

Dead leaves and the dirty ground: Low-level image statistics in transmissive and occlusive imaging environments

Joel Zylberberg

*Department of Applied Mathematics, University of Washington, Seattle, Washington 98195, USA,
Department of Physics, University of California, Berkeley, California 94720, USA, and
Redwood Center for Theoretical Neuroscience, University of California,
Berkeley, California 94720, USA*

David Pfau

*Redwood Center for Theoretical Neuroscience, University of California, Berkeley, California 94720, USA and
Center for Theoretical Neuroscience, Columbia University, New York, New York 10032, USA*

Michael Robert DeWeese*

*Department of Physics, University of California, Berkeley, California 94720, USA,
Redwood Center for Theoretical Neuroscience, University of California, Berkeley, California 94720, USA, and
Helen Wills Neuroscience Institute, University of California, Berkeley, California 94720, USA*

(Received 8 June 2012; published 14 December 2012)

The opacity of typical objects in the world results in occlusion, an important property of natural scenes that makes inference of the full three-dimensional structure of the world challenging. The relationship between occlusion and low-level image statistics has been hotly debated in the literature, and extensive simulations have been used to determine whether occlusion is responsible for the ubiquitously observed power-law power spectra of natural images. To deepen our understanding of this problem, we have analytically computed the two- and four-point functions of a generalized “dead leaves” model of natural images with parameterized object transparency. Surprisingly, transparency alters these functions only by a multiplicative constant, so long as object diameters follow a power-law distribution. For other object size distributions, transparency more substantially affects the low-level image statistics. We propose that the universality of power-law power spectra for both natural scenes and radiological medical images, formed by the transmission of x-rays through partially transparent tissue, stems from power-law object size distributions, independent of object opacity.

DOI: [10.1103/PhysRevE.86.066112](https://doi.org/10.1103/PhysRevE.86.066112)

PACS number(s): 89.75.Da, 42.30.Va, 87.57.–s, 89.75.Kd

I. INTRODUCTION

Natural images are surprisingly statistically uniform. The autocorrelation function, a measure of how similar nearby pixels tend to be, is virtually universal for natural images [1–7] (Fig. 1). This is typically quantified by measuring image power spectra (Fourier transform of the autocorrelation function), which are well described by scale-invariant power-law functions with power \mathcal{P} and spatial frequency k related by $\mathcal{P}(k) \propto k^{-\alpha}$, with exponents $\alpha \approx 2$. The exponents α vary slightly from image to image, and there are small differences in average exponent α between terrestrial [2,4,7] and aquatic [8] environments, and between natural and man-made ones [6].

Intriguingly, even radiological images like mammograms have power-law power spectra [9,10], typically with larger α values, despite the fact that the physics of image formation are very different for radiological and natural images. In natural images, formed by reflection of light off of surfaces, objects tend to be opaque, and thus they occlude one another, whereas in mammograms, formed by the transmission of x-rays through breast tissue, objects are more transmissive and do not completely occlude one another. The statistics of radiological images have received less attention and are less

well understood. Interestingly, however, the powers α typically vary between mammogram images of patients with low versus high risk of developing breast cancer [10] and vary as a function of the density of the breast tissue [11], highlighting the potential clinical importance of these image statistics.

The statistical regularity of natural scenes implies that engineers can design, and evolution might have selected for, coding schemes that exploit this structure [5,12]. Indeed, the peripheral mammalian visual system appears to exploit this homogeneity by using simple filters to decorrelate the incoming signal [13–15] and more complex feature dictionaries to efficiently encode the decorrelated signal [5,16,17].

Using the intuition that the environment is composed of distinct objects, Ruderman studied a “dead leaves” model [18,19] for natural scenes, in which images are created by sequentially placing opaque, potentially overlapping circles of random brightness in random locations on a two-dimensional image plane [20] [Fig. 2(a)]. Ruderman modeled correlations between pixels by assuming a different correlation function for points falling within a visible circle than for points falling in different visible circles. Using analytical calculations he demonstrated that, so long as the diameters s of the circles follow a power-law distribution with probabilities $p(s) \propto s^{-(3+\eta)}$, the images exhibit power-law correlation functions, $C(q) \propto q^{-\eta}$, where q is the separation between pixels, and power-law power spectra, $\mathcal{P}(k) \propto k^{-(2-\eta)}$. If the circle sizes

*deweese@berkeley.edu

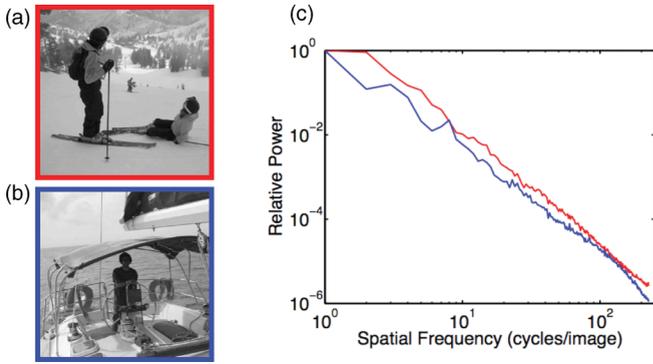


FIG. 1. (Color online) Natural images have nearly identical scale-invariant power spectra. Even very different natural images (a), (b) have similar rotation-averaged power spectra that each follow a power law (c). Line colors in panel c match the borders of corresponding panels a and b. The upper curve in panel c corresponds to the image in panel a, while the lower one corresponds to b.

are drawn from other distributions, Ruderman’s analytical calculations suggest that the power spectra that could be made to differ from a power law, contrary to the old notion [21] that the $1/k^2$ power spectra result from the mere presence of edges, each of which has a $1/k^2$ one-dimensional power spectrum (cf. Balboa *et al.* [22]). More recently, Balboa *et al.* [22] simulated the analytical examples presented by Ruderman [20], including images with the exponential distribution of object sizes that was claimed [20] to yield non-power-law power spectra. They found that these images had nearly power-law power spectra, and subsequently they reiterated the previous claim that occlusion, and not object size distributions, are the cause of power-law power spectra in natural images.

This “edges versus size distributions” debate was subsequently resolved when Hsiao and Milane demonstrated, via numerical simulations, that dead leaf models with partially transparent objects (and thus only partial occlusion) whose sizes follow a power-law distribution yield power-law power spectra, and that dead leaf models with opaque objects from other size distributions can have non power-law power spectra [23]. In other words, occlusion is neither necessary, nor sufficient, to yield power-law image power spectra. In the same paper, Hsiao and Milane computed the power spectrum of a simplified ensemble of images formed by summing the intensities of different randomly placed disks. This model was simpler than the images with partially occluding leaves that they simulated. The linearity of this model makes it relatively straightforward to compute the Fourier transform of the model images, and thus to estimate the power spectra.

Thus, to date, the two-point statistics of dead leaf image models have been analytically calculated both for fully opaque leaves [20] and for fully transmissive leaves [23]. What remains is to solve for the two-point function of images with partial occlusion, which will deepen our understanding of how opacity and image statistics interrelate along this continuum of object properties. Thus motivated, we studied a generalized dead leaves model, in which the leaves have variable transparency. While general feature probabilities have been solved exactly for the fully opaque dead leaves model [24], our transparent generalization requires other methods

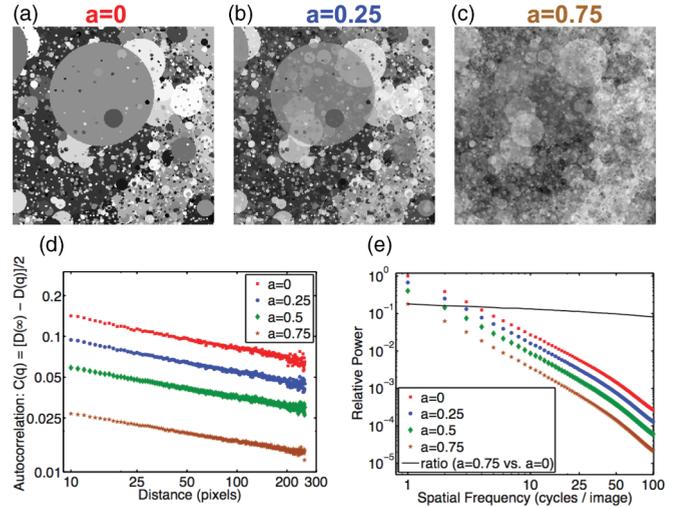


FIG. 2. (Color online) For power-law object size distributions, the two-point statistics of opaque and transmissive dead leaves images differ by a multiplicative constant. (a) A representative image from the opaque ($a = 0$) dead leaves model with circle diameters drawn from the distribution $p(s) \propto s^{-3.2}$ for $s > s_0 = 1$ pixel and circle brightnesses drawn uniformly within $b \in [-1, 1]$. (b) When the circles are partially transparent ($a = 0.25$ for all circles), but all other parameters are the same, previously occluded circles are partially visible. (c) A higher level of transparency ($a = 0.75$) results in an image that begins to approximate Gaussian pink noise, as expected from the central limit theorem [24]. (d) Autocorrelation functions of dead leaves image ensembles of different opacity levels differ only by a multiplicative constant for power-law object size distributions. The two-point functions are power-law functions of distance, with power ~ -0.2 , in good agreement with our analytical calculation. (e) Similarly the power spectra of these image ensembles are roughly power-law functions and are all the same up to a multiplicative constant. The ratio of the opaque and most transparent power spectra is nearly flat. At relatively high spatial frequencies (above ~ 20 cycles/image), corresponding to small length scales, the $q \gg s_0$ approximation in our analytical calculation fails, and slight deviations from power-law power spectra can be observed, as can deviations from constancy in the ratio.

and has not previously been systematically explored. We show herein that, so long as leaf sizes follow a power-law distribution, transparency results in an overall multiplicative factor in the two- and four-point functions but does not change their functional (power-law) form. For other size distributions, transparency does change the form of the autocorrelation function, suggesting that power-law size distributions unify the observed power spectra of natural and radiological images.

II. ANALYTICAL CALCULATION OF THE TWO-POINT FUNCTION IN THE TRANSMISSIVE DEAD LEAVES MODEL

We begin by analytically computing the two-point functions of images in our “transmissive dead leaves” environment. For image pixels values $I(\vec{x})$, the two-point function is given by $C(\vec{x}, \vec{x}') = \langle I(\vec{x})I(\vec{x}') \rangle = C(|\vec{x} - \vec{x}'|)$, where the angle brackets denote averaging over images drawn from this ensemble and the second step stems from the fact that, since our model

world is invariant under both translations and rotations, the two-point function depends only on the distance $|\vec{x} - \vec{x}'| = q$ between sample points.

The image is formed by randomly placing a circle whose diameter s is drawn from some distribution, with brightness value b , and transparency a , on a surface of diameter L . The brightnesses b will be drawn from a zero-mean distribution, and the transparencies $a \in [0, 1]$ can also be random. A value $a = 1$ specifies a fully transparent (invisible) circle, while a value of $a = 0$ specifies a fully opaque circle, as in Ruderman's model [20]. When a new circle is added, the pixel value $I(\vec{x})$ at a point \vec{x} that falls within the circle undergoes the transformation

$$I(\vec{x}) \rightarrow (1 - a)b + aI(\vec{x}). \quad (1)$$

Pixels not lying under the circle are unaffected by its addition. This process is continued *ad infinitum* to create model images (Fig. 2).

We will compute $\langle I(\vec{x})^2 \rangle$ and $C(q)$ recursively by noting that adding another leaf to an image creates a new image from the same transmissive dead leaves ensemble, and thus the (average) statistical properties must remain unchanged by this transformation [20].

Using Eq. (1), we can compute the pixel variance

$$\begin{aligned} \langle I^2(\vec{x}) \rangle &= (1 - P_{\text{in}})\langle I^2(\vec{x}) \rangle + P_{\text{in}}\langle [aI(\vec{x}) + (1 - a)b]^2 \rangle \\ \Rightarrow \langle I^2(\vec{x}) \rangle &= \frac{\langle b^2 \rangle \langle (1 - a)^2 \rangle}{1 - \langle a^2 \rangle}, \end{aligned} \quad (2)$$

where P_{in} is the probability that the point in question falls within the newly added circle. The quantity P_{in} , and thus the distribution of circle sizes, does not affect the pixel variance. It will, however, affect the spatial properties of the image, including $C(q)$.

To compute $C(q)$, consider how the pixel values of a pair of points with separation q are affected by the addition of a new leaf. After adding the leaf, either one, both, or neither of the sample points lie under the leaf, resulting in three different possible modifications to the pixel values [Eq. (1)]. These outcomes occur with probabilities $P_1(q)$, $P_2(q)$, or $P_0(q)$, respectively, which we will later compute. Equating the two-point functions before and after the addition of a new leaf, we obtain

$$\begin{aligned} C(q) &= P_0(q)C(q) + P_1(q)\langle [aI(\vec{x}) + (1 - a)b]I(\vec{x}') \rangle \\ &\quad + P_2(q)\langle [aI(\vec{x}) + (1 - a)b][aI(\vec{x}') + (1 - a)b] \rangle. \end{aligned} \quad (3)$$

Recalling the definition of the autocorrelation function and the normalization $P_0(q) + P_1(q) + P_2(q) = 1$, we find

$$C(q) = \frac{\langle b^2 \rangle \langle (1 - a)^2 \rangle P_2(q)}{P_1(q)\langle 1 - a \rangle + P_2(q)\langle 1 - a^2 \rangle}. \quad (4)$$

The quantities $\langle b^2 \rangle$, $\langle a^2 \rangle$, and $\langle a \rangle$ depend on the distributions of circle brightnesses and opacities.

To calculate $P_1(q)$, we first define $P^* = \langle s^2 \rangle / L^2$, which is the probability that any given point in the image falls within a newly deposited leaf. Here L is the diameter of the circular image area, s is the diameter of the newly added circle, and we assume $\langle s^2 \rangle \ll L^2$. The probability $P_1(q)$ that either point, but not both, falls within the circle is then $P_1(q) = 2[P^* - P_2(q)]$,

where the factor of 2 comes in because there are two such points to consider.

To determine the probability $P_2(q)$, note that, for a circle of diameter s , given that one particular point \vec{x} is within the circle (which occurs with probability s^2/L^2), the probability that another point, a distance q away, is also within the circle, is given by Ref. [20] $g(q/s \in [0, 1]) = \frac{2}{\pi} [\cos^{-1}(q/s) - (q/s)\sqrt{1 - (q/s)^2}]$, and thus

$$P_2(q) = \int_0^\infty \frac{s^2}{L^2} g(q/s) p(s) ds. \quad (5)$$

For a power-law size distribution $p(s) = (A/s_0)(s/s_0)^{-\alpha}$, where $\alpha > 3$, A is a unitless normalization constant, and s_0 is the small-size cutoff, the change of variables $u = s/q$ in the above integral yields

$$P_2(q) = A \left(\frac{s_0}{L} \right)^2 \left(\frac{q}{s_0} \right)^{-(\alpha-3)} \int_1^\infty g(1/u) u^{2-\alpha} du. \quad (6)$$

Define the integral to be $B(\alpha)$. For pixel separations much larger than the small-size cutoff of our leaf diameter distribution, $q \gg s_0$ [in which case $P^* = \frac{A}{\alpha-3} \left(\frac{s_0}{L} \right)^2 \gg P_2(q)$], Eq. (4) becomes

$$C(q) = \frac{B(\alpha)(\alpha - 3)\langle b^2 \rangle \langle (1 - a)^2 \rangle}{2\langle 1 - a \rangle} \left(\frac{q}{s_0} \right)^{-(\alpha-3)}, \quad (7)$$

yielding an image power spectrum [20] $\mathcal{P}(k) \propto \frac{\langle b^2 \rangle \langle (1 - a)^2 \rangle}{(1 - a)} k^{-(5-\alpha)}$ in which the opacity affects the power spectrum only as a multiplicative prefactor. When $a = 0$ for all circles (opaque limit), our result is equal to that of Ruderman [20], as it must be. Also note that, as one might expect, the two-point function does not depend on the size L of the image surface.

To demonstrate that leaf opacity can affect the functional form of the two-point function, we repeat the above calculations but now have all leaves be the same size s^* . The size distribution is thus $p(s) = \delta(s - s^*)$, in which case the correlation function is

$$C_\delta(q) = \frac{\langle b^2 \rangle \langle (1 - a)^2 \rangle g(q/s^*)}{2\langle 1 - a \rangle - \langle (1 + a)^2 \rangle g(q/s^*)}, \quad (8)$$

which depends nontrivially on a : for $q > s^*$, $g(q/s^*) = 0$ and the correlation function vanishes, so the large- q limit in which Eq. (7) was derived is irrelevant for δ function size distributions. Furthermore, even for fully opaque leaves, it is clear that this correlation function, which is identically zero for $q > s^*$, is not described by a power-law function of distance.

A comparison of Eqs. (2) and (7) shows that the pixel variance, and the image autocorrelation function, are multiplied by different opacity dependent prefactors. For $q = 0$, the variance and the two-point function are equal, so the fact that for $q \gg s_0$, they scale differently with changing opacity highlights that there is a qualitative change in the two-point function near the $q \sim s_0$ boundary. For natural images, the minimum object size is much smaller than our cameras can resolve, so this boundary is never encountered in practice. Furthermore, this comparison demonstrates that not all image statistics vary in the same way with changing leaf opacity.

III. ANALYTICAL CALCULATION OF THE FOUR-POINT FUNCTION FOR COLLINEAR POINTS IN THE TRANSMISSIVE DEAD LEAVES MODEL

As we have seen, the form of the two-point function is independent of leaf opacity for power-law object size distributions. At the same time, the images generated with different leaf opacities (Fig. 2) are visibly different, so there must be some difference in the image statistics (aside from the overall pixel variance) from ensembles with different object opacities. To understand this difference, we consider higher-order statistics beyond the two-point function. If the leaf brightnesses $\langle b \rangle$ are symmetrically distributed about zero, then the three-point function will vanish, and so the next possible candidate beyond the two-point function is the four-point function.

In this section, we will compute the four-point function $C_4^{\text{coll}}(\vec{x}, \vec{x}', \vec{x}'', \vec{x}''') = \langle I(\vec{x})I(\vec{x}')I(\vec{x}'')I(\vec{x}''') \rangle$ for equidistant

collinear points: $|\vec{x} - \vec{x}'| = |\vec{x}' - \vec{x}''| = |\vec{x}'' - \vec{x}'''| = q$ and $|\vec{x} - \vec{x}''| = |\vec{x}' - \vec{x}'''| = 2q$, for the dead leaves model with power-law leaf size distribution. We chose this arrangement of points because it considerably simplifies the analysis of the four-point function, for reasons that will become apparent during the calculation. Nevertheless, the calculation itself is still somewhat tedious, so some readers may wish to skip to the result at the end of this section.

As in the case of the two-point function described above, since our image ensemble is invariant under translations and rotations, the result depends only on the pixel spacing q : $C_4^{\text{coll}}(\vec{x}, \vec{x}', \vec{x}'', \vec{x}''') = C_4^{\text{coll}}(q)$. We apply the same recursive logic that we used for computing the two-point function in order to infer the four-point function, and start by enumerating all of the possible modifications to the four-point function upon the addition of a new circle. We will number the points from left to right. Thus,

$$\begin{aligned}
 C_4^{\text{coll}}(q) = & P_{\emptyset}^{\text{coll}}(q)C_4^{\text{coll}}(q) + P_{1,\emptyset}^{\text{coll}}(q)\langle [aI(\vec{x}) + (1-a)b]I(\vec{x}')I(\vec{x}'')I(\vec{x}''') \rangle + P_{2,\emptyset}^{\text{coll}}(q)\langle I(\vec{x})[aI(\vec{x}') + (1-a)b]I(\vec{x}'')I(\vec{x}''') \rangle \\
 & + P_{3,\emptyset}^{\text{coll}}(q)\langle I(\vec{x})I(\vec{x}')[aI(\vec{x}'') + (1-a)b]I(\vec{x}''') \rangle + P_{4,\emptyset}^{\text{coll}}(q)\langle I(\vec{x})I(\vec{x}')I(\vec{x}'')[aI(\vec{x}''') + (1-a)b] \rangle \\
 & + P_{1,2}^{\text{coll}}(q)\langle [aI(\vec{x}) + (1-a)b][aI(\vec{x}') + (1-a)b]I(\vec{x}'')I(\vec{x}''') \rangle \\
 & + P_{2,3}^{\text{coll}}(q)\langle I(\vec{x})[aI(\vec{x}') + (1-a)b][aI(\vec{x}'') + (1-a)b]I(\vec{x}''') \rangle \\
 & + P_{3,4}^{\text{coll}}(q)\langle I(\vec{x})I(\vec{x}')[aI(\vec{x}'') + (1-a)b][aI(\vec{x}''') + (1-a)b] \rangle \\
 & + P_{1,2,3}^{\text{coll}}(q)\langle [aI(\vec{x}) + (1-a)b][aI(\vec{x}') + (1-a)b][aI(\vec{x}'') + (1-a)b]I(\vec{x}''') \rangle \\
 & + P_{2,3,4}^{\text{coll}}(q)\langle I(\vec{x})[aI(\vec{x}') + (1-a)b][aI(\vec{x}'') + (1-a)b][aI(\vec{x}''') + (1-a)b] \rangle \\
 & + P_{1,2,3,4}^{\text{coll}}(q)\langle [aI(\vec{x}) + (1-a)b][aI(\vec{x}') + (1-a)b][aI(\vec{x}'') + (1-a)b][aI(\vec{x}''') + (1-a)b] \rangle, \tag{9}
 \end{aligned}$$

where $P_{\emptyset}^{\text{coll}}$ is the probability that none of the four collinear points fall under the newly deposited circle, $P_{i,\emptyset}^{\text{coll}}$ is the probability that only the i th point falls under the newly deposited circle, $P_{i,j}^{\text{coll}}$ is that probability that only the i th and j th collinear points fall under the newly deposited circle, and so on. Because the points are collinear, it is impossible for non-neighboring pixels to fall under a given circle unless all of the pixels between them also fall under that circle. Hence, there are no terms like $P_{1,3}^{\text{coll}}$ or $P_{1,2,4}^{\text{coll}}$ in the above equation, since they would require there to be “gaps” between neighboring pixels. Alternatively, one can include those terms but note that the probabilities associated with them are zero.

To simplify Eq. (9) to the point that we can easily solve for $C_4^{\text{coll}}(q)$, we will first expand and simplify all of the average products $\langle \cdot \rangle$, then compute all of the probabilities $P_{\{i\}}^{\text{coll}}$, and finally assemble all of these pieces.

A. Expanding and simplifying the average pixel-value products

Since the circle brightnesses b are zero-mean and independently drawn, each of the terms in which a single pixel is modified [the second through fifth terms in Eq. (9)] reduces to $\langle a \rangle P_{i,\emptyset}^{\text{coll}} C_4^{\text{coll}}(q)$. Similarly, expanding the terms in which two points fall under the circle [the sixth through eighth terms in

Eq. (9)], recalling that $\langle b \rangle = 0$, and performing a bit of algebra, each of those terms can be simplified to

$$P_{i,j}^{\text{coll}}(q) [\langle a^2 \rangle C_4^{\text{coll}}(q) + \langle (1-a)^2 \rangle \langle b^2 \rangle C_2(|k-m|q)], \tag{10}$$

where $k \neq m$, $k, m \in \{1, 2, 3, 4\} \setminus \{i, j\}$, $C_2(\cdot)$ is the two-point function that we calculated in the previous section [Eqs. (4) and (7) for power-law object size distributions], and we now denote it with a subscript 2 to avoid confusion with the four-point function.

Assuming that the circle brightnesses are symmetrically distributed about zero (and thus $\langle b^3 \rangle = 0$), the $P_{i,j,k}^{\text{coll}}$ terms in which three points fall under the circle reduce to

$$\begin{aligned}
 P_{i,j,k}^{\text{coll}}(q) [\langle a^3 \rangle C_4^{\text{coll}}(q) + \langle a(1-a)^2 \rangle \langle b^2 \rangle \\
 \times [C_2(q) + C_2(2q) + C_2(3q)]]. \tag{11}
 \end{aligned}$$

Finally, the last term in Eq. (9), in which all four points fall under the new circle, simplifies to

$$\begin{aligned}
 P_{1,2,3,4}^{\text{coll}}(q) [\langle a^4 \rangle C_4^{\text{coll}}(q) + \langle a^2(1-a)^2 \rangle \langle b^2 \rangle \\
 \times [3C_2(q) + 2C_2(2q) + C_2(3q)] + \langle (1-a)^4 \rangle \langle b^4 \rangle]. \tag{12}
 \end{aligned}$$

B. Computing the probabilities $P_{\{i\}}^{\text{coll}}$

We now require the probabilities $P_{\emptyset}^{\text{coll}}, P_{1,\emptyset}^{\text{coll}}, P_{2,\emptyset}^{\text{coll}}, P_{1,2}^{\text{coll}}, P_{2,3}^{\text{coll}}, P_{1,2,3}^{\text{coll}}$, and $P_{1,2,3,4}^{\text{coll}}$. The remaining probabilities in Eq. (9) are equivalent to these because of the symmetry of the arrangement of points (and of the image ensemble).

Because all intervening pixels must lie under the circle if the bounding ones do, $P_{1,2,3,4}^{\text{coll}}(q) = P_2(3q)$, where $P_2(\cdot)$ is the probability that two pixels of a given separation lie under the same circle and is calculated in the previous section [Eq. (5) for power-law distributions of circle sizes]. We will use similar arguments to obtain the other six probability functions that we require.

The ‘‘triplet’’ probability $P_{1,2,3}^{\text{coll}}(q)$ is thus given by the probability that three of the (adjoining) pixels fall under the circle, minus the probability that all four pixels fall under it: $P_{1,2,3}^{\text{coll}} = P_2(2q) - P_2(3q)$. And by the same logic, $P_{1,2}^{\text{coll}} = P_2(q) - P_2(2q)$.

For the ‘‘inner’’ pairs, we compute the probability of the two ‘‘inner’’ points falling under the circle minus the probability that those two points *and any adjoining ones* all fall under the circle. Thus,

$$\begin{aligned} P_{2,3}^{\text{coll}}(q) &= P_2(q) - P_{1,2,3}^{\text{coll}} - P_{2,3,4}^{\text{coll}} - P_{1,2,3,4}^{\text{coll}} \\ \Rightarrow P_{2,3}^{\text{coll}}(q) &= P_2(q) - 2P_2(2q) + P_2(3q). \end{aligned} \quad (13)$$

Similarly, $P_{1,\emptyset}^{\text{coll}}(q) = P^* - P_2(q)$, where $P^* = \langle s^2 \rangle / L^2$ is the probability of any given point falling under the newly deposited circle, and

$$\begin{aligned} P_{2,\emptyset}^{\text{coll}}(q) &= P^* - P_{1,2}^{\text{coll}}(q) - P_{2,3}^{\text{coll}}(q) - P_{1,2,3}^{\text{coll}}(q) \\ &\quad - P_{2,3,4}^{\text{coll}}(q) - P_{1,2,3,4}^{\text{coll}}(q) \\ \Rightarrow P_{2,\emptyset}^{\text{coll}}(q) &= P^* - 2P_2(q) + P_2(2q). \end{aligned} \quad (14)$$

Finally,

$$\begin{aligned} P_{\emptyset}^{\text{coll}}(q) &= 1 - \sum_i P_{i,\emptyset}^{\text{coll}}(q) - \sum_{i,j \neq i} P_{i,j}^{\text{coll}}(q) \\ &\quad - \sum_{i,j \neq i, k \neq i,j} P_{i,j,k}^{\text{coll}}(q) - P_{1,2,3,4}^{\text{coll}}(q) \\ \Rightarrow P_{\emptyset}^{\text{coll}}(q) &= 1 - 4P^* + 3P_2(q). \end{aligned} \quad (15)$$

C. Assembling the pieces to find $C_4^{\text{coll}}(q)$

Before substituting all of our results into Eq. (9) and solving for $C_4^{\text{coll}}(q)$, it will be useful to first consider the $q \gg s_0$ limit, in which we derived the two-point function. In that limit [Eq. (6)],

$$P_2(q) = AB(\alpha) \left(\frac{s_0}{L} \right)^2 \left(\frac{q}{s_0} \right)^{-(\alpha-3)} \ll 1 \quad (16)$$

and

$$C_2(q) = \frac{B(\alpha)(\alpha-3)\langle b^2 \rangle \langle (1-a)^2 \rangle \left(\frac{q}{s_0} \right)^{-(\alpha-3)}}{2(1-a)} \ll 1, \quad (17)$$

so only the lowest-order terms in these quantities need to be considered. Because of the power-law nature of these functions, $C_2(2q)$ and $P_2(2q)$ have the same dependence on distance q as do the $C_2(q)$ and $P_2(q)$ terms, but are smaller by a factor of $2^{-(\alpha-3)}$, and similarly for the $f(3q)$ -type terms.

Substituting all of the products and probabilities derived in the preceding subsections into Eq. (9), keeping only the lowest-order terms in $(q/s_0)^{-(\alpha-3)}$, which dominate for $q \gg s_0$, and solving for $C_4^{\text{coll}}(q)$, we find that

$$C_4^{\text{coll}}(q) \approx \frac{B(\alpha)(\alpha-3)\langle b^4 \rangle \langle (1-a)^4 \rangle \left(\frac{3q}{s_0} \right)^{-(\alpha-3)}}{4(1-a)}. \quad (18)$$

Thus, the four-point function for this arrangement of points (in the $q \gg s_0$ limit) has the same power-law form as does the two-point function [Eq. (7)], and it also only depends on opacity by a multiplicative prefactor. Given that this (collinear) arrangement of points is so similar to the arrangement of points in the two-point function (two points will always be collinear), this result is perhaps unsurprising. To test the generality of this result, we will compute the four-point function for a square arrangement of points in the next section.

IV. ANALYTICAL CALCULATION OF THE FOUR-POINT FUNCTION FOR A SQUARE ARRANGEMENT OF POINTS IN THE TRANSMISSIVE DEAD LEAVES MODEL

In this section we calculate the four-point function for our transmissive dead leaves ensemble, for the case in which the four points lie on the vertices of a square with edge length q . Similar to the collinear arrangement of points, the symmetry in this arrangement will greatly simplify our calculations and, since it has nontrivial geometry when compared to the collinear arrangement, there is a possibility for interesting features to arise in this four-point function that are not apparent in either the two-point function, or the four-point function for collinear points.

We will label these points 1, 2, 3, 4, going clockwise, and beginning in the upper left-hand corner. Similar to the calculation for the collinear case, we first list all of the possible modifications to the four-point function, and the probabilities with which they occur. We will then simplify this expression, calculate the relevant probabilities, and use recursion to solve for the four-point function. Similar to the previous calculations, the translation and rotation invariance of our image ensemble means that this four-point function will depend only on the edge length of the square: $C_4^{\text{square}}(\vec{x}_1, \vec{x}_2, \vec{x}_3, \vec{x}_4) = C_4^{\text{square}}(q)$.

Enumerating all possible modifications caused by the addition of a new circle, we find that

$$\begin{aligned} C_4^{\text{square}}(q) &= P_{\emptyset}^{\text{square}}(q)C_4^{\text{square}}(q) + 4P_{1,\emptyset}^{\text{square}}(q)[aI(\vec{x}_1) + (1-a)b]I(\vec{x}_2)I(\vec{x}_3)I(\vec{x}_4) + 4P_{1,2}^{\text{square}}(q)[aI(\vec{x}_1) + (1-a)b] \\ &\quad \times [aI(\vec{x}_2) + (1-a)b]I(\vec{x}_3)I(\vec{x}_4) + 4P_{1,2,3}^{\text{square}}(q)[aI(\vec{x}_1) + (1-a)b][aI(\vec{x}_2) + (1-a)b][aI(\vec{x}_3) + (1-a)b]I(\vec{x}_4) \\ &\quad + P_{1,2,3,4}^{\text{square}}(q)[aI(\vec{x}_1) + (1-a)b][aI(\vec{x}_2) + (1-a)b][aI(\vec{x}_3) + (1-a)b][aI(\vec{x}_4) + (1-a)b], \end{aligned} \quad (19)$$

where $P_\emptyset^{\text{square}}(q)$ is the probability that none of the four corners of the square fall under the newly deposited circle, $P_{i,\emptyset}^{\text{square}}$ is the probability that only the i th corner falls under the newly deposited circle, $P_{i,j}^{\text{square}}$ is that probability that only the i th and j th corners fall under the newly deposited circle, and so on. The symmetries in the square configuration (all edges are equivalent, and all corners are equivalent) allow us to collapse the (equivalent) $P_{i,\emptyset}^{\text{square}}$ terms, and similarly for the $P_{i,j}^{\text{square}}$ terms and the $P_{i,j,k}^{\text{square}}$ terms. We further note that terms like

$$\begin{aligned} C_4^{\text{square}}(q) = & P_\emptyset^{\text{square}}(q)C_4^{\text{square}}(q) + 4\langle a \rangle P_{1,\emptyset}^{\text{square}}(q)C_4^{\text{square}}(q) + 4P_{1,2}^{\text{square}}(q)[\langle a^2 \rangle C_4^{\text{square}}(q) + \langle (1-a)^2 \rangle \langle b^2 \rangle C_2(q)] \\ & + 4P_{1,2,3}^{\text{square}}(q)[\langle a^3 \rangle C_4^{\text{square}}(q) + 2\langle a(1-a)^2 \rangle \langle b^2 \rangle C_2(q) + \langle a(1-a)^2 \rangle \langle b^2 \rangle C_2(\sqrt{2}q)] \\ & + P_{1,2,3,4}^{\text{square}}(q)[\langle a^4 \rangle C_4^{\text{square}}(q) + \langle a^2(1-a)^2 \rangle \langle b^2 \rangle (4C_2(q) + 2C_2(\sqrt{2}q)) + \langle (1-a)^4 \rangle \langle b^4 \rangle], \end{aligned} \quad (20)$$

where the function $C_2(\cdot)$ is the two-point function we discuss in Eq. (7).

A. Computing the probabilities $P_{i,j}^{\text{square}}$

To finish our calculation of the four-point function for square geometries, we require the probabilities $P_\emptyset^{\text{square}}(q)$, $P_{1,\emptyset}^{\text{square}}(q)$, $P_{1,2}^{\text{square}}(q)$, $P_{1,2,3}^{\text{square}}(q)$, and $P_{1,2,3,4}^{\text{square}}(q)$.

For the calculation of $P_{1,2,3,4}^{\text{square}}(q)$, we first note that, given that one of the corners of the square falls under a newly deposited circle (with diameter s), the probability that all four points fall under it is $g_4(q/s) \in [0, 1/\sqrt{2}] = \frac{4}{\pi}[\cos^{-1}(q/s) - (\pi/4) + (q/s)^2 - (q/s)\sqrt{1 - (q/s)^2}]$.

Using the same logic (and variable substitution) as in Eq. (6), we find that

$$P_{1,2,3,4}^{\text{square}}(q) = \int_0^\infty \frac{s^2}{L^2} g_4(q/s) p(s) ds \quad (21)$$

$$= A \left(\frac{s_0}{L} \right)^2 \left(\frac{q}{s_0} \right)^{-(\alpha-3)} B_4(\alpha), \quad (22)$$

where $B_4(\alpha) = \int_{\sqrt{2}}^\infty g_4(1/u) u^{2-\alpha} du$.

To derive $P_{1,2,3}^{\text{square}}(q)$, we seek the probability that three of the points, but not all four, lie under the newly deposited circle. If the two diagonal points are under the circle, so will at least one of the corners, and thus $P_{1,2,3}^{\text{square}}(q) = [P_2(\sqrt{2}q) - P_{1,2,3,4}^{\text{square}}(q)]/2$, where $P_2(x)$ is the probability that two points a distance x apart lie under a newly deposited circle, and is calculated in Eqs. (5) and (6).

The ‘‘doublet’’ probability $P_{1,2}^{\text{square}}(q)$ is the probability that two, but not three or four, of the points fall under the circle and thus is given by $P_{1,2}^{\text{square}}(q) = P_2^{\text{square}}(q) - P_{1,2,3}^{\text{square}}(q) - P_{1,2,4}^{\text{square}}(q) - P_{1,2,3,4}^{\text{square}}(q) = P_2(q) - P_2(\sqrt{2}q)$.

The ‘‘singlet’’ probability $P_{1,\emptyset}^{\text{square}}(q)$ is the probability that one, but not more, of the points fall under the circle, and is thus given by $P_{1,\emptyset}^{\text{square}}(q) = P^* - P_{1,2}^{\text{square}}(q) - P_{1,4}^{\text{square}}(q) - P_{1,2,3}^{\text{square}}(q) - P_{1,3,4}^{\text{square}}(q) - P_{1,2,4}^{\text{square}}(q) - P_{1,2,3,4}^{\text{square}}(q)$, where $P^* = \langle s^2 \rangle / L^2$ is the probability that a newly deposited circle covers

$P_{1,3}^{\text{square}}$ and $P_{2,4}^{\text{square}}$, which contain opposite corners of the circle, are omitted because it is impossible for a circle to cover diagonally opposite corners of the square without covering at least one other corner. The factors of four in the above equation come in because there are four corners to a square, and four edges to a square, and $\binom{4}{3} = 4$ different ways to choose groupings of three of the four corners.

We can expand and simplify the averages of the products of the pixel values, as in the previous section, to find

any given point, and is calculated in the previous sections. Simplifying this expression using our previously derived results, we find that $P_{1,\emptyset}^{\text{square}}(q) = P^* - 2P_2(q) + \frac{1}{2}P_2(\sqrt{2}q) + \frac{1}{2}P_{1,2,3,4}^{\text{square}}(q)$.

Finally, the probability that none of the points falls under a newly deposited circle is given by $P_\emptyset^{\text{square}}(q) = 1 - \sum_i P_{i,\emptyset}^{\text{square}}(q) - \sum_{i,j \neq i} P_{i,j}^{\text{square}} - \sum_{i,j \neq i, k \neq i,j} P_{i,j,k}^{\text{square}} - P_{1,2,3,4}^{\text{square}} = 1 - 4P^* + 4P_2(q) - P_{1,2,3,4}^{\text{square}}(q)$.

B. Combining the pieces to find $C_4^{\text{square}}(q)$

As in our calculation of the four-point function for collinear points, we again consider the $q/s_0 \gg 1$ limit, in which we need only consider the lowest-order terms in $(q/s_0)^{-(\alpha-3)}$. In that limit, we find that

$$C_4^{\text{square}}(q) \approx \frac{B_4(\alpha)(\alpha-3)\langle b^4 \rangle \langle (1-a)^4 \rangle}{4\langle 1-a \rangle} \left(\frac{q}{s_0} \right)^{-(\alpha-3)}. \quad (23)$$

Like the other n -point functions computed thus far, the four-point function for square geometries is a power law with power $-(\alpha-3)$, and it depends on opacity only as a multiplicative prefactor. We note that, for $\alpha = 3.2$, $B_4(\alpha) \approx 4.014$, while $B_4(\alpha) \approx 3.581$, where these values come from numerical integration using Simpson’s method. These values are similar in magnitude, and thus the four-point function is not inherently much smaller than the two-point function.

Finally, we note that the two- and four-point functions depend differently on object opacity, and thus the visible difference in the different image ensembles likely arises from the relative amplitudes of these (power-law) functions, and not any difference in their functional forms.

V. NUMERICAL ANALYSIS OF THE TRANSMISSIVE FALLEN-LEAF IMAGES

To confirm our analytical calculations of the two-point functions, we simulated 500-frame ensembles of 256×256 pixel images, using the procedure described in Eq. (1): Circles of random size (following a power-law distribution

$p(s) \propto s^{-3.2}$ above the cutoff of $s_0 = 1$ pixel), brightness, and position were iteratively placed on the image frame to build up the images. For each frame, 10^6 circles were deposited, which is the number required to cover the image surface ~ 100 times.

To avoid edge effects, circle centers were allowed to fall up to $256 + s/2$ pixels away from the center of the image frame, where s is the circle diameter in pixels. We used a large maximum circle size, $s_{\max} = 10^8$ pixels, because prior work [25] on dead leaves models found that the functional form of the measured autocorrelation function approaches the analytically calculated curve only in the $s_{\max} \rightarrow \infty$ limit. The heavy tail of the power-law distribution contains a non-negligible number of very large leaves, which contribute to the long-range correlations in the images.

We then measured the difference functions $D(q) = \langle |I(\vec{x}) - I(\vec{x}')|^2 \rangle = 2\langle I(\vec{x})^2 \rangle - 2C(q)$ for the image ensembles. $D(q)$ is clearly related to the autocorrelation function $C(q)$, but is easier to measure [20] as it is unaffected by the mean values of the individual images. We fit the measured difference functions to power-law functions of the form $D(q) = \eta \times q^\mu + \nu$, as is suggested by our analytical calculations [Eq. (7)]. The best-fit parameters (η, μ, ν) for the image ensembles with $a = \{0, 0.25, 0.5, 0.75\}$ were $(-0.48 \pm 0.01, -0.24 \pm 0.04, 0.69 \pm 0.03)$, $(-0.32 \pm 0.01, -0.23 \pm 0.03, 0.41 \pm 0.02)$, $(-0.191 \pm 0.004, -0.22 \pm 0.03, 0.23 \pm 0.01)$, $(-0.086 \pm 0.002, -0.21 \pm 0.02, 0.098 \pm 0.005)$, respectively, where the uncertainties represent 95% confidence intervals. These values are in good agreement with the analytical calculations that predict $\mu = -0.2$ for all ensembles, and $\nu = \{0.66, 0.396, 0.22, 0.094\}$ for the ensembles with $a = \{0, 0.25, 0.5, 0.75\}$, respectively. The correlation functions shown [Fig. 2(d)] are the measured difference functions subtracted from the constants ν measured in the fit: $C(q) = [\nu - D(q)]/2$. These correlation functions are power-law functions of distance (linear on the log-log plot) and differ by a multiplicative constant. Similarly, the power spectra of the image ensembles [Fig. 2(e)] differ only by a multiplicative constant for low spatial frequencies, where the $q \gg s_0$ approximation holds.

Figure 3 demonstrates that the two-point function is affected substantially by leaf opacity for δ function size distributions. In particular, the modulation depth of the “ripples” in the power spectra depend on the leaf opacity, and thus the opacity does not modify the power spectra simply by a multiplicative factor. The procedures used to generate the data shown in Fig. 3 were the same as for the power-law object size distribution, except for the different distribution of object sizes.

VI. A MORE REALISTIC MODEL OF RADIOLOGICAL IMAGES

Our transmissive dead leaves model is not a perfect model for radiological images. Image formation in mammograms and other projectional radiographs results from the partial blockage of a roughly uniform illumination of x-rays due to local regions of dense tissue, unlike our dead leaves model. Moreover, imaged tissue is typically much thinner than the path length required to fully block the x-rays throughout the image, unlike the effectively infinite optical depth of our “additive” transparent dead leaves model [Eq. (1)].

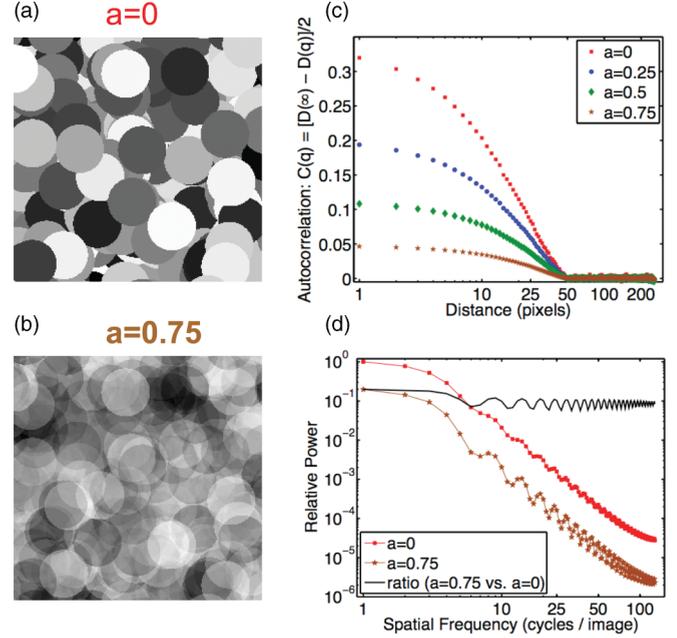


FIG. 3. (Color online) For δ function object size distributions, opaque and transmissive dead leaves images yield different two-point statistics. (a), (b) Sample images in which the leaves are all the same size ($s^* = 25$ pixels), from opaque ($a = 0$, a) and transmissive ($a = 0.75$, b) ensembles. (c) The autocorrelation functions of these image ensembles do not follow power laws, and they differ from one another. (d) Their power spectra also differ nontrivially: The ratio between the power spectra is not constant. The ripples are at multiples of the $256/25 \approx 10$ cycles/image frequency imposed by the uniform circle size.

It is thus natural to ask whether our conclusions about variable object opacity generalize to these types of images. Analytically computing the two-point statistics for this radiographic model is more involved than for the infinite depth models, since recursion is more complex in this case. For this reason, we chose to verify via simulation that the qualitative results from our analytical calculations hold for these types of images.

Figure 4(a) shows a typical image from a shadowing dead leaves model with finite optical depth and the same power-law leaf size distribution as in the previous models. To generate these model images, a uniform background illumination (of 1) was imposed across the whole image. Randomly sized and located circles were then deposited onto the image plane, with each leaf multiplying the brightness of the pixels it subtends by a factor drawn uniformly within $[0.5, 1]$. The circle sizes were drawn from the same power-law distribution as in the previous simulations, and the simulation code was thus very similar.

For an ensemble of these “radiographic” images, the empirically measured two-point function [Fig. 4(b)] and power spectrum [Fig. 4(c)] exhibit the same power laws as we found for our previous models (Figs. 2 and 3), suggesting that our calculation holds more generally than for the specific model for which we performed the analytical calculations.

Intuitively, one might expect the same scale-invariant two-point function for this model as for the previous one since no new length scale has been introduced.

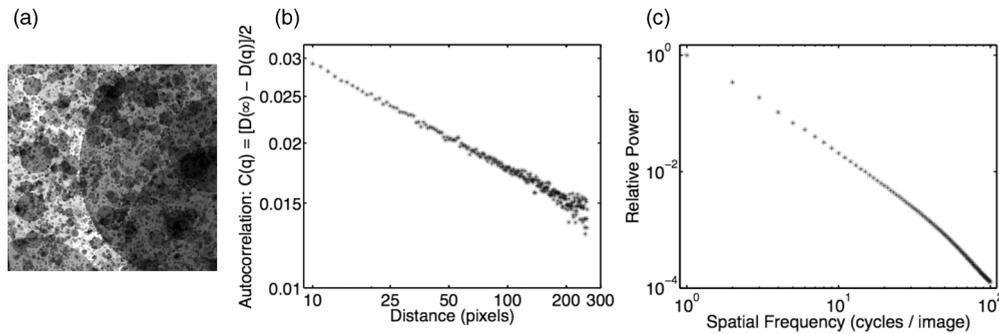


FIG. 4. A shadowing dead leaves model with finite optical depth also exhibits scale-invariant two-point statistics for power-law object size distributions. (a) An example image from a dead leaves model with the same power-law distributed leaf sizes as in Fig. 2, in which each leaf leaves a shadow by multiplying the brightness of the pixels it subtends by a factor no greater than one, drawn uniformly within $[0.5, 1]$. Unlike the previous models, each pixel starts out at full brightness, and only a finite number of circles is added to generate the image. The autocorrelation function (b) and power spectrum (c) of this ensemble show scale invariance (for relatively low frequencies, which corresponds to $q \gg s_0$), just like the previous models.

VII. CONCLUSIONS

For the special case of power-law object size distributions, object opacity does not affect the form of either the two- or four-point functions, or the power spectrum of images: It is manifest only by a multiplicative constant in these power-law functions. Ours is the first analytic calculation that demonstrates these facts and thus deepens our understanding of image statistics.

For object size distributions other than power law, object opacity can (potentially dramatically) alter the low-level image statistics. Occlusion is important for natural image formation, but we find that it does not change the form of the power spectrum. Since images formed by opaque leaves that are all the same size have oscillatory, non-power-law, power spectra (Fig. 3), and transmissive leaves can yield power-law power spectra (Figs. 2 and 4), occlusion is likely not

responsible for scale invariance of images. We propose that the universality of power-law power spectra in both occlusive imaging environments, such as natural photographic images, and transmissive ones, such as mammography, is likely due to power-law object size distributions in both settings.

ACKNOWLEDGMENTS

J.Z.'s contribution to this work was supported by an international student research fellowship from the Howard Hughes Medical Institute (HHMI). This material is based upon work supported by a National Science Foundation Graduate Research Fellowship to D.P. under Grant No. DGE 11-44155. M.R.D. thanks the Hellman Family Foundation, the James S. McDonnell Foundation, and the McKnight Foundation for support.

-
- [1] G. J. Stephens, T. Mora, G. Tkačik, and W. Bialek, arXiv:0806.2694 (2008).
 - [2] D. L. Ruderman and W. Bialek, *Phys. Rev. Lett.* **73**, 814 (1994).
 - [3] D. W. Dong and J. J. Atick, *Network: Comput. Neural Syst.* **6**, 345 (1995).
 - [4] D. J. Field, *J. Opt. Soc. Am. A* **4**, 2379 (1987).
 - [5] E. P. Simoncelli and B. A. Olshausen, *Annu. Rev. Neurosci.* **24**, 1193 (2001).
 - [6] A. Torralba and A. Oliva, *Network: Comput. Neural Syst.* **14**, 391 (2003).
 - [7] A. van der Schaff and J. van Hateren, *Vis. Res.* **36**, 2759 (1996).
 - [8] R. M. Balboa and N. M. Grzywacz, *Vis. Res.* **43**, 2527 (2003).
 - [9] J. J. Heine and R. P. Velthuisen, *Med. Phys.* **29**, 647 (2002).
 - [10] H. Li, M. L. Giger, O. I. Olopade, and M. R. Chinander, *J. Digit. Imaging* **21**, 145 (2008).
 - [11] K. G. Metheany, C. K. Abbey, N. Packard, and J. M. Boone, *Med. Phys.* **35**, 4685 (2008).
 - [12] H. B. Barlow, in *Sensory Communication*, edited by W. A. Rosenblith (MIT Press, Cambridge, MA, 1961), pp. 217–234.
 - [13] D. W. Dong and J. J. Atick, *Network: Comput. Neural Syst.* **6**, 159 (1995).
 - [14] J. J. Atick and A. N. Redlich, *Neural Comput.* **4**, 196 (1992).
 - [15] Y. Dan, J. J. Atick, and R. C. Reid, *J. Neurosci.* **16**, 3351 (1996).
 - [16] J. Zylberberg, J. T. Murphy, and M. R. DeWeese, *PLoS Comput. Biol.* **7**, e1002250 (2011).
 - [17] M. Rehn and F. T. Sommer, *J. Comput. Neurosci.* **22**, 135 (2007).
 - [18] G. Matheron, Technical Report, Centre de Morphologie Mathématique, Fontainebleau (unpublished).
 - [19] C. Bordenave, Y. Gousseau, and F. Roueff, *Adv. Appl. Prob.* **38**, 31 (2006).
 - [20] D. L. Ruderman, *Vis. Res.* **37**, 3385 (1997).
 - [21] C. R. Carlson, *Photogr. Sci. Eng.* **22**, 69 (1978).
 - [22] R. M. Balboa, C. W. Tyler, and N. M. Grzywacz, *Vis. Res.* **41**, 955 (2001).
 - [23] W. H. Hsiao and R. P. Milane, *J. Opt. Soc. Am. A* **22**, 1789 (2005).
 - [24] X. Pitkow, *J. Vis.* **10**, 42 (2010).
 - [25] A. B. Lee, D. Mumford, and J. Huang, *Int. J. Comp. Vis.* **41**, 35 (2001).