Implications of triple correlation uniqueness for texture statistics and the Julesz conjecture

John I. Yellott, Jr.

Department of Cognitive Sciences, University of California, Irvine, Irvine, California 92717

Received May 15, 1992; revised manuscript received September 21, 1992; accepted November 11, 1992

Triple correlation uniqueness refers to the fact that every monochromatic image of finite size is uniquely determined up to translation by its triple correlation function. Here that fact is used to prove that every finite image composed of discrete colors is determined up to translation by its third-order statistics. Consequently, if two texture samples have identical third-order statistics, they must be physically identical images and thus visually nondiscriminable by definition. It follows that a third-order version of Julesz's long-abandoned conjecture about spontaneous texture discrimination is necessarily true (notwithstanding such well-known counterexamples as the odd and even textures). The second-order (i.e., original) version of that conjecture is not necessarily true: physically distinct finite images with identical second-order statistics can be constructed, so counterexamples are possible. However, the counterexamples that one finds in the literature are either nonexact or difficult to reconstruct. A new principle is described that permits the easy construction of discriminable black and white texture samples that have strictly identical second-order statistics and thus provide exact counterexamples to the Julesz conjecture.

1. INTRODUCTION

A. The Julesz Conjecture

In a seminal paper in 1962, Julesz[1] introduced a program of research on visual texture perception motivated by the observation that some pairs of distinct textures are instantly seen as different (spontaneously discriminated) while others can be distinguished only after careful scrutiny. He asked whether that perceptual dichotomy can be predicted on the basis of global image statistics corresponding to the joint probability distributions that characterize stationary stochastic processes. A decade later[2] he and co-workers proposed a specific hypothesis along those lines that came to be known as the Julesz conjecture: the hypothesis that textures cannot be spontaneously discriminated if they have the same first-order and second-order statistics and differ only in their third-order or higher-order statistics. (Those statistics are defined here in Section 2.) For a period in the 1970's that conjecture survived many experimental tests and seemed to have captured a fundamental property of vision.[3] By 1981, however, the Julesz conjecture appeared to be disproved by results from his own laboratory[4–6] and elsewhere,[7] and Julesz himself abandoned it in favor of a theoretical approach based on local image features called textons.[8] Today both the conjecture itself and the global statistical approach that it represented are generally regarded as a closed chapter in the history of texture discrimination research.[9]

In the present paper that chapter is reexamined in light of recent mathematical discoveries about higher-order autocorrelation functions,[10,11] which prove to impose strict constraints on image statistics. The strongest evidence against the Julesz conjecture took the form of a pair of easily discriminable black and white textures devised by Julesz *et al.*,[6] the odd and even textures, which were said to have identical third-order statistics. (Figure 1 shows samples of odd and even texture. Section 4 explains their construction.) Identical third-order statistics imply identical second-order statistics, so the odd and even textures apparently provided a strong counterexample to Julesz's second-order conjecture. But they also had broader implications. Underlying that specific hypothesis were ideas about computational complexity and limited processing capacity. Higher-order statistics entail more calculations, and the visual mechanism responsible for spontaneous texture discrimination was thought to be limited to the computation of at most the second-order statistics of images. So the easy discriminability of the odd and even textures, which ostensibly differed only at the level of fourth-order statistics, seemed to discredit not only the second-order conjecture itself but also the general idea that spontaneous texture discrimination is based on a global statistical computation performed by a mechanism with limited capacity. That conclusion was reinforced by the fact that both the odd and the even textures can be easily discriminated from a purely random texture created by coloring all the squares of a checkerboard by independent fair-coin tosses—a texture that was said to have the same third-order statistics as those of the two other textures.[8] (Figure 1 shows samples of coin-toss texture compared with samples of odd and even texture.)

B. Triple Correlation and Third-Order Statistics

What prompts a reexamination of that episode now is the realization that third-order image statistics have unexpectedly strong characterization properties. Using a fairly recent uniqueness theorem for triple correlations,[10,11] one can show that the third-order statistics of any image of finite size uniquely determine that image up to a translation. (The proof is given in Section 3.) In other words, two pictures with identical third-order statistics must be physically identical. Consequently, odd and even texture

0740-3232/93/050777-17$05.00

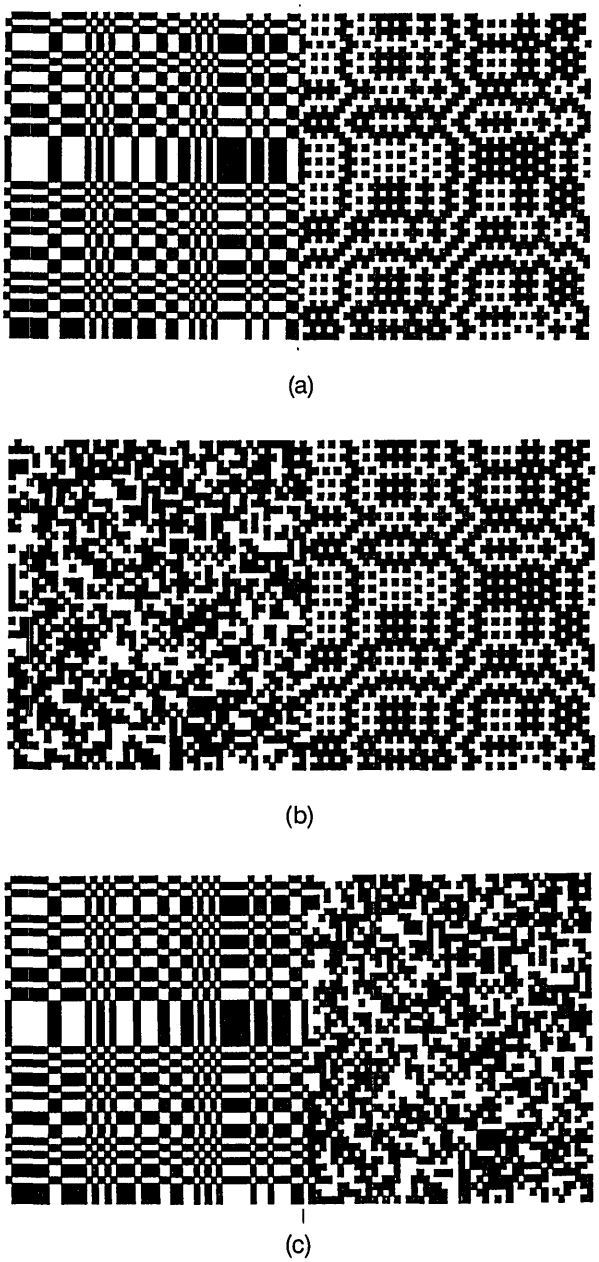

Fig. 1. (a) 50 × 50 pixel samples of odd (right side) and even (left side) texture. Tick marks indicate the boundary between the two. (These are matched samples, as explained in Section 4.) (b) 50 × 50 samples of odd (right side) versus coin-toss (left side) texture. (c) 50 × 50 samples of even (left side) versus coin-toss (right side) texture.

samples (such as the images on the right-hand and left-hand sides, respectively, in Fig. 1) can never have identical third-order statistics, and the same is true of odd or even texture samples compared with physically distinct samples of the coin-toss texture. And, in general, visual discrimination of images with identical third-order statistics is impossible in principle because it would mean discrimination between physically identical objects. Thus a third-order version of the Julesz conjecture is necessarily true because it is impossible to construct a counterexample to the claim that texture samples cannot be discriminated if their third-order statistics are identical. (Subsection 1.D discusses possible ways in which one can reinterpret the Julesz conjecture to make its third-order version nontautological. The conclusion is that no logical way to do this exists.)

C. Image Statistics Versus Ensemble Statistics

If odd and even texture samples cannot have identical third-order statistics, how could Julesz *et al.* say that they had demonstrated "visual discrimination of textures with identical third-order statistics"?[6] The answer is a matter of definitions, which revolves around the distinction between the ensemble statistics and the sample statistics of a stochastic process. The names "odd and even textures" refer not to specific images but rather to probabilistic algorithms for the construction of images—in other words, to stochastic processes whose samples are specific images. (The odd and even processes are described and analyzed in Section 4.) What Julesz *et al.*[6] actually proved is that the odd and even stochastic processes have identical third-order ensemble statistics: the expected values of the third-order statistics of sample images are the same for both processes. In other words, if one causes both processes to generate sequences of images, the averages of the third-order statistics of those images across the odd sequence and the averages of the same statistics across the even sequence should converge to common values as the sequences become infinitely long. That result does not imply that the third-order statistics of any specific pair of odd and even sample images will be identical (which the uniqueness theorem proved in Section 3 shows to be impossible), so there is no conflict between the mathematical results of Julesz *et al.*[6] and those reported here.

D. What Was the Julesz Conjecture?

What the odd and even textures actually demonstrate, then, is visual discrimination of texture samples drawn from stochastic processes with identical third-order ensemble statistics. (The same is true of the odd and even textures versus the coin-toss texture.) Consequently, the hypothesis that is actually refuted by the odd and even textures is not

H1: Texture images cannot be spontaneously discriminated if they have the same third-order statistics,

i.e., a hypothesis that predicts whether two specific texture images will be discriminable on the basis of statistics computed from the images themselves. Instead, what is refuted is the hypothesis

H2: Texture images cannot be spontaneously discriminated if they are samples from stochastic processes whose third-order ensemble statistics are identical,

i.e., a hypothesis that makes the discriminability of texture images depend on the statistical properties of the ensembles to which they belong. Since those properties are not directly visible to a naïve observer confronted with a specific pair of texture images, that idea clearly poses some conceptual difficulties. Nevertheless, in the paper of Julesz *et al.* on the odd and even textures,[6] hypothesis H2 was evidently the interpretation implicitly given to the Julesz conjecture (with "third-order" replaced by "second-order" in statement H2 to form the original conjecture).

Under the ensemble-statistics interpretation H2, a third-order version of the Julesz conjecture is not rendered tautological by the fact that every image is uniquely determined by its third-order statistics. But, without further amendment, hypothesis H2 is logically defective because its predictions about the discriminability of specific texture samples are intrinsically ambiguous: any pair of images A and B can always be construed as samples from two stochastic processes that have identical third-order (or nth-order, for any n) ensemble statistics, which implies nondiscriminability, and also as samples from two processes whose third-order ensemble statistics are different, which implies the opposite. (To satisfy the identical statistics condition, one construes A and B as samples from a stochastic process whose sample images are only A and B, with arbitrary probabilities assigned to each. Then A and B are samples from the same process, so they share the same nth-order ensemble statistics for all n. To satisfy the different-statistics condition when the images A and B themselves have nonidentical nth-order statistics, one construes A as a sample from a process that generates only A, and B as a sample from a process that generates only B. If A and B have identical nth-order statistics, one construes A as before and B as a sample from a process whose sample images are B and any other image C whose nth-order statistics are different from those of B, where B and C each are assigned probability 0.5.)

To remove the ambiguity in hypothesis H2, a reviewer proposed that it should be regarded as applying to predefined stochastic processes in a form such as

H3: If two stochastic image-generating processes S and S' have identical nth-order ensemble statistics, sample images from S will not be spontaneously discriminable from sample images from S'.

But hypothesis H3 is trivially refutable by the construction of a 2-image stochastic process S whose samples are any two easily discriminable images A and B. Then S and $S' = S$ form a counterexample. To overcome that difficulty, the reviewer suggested the addition of the proviso that S and S' must be ergodic processes (the intuition being that all the sample images of a process should be statistically similar). In that case the nth-order version of the conjecture would become

H4: If S and S' are nth-order ergodic stochastic processes whose nth-order ensemble statistics are identical, sample images from S will not be spontaneously discriminable from sample images from S'.

But nth-order ergodic literally means that, with probability 1, the nth-order statistics of every sample image generated by a stochastic process equal their expectations (i.e., the ensemble values). So, under this interpretation of an ensemble statistics version of the Julesz conjecture, the ensemble statistics themselves become irrelevant: the hypothesis

H5: If S and S' are stochastic processes whose sample images all have identical nth-order statistics, images from S will not be spontaneously discriminable from images from S'

is equivalent to the hypothesis

H6: If two images A and B have the same nth-order statistics, they will not be spontaneously discriminable,

which is the nth-order version of hypothesis H1.

(Note too that for stochastic processes whose samples are images of finite size, the uniqueness theorem for third-order statistics implies that, if a process is nth-order ergodic for $n \geq 3$, all its sample images must be physically identical. Since only finite images are possible in practice, it follows that two stochastic processes S and S' can satisfy the conditions of hypothesis H4 for $n = 3$ if and only if both always generate exactly one and the same image. So a third-order version of the Julesz conjecture is still tautological under this ensemble-statistics interpretation.)

The conclusion of the preceding analysis is that it does not make sense conceptually to interpret the Julesz conjecture as a hypothesis linking the discriminability of texture samples directly to the ensemble statistics of stochastic algorithms that generate those samples. Instead, the conjecture's only logical interpretation would seem to be as a hypothesis that predicts the discriminability of specific texture images on the basis of the nth-order statistics of the images themselves, i.e., as hypothesis H6, with $n = 2$ for the original conjecture. In the remainder of this paper that interpretation is assumed. In that case the uniqueness theorem proved in Section 3 implies that for $n \geq 3$ the Julesz conjecture is irrefutable because an exact counterexample cannot be constructed. But the original $n = 2$ conjecture is potentially refutable, since it is possible to construct pairs of physically distinct texture images whose second-order statistics are exactly identical. (That point is pursued in Subsection 1.F.)

E. Third-Order Statistics of Odd, Even, and Coin-Toss Images

The uniqueness theorem for third-order statistics proved in Section 3 prohibits only an exact identity between the third-order statistics of physically distinct images. It leaves open the possibility that the third-order statistics of odd, even, and coin-toss texture samples might be close enough to be regarded as equal for all practical purposes. However, a direct examination shows differently. Figure 2 shows scatterplots that compare the third-order statistics of the odd, even, and coin-toss texture samples from Fig. 1. In these graphs the x coordinate of each data point is the value of a given third-order statistic for one member of an image pair, and the y coordinate is the value of the same third-order statistic for the other member of the pair. [For example, in Fig. 2(a) the x coordinate is the value of a third-order statistic for the even member of the pair of images shown in Fig. 1(a) and the y coordinate is the value of the same statistic for the odd member. Third-order statistics were computed with Eq. (9′) below.] If the texture pairs in Fig. 1 had identical third-order statistics, all data points in the graphs in Fig. 2 would lie on a 45° diagonal line through the origin; and if all the third-order statistics of the images equaled their corresponding ensemble values (i.e., their expectations), all the data points would lie at (1/2, 1/2), (1/4, 1/4), or (1/8, 1/8). It can be seen that neither condition is close to being satisfied. These results are typical of what one always finds for odd, even, and coin-toss images of this size (50 × 50 pixels): overall, their third-order statistics are far from identical.

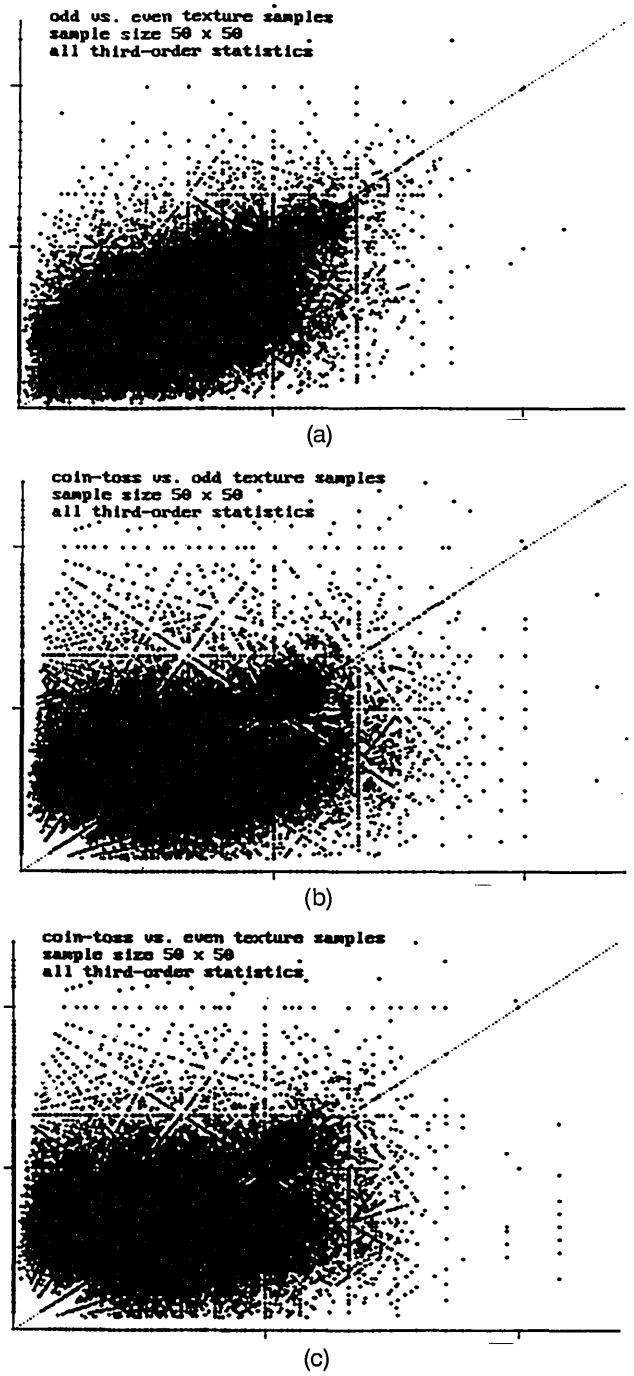

Fig. 2. (a) Scatterplot that compares all the third-order statistics [Eq. (9′)] of the odd and even texture samples shown in Fig. 1. The x axis corresponds to even texture statistic values, and the y axis corresponds to odd texture values. Both axes run from 0 to 0.6, with tick marks at 0.25 intervals. If the odd and even samples had identical third-order statistics, all the points would lie on the diagonal line. If all the statistics equaled their expected values (i.e., if the sample statistics equaled the ensemble statistics), all the points would fall at either (1/2, 1/2), (1/4, 1/4), or (1/8, 1/8). The two other scatterplots compare the third-order statistics of the odd versus coin-toss texture samples (b) and the even versus coin-toss samples (c) from Fig. 1. In (b) and (c) the y axis corresponds to coin-toss statistic values.

But the matter does not quite end there. The graphs in Fig. 2 compare all the third-order statistics of the texture samples of Fig. 1. It can be argued on both statistical and perceptual grounds that a more meaningful comparison would focus on the small-triangle statistics: third-order statistics that correspond to small distances across an image. The odd, even, and coin-toss algorithms all have the property of third-order ergodicity for infinitely large images: with probability 1 the third-order statistics of infinitely large sample images generated by all three processes must equal their expected values. (The uniqueness theorem for third-order statistics does not hold for infinitely large images, for reasons explained in Section 3.) So, as image size increases, one can expect odd, even, and coin-toss sample images to become statistically more alike, and the convergence should be more evident for the smaller-triangle statistics. The practical question is, How rapidly does that convergence occur? Section 4 examines that issue. Its conclusion is that the convergence is too slow to be perceptually meaningful; for the images in Fig. 1, even the small-triangle statistics show large differences by comparison with pairs of independent coin-toss images. Statistically, that result is not surprising, because the variance of the statistics of odd and even images decreases not in proportion to the number of pixels but more nearly in proportion to the square root of that number.

(The analysis in Section 4 also considers various quirks of the odd and even textures, such as the fact that odd and even sample images can be either matched or unmatched pairs, which have quite different statistical properties. An interesting conclusion from that analysis is that the eye appears to be insensitive to those large statistical differences.)

F. Second-Order Statistics and the Julesz Conjecture

The fact that odd and even texture samples cannot have identical third-order statistics does not automatically rule out the possibility that they might have identical second-order statistics and thus still provide a counterexample to Julesz's second-order conjecture. However, direct examination removes this possibility: the second-order statistics of odd and even texture samples are positively correlated (and contain a strong hidden correlation, described in Section 4) but clearly are not identical, as is demonstrated by the scatterplots in Fig. 3, which compare the second-order statistics of the texture samples of Fig. 1. Consequently, the odd and even textures do not provide counterexamples to the Julesz conjecture. An earlier analysis by Gagalowicz[12] arrived at the same conclusion.

That realization prompts a second look at all the original evidence against the Julesz conjecture. A decisive counterexample to the conjecture would take the form of a pair of texture samples that have identical second-order statistics and that are instantly discriminable. Most of the counterexamples found in the literature do not come close to meeting that description, because the texture samples involved are created by probabilistic algorithms that are guaranteed only to produce identical second-order statistics in the limit, as the images become infinitely large. That objection applies not only to the odd and even textures but also to other textures created by the floater technique of Gilbert[13,14] and by the ϕ-matrix technique of Diaconis and Freedman,[7] and also to the counterexamples created by Caelli and Julesz[4,5] with the four-disk method. The only exceptions appear to be the images of Gagalowicz,[12] who recognized the defects of previous

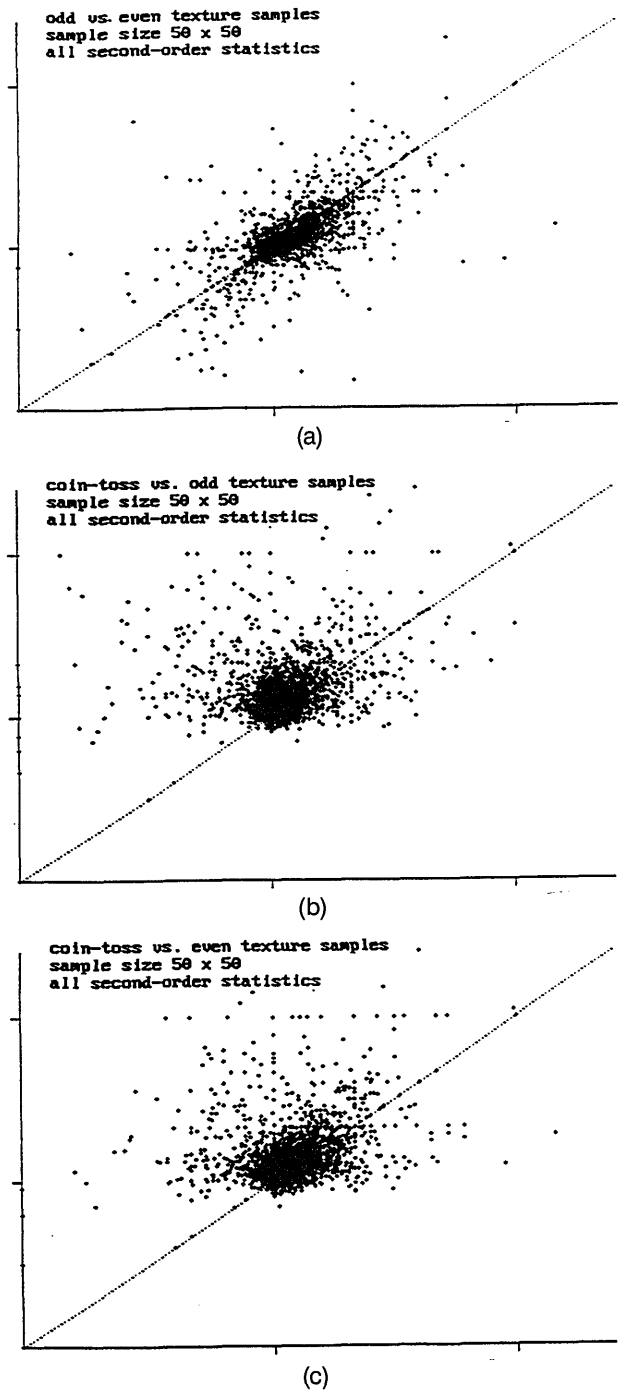

Fig. 3. (a) Scatterplots that compare all the second-order statistics [Eq. (8′)] of the texture samples from Fig. 1. The axes are the same as those in Fig. 2. (a) Odd versus even, (b) odd versus coin-toss, (c) even versus coin-toss.

counterexamples and used linear programming methods in an attempt to create better ones. He succeeded in constructing quite easily discriminable texture samples whose second-order statistics differ by only approximately 2%—identical for all practical purposes and certainly close enough for one to justify the conclusion that the second-order Julesz conjecture is false. However, Gagalowicz's texture construction algorithm is quite complicated and does not guarantee strict identity between second-order image statistics. It seems worthwhile to seek a simpler method that achieves exact results. Section 5 describes an easily implemented principle for the construction of pairs of black and white texture images that appear to be spontaneously discriminable and whose second-order statistics are guaranteed to be exactly identical. Pairs of one-dimensional textures (bar codes) and of two-dimensional textures created by that technique are shown in Figs. 4 and 5, respectively.

This construction principle may prove useful in other contexts because it provides an easy way by which one can create pairs of black and white images whose Fourier power spectra are exactly identical (and of course remain identical if both images are blurred by a common point-spread function, like that of the eye).

G. Texture Discrimination and Third-Order Statistical Similarity

Julesz's second-order conjecture appears to be false empirically, and one can now see that a third-order version of the conjecture is irrefutable in principle and thus true but vacuous. However, there is a variant of the third-order conjecture that is not tautological and can be tested: the hypothesis that discrimination between textures becomes increasingly difficult as their third-order statistics become more similar. Offhand, that idea does not seem promising. One type of negative evidence is provided by the odd and coin-toss textures. If the odd and coin-toss algorithms are used in the generation of texture samples, the third-order statistics of those samples are not at all similar when the samples are small, but the same statistics become more nearly alike as the size of the images (the number of pixels) increases. Perceptually, however, things go the other way: small samples of odd and coin-toss textures are weakly discriminable, and discrimination becomes easier as the samples become larger.

Figure 6, in conjuction with Fig. 1, illustrates that point. Figure 6(a) shows small odd-versus-coin-toss images (10 × 10 pixels) and a scatterplot that compares their third-order statistics. For images as small as those, the odd and coin-toss textures are not spontaneously discriminable, but the third-order statistics of the images differ greatly. Figure 6(b) compares the same third-order statistics for the 50 × 50-pixel samples of odd and coin-toss texture shown in Fig. 1. For images of that size the odd and coin-toss textures are easily discriminable, but their third-order statistics are now much more similar. In other words, as their global third-order statistics become more nearly alike, odd and coin-toss texture samples look less alike. That result is the opposite of what one would expect if spontaneous discriminability depended on a comparison of global third-order image statistics. (Section 4 describes another counterargument based on matched-versus-nonmatched samples of odd and even texture.)

H. Contents

Section 2 defines the image statistics that figured in the Julesz conjecture. Section 3 describes triple correlation functions and their uniqueness properties and uses them to prove that every finite-sized image composed of discrete colors is uniquely determined by its third-order statistics, which is the main new result of this paper. Section 4 examines the statistical properties of odd, even, and coin-toss texture samples. Section 5 deals with the problem of

Fig. 4. Exact one-dimensional counterexample to the Julesz conjecture, based on the construction principle explained in Section 5. The texture samples in (a) and (b) have identical second-order statistics. Sample (a) is composed of four replicas of the micropattern shown below in Fig. 12(a); sample (b) is composed of four replicas of the micropattern shown in Fig. 12(b). In (c) the two texture samples are displayed side by side.

devising counterexamples to the Julesz conjecture and explains the construction of the counterexample images shown in Figs. 4 and 5. The paper is mostly self-contained, but a detailed proof of the key uniqueness theorem for triple correlations seems too long for reproduction here; the basic ideas are sketched in Section 3, and the full argument can be found in Ref. 11, along with an analysis of the general uniqueness properties of triple correlation functions and other higher-order relatives of the ordinary autocorrelation function. Reference 15 is an excellent introduction to physical applications of triple correlation, and Ref. 16 includes a broad spectrum of recent papers on higher-order autocorrelation functions. The first application of those functions to vision research is due to Klein and Tyler,[17] who referred to them as generalized autocorrelation functions.

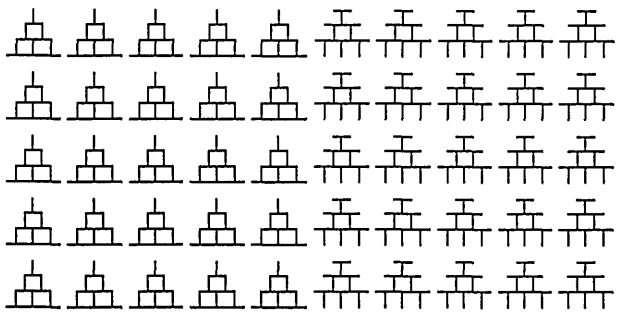

Fig. 5. Two-dimensional counterexample to the Julesz conjecture, based on the construction principle described in Section 5. The texture samples on the left and the right have identical second-order statistics.

2. IMAGE STATISTICS

A. Black and White Images

The statistics of various orders that figured in the Julesz conjecture were defined in terms of a hypothetical experiment performed on an image composed of discrete colors: "The n-gon statistic [or nth-order joint probability distribution] of an image can be obtained by randomly throwing n-gons of all possible shapes on the image and observing the probabilities that their n vertices fall on certain color combinations."[8] The n-gons here are geometrical objects: points (1-gons), line segments (2-gons, or dipoles), triangles (3-gons), etc. Random throwings of these objects lead to the first-, second-, and third-order statistics of an image. Most research has focused on two-color images, e.g., black and white. I begin there, taking black to be the foreground color and white to be the background color. Suppose that the function $f:\mathcal{R}^2 \rightarrow \{0,1\}$ represents such an image, with $f(x,y) = 1$ if image point (x,y) is black and $f(x,y) = 0$ otherwise (f is assumed to be locally integrable). One can think of the image as occupy-

ing some finite rectangular region of the plane that has area A and imagine that outside that region the rest of the plane is white [i.e., $f(x, y)$ is defined to be zero outside the image region]. "Randomly throwing" a single point onto the image means picking a point (x, y) at random in the image region, so the point is a random vector $(\mathbf{x}, \mathbf{y})$ with probability-density function $1/A$ throughout that region and with zero value outside it. The first-order statistic of image f, $s_{1,f}$, is the probability that the randomly thrown point lands on black, $P[f(\mathbf{x}, \mathbf{y}) = 1]$, which is simply the proportion of the image region that is black, i.e.,

$$s_{1,f} = \frac{1}{A}\iint_{\mathcal{R}^2} f(x, y)\mathrm{d}x\mathrm{d}y. \tag{1}$$

[The integration here could be carried out over just the image region instead of the entire plane, but because $f \equiv 0$ outside the image region, Eq. (1) gives the same result.]

The second-order statistics of the image are values of a function $s_{2,f}: \mathcal{R}^2 \rightarrow [0, 1]$ whose arguments represent horizontal and vertical separations between pairs of image points. The second-order statistic for a given separation vector (h, v) is the probability that when a point $(\mathbf{x}, \mathbf{y})$ is chosen at random within the image, both $(\mathbf{x}, \mathbf{y})$ and the point $(\mathbf{x} + h, \mathbf{y} + v)$ are black. Thus

$$s_{2,f}(h, v) = \frac{1}{A}\iint_{\mathcal{R}^2} f(x, y)f(x + h, y + v)\mathrm{d}x\mathrm{d}y. \tag{2}$$

[Note that even in an all-black image these probabilities will always be less than 1 (except when $h = v = 0$) because the point $(\mathbf{x} + h, \mathbf{y} + v)$ sometimes falls outside the image region.]

Third-order statistics are values of a function $s_{3,f}: \mathcal{R}^4 \rightarrow [0, 1]$ whose arguments represent pairs of separation vectors, (h_1, v_1) and (h_2, v_2). For a given pair the third-order statistic is the probability that when a point $(\mathbf{x}, \mathbf{y})$ is chosen at random in the image, the three points $(\mathbf{x}, \mathbf{y})$, $(\mathbf{x} + h_1, \mathbf{y} + v_1)$, and $(\mathbf{x} + h_2, \mathbf{y} + v_2)$ are all black. Thus

$$s_{3,f}(h_1, v_1, h_2, v_2) = \frac{1}{A}\iint_{\mathcal{R}^2} f(x, y)f(x + h_1, y + v_1)f(x + h_2, y + v_2)\mathrm{d}x\mathrm{d}y. \tag{3}$$

The third-order statistics can be thought of as the probabilities, for all triangles, that when a given triangle is dropped at random onto the image (preserving its orientation), all three vertices land on black. Higher-order statistics are defined analogously: fourth-order statistics have arguments that are triples of separation vectors, and so on. In general, statistics of order n include all statistics of lower orders; e.g., the third-order statistics include the second-order statistics as the special case in which one of the two separation vectors is $(0, 0)$.

In addition to nth-order statistics defined in terms of all-black coincidences, one could also define statistics based on white–white and black–white combinations. However, those statistics would be redundant because their values are already determined by the all-black statistics. For example, the black–white second-order statistic for argument (h, v) would be

$$\frac{1}{A}\iint_{\mathcal{R}^2} f(x, y)[1 - f(x + h, y + v)]\mathrm{d}x\mathrm{d}y = s_{1,f} - s_{2,f}(h, v).$$

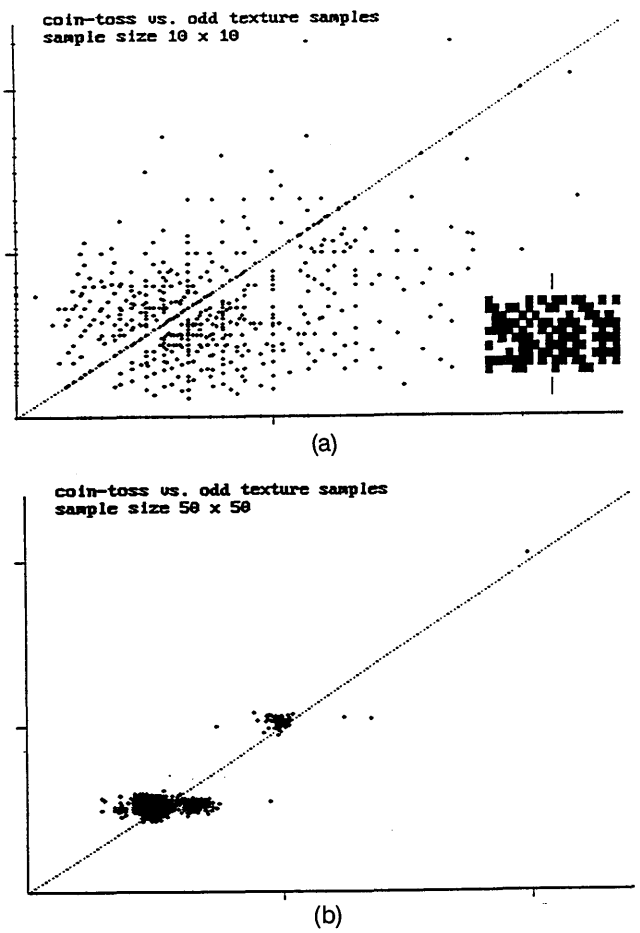

Fig. 6. (a) Scatterplot that compares all third-order statistics of small (10 × 10) samples of odd-versus-coin-toss texture (shown in the inset). (b) Scatterplot that compares the same third-order statistics for 50 × 50 samples of odd-versus-coin-toss texture. The axes are the same as those in Fig. 2, with coin-toss statistic values on the y axis.

B. Multicolored Images

For images composed of more than two colors, nth-order statistics are defined by a natural extension of the two-color definitions. Suppose that a finite-sized rectangular image $f(x, y)$ with area A is composed of L discrete colors or gray levels, labeled $0, 1, \ldots, L - 1$, where 0 is the background color. (Outside the image region the assumption is made that the plane has the background color.) The image can be thought of as the sum of L indicator functions $f_i: \mathcal{R}^2 \rightarrow \{0, 1\}$, $i = 0, \ldots, L = 1$, where $f_i(x, y) = 1$ if point (x, y) has color i and $f_i(x, y) = 0$ otherwise. Then $f(x, y) = \sum_{i=0}^{L-1} w_i f_i(x, y)$, where w_i is the numerical value of color i ($w_0 = 0$). There are L first-order statistics $s_{1,i,f}$, $i = 0, 1, \ldots, L - 1$, which correspond to the probabilities that a point chosen at random within the image has each of the L possible colors. For the $L - 1$ foreground colors these probabilities are

$$s_{1,i,f} = \frac{1}{A}\iint_{\mathcal{R}^2} f_i(x, y)\mathrm{d}x\mathrm{d}y, \tag{4}$$

and $s_{1,0,f}$ is equal to 1 minus the sum of the rest. Second-order statistics must now be defined for all possible colorings of a pair of points, (x, y) and $(x + h, y + v)$. For example, the second-order statistic $s_{2,i,j,f}(h,v)$ is the probability that when a point $(\mathbf{x}, \mathbf{y})$ is chosen at random in the image, $(\mathbf{x}, \mathbf{y})$ has color i and $(\mathbf{x} + h, \mathbf{y} + v)$ has color j. Except for the case in which both i and j are the background color 0, that statistic is

$$s_{2,i,j,f}(h,v) = \frac{1}{A}\iint_{\mathcal{R}^2} f_i(x,y) f_j(x + h, y + v)\mathrm{d}x\mathrm{d}y, \quad (5)$$

and $s_{2,0,0,f}(h,v)$ is defined to be 1 minus the sum of all the other $s_{2,i,j,f}(h,v)$. Third-order statistics $s_{3,i,j,k,f}$ are defined for all possible colorings of triples of points; they are the probabilities that when a point $(\mathbf{x}, \mathbf{y})$ is chosen at random within the image, $(\mathbf{x}, \mathbf{y})$, $(\mathbf{x} + h_1, \mathbf{y} + v_1)$, and $(\mathbf{x} + h_2, \mathbf{y} + v_2)$ have the colors i, j, and k, respectively. Except for the special case $i = j = k = 0$, those probabilities are given by

$$s_{3,i,j,k,f}(h_1,v_1,h_2,v_2) = \frac{1}{A}\iint_{\mathcal{R}^2} f_i(x,y) f_j(x + h_1, y + v_1) f_k(x + h_2, y + v_2)\mathrm{d}x\mathrm{d}y, \quad (6)$$

and $s_{3,0,0,0,f}(h_1,v_1,h_2,v_2)$ is equal to 1 minus the sum of all the other third-order statistics for the same set of arguments.

C. Discrete Image Statistics

Statistics of order n are defined above in terms of integrals of functions with continuous arguments—functions that describe colorings of a continuous surface. Such definitions are appropriate for physical images viewed by the eye. Computationally, images are represented by discrete arrays of pixel values: arrays of the form $\langle p(c,r);$ $c = 0, 1, \ldots, C - 1;$ $r = 0, 1, \ldots, R - 1\rangle$. The nth-order statistics of such arrays are defined by sums; e.g., for a C(olumn) $\times$ R(ow) array of binary pixels, with $p(c,r) = 1$ for the foreground color, the discrete first-, second-, and third-order statistics are defined as

$$s_{1,p} = \frac{1}{CR}\sum_{c=0}^{C-1}\sum_{r=0}^{R-1} p(c,r), \quad (7)$$

$$s_{2,p}(n,m) = \frac{1}{CR}\sum_{c=0}^{C-1}\sum_{r=0}^{R-1} p(c,r)p(c + n, r + m), \quad (8)$$

$$s_{3,p}(n_1,m_1,n_2,m_2) = \frac{1}{CR}\sum_{c=0}^{C-1}\sum_{r=0}^{R-1} p(c,r)p(c + n_1, r + m_1) \times p(c + n_2, r + m_2), \quad (9)$$

respectively, where the arguments are all integers. [These definitions assume that $p(c,r) \equiv 0$ unless $0 \le c \le C - 1$ and $0 \le r \le R - 1$. The upper limits of the sum in Eq. (8) could equivalently be set to $C - 1 - n$ and $R - 1 - m$, and in Eq. (9) they could be set to $C - 1 - \max(n_1, n_2)$ and $R - 1 - \max(m_1, m_2)$.]

Sometimes it is useful to renormalize the second- and third-order statistics so that their expected values for sample images of stochastic textures become independent of their arguments. For that purpose the statistics

$$S_{2,p}(n,m) = [CR/(C - n)(R - m)]s_{2,p}(n,m), \quad (8')$$

$$S_{3,p}(n_1,m_1,n_2,m_2) = \{CR/[(C - \max(n_1,n_2)] \times [R - \max(m_1,m_2)]\} \times s_{3,p}(n_1,m_1,n_2,m_2), \quad (9')$$

are defined. The statistical comparisons in Figs. 2 and 3 are based on the renormalized statistics defined by Eqs. (8′) and (9′), as are the comparisons made in Section 4. For odd, even, and coin-toss texture samples, the expected values of those statistics are always 1/2 (when all the arguments equal zero), 1/4 (for statistics based on pairs of distinct pixels), or 1/8 (for statistics based on triples of distinct pixels).

The relationship between the discrete-domain statistics of a binary pixel array $\langle p(c,r)\rangle$ and the continuous-domain statistics of a black and white physical image $f(x,y)$ can be expressed by means of delta functions. Mathematically, a display device converts an array of binary pixel values $\langle p(c,r)\rangle$ into a black and white image $f(x,y)$ by an operation of the form

$$\langle p(c,r)\rangle \rightarrow f(x,y) = d(x,y) * \sum_{c=0}^{C-1}\sum_{r=0}^{R-1} p(c,r)\delta(x - c)\delta(y - r),$$

where $*$ denotes a convolution, δ is the Dirac delta, and $d(x,y)$ is a 0–1 function representing a single screen pixel; e.g., $d = 1$ inside an open unit square centered at the origin (screen distances are measured in units of pixel width and height), and $d = 0$ otherwise. The discrete second- and third-order statistics can be translated into functions in continuous space by the operations

$$\langle s_{2,p}(n,m)\rangle \rightarrow s_{2,f}(h,v) = \sum_n\sum_m s_{2,p}(n,m)\delta(h - n)\delta(v - m),$$

$$\langle s_{3,p}(n_1,m_1,n_2,m_2)\rangle \rightarrow s_{3,f}(h_1,v_1,h_2,v_2) = \sum_{n_1}\sum_{m_1}\sum_{n_2}\sum_{m_2} s_{3,p}(n_1,m_1,n_2,m_2) \times \delta(h_1 - n_1)\delta(v_1 - m_1) \times \delta(h_2 - n_2)\delta(v_2 - m_2),$$

respectively.

The relationship between the discrete statistics of the pixel array and the continuous statistics of the image can now be stated concisely: $s_{2,f} = a_d * s_{2,p}$, where $a_d(h,v)$ is the autocorrelation function of the pixel shape function $d(x,y)$, and $s_{3,f} = t_d * s_{3,p}$, where t_d is the triple correlation function of $d(x,y)$. (Section 3 defines the autocorrelation and triple correlation functions.) Thus, when two binary pixel-value arrays with the same second or third-order discrete statistics are displayed as viewable images, those statistical identities are preserved. And the second- and third-order statistics of an image contain the discrete statistics of its pixel array (as their values for integer arguments), so identical continuous statistics for two images in the physical domain imply identical discrete statistics inside the computer.

Multicolored pixel arrays can be decomposed into binary arrays, one for each color value, and the relationship between their nth-order statistics and those of the corresponding multicolored physical image can be expressed as $s_{2,i,j,k,f} = a_d * s_{2i,j,p}$ and $s_{3,i,j,k,f} = t_d * s_{i,j,k,p}$.

3. AUTOCORRELATION, TRIPLE CORRELATION, AND IMAGE STATISTICS

A. Autocorrelation and Second-Order Statistics

It has always been recognized[1] that there is a close connection between second-order image statistics and the ordinary autocorrelation function—the function $a_f:\mathcal{R}^2 \rightarrow \mathcal{R}$ defined for any (integrable) function $f:\mathcal{R}^2 \rightarrow \mathcal{R}$ by

$$a_f(h,v) = \iint_{\mathcal{R}^2} f(x,y)f(x+h, y+v)\mathrm{d}x\mathrm{d}y. \quad (10)$$

If one compares a_f in Eq. (10) with the function $s_{2,f}$ defined in Subsection 2.A by Eq. (2), it is apparent that the second-order statistic function of a black and white image is its autocorrelation function divided by its area. The area A can be recovered from the second-order statistics [$A = S_{2,f}(0,0)/s_{1,f}{}^2$, where $S_{2,f}$ is the Fourier transform of $s_{2,f}$], so the second-order statistics of a black and white image completely determine its autocorrelation function and thus its power spectrum [since the Fourier transform of a_f is $|F(\alpha,\beta)|^2$, where F is the transform of f]. Conversely, if the autocorrelation function of a black and white image is known, its second-order statistics are completely determined up to the (somewhat arbitrary) area parameter A. In this sense the second-order statistics of a black and white image are equivalent to its autocorrelation function. For images containing more than two colors, that equivalence breaks down: in that case two images can have the same autocorrelation function but different second-order statistics. A simple example is that in which one image consists of three adjacent pixels with luminance levels 4, 4, and 1 and another consists of three pixels with luminances 2, 5, and 2. Calculation shows that the autocorrelation functions of these images are the same, but, since the images have no luminances in common, their second-order statistics are obviously different.

{The Fourier transform $F(\alpha,\beta)$ of an integrable function $f(x,y)$ is defined here as $F(\alpha,\beta) = \int\int_{\mathcal{R}^2} f(x,y)\exp[-i2\pi(\alpha x + \beta y)]\mathrm{d}y\mathrm{d}x$.}

B. Triple Correlation and Third-Order Statistics

In exactly the same sense, the third-order statistics of a black and white image are equivalent to a natural generalization of the ordinary autocorrelation function, called the triple correlation, which is not yet widely known in visual science but over the past decade and a half has attracted considerable attention in other fields.[15,16] The triple correlation function of an integrable function $f:\mathcal{R}^2 \rightarrow \mathcal{R}$ is the function $t_f:\mathcal{R}^4 \rightarrow \mathcal{R}$, defined by

$$t_f(h_1,v_1,h_2,v_2) = \iint_{\mathcal{R}^2} f(x,y)f(x+h_1, y+v_1)f(x+h_2, y+v_2)\mathrm{d}x\mathrm{d}y. \quad (11)$$

A straightforward calculation shows that the Fourier transform T_f of the triple correlation t_f (commonly known as the bispectrum) is related to the image transform F through the equation

$$T_f(\alpha_1,\beta_1,\alpha_2,\beta_2) = F(\alpha_1,\beta_1)F(\alpha_2,\beta_2)F(-\alpha_1-\alpha_2,-\beta_1-\beta_2), \quad (12)$$

which proves endlessly useful. For example, setting all arguments to zero in Eq. (12) gives $[F(0,0)]^3$, and setting $\alpha_2 = \beta_2 = 0$ yields $F(0,0)|F(\alpha_1,\beta_1)|^2$. So for nonnegative functions, such as images [where $F(0,0)$ cannot vanish unless $f = 0$], the bispectrum determines the power spectrum, and thus the triple correlation determines the autocorrelation. As another example, Eq. (12) shows immediately that $t_{f*g} = t_f * t_g$; i.e., the triple correlation of the convolution of two functions is the convolution of their triple correlations. Thus, if two images have the same triple correlation, that identity is preserved when both are blurred by the same point-spread function.

Comparing Eqs. (3) and (11), one can see that for a black and white image [$f(x,y) = 0$ or 1], the third-order statistic function $s_{3,f}$ is the triple correlation function t_f divided by the image area A. The area can be recovered from the third-order statistics by means of the facts that $s_{3,f}(0,0,0,0) = (1/A)F(0,0)$ (since $f^3 = f$) and $S_{3,f}(0,0,0,0) = (1/A)[F(0,0)]^3$ [from Eq. (12)]. Consequently, the third-order statistics of such an image completely determine its triple correlation function, and its triple correlation determines its third-order statistics up to the factor $1/A$. Thus two black and white images that have identical third-order statistics have identical triple correlation functions.

C. Uniqueness Properties of Triple Correlation Functions

The triple correlation function has aroused interest in large part because of its potential for the recovery of the phase spectrum of images. Ordinary autocorrelation is, of course, phase blind: the autocorrelation function of an image determines the amplitudes of all its Fourier components but none of their phases. A consequence is that the autocorrelation function of an image is unaffected by translations [i.e., $f(x,y) \rightarrow f(x+a, y+b)$]. Triple correlations have the attractive property of also being invariant under image translations but otherwise not phase blind: the triple correlation function of any finite-size image completely determines both its amplitude spectrum and (except for an absolute location term) its phase spectrum. Consequently, the triple correlation function of a finite image uniquely determines that image up to a translation; that is, two finite-sized images have identical triple correlations if and only if they are identical images, which differ at most in their absolute location in the plane.

The uniqueness theorem for triple correlations is valid for all functions that one can reasonably regard as representing finite-sized monochromatic images. In Ref. 11 the theorem is proved (in two different ways) for the class of all image functions with bounded support, where an image function is any nonnegative locally integrable function $f:\mathcal{R}^2 \rightarrow \mathcal{R}$. [Such a function is thought of as representing an image in the sense that its integral over any region of the plane specifies the total amount of light in that region (hence the requirement of local integrability). Two image functions are equal if their integrals agree for all regions, so that a photometer cannot distinguish them.]

An image function has bounded support (represents an image of finite size) if it is identically zero outside some finite rectangle. Every image function f with bounded support has a well-defined triple correlation function given by Eq. (11), and using Eq. (12) one can prove the following theorem (Ref. 11, Theorem 1′):

Triple Correlation Uniqueness [TCU] Theorem If f is an image function with bounded support, and another image g has the same triple correlation function as that of f, then $g(x, y) = f(x + a, y + b)$ for some pair of constants a and b.

The proof of this theorem seems too long to be given here in complete detail, but the basic ideas can be quickly sketched as an indication of what is involved. To simplify matters, one can consider one-dimensional image functions (the extension to two dimensions is straightforward). The proof makes use of well-known facts from probability theory. Suppose that two image functions $f(x)$ and $g(x)$ have the same triple correlation. Then they have the same bispectrum, and Eq. (12) implies that their Fourier transforms F and G satisfy the relationship

$$F(\alpha)F(\beta)F(-\alpha - \beta) = G(\alpha)G(\beta)G(-\alpha - \beta). \quad (13)$$

I show that Eq. (13) implies that $G(\alpha) = \exp(i2\pi c\alpha)F(\alpha)$ for some real constant c, and thus $g(x) = f(x + c)$. Note first that Eq. (13) implies that $F(0) = G(0)$. If $F(0) = 0$, both f and g are the zero image ($f = g = 0$ a.e.); if not, there is no loss of generality in the assumption that both $F(0)$ and $G(0)$ equal 1.0. {One can always divide both sides of Eq. (13) by $[F(0)]^3$ and show that $G(\alpha)/F(0) = \exp(i2\pi c\alpha)F(\alpha)/F(0)$, which proves the claim.} Then the image functions f and g are probability-density functions, and F and G are their characteristic functions. All characteristic functions are continuous[19] and equal 1.0 at the origin, so there is a neighborhood of the origin in which both F and G are nonvanishing. Within that neighborhood the functions in Eq. (13) can be divided freely, and after appropriate divisions [and with the use of the fact that Eq. (7) implies that $|F(\alpha)|^2 = |G(\alpha)|^2$ for all α], Eq. (13) becomes

$$G(\alpha)G(\beta)/F(\alpha)F(\beta) = G(\alpha + \beta)/F(\alpha + \beta), \quad (14)$$

which is valid for all α and β in some neighborhood of the origin. With the substitution $H = G/F$, Eq. (14) is a complex functional equation:

$$H(\alpha)H(\beta) = H(\alpha + \beta), \quad (15)$$

where the function H is continuous and equals 1.0 at the origin (so it is nonvanishing in a neighborhood of the origin). Using ideas from functional equation theory,[20] one can easily show that all solutions to Eq. (15) take the form $H(\alpha) = \exp(i2\pi c\alpha)$, where c is a real constant. So, in a neighborhood of the origin, $G(\alpha) = \exp(i2\pi c\alpha)F(\alpha)$; i.e., the characteristic function of $g(x)$ agrees with that of $f(x + c)$ for some c. But for every c, $f(x + c)$ is a probability-density function with bounded support, and consequently its characteristic function is uniquely determined for all α by its values in a neighborhood of the origin (specifically, by its derivatives at the origin, which determine the function everywhere through a power series[19]). It follows that $G(\alpha) = \exp(i2\pi c\alpha)F(\alpha)$ for all α, which proves the theorem for one-dimensional image functions. The proof for two-dimensional images is exactly analogous. Two more proofs based on different arguments can be found in Refs. 10 and 11. One proof in Ref. 11 is constructive, in the sense that it shows how, in principle, any finite image can be reconstructed from its triple correlation function.

The TCU theorem does not hold for images of infinite size, and counterexamples can be quite simple; e.g., the band-limited one-dimensional integrable functions $\mathrm{sinc}^2(x)(1 \pm \cos 6\pi x)$ have identical triple correlations[11] [$\mathrm{sinc}(x) = (\sin \pi x)/\pi x$]. (The critical difference between the finite and infinite cases is that the transforms of finite images have only isolated zeros, while the transforms of infinite images can vanish over intervals. Reference 11 discusses that difference in detail.) But for images of finite size the scope of the TCU theorem can be readily extended from the specific class of image functions defined above (functions proportional to probability-density functions on $\mathcal{R}^2$) to the more general class of all measures on $\mathcal{R}^2$ that are proportional to probability distributions with bounded support in the plane. (The characteristic function argument given above goes through in that case in exactly the same way.) That class would appear to be sufficiently general that it could describe all images of finite size, so it seems fair to say that all such images are uniquely determined up to translation by their triple correlation functions. (In particular, images defined on discrete domains, such as the lattice of two-dimensional integers, can be identified with probability distributions concentrated on discrete points in the plane, and the TCU theorem holds for all finite images of that sort.)

D. Uniqueness Theorem for Third-Order Statistics

The TCU theorem implies that two physically distinct black and white texture samples (for example, a pair of images generated by the odd and even texture algorithms, such as the left and right sides of Fig. 1) can never have identical third-order statistics: if two finite-size black and white images have identical third-order statistics, they must have identical triple correlation functions and thus be physically identical. Consequently, the odd and even textures cannot provide a demonstration of visual discrimination of textures with identical third-order statistics. And, in general, no such demonstration is possible in principle.

The truth of the last assertion for black and white images is apparent from what has been said so far, but to prove it for images composed of more than two colors requires a bit more argument. Suppose that f and g are two finite images composed of L discrete colors labeled $0, 1, \ldots, L - 1$, where 0 is the background color. Each image is described by L indicator functions f_i and g_i, $i = 0, 1, \ldots, L - 1$, where $f_i(x, y) = 1$ if the point (x, y) has color i in image f and $f_i(x, y) = 0$ otherwise, and g_i is defined analogously for image g. It is assumed that all these indicator functions are locally integrable. The third-order statistics of f and g are defined by Eq. (6). Suppose that all the third-order statistics of g are the same as those of f. Comparing Eqs. (6) and (11) (and recalling that image areas are uniquely determined by third-order statistics), one can see that the identities $s_{3,i,i,i,g} = s_{3,i,i,i,f}$, $i = 1, 2, \ldots, L - 1$,

imply that each pair of foreground color indicator functions, g_i and f_i, has identical triple correlation functions, and the TCU theorem thus implies that for each i, $g_i(x, y) = f_i(x + a_i, y + b_i)$ for some pair of constants a_i, b_i. In other words, each of the indicator functions g_i for image g is identical to the corresponding indicator function f_i for image f, except for a possible translation (a_i, b_i). It remains to show that the vectors (a_i, b_i) are all the same, so that the entire image g is a bodily translation of image f. The identity $s_{3,i,i,1,g} = s_{3,i,i,1,f}$ implies that

$$\iint_{\mathcal{R}^2} g_i(x,y)g_i(x+h_1,y+v_1)g_1(x+h_2,y+v_2)\mathrm{d}x\mathrm{d}y = \iint_{\mathcal{R}^2} f_i(x,y)f_i(x+h_1,y+v_1)f_1(x+h_2,y+v_2)\mathrm{d}x\mathrm{d}y \tag{16}$$

for each $i = 2, \ldots, L - 1$. Taking Fourier transforms of both sides of Eq. (16) yields

$$G_i(\alpha_1,\beta_1)G_i(\alpha_2,\beta_2)G_1(-\alpha_1-\alpha_1,-\beta_1-\beta_2) = F_i(\alpha_1,\beta_1)F_i(\alpha_2,\beta_2)F_1(-\alpha_1-\alpha_2,-\beta_1-\beta_2), \tag{17}$$

where G_i and G_1 are the transforms of g_i and g_1, respectively, and F_i and F_1 are the transforms of f_i and f_1, respectively. Because the indicator functions f_j are nonnegative, $F_j(0,0) \neq 0$ unless $f_j = 0$ (in which case g_j does also, and both can be ignored); and, because they are locally integrable functions with bounded support, their transforms are continuous and hence nonzero in some neighborhood of the origin. Since $g_j(x, y) = f_j(x + a_j, y + b_j)$ for each j, $G_j(\alpha, \beta) = \exp[i2\pi(a_j\alpha + b_i\beta)]F_j(\alpha, \beta)$, and substituting those expressions for G_1 and G_i on the left-hand side of Eq. (17) and dividing out the common F_1 and F_i factors, one finds that, within a neighborhood of the origin,

$$\exp\{i2\pi[(a_i - a_1)(\alpha_1 + \alpha_2) + (b_i - b_1)(\beta_1 + \beta_2)]\} = 1 \tag{18}$$

for all values of α_1, α_2, β_1, and β_2. That relation is possible only if $a_i = a_1$ and $b_i = b_1$. Consequently, all the translation vectors (a_i, b_i) are the same, and $g(x, y) = f(x + a, y + b)$. Hence the following is proved:

Uniqueness Theorem for Third-Order Statistics Two finite-sized images $f(x, y)$ and $g(x, y)$ composed of the same set of discrete colors have identical third-order statistics if and only if f and g are identical up to a translation [i.e., $g(x, y) = f(x + a, y + b)$ for some pair of constants a, b].

4. STATISTICAL COMPARISONS OF STOCHASTIC TEXTURES

The uniqueness theorem of Section 3 guarantees that physically different texture samples cannot have identical third-order statistics, but it says nothing about how close those statistics can be. The relationship between physical similarity and third-order statistical similarity (that is, triple correlation similarity) is not at all straightforward. For example, rather small finite segments of the one-dimensional image functions $\mathrm{sinc}^2(x)(1 + \cos 6\pi x)$ and $\mathrm{sinc}^2(x)(1 - \cos 6\pi x)$ have triple correlation functions that differ by only a tiny amount, though the functions themselves are obviously quite different. So while we know that odd and even texture samples cannot have identical third-order statistics, it is still possible that their third-order statistics might be similar—close enough that one may justifiably call them identical on some scale. The same possibility exists for odd and even samples compared with samples of the coin-toss texture. Of course, the scatterplots in Fig. 2 show that when the complete set of third-order statistics of sample odd, even, and coin-toss images are compared, there is vast disagreement overall. But that comparison is imprecise because it lumps together large-triangle statistics, which depend on only a few pixels and consequently can be expected to have a large variance, and small-triangle statistics, which depend on many pixels and should provide a fairer picture of the actual structural similarity of sample images. This section takes a closer look at the statistics of sample images generated by the odd, even, and coin-toss processes. All these processes are probabilistic algorithms that color a rectangular array of square cells (pixels) with R rows $r = 0, 1, \ldots, R - 1$ and C columns $c = 0, 1, \ldots, C - 1$. Each pixel is black or white; the pixel value $p(c, r) = 1$ if (c, r) is black, and $p(c, r) = 0$ if (c, r) is white. A coloring algorithm, such as the odd texture, defines a stochastic process $\langle \mathbf{p}(c, r);\ c = 0, 1, \ldots, C - 1;\ r = 0, 1, \ldots, R - 1 \rangle$ whose random variables $\mathbf{p}(c, r)$ are 0–1. The natural sample space of such a process is the set of all 2^{CR} possible colorings of a $C \times R$ array; a given process is an assignment of probabilities to each of those colorings. A specific coloring will be called an image. As usual, an event such as $\{\mathbf{p}(0,0) = 1\}$ is the subset of all images in the sample space that have the property in question; the ensemble probability of that event, $P\{\mathbf{p}(0, 0) = 1\}$, is the sum of the probabilities of the images in that subset.

Our concern is the relationship between ensemble probabilities and the statistics of specific sample images generated by a process. First-, second-, and third-order statistics for discrete images were defined in Subsection 2.C. For a stochastic process $\langle \mathbf{p}(c, r) \rangle$ those statistics are random variables whose values vary from image to image. Interest centers on their means (the ensemble statistics) and variances. Here it is natural to focus on the renormalized statistics defined by Eqs. (8′) and (9′), whose means have a simple relationship to ensemble probabilities. The third-order statistic function $\mathbf{s}_{3,p}$ includes the first-order statistic $\mathbf{s}_{1,p}$ (when all its arguments are zero) and the second-order statistic function $\mathbf{s}_{2,p}$ [when one of the n_i, m_i) is $(0, 0)$], so comparisons of the third-order statistics of sample images include the first- and second-order statistics as well. All the textures considered here have identical third-order ensemble probabilities, and all three are third-order ergodic in the limit, as image size becomes infinite. The question is, How close are their third-order sample statistics for images of visually meaningful sizes? For the coin-toss texture it is clear in advance what will happen, since the mathematics is straightforward. But it is still useful to take an empirical look at the statistics of coin-toss samples, as a baseline for the comparison of samples of the odd and even textures, where it is not so obvious what to expect.

A. Coin-Toss Textures

Here one creates an image by coloring each pixel independently black or white, each with probability 0.5. All 2^{CR} possible colorings can occur and all are equally likely, so each image in the sample space has probability $1/2^{CR}$. The first-, second-, and third-order ensemble probabilities are, respectively, $P[\mathbf{p}(c,r) = 1] = 1/2$ if (c,r) is inside the array and $P[\mathbf{p}(c,r)] = 0$ otherwise; $P[\mathbf{p}(c,r) = i, \mathbf{p}(c+n, r+m) = j] = 1/4$ for all possible combinations $i,j = 0,1$, as long as the pixels are distinct and inside the array, and that probability equals zero otherwise [or 1/2 if $(n,m) = (0,0)$ and (c,r) is in the array]; and $P[\mathbf{p}(c,r) = i, \mathbf{p}(c+n_1, r+m_1) = j, \mathbf{p}(c+n_2, r+m_2) = k] = 1/8$ for all values of $i,j,k = 0,1$ when all three pixels are in the array and that probability equals zero otherwise (or 1/4 or 1/2 in the obvious special cases). Consequently, the expected value of the renormalized third-order statistic, $E[\mathbf{S}_{3,p}(n_1,m_1,n_2,m_2)]$, is 1/2 (when all the arguments are zero), 1/4 [when one of the (n_i,m_i) is $(0,0)$, or $(n_1,m_1) = (n_2,m_2) \neq (0,0)$], or 1/8 (for all other arguments). The third-order statistics of specific images will vary around those expected values, but as array size increases, that variance should decrease for fixed values of the arguments n_i, m_i because the coin-toss process is third-order ergodic in the limit: with probability 1,

$$\lim_{R\to\infty, C\to\infty} \mathbf{S}_{3,p}(n_1,m_1,n_2,m_2) = E[\mathbf{S}_{3,p}(n_1,m_1,n_2,m_2)]. \tag{19}$$

In other words, for every fixed set of arguments n_i, m_i, the variance of the sample statistic $\mathbf{S}_{3,p}(n_1,m_1,n_2,m_2)$ shrinks to zero as sample images become infinitely large. (Diaconis and Freedman[7] show that all their ϕ-matrix textures are third-order ergodic; the coin-toss texture is one special case, and the even texture is another.) However, an image size increases, so does the number of possible third-order statistics, since $\mathbf{S}_{3,p}$ can be nonzero for larger arguments. The variance of the newly arriving statistics is always large, since they depend on only a few pixels. Consequently, comparisons of the complete statistics of two sample images always show considerable disagreement. Figure 7 illustrates those points. It shows scatterplots that compare the third-order statistics of two sample coin-toss images. [If a given statistic has the same value for both images, its data point falls on the diagonal line. If all the statistics equaled their expected values, all the data points would fall at (1/2,1/2), (1/4,1/4), or (1/8,1/8).] In Fig. 7(a) the array sizes are 10×10, and the graph compares all the third-order statistics of the two images. Figure 7(b) compares the same set of statistics ($0 \le n_i, m_i \le 9$) for 50×50 arrays. The law of large numbers is clearly doing its job. Figure 7(c) compares the complete set of third-order statistics for the 50×50 images.

B. Odd and Even Textures

The odd and even algorithms[6,13] both begin by coloring all the pixels of row 0 and column 0 by independent tosses of a fair coin. For the creation of an odd image the remaining pixels are then colored recursively, under the rule that every 2×2 block of four adjacent pixels must contain an odd number of blacks. For the creation of an even image, the remaining pixels are colored so that every 2×2 block contains an even number of blacks. Equivalently, for the even texture, one could color each pixel (c,r) outside the zero row and column so that it satisfies the constraint

$$p(c,r) = p(c,0) + p(0,r) + p(0,0) \pmod 2, \tag{20}$$

and for the odd texture, one could use the constraint

$$p(c,r) = p(c,0) + p(0,r) + p(0,0) + cr \pmod 2. \tag{21}$$

Every initial coloring of row 0 and column 0 creates a

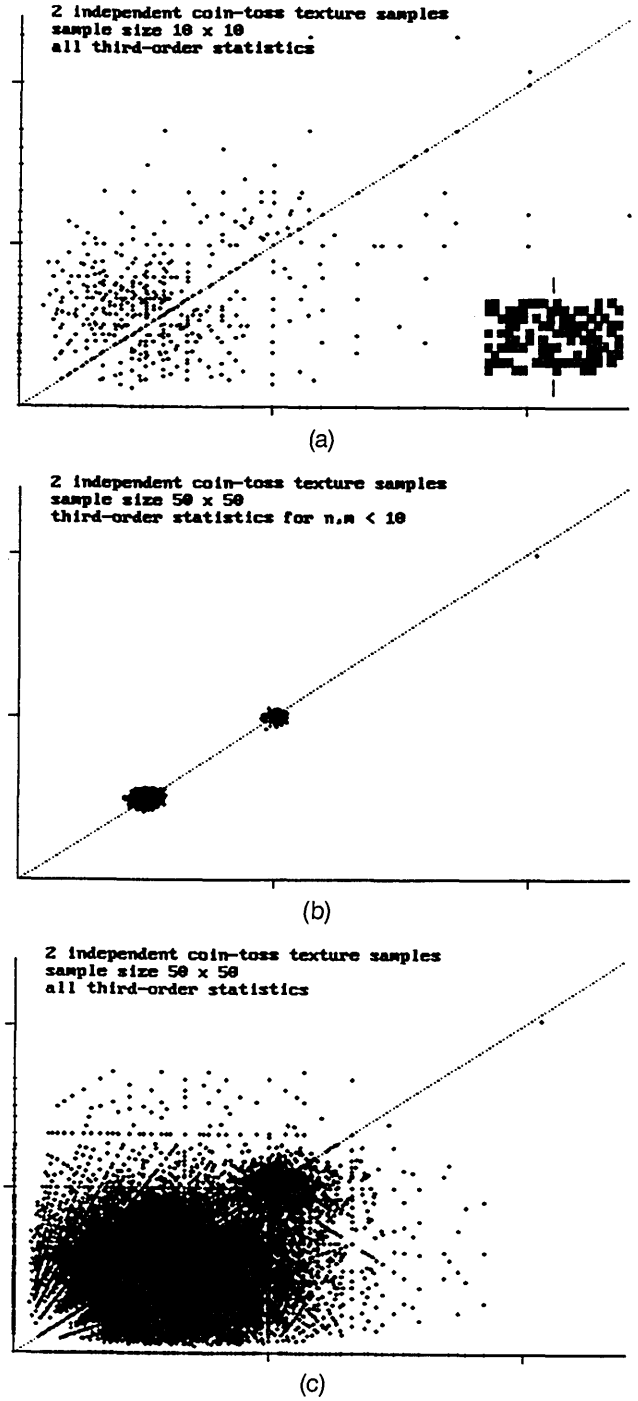

Fig. 7. (a) Scatterplot that compares all third-order statistics for two independent 10×10 samples of coin-toss texture. The axes are the same as those in Fig. 2. (b) Comparison of the same statistics for independent 50×50 coin-toss samples. (c) Comparison of all third-order statistics for the same 50×50 samples.

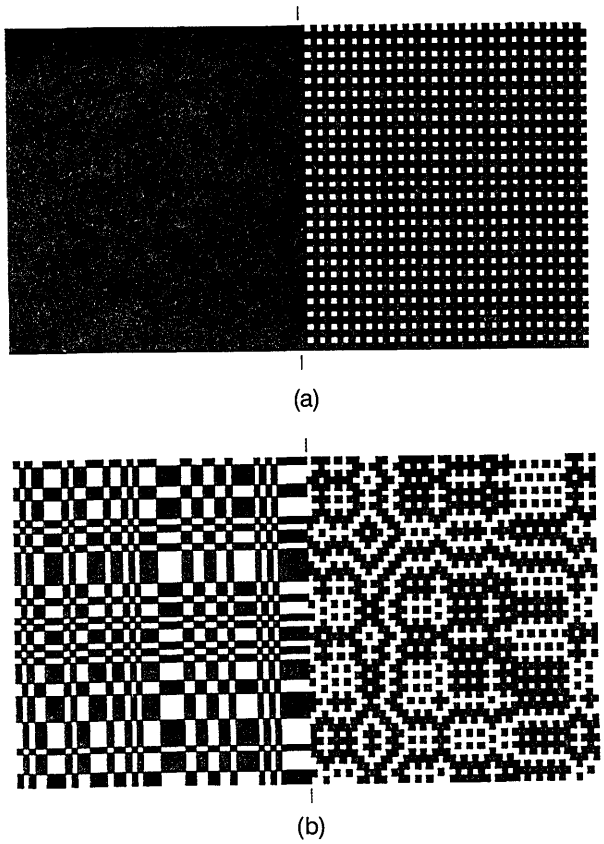

Fig. 8. (a) Solid black even image and its odd match. (b) Nonmatched 50 × 50 samples of odd (right side) and even (left side) texture. The two images are created with different initial colorings for the zeroth row and column. (The odd and even images in Fig. 1 are a matched pair; the same initial coloring generated both images.)

unique odd and even texture sample. All 2^{C+R-1} such colorings are equally probable, so every odd image and every even image has probability $\frac{1}{2}^{C+R-1}$. (When the initial coloring of row zero and column zero is all black or all white, the resulting even image is solid black or solid white. Intuitively, those two images are obviously nongeneric even images, but their ensemble probabilities are the same as all the others.)

The definitions in Eqs. (20) and (21) can be restated verbally in a helpful way: to create an even texture, one should color row 0 and column 0 by coin tosses and then, for each row $r > 0$, color the entire row the same as row 0 or the opposite, according to whether the column 0 pixel in row r has the same color as the column 0 pixel in row 0 or the opposite. To create the odd texture for the same initial coloring, one should reverse the even image coloring in every pixel whose coordinates are both odd numbers. That description makes apparent the one-to-one correspondence between odd and even texture samples: each even image is matched to a unique odd image, which is created by a reversal of the coloring of all its odd–odd pixels.[21] [As a curiosity, the top of Fig. 8(a) shows the solid black even image and its odd match.] Statistical and visual comparisons between odd and even texture samples need to take that point into account; they can involve either matched or unmatched pairs.

Visually, matched odd and even samples seem just as highly discriminable as unmatched samples. [The odd and even samples in Fig. 1 are a matched pair. Figure 8(b) shows an unmatched pair.] That observation is interesting, because, in a matched pair, three fourths of the pixels are colored identically. Statistically, matched pairs have a strong odd–even correlation between statistics whose arguments are all even numbers [shown below in Fig. 9(c)]—a correlation that is absent for unmatched pairs (cf. Fig. 10 below). The presence versus absence of this correlation apparently goes unnoticed by the eye, thus providing another bit of evidence (in addition to the point about the odd and coin-toss textures mentioned in Section 1) that texture discriminability is not governed by global third-order statistical similarity.

The strong correlation between even-argument statistics of matched odd and even texture samples can be understood from the fact that if $\langle p_o(c,r)\rangle$ and $\langle p_e(c,r)\rangle$ are a matched pair of odd and even images, then $2p_o(c,r) - 1 = (-1)^{cr}(2p_e - 1)$ for all pixels. If one sets $o(c,r) = 2p_o(c,r) - 1$ and $e(c,r) = 2p_e(c,r) - 1$, that relationship implies that for even-numbered arguments the second-order statistics of the arrays $\langle o\rangle$ and $\langle e\rangle$ are perfectly correlated, since

$$\sum_{c=0}^{C-1}\sum_{r=0}^{R-1} o(c,r)o(c+2n, r+2m) = \sum_{c=0}^{C-1}\sum_{r=0}^{R-1}(-1)^{2cr+2cm+2rn}e(c,r)e(c+2n,r+2m)$$

and $(-1)^{2cr+2cm+2rn} \equiv 1$. Consequently, for even arguments the difference between the second-order statistics of matched odd and even images depends only on the difference between the total number of black pixels in the two images, which tends to be small, since three fourths of the pixels are always initially identical, and the law of large numbers averages out the rest. Similar considerations explain why the other (i.e., the truly third-order) statistics will be highly correlated for all-even arguments.

Using a recursive argument based on Eqs. (20) and (21), Julesz *et al.*[6] proved that both the odd and even textures have the same third-order ensemble probabilities as those of coin-tossing. One could also reach that conclusion by using the verbal restatement of the even algorithm to show that its third-order probabilities are the same as those of coin tossing (by considering the possibilities for representative triples of pixels) and then appealing to the matched-pair property to show the same for the odd process. One could also use that same property to show that since the even texture process is third-order ergodic in the limit, as image size becomes infinite the odd process is also. (The argument is straightforward but long winded; it is omitted here.)

Figure 9(a) compares the complete set of third-order statistics of a matched pair of small (10 × 10) odd and even texture samples (shown in the inset). Figure 9(b) compares the same statistics ($0 \le n_i, m_i \le 9$) for matched 50 × 50 odd and even images (the images in Fig. 1), and Fig. 9(c) compares their statistics for arguments in the same range but now for even-numbered arguments only (showing the strong hidden correlation mentioned above). Figure 10 shows the same comparisons for the pair of unmatched odd and even samples from Fig. 8. Apart from the absence of the strong correlation for statistics with

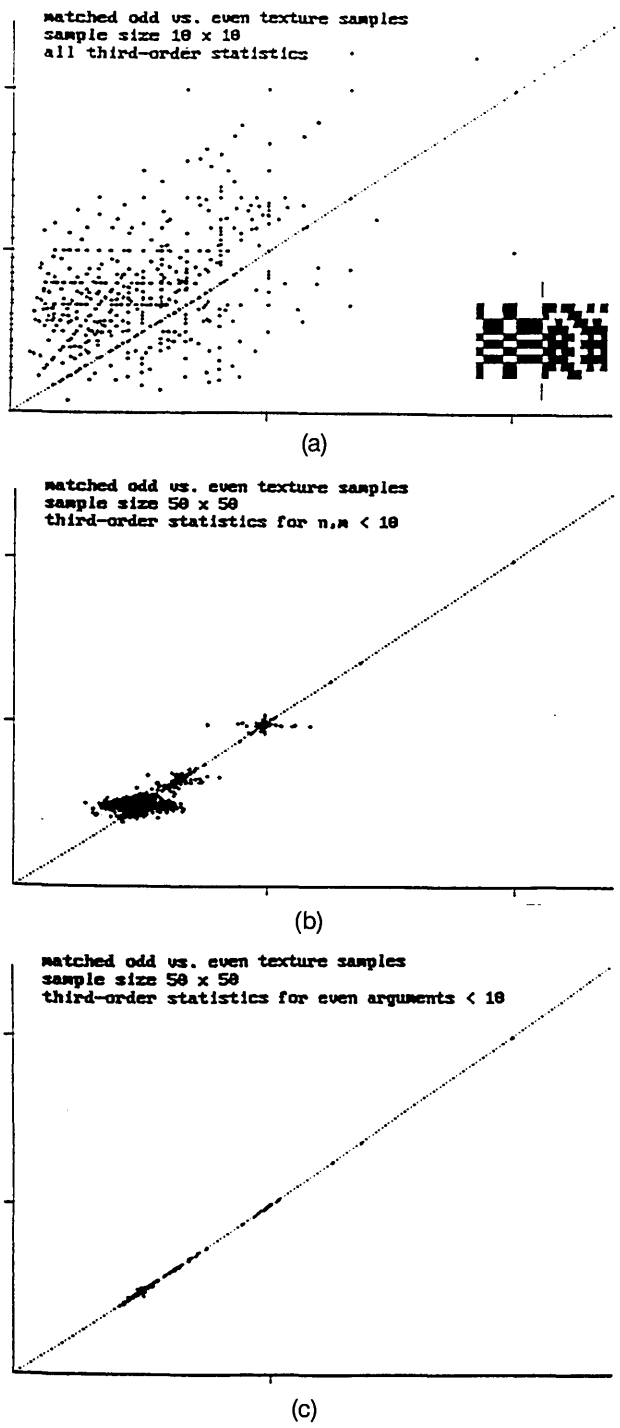

Fig. 9. (a) Scatterplot that compares all third-order statistics of 10×10 matched samples of odd and even texture (shown in the inset). The axes are the same as those in Fig. 2. (b) Comparison of the same third-order statistics ($0 \leq n_i, m_i \leq 9$) for a matched pair of 50×50 odd and even samples (the samples in Fig. 1). (c) Comparison of third-order statistics in the same range as that in (b) but only for statistics whose arguments are all even numbers. For this subset of statistics, matched odd and even samples are strongly correlated.

all-even arguments [compare Fig. 10(c) with Fig. 9(c)], the overall statistical similarity of unmatched odd–even pairs is nearly the same as that for matched pairs. In both cases, odd and even sample images are statistically quite a bit less similar than are independent coin-toss samples, though not in any qualitatively revealing way. The main overall difference is that the variance of odd and even sample statistics is greater than that of coin-toss statistics.

Direct comparisons of odd and even images with coin-toss images show this point clearly. Figure 6 compared the small-triangle ($0 \leq n_i, m_i \leq 9$) statistics of 50×50 odd and coin-toss sample images, and Fig. 11 does the same for 50×50 even-versus-coin-toss samples (the images are those in Fig. 2). The horizontal elongation of the data point clouds in Figs. 6 and 11 indicates the greater variance of the odd and even image statistics. That fact is not surprising when one considers that the statistics of

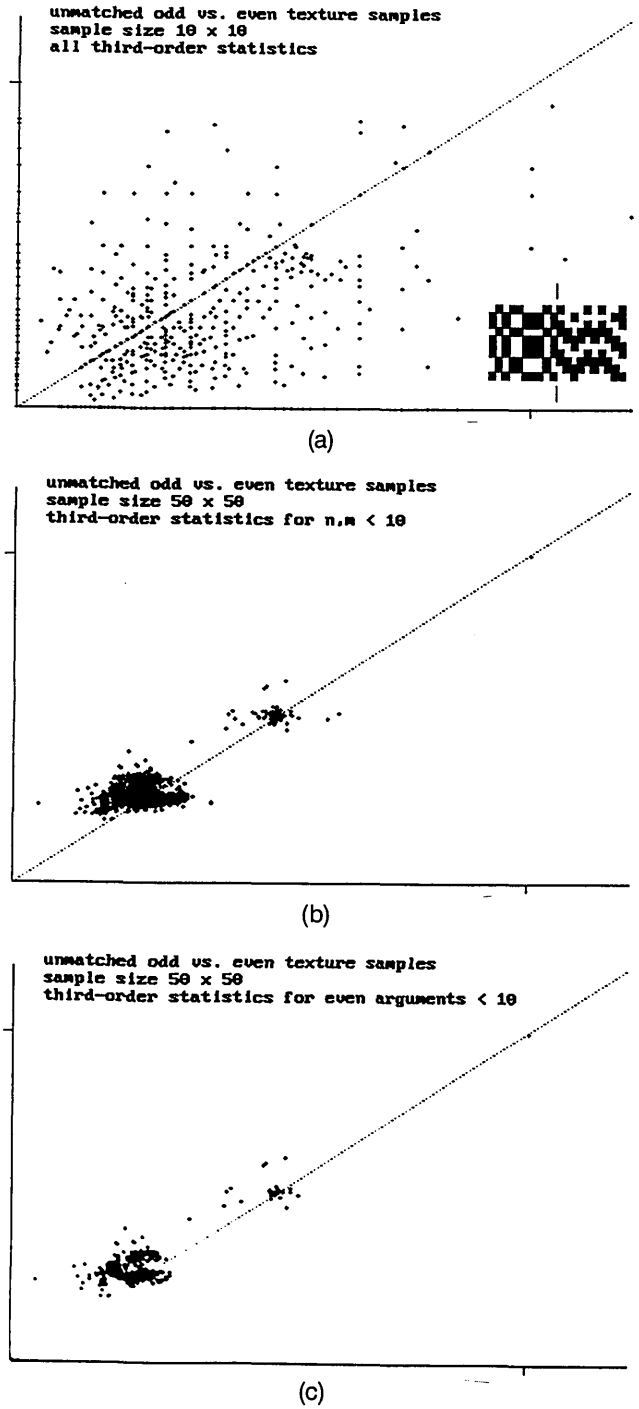

Fig. 10. Same comparisons as those in Fig. 9 but for the non-matched odd and even samples shown in Fig. 8.

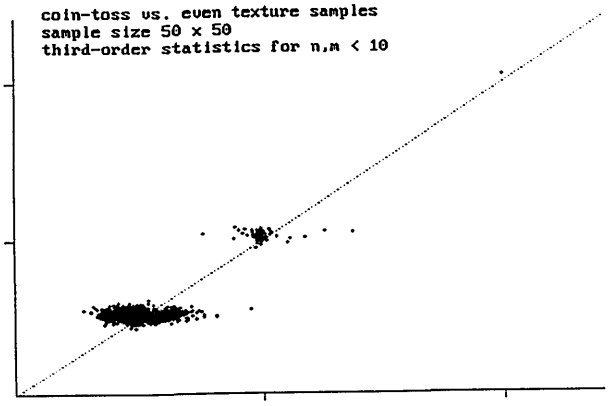

Fig. 11. Comparison of small-argument ($0 \le n_i, m_i \le 9$) third-order statistics of the 50×50 even and coin-toss texture samples fom Fig. 1. Coin-toss values are on the y axis.

a coin-toss sample depend on CR independent random variables, so one can expect that their variance will decrease as $1/CR$, while in odd and even samples there are only $C + R - 1$ independent random variables (the pixels in row 0 and column 0), so the variance of their statistics will decrease at the much slower rate $1/(C + R - 1)$. That point suggests that in scatterplots comparing the small-triangle statistics of coin-toss versus odd or even texture samples containing CR pixels, the data-point clouds should be roughly elliptical, with aspect ratios of the order of $[CR/(C + R - 1)]^{1/2}$. As a rule of thumb, that is a fairly accurate description of what one finds (e.g., for 50×50 images that ratio would be 5:1).

The overall conclusion of the preceding analysis is that, apart from the strong correlation between statistics with all-even arguments (shown in Fig. 9), the third-order statistics of sample odd, even, and coin-toss images as large as 50×50 pixels cannot be described as identical. (Extending the comparisons to 100×100 arrays still leads to the same conclusion.) If all the third-order statistics of matched odd and even texture samples showed the same strong agreement as those with all-even arguments, it might reasonably be argued that the odd and even algorithms create viewable images whose third-order statistics are essentially identical. But evidently that is not the case.

5. COUNTEREXAMPLES TO THE JULESZ CONJECTURE

A. Previous Examples

Attempts to create counterexamples to Julesz's second-order conjecture have involved two types of texture. One is exemplified by the odd and even textures: probabilistic algorithms that color the squares of a checkerboard black or white. (No counterexamples that involve multicolored images seem to have been proposed.) This class includes other textures based on Gilbert's[13] floater (or glider[14]) algorithms, Diaconis and Freedman's[7] ϕ-matrix textures, and the textures devised by Gagalowicz.[12] The floater and ϕ-matrix techniques create pairs of stochastic processes that have identical second-order ensemble probabilities, but the second-order statistics of sample images produced by different processes are not constrained to be identical, and, as previous sections have demonstrated, they will generally be quite different. (The odd and even textures are both floater textures; the even and coin-toss textures are ϕ-matrix textures.) Consequently, Gagalowicz[12] sought to create discriminable stochastic texture samples with second-order statistics constrained to be as close as possible, using linear programming. He succeeded in creating texture samples that are quite easily discriminable despite having second-order statistics that differ by only approximately 2%. That is not quite an identity, but it is surely close enough to dispose of the Julesz conjecture for all practical purposes. The only real drawback of Gagalowicz's counterexamples is the complexity of their construction: one would prefer images that are easy to create, and, of course, exact second-order statistical identity is better than near identity.

The other class of counterexamples in the literature involves textures composed of micropatterns: a small pattern (e.g., a letter or an abstract shape) is replicated at numerous sites in the plane to create a texture sample. (Here again, all the examples involved black and white images.) The mathematics underlying this approach can be understood in terms of autocorrelation. Recall that if two black and white images have the same area and the same autocorrelation function (the same power spectrum), their second-order statistics are the same. Suppose that $\phi(x, y)$ is a 0–1 function that describes a black ($\phi = 1$) micropattern on a white ($\phi = 0$) background. Replicating that pattern at N sites in the plane corresponds to convolving ϕ with a function of the form $\rho(x, y) = \Sigma_{i=1}^{N} \delta(x - x_i)\delta(y - y_i)$, where the locations (x_i, y_i) are chosen so that none of the replicas overlap. The result is a black and white image function $f: \mathcal{R} \rightarrow [0, 1]$ given by

$$f(x, y) = \rho(x, y) * \phi(x, y) = \sum_{i=1}^{N} \phi(x - x_i, y - y_i).$$

The autocorrelation function of $f(x, y)$, $a_f(h,v)$, is the convolution of the autocorrelation functions of the micropattern ϕ and the replication function ρ: $a_f(h,v) = a_\rho(h,v) * a_\phi(h,v)$. If another micropattern $\gamma(x, y)$ can be found whose area and autocorrelation function are the same as those of pattern $\phi(x, y)$, and if one replicates that pattern at the same set of locations (x_i, y_i), creating the texture sample $g(x, y) = \rho(x, y) * \gamma(x, y)$, then, since $a_\gamma(h,v) = a_\phi(h,v)$, it follows that $a_g(h,v) = a_f(h,v)$; i.e., the two samples have the same autocorrelation function. Since their areas are also the same, the two texture samples have identical second-order statistics.

The problem is to find micropatterns with identical areas and autocorrelation functions that yield spontaneously discriminable textures. The simplest way in which one can create such micropatterns is to take any pattern $\phi(x, y)$ and rotate it 180°, turning it into $\phi(-x, -y)$. Since $\phi(-x, -y)$ and $\phi(x, y)$ have the same power spectrum [i.e., $\Phi(\alpha, \beta)\Phi(-\alpha, -\beta)$, where Φ is the transform of ϕ], they have the same autocorrelation function. So a texture sample composed of, say, inverted letters A has the same second-order statistics as those of a sample composed of upright letters A located at the same positions. One might expect that such textures would be easily discriminable for micropatterns that look quite different after a 180° rotation. But typically that is not the case: one of the surprising facts about texture percep-

tion that emerged from Julesz's research[2,3] is that pairs of texture samples created in that way are usually not spontaneously discriminable. That fact provided strong evidence in favor of the Julesz conjecture.

In their search during the 1970's for counterexamples to this conjecture, Julesz and his co-workers did not succeed in finding pairs of micropatterns with identical autocorrelation functions that produced discriminable textures. However, they did present examples of discriminable textures based on micropatterns that do not themselves have identical autocorrelations but that would yield distinct textures with identical autocorrelation functions if they were replicated infinitely often throughout the plane, with each replica independently rotated by a random amount.[4,5,8] (Random rotation forces the power spectra of the textures to be circularly symmetric, erasing the angular differences that distinguish the spectra of the individual micropatterns.) That technique was referred to as the four-disk method (named after the prototypical micropatterns). Its drawback was pointed out by Gagalowicz[12]: one cannot expect finite samples of texture that correspond to distinct micropatterns to have identical power spectra, because they do not contain enough randomly rotated micropatterns—the actual power spectra of samples do not achieve perfect circular symmetry, so, in practice, the spectral differences between samples composed of distinct micropatterns will not be erased.

B. New Counterexamples to the Julesz Conjecture

This subsection describes a simple technique for the creation of discriminable texture samples (e.g., Figs. 4 and 5) whose second-order statistics are exactly identical. The textures are created by the replication of black and white micropatterns that themselves have identical second-order statistics, so they can be said to have identical second-order statistics locally as well as globally. The construction principle is the same in one and two dimensions. In the one-dimensional case (for the creation of bar-code textures like those in Fig. 4) the idea is as follows. Suppose that $p(x)$ and $q(x)$ are arbitrary real functions on the line, with Fourier transforms $P(\alpha)$ and $Q(\alpha)$, respectively. The transforms of the convolutions $p(x) * q(x)$ and $p(x) * q(-x)$ are $P(\alpha)Q(\alpha)$ and $P(\alpha)Q(-\alpha)$, respectively, so the power spectra of $p(x) * q(x)$ and $p(x) * q(-x)$ both equal $|P(\alpha)|^2|Q(\alpha)|^2$. Thus $p(x) * q(x)$ and $p(x) * q(-x)$ have the same autocorrelation function. Then, if $h(y)$ is any third function, and p, q, and h are all nonnegative, the functions $h(y)[p(x) * q(x)]$ and $h(y)[p(x) * q(-x)]$ represent two-dimensional images with identical autocorrelation functions, since both have the power spectrum $|H(\beta)|^2|P(\alpha)|^2|Q(\alpha)|^2$. To use this principle to create black and white micropatterns, one supposes that p and q are both composed entirely of unit-amplitude delta functions located at selected integers, e.g.,

$$p(x) = \sum_{n=0}^{4} \delta(x - 4n^2), \tag{22}$$

$$q(x) = \sum_{n=0}^{4} \delta[x - (4n^2 + n + 1)]. \tag{23}$$

In that case the convolutions $p(x) * q(x)$ and $p(x) * q(-x)$ both consist entirely of unit-amplitude deltas located at integers. (The choices $4n^2$ and $4n^2 + n + 1$, $n = 0, 1, \ldots, 4$, are motivated by the need to prevent the deltas from piling up in the convolutions. They are not unique in that respect.) If the function $h(y) = 1$ for $|y| <$ some constant L, then $\phi(x, y) = h(y)[p(x) * q(x)]$ and $\gamma(x, y) = h(y)[p(x) * q(-x)]$ both describe bar-code micropatterns composed of 25 vertical black lines of length L, distributed over 133 spaces (so that their areas are equal), and those micropatterns have identical autocorrelation functions. Figure 12 shows those micropatterns and also the patterns created by Eqs. (22) and (23) when the sums run from $n = 0$ to 3 instead of 4. Texture samples that one constructs by locating replicas of micropatterns ϕ and γ at the same sites inherit that identity and consequently have identical second-order statistics.

To extend the principle to two-dimensional micropatterns, one supposes $p(x, y)$ and $q(x, y)$ are arbitrary functions on the plane. Then $p(x, y) * q(x, y)$ and $q(x, y) * q(-x, -y)$ have the same autocorrelation function, since each has the power spectrum $|P(u,v)|^2|Q(u,v)|^2$. Now let $p(x, y)$ be a sum of two-dimensional delta functions that represent a set of locations in the plane, say $(x_1, y_1), \ldots, (x_n, y_n)$, and let $q(x, y)$ be a 0–1 function that represents a black and white pattern. The convolution $p(x, y) * q(x, y)$ creates replicas of $q(x, y)$ centered at the (x_i, y_i) locations, and $p(x, y) * (q - x - y)$ creates replicas of 180° rotations of $q(x, y)$ centered at the same locations. If none of the replicas overlap, the convolution patterns are black and white and have identical autocorrelations and hence identical second-order statistics. That identity will then be inherited by textures formed by the replication of those patterns at the same sites in the plane. Victor[22] suggested the application of that principle illustrated in Fig. 5: take $p(x, y)$ to be a set of two-dimensional delta functions arranged in a triangle (the micropatterns in Fig. 5 use six deltas), and let $q(x, y)$ be

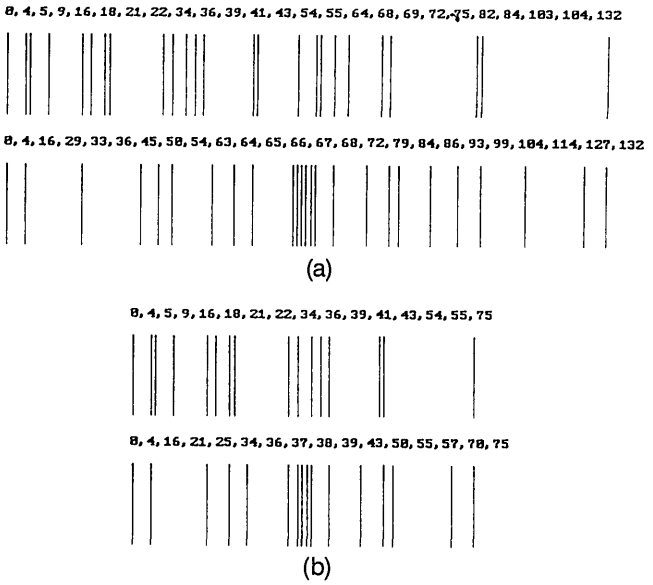

Fig. 12. Pairs of bar-code micropatterns with identical second-order statistics. The numbers are the positions of the bars. Pair (a) and (b) are $p(x) * q(x)$ and $p(x) * q(-x)$, respectively, where p and q are defined by Eqs. (22) and (23), respectively. Pair (c) and (d) result from a change in the upper limits of the sums from 4 to 3. (The functions have been translated here so that the first bars of all the patterns fall at position 0.)

some other shape that is altered by a 180° rotation, such as the letter T. Textures created in that way appear to be quite easily discriminable.

ACKNOWLEDGMENTS

Assistance and advice from A. Ahumada, M. D'Zmura, G. Iverson, and R. Kakarala are gratefully acknowledged. The author is affiliated with the Institute for Mathematical Behavioral Sciences at the University of California, Irvine (UCI) as well as with its Department of Cognitive Sciences. This research was supported in part by a UCI faculty research grant.

REFERENCES AND NOTES

1. B. Julesz, "Visual pattern discrimination," IRE Trans. Inf. Theory **IT-8,** 84–92 (1962).
2. B. Julesz, E. N. Gilbert, L. A. Shepp, and H. L. Frisch, "Inability of humans to discriminate between visual textures that agree in second-order statistics—revisited," Perception **2,** 391–405 (1973).
3. B. Julesz, "Experiments in the visual perception of texture," Sci. Am. **232,** 34–43 (1975).
4. T. M. Caelli and B. Julesz, "On perceptual analyzers underlying visual texture discrimination: Part 1," Biol. Cybernet. **28,** 167–175 (1978).
5. T. M. Caelli, B. Julesz, and E. N. Gilbert, "On perceptual analyzers underlying visual texture discrimination: Part 2," Biol. Cybernet. **29,** 201–214 (1978).
6. B. Julesz, E. N. Gilbert, and J. D. Victor, "Visual discrimination of textures with identical third-order statistics," Biol. Cybernet. **31,** 137–140 (1978).
7. P. Diaconis and D. Freedman, "On the statistics of vision: the Julesz conjecture," J. Math. Psychol. **24,** 112–138 (1981).
8. B. Julesz, "Textons, the elements of texture perception, and their interactions," Nature (London) **290,** 91–97 (1981).
9. J. R. Bergen, "Theories of visual texture perception," in *Spatial Vision,* Vol. 10 of Vision and Visual Dysfunction, D. M. Regan, ed. (Macmillan, New York, 1991).
10. H. Bartelt, A. W. Lohmann, and B. Wirnitzer, "Phase and amplitude recovery from bispectra," Appl. Opt. **23,** 3121–3129 (1984).
11. J. I. Yellott, Jr. and G. J. Iverson, "Uniqueness properties of higher-order autocorrelation functions," J. Opt. Soc. Am. A **9,** 388–404 (1992).
12. A. Gagalowicz, "A new method for texture fields synthesis: some applications to the study of human vision," IEEE Trans. Pattern Anal. Mach. Intell. **PAMI-3,** 520–533 (1981).
13. E. N. Gilbert, "Random colorings of a lattice of squares in the plane," SIAM J. Algebra Discrete Meth. **1,** 152–159 (1980).
14. J. D. Victor and M. M. Conte, "Spatial organization of nonlinear interactions in form perception," Vision Res. **31,** 1457–1488 (1991).
15. A. W. Lohmann and B. Wirnitzer, "Triple correlation," Proc. IEEE **72,** 889–901 (1984).
16. C. L. Nikias and J. M. Mendl, eds., *Workshop on Higher-Order Spectral Analysis,* Cat. No. 89TH0267-5 (Institute of Electrical and Electronics Engineers, New York, 1989).
17. S. A. Klein and C. W. Tyler, "Phase discrimination of compound gratings: generalized autocorrelation analysis," J. Opt. Soc. Am. A **3,** 868–879 (1986). That paper was apparently written without awareness of the fact that finite-sized images are uniquely determined up to translation by their triple correlation functions. It focused on the higher-order autocorrelation functions of infinitely extended periodic images, which are not uniquely determined by their triple correlations, since they do not have bounded support: two periodic functions can have identical triple correlations without being translations of each other. (A one-dimensional example pointed out by Klein and Tyler is the pair $2 + \cos x \pm \cos 3x$. Because infinitely supported periodic images have impulsive spectra, the uniqueness properties of their higher-order autocorrelations are quite different from those of finite images. Those properties are discussed in Ref. 11, and the complete theory is worked out in Ref. 18.)
18. R. Kakarala and G. J. Iverson, "Characterization of periodic functions by higher-order spectra," Tech. Rep. MBS 91-12 (Institute for Mathematical Behavioral Sciences, University of California, Irvine, Irvine, Calif., 1991).
19. W. Feller, *An Introduction to Probability Theory and Its Applications* (Wiley, New York, 1966), Vol. II.
20. J. Aczel, *Lectures on Functional Equations* (Academic, New York, 1966).
21. J. D. Victor, "Complex visual textures as a tool for studying the VEP," Vision Res. **25,** 1811–1828 (1985).
22. J. D. Victor, Department of Neurology and Neuroscience, Cornell University Medical College, Ithaca, New York 10021 (personal communication).