phi}$ and compress the deviations in $P_{U-\phi}$, thereby imbuing $P_{U+\phi}$ with higher perceived variance than $P_{U-\phi}$. Thus dark texels could operate in two ways to increase perceived variance: first by deviating strongly from mean luminance, but also second by driving down local mean luminance, thereby tending to amplify all other deviations from mean luminance in the nearby vicinity.

We performed a simulation to decide whether such a gain control mechanism might account for the obtained impact functions (Fig. 1). For each of the experimental conditions i used for observers CC and JN in experiment 1, we computed

$$\begin{aligned} \Delta_i = & \sum_{v \in \Omega} \left(\frac{v - \mu_{U+\phi_i}}{\mu_{U+\phi_i}} \right)^2 (U + \phi_i)(v) \\ & - \sum_{v \in \Omega} \left(\frac{v - \mu_{U-\phi_i}}{\mu_{U-\phi_i}} \right)^2 (U - \phi_i)(v). \end{aligned} \quad (12)$$

This is analogous to computing the difference $\text{SS}_{U+\phi_i} - \text{SS}_{U-\phi_i}$; however, in the computation that yields Δ_i , differences of texel luminances from their patch means are divisibly normalized by their patch means. We proceed to interpret Δ_i , $i = 1, 2, \dots, 15$, as d' values. We can process these values in essentially the same way as we process our data, measuring the impact exerted on Δ_i by different luminances.

The resulting impact function (Fig. 6) shows us what we might expect from a mechanism that preceded a straightforward variance computation by a luminance normalization. This function differs very little from the parabolic function to be expected if the observer were veridically sensing luminance variance. We infer that divisive normalization is insufficient to explain the form of the texture variance impact functions obtained from our observers (Figs. 1, 3 and 4).

We submit that the perceived variance impact functions f of Fig. 1, deviate from the quadratic parabola not because texel luminances have been compressively distorted, nor because of divisive contrast normalization. Rather, the impact functions of Fig. 1 take the form they do because the mechanism resident in human vision that is best suited to the task of sensing texture variance senses not texture variance per se, but the average of the negative half-wave rectified stimulus.

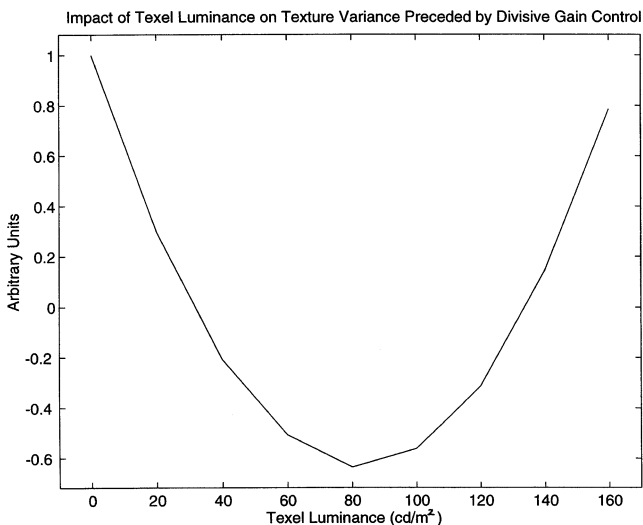


Fig. 6. The impact function predicted if observers were actually computing texture variance following normalization due to contrast gain control. This function differs very little from the quadratic function to be expected if the observer were veridically sensing luminance variance. We infer that contrast normalization due to retinal gain control is insufficient to explain the form of the texture variance impact functions in Fig. 1.

5. Experiment 2

On reflection, there is some sense to the results of experiments 1, 1.1 and 1.2. Plausibly, our observers make their texture variance judgments using the mechanism that is normally used to sense contrast in a scene. From an ecological perspective, it is reasonable for such a contrast-sensing mechanism to take special note of luminances near 0. Luminances in the scene can increase without bound; however, luminances cannot decrease below 0. Thus the presence of luminances near 0 in a scene is highly significant for assessing the contrast of that scene. For example, there may be highly luminant elements in a foggy scene; however, a foggy scene is necessarily devoid of black. Thus, for the specific purpose of assessing the transparency of a visual medium, the presence and/or absence of black elements in the scene is a potent indicator.

These considerations suggest that it may be the presence of very low luminances in the stimuli of experiment 1 that lead to the observed performance.

Table 7
Data from Experiment 2 for observers CC and JN^a

Modulator									CC		JN	
$\phi(0)$	$\phi(1)$	$\phi(2)$	$\phi(3)$	$\phi(4)$	$\phi(5)$	$\phi(6)$	$\phi(7)$	$\phi(8)$	C	I	C	I
30	7	-9	-18	-20	-18	-9	7	30	76	24	45	15
0	-15	-24	-26	-20	-11	6	30	60	58	42	48	12
0	22	19	1	-22	-37	-36	-7	60	74	26	38	22
50	-23	-24	-5	3	-5	-24	-22	50	83	17	42	18
19	37	-19	-43	-20	6	2	-23	41	67	33	41	19
35	-16	21	-17	-46	-17	21	-16	35	67	33	42	18
28	20	-39	12	-21	-48	21	-5	32	69	31	43	17
30	4	3	-42	10	-42	3	4	30	75	25	50	10
60	30	6	-11	-20	-26	-24	-15	0	71	29	44	16
60	-7	-36	-37	-22	1	19	22	0	75	25	47	13
10	38	7	-31	-47	-31	7	37	10	78	22	48	12
41	-22	2	6	-21	-43	-19	37	19	76	24	50	10
25	31	-39	-20	6	-20	-39	31	25	75	25	38	22
32	-5	21	-48	-21	12	-39	20	28	69	31	45	15
30	11	-21	6	-52	6	-21	11	30	68	32	43	17

^a The first nine columns of row i give the modulator ϕ_i used in the i th condition. Column 10 (11) of row i gives the number of trials on which CC responded correctly (incorrectly) in condition i . Column 12 (13) of row i gives the number of trials on which JN responded correctly (incorrectly) in condition i .

Luminances near 0 may provide a cue to texture contrast that overrides other sources of information.

Alternatively, it may not be the presence of black texels per se that leads to the observed neglect of information carried by luminances greater than the mean. Perhaps the same negative half-wave characteristic will predominate even when all stimulus contrasts are small. The purpose of experiment 2 was to investigate this possibility.

5.1. Procedure

The procedure was identical to that of experiment 1, except that the nine luminances used were 60, 65, 70, 75, 80, 85, 90, 95 and 100 cd/m². Thus the range of stimulus contrasts was reduced from 100% in experiment 1 to 25% in experiment 2.

5.2. Results and discussion

The raw data for observers CC and JN are given in Table 7 and the data for EF are given in Table 8. JN performed 60 trials per condition, and CC performed 100 trials per condition. EF initially ran 90 trials per condition (for each of the first 15 conditions listed in Table 8). However, there was a discrepancy between his results versus those of the other two observers. He therefore conducted 100 additional trials in each of a set of 15 conditions of increased difficulty. The impact functions resulting from these two sets of data from EF were very similar. Accordingly, the data were pooled for EF, and a single impact function computed.

5.2.1. Model fits

As in experiment 1, fits of the model were assessed using likelihood ratio tests. The results are shown in Table 9. The p -values for EF and JN are acceptable. However, as was found in experiment 1, the p -value for CC is low, 0.030. This suggests that his data might be better fit by elaborating the basic model to take account of interactions between different luminances. We must acknowledge the possibility that such an elaboration might alter the form of CC's impact function (Fig. 7).

5.2.2. Perceived variance impact functions

The resulting perceived variance impact functions are shown in Fig. 7. We had hoped to discover whether or not the striking form taken by the perceived variance impact function in experiment 1 would persist when the gamut of stimulus contrasts was reduced from 100 to 25%. What we find are distinct individual differences, precluding any firm conclusions about this point.

With the restricted gamut of texel contrasts used in experiment 2, the impact functions for CC and JN are crudely parabolic, suggesting that they are indeed using a different system and/or strategy to perform the variance judgments than they used in experiment 1. By contrast, EF continues to use a negative, half-wave rectifying system.

It may be relevant to note that in experiment 1, EF's perceived variance impact function was differentially sensitive to all texel luminances up to 80 cd/m², whereas the impact functions of JN and CC were differentially sensitive to luminances only up to 60 cd/m², and were flat across all luminances greater than or equal to 60 cd/m². Perhaps EF's larger range of differential sensi-

Table 8
Data from Experiment 2 for observer EF^a

Modulator									EF	
$\phi(0)$	$\phi(1)$	$\phi(2)$	$\phi(3)$	$\phi(4)$	$\phi(5)$	$\phi(6)$	$\phi(7)$	$\phi(8)$	C	I
40	10	-11	-24	-30	-24	-11	10	40	62	28
0	-20	-31	-34	-30	-14	9	40	80	8	82
0	30	26	1	-28	-50	-49	-10	80	68	22
67	-30	-32	-7	4	-7	-32	-30	67	56	34
25	50	-26	-57	-28	8	3	-30	55	61	29
47	-21	29	-22	-66	-22	29	-21	47	67	23
37	27	-51	16	-30	-64	29	-7	43	62	28
41	5	5	-56	11	-56	4	5	41	63	27
80	40	9	-14	-30	-34	-31	-20	0	90	0
80	-10	-49	-50	-28	1	26	30	0	56	34
13	50	10	-41	-64	-41	10	50	13	63	27
55	-30	3	8	-28	-57	-26	50	25	69	21
33	41	-51	-26	6	-26	-51	41	33	66	24
43	-7	29	-64	-30	16	-51	27	37	64	26
39	15	-28	8	-69	8	-27	15	39	65	25
30	7	-9	-18	-20	-18	-9	7	30	61	39
0	-15	-24	-26	-20	-11	6	30	60	19	81
0	22	19	1	-22	-37	-36	-7	60	62	38
50	-23	-24	-5	3	-5	-24	-22	50	62	38
19	37	-19	-43	-20	6	2	-23	41	69	31
35	-16	21	-17	-46	-17	21	-16	35	66	34
28	20	-39	12	-21	-48	21	-5	32	71	29
30	4	3	-42	10	-42	3	4	30	53	47
60	30	6	-11	-20	-26	-24	-15	0	94	6
60	-7	-36	-37	-22	1	19	22	0	66	34
10	38	7	-31	-47	-31	7	37	10	70	30
41	-22	2	6	-21	-43	-19	37	19	62	38
25	31	-39	-20	6	-20	-39	31	25	77	23
32	-5	21	-48	-21	12	-39	20	28	65	35
30	11	-21	6	-52	6	-21	11	30	63	37

^a The first nine columns of row i give the modulator ϕ_i used in the i th condition. Column 10 (11) of row i gives the number of trials on which EF responded correctly (incorrectly) in condition i .

tivity, which encompassed the lower 4 texel luminances used in experiment 2, inclined him to use the same system/strategy in experiment 2 that he used in experiment 1. On the other hand, the stimulus contrasts used in experiment 2 were all drawn from a region across which the previously measured impact functions of CC and JN were flat. Perhaps, then, CC and JN were forced in experiment 2 to adopt a new strategy involving other visual subsystems.

6. Final remarks

The results of experiment 1, sustained by control experiments 1.1 and 1.2 suggest a surprising possibility. It has sometimes been proposed that the signals carried by the on- and off-center channels are combined in two distinct ways in the cortex (e.g. Derrington, 1990; Wilson, 1994). It is hypothesized that

1. Judgments requiring sensitivity to luminance modulations (first-order judgments) use cortical neurons

that combine on-center and off-center outputs in opponent (push–pull) fashion, thereby achieving a quasi-linear sensitivity to different luminances in the retinal input.

2. Judgments requiring sensitivity to variations in texture contrast (second-order judgments) use cortical neurons that combine on-center and off-center outputs in additive (push–push) fashion, thereby achieving sensitivity to a full-wave rectified transformation of the retinal input.

Table 9

Results of a likelihood ratio test of the model of Section 2.2.1 as applied to the data of Experiment 2^a

	Deg. of freedom	χ^2	P
CC	7	15.493	0.030
JN	7	5.606	0.586
EF	22	24.308	0.331

^a The model fits adequately for EF and JN, but the fit is poor for CC.

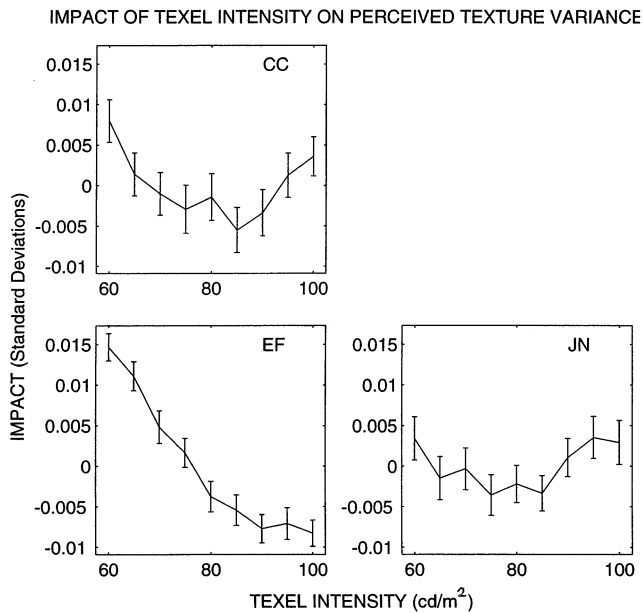


Fig. 7. Perceived variance impact functions for three observers, obtained using a restricted range of texel luminances. Texel luminances used were 60, 65, 70, 75, 80, 85, 90, 95 and 100 cd/m^2 . As in Figs. 1–4, the units of the vertical axis are standard deviations of the total noise compromising the observer’s judgments.

The current results, however, run contrary to this a priori plausible scheme. We find — at least for textures comprising a broad gamut of luminances, including luminances near 0 — that judgments of texture contrast (second order judgments) are mediated not by full-wave rectification but rather by negative half-wave rectification. This leads us to conjecture that such judgments are mediated exclusively by the off-center system.

Acknowledgements

Much of the research reported here was supported by National Science Foundation Human Cognition and Perception Program grant DBS-9203291. Additional support was provided by AFOSR Life Sciences Visual Information Processing Program grant F49620-94-1-0345. Finally, the authors are grateful to an anonymous reviewer for prompting experiments 1.1 and 1.2.

Appendix A

Here we show that for any U -modulator ϕ , $\text{SS}_{U+\phi} - \text{SS}_{U-\phi} = 2\psi \cdot \phi$, for ψ defined by Eq. (7). For all $v \in \Omega$, it will be convenient to set

$$\lambda(v) = v - 4, \quad (\text{A1})$$

$$\tilde{U}(v) = \frac{1}{9} = \frac{U(v)}{1800}, \quad (\text{A2})$$

$$\tilde{\phi}(v) = \frac{\phi(v)}{1800}. \quad (\text{A3})$$

We consider a constant k to be the function mapping every $v \in \Omega$ onto the value k . Thus

$$k \cdot \tilde{U} = \sum_{v \in \Omega} k \tilde{U}(v) = k, \quad (\text{A4})$$

whereas

$$k \cdot \tilde{\phi} = k \cdot \phi = 0. \quad (\text{A5})$$

For ψ defined by Eq. (7), the reader may verify that

$$\psi(v) = \lambda^2(v) - K, \quad \text{for } K = \frac{\lambda \cdot \lambda}{9}. \quad (\text{A6})$$

Note that ψ , \tilde{U} , and λ are mutually orthogonal. That is,

$$\psi \cdot \tilde{U} = \psi \cdot \lambda = \tilde{U} \cdot \lambda = 0. \quad (\text{A7})$$

For $\mu_{U+\phi}$ defined by Eq. (3) we observe that

$$\begin{aligned} \mu_{U+\phi} &= \frac{1}{1800} \sum_{v \in \Omega} v(U+\phi)(v) \\ &= \sum_{v \in \Omega} v(\tilde{U} + \tilde{\phi})(v) \\ &= \sum_{v \in \Omega} (\lambda(v) + 4)(\tilde{U} + \tilde{\phi})(v) \\ &= 4 + \lambda \cdot \tilde{\phi}. \end{aligned} \quad (\text{A8})$$

Similarly, we find

$$\mu_{U-\phi} = 4 - \lambda \cdot \tilde{\phi}. \quad (\text{A9})$$

We next observe that

$$\begin{aligned} \text{SS}_{U+\phi} &= \sum_{v \in \Omega} (v - (4 + \lambda \cdot \tilde{\phi}))^2 (U + \phi)(v) \\ &= 1800 \sum_{v \in \Omega} (v - (4 + \lambda \cdot \tilde{\phi}))^2 (\tilde{U} + \tilde{\phi})(v) \\ &= 1800 \sum_{v \in \Omega} (\lambda(v) - \lambda \cdot \tilde{\phi})^2 (\tilde{U} + \tilde{\phi})(v) \\ &= 1800 \sum_{v \in \Omega} (\lambda^2(v) - 2(\lambda \cdot \tilde{\phi})\lambda(v) + (\lambda \cdot \tilde{\phi})^2) (\tilde{U} + \tilde{\phi})(v) \\ &= 1800 \sum_{v \in \Omega} (\psi(v) + K - 2(\lambda \cdot \tilde{\phi})\lambda(v) + (\lambda \cdot \tilde{\phi})^2) (\tilde{U} + \tilde{\phi})(v) \\ &= 1800(K + \psi \cdot \tilde{\phi} - (\lambda \cdot \tilde{\phi})^2). \end{aligned} \quad (\text{A10})$$

A similar argument shows that

$$\text{SS}_{U-\phi} = 1800(K - \psi \cdot \tilde{\phi} - (\lambda \cdot \tilde{\phi})^2). \quad (\text{A11})$$

It follows that

$$\text{SS}_{U+\phi} - \text{SS}_{U-\phi} = 3600\psi \cdot \tilde{\phi} = 2\psi \cdot \phi. \quad (\text{A12})$$

References

- Barlow, H. B. (1965). Optic nerve impulses and Weber's law. *Cold Spring Harbour Symposium on Quantitative Biology*, 30, 539–546.
- Barlow, H. B., & Levick, W. R. (1969). Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. *Journal of Physiology*, 200, 1–24.
- Bergen, J. R., & Adelson, E. H. (1988). Visual texture segmentation based on energy measures. *Journal of the Optical Society of America A*, 3(13), 98–101.
- Bergen, J. R., & Landy, M. S. (1991). In J. A. Movshon, & Michael S. Landy, *Computational modeling of visual texture segregation. Computational models of visual processing* (pp. 253–271). Cambridge, MA: MIT Press.
- Chen, B., Makous, W., & Williams, D. R. (1993). Serial spatial filters in vision. *Vision Research*, 33(3), 413–427.
- Chubb, C. (1999). Texture based methods for analyzing elementary visual substances. *Journal of Mathematical Psychology*, 43, 539–567.
- Cornsweet, T. N., & Teller, D. Y. (1965). Relation of increment thresholds to brightness and luminance. *Journal of the Optical Society of America A*, 55(10 pt. 1), 1303–1308.
- Derrington, A. (1990). Mechanisms for coding luminance patterns: are they really linear? In C. Blakemore, *Vision: coding and efficiency*. Cambridge: Cambridge University Press.
- Freeman, A. W., & Badcock, D. R. (1996). Visual adaptation is highly localized in the human retina. *Investigative Ophthalmology and Visual Science*, 37(3), S726.
- Fechner, G. T. (1966). *Elements of psychophysics*. New York: Holt, Reinhardt and Winston (originally published in 1860).
- He, S., & MacLeod, D. I. A. (1996). Local luminance nonlinearity and receptor aliasing in the detection of high-frequency gratings. *Journal of the Optical Society of America A*, 13, 6.
- He, S., & MacLeod, D. I. A. (1998). Contrast-modulation flicker: dynamics and spatial resolution of the light adaptation process. *Vision Research*, 38(7), 985–1000.
- Hoel, P. G., Port, S. C., & Stone, C. J. (1971). *Introduction to statistical theory*. Boston: Houghton Mifflin.
- Kingdom, F., & Moulden, B. (1991). A model for contrast discrimination with incremental and decremental test patches. *Vision Research*, 31(5), 851–858.
- Knutsson, H., & Granlund, G. H. (1983). *Texture analysis using two-dimensional quadrature filters*. IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management.
- Landy, M. S., & Bergen, J. R. (1991). Texture segregation and orientation gradient. *Vision Research*, 31(4), 679–691.
- Leshowitz, B., Taub, H. B., & Raab, D. H. (1968). Visual detection of signals in the presence of continuous and pulsed backgrounds. *Perception and Psychophysics*, 4(4), 20–213.
- MacLeod, D. I. A., Williams, D. R., & Makous, W. (1992). A visual nonlinearity fed by single cones. *Vision Research*, 32(2), 347–363.
- MacLeod, D. I. A., & He, S. (1993). Visible flicker from invisible patterns. *Nature*, 361, 256–258.
- Makous, W. L. (1997). Fourier models and the loci of adaptation. *Journal of the Optical Society of America A*, 14, 2323–2345.
- Maxwell, J. C. (1860). On the theory of compound colors and the relation of the colors of the spectrum. *Philosophical Transactions of the Royal Society*, 150, 57–84.
- Nam, J.-H., & Chubb, C. (2000). Texture luminance judgments are approximately veridical. *Vision Research*, 40, 1695–1709.
- Robson, J. G. (1980). Neural images: the physiological basis of spatial vision. In C. S. Harris, *Visual coding and adaptability* (pp. 177–214). Erlbaum, NJ: Hillsdale.
- Sekiguchi, N., Williams, D. R., & Packer, O. (1991). Nonlinear distortion of gratings at the foveal resolution limit. *Vision Research*, 31(5), 815–831.
- Shapley, R., & Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. *Progress in Retinal Research*, 3, 263–346.
- Whittle, P. (1986). Increments and decrements: luminance discrimination. *Vision Research*, 26(10), 1493–1507.
- Wilks, S. S. (1944). *Mathematical statistics*. Princeton: Princeton University Press.
- Wilson, H. R. (1994). The role of second-order motion signals in coherence and transparency. In J. A. G. Gregory, & R. Bock, *Higher-order processing in the visual system. Ciba Foundation symposium*, 184. Chichester, UK: Wiley.