



# Scale invariance is driven by stimulus density

Stéphane J.M. Rainville<sup>a,\*</sup>, Frederick A.A. Kingdom<sup>b</sup>

<sup>a</sup> Center for Visual Science, University of Rochester, 274 Meliora Hall, Rochester, NY 14627, USA

<sup>b</sup> McGill Vision Research Unit, Room H4-14, 687 Pine Ave W., Montreal, Qu., Canada H3A 1A1

Received 29 October 1999; received in revised form 20 August 2001

---

## Abstract

Scale invariance refers to aspects of visual perception that remain constant with changes in viewing distance. Previously, Dakin and Herbert [Proc. Roy. Soc. B. 265 (1397) (1998) 659] reported that the spatial integration region (IR) for mirror symmetry in bandpass noise is scale invariant because its dimensions scale with the inverse of peak spatial frequency. In bandpass noise, however, peak spatial frequency covaries with stimulus numerosity (i.e. the total number of information samples) and density (i.e. the total number of information samples per unit area). In this study, we report four experiments that decoupled properties of the retinal image affected by viewing distance—spatial frequency, numerosity, size, and density—and measured their effect on IR size. Stimuli consisted of bandpass microelements with vertically mirror-symmetric but otherwise random positions, and we measured observer resistance to random jitter imposed on microelement position. Results show that jitter resistance and IR size vary with the inverse of stimulus *density* but are unaffected by changes in stimulus spatial frequency, numerosity, or size. We found the IR has a 2:1 height-to-width aspect ratio and integrates information from  $\approx 18$  microelements regardless of their spatial separation. Our results reveal that stimulus density plays a central role in the visual system's implementation of scale invariance. Using an ideal-observer, we demonstrate that scale invariance reflects genuine neural scale selection rather than a physical limitation on the stimulus' information content. Our findings that jitter resistance and IR size vary with the inverse of density challenge current models of spatial vision but can be reconciled with a model that compares the output of bandpass non-Fourier mechanisms to select spatial scales that match stimulus density. © 2002 Elsevier Science Ltd. All rights reserved.

*Keywords:* Density; Scale invariance; Scale selection; Symmetry; Texture; Spatial filtering; Correspondence problem; Grouping; Second-order; Non-Fourier; Ideal observer; Information theory

---

## 1. Introduction

Scale invariance refers to those aspects of visual perception that remain constant with changes in viewing distance. Examples of scale invariance include our ability to recognize a face or read text whether these objects are, within limits, near or far. In the present paper, we exploit mirror symmetry as a psychophysical tool to investigate the stimulus properties (and corresponding neural mechanisms) that human vision uses to implement scale invariance. The remainder of this Section describes our rationale in more detail.

### 1.1. Mirror symmetry and the spatial correspondence problem

Mirror symmetry is a visual property with ecological significance. For instance, the anatomy of most animals exhibits approximate bilateral symmetry—a potent visual cue that can signal the presence of friend, foe, or food otherwise camouflaged by an appropriate background (e.g. Enquist & Arak, 1994; Horridge, 1996; Møller, 1995; Swaddle & Cuthill, 1994). However, mirror symmetry should also be considered as a particularly valuable psychophysical tool that provides insight into the way in which human vision solves the more general and more fundamental *grouping* problem, otherwise known as the *spatial correspondence* problem.

The problem of spatial correspondence arises when an observer must determine whether visual elements, say two dots separated by some distance, are dependent or independent of each other. To illustrate this, consider

---

\* Corresponding author. Tel.: +1-716-275-8443; fax: +1-716-271-3043.

E-mail address: stephane@cvs.rochester.edu (S.J.M. Rainville).

an array of dots positioned in a mirror-symmetric but otherwise random fashion. Two non-symmetric dots are independent since their positions are determined by different causes (i.e. knowledge of one dot's position provides no information as to the position of the other). Conversely, two symmetric dots are dependent since the position of one dot is constrained by the position of the other. In Barlow's terminology (Barlow, 1959; Barlow & Reeves, 1979), symmetric dots are said to be spatially *redundant* whereas non-symmetric dots are not. Solving the spatial correspondence problem—or the grouping problem—amounts to detecting redundancies between visual elements.

In principle, all images (including natural scenes) can be described in terms of spatial redundancies, but in practice natural scenes often make poor experimental stimuli since their spatial structure is complex and only partially understood (Field, 1987; Olshausen & Field, 1996; Simoncelli & Portilla, 1998). In this context, mirror symmetry emerges as a better stimulus because the correspondence problem is simple and well defined. However, what sets symmetry detection apart from most psychophysical tasks is that the correspondence problem remains the same regardless of whether the statistical structure of a stimulus class is simple (e.g. symmetric white-noise patterns) or highly complex (e.g. symmetric human faces). Symmetry detection is a valuable psychophysical tool because it allows one to study how the correspondence problem is solved in the context of images with variable complexity.

### 1.2. Scale invariance and the spatial integration region for mirror symmetry

Because it factors out the effects of viewing distance on the retinal image, scale invariance can be understood as a transformation from a retino-centric to an object-centric frame of reference; retino-centric properties vary with viewing distance whereas object-centric properties do not. Thus, the signature of scale-invariant perception is characterized by variable performance when expressed in retino-centric units (e.g. degrees of arc) but constant performance when expressed in object-centric units (e.g. centimeters).

In a study that serves as the starting point for the present paper, Dakin and Herbert (1998) have shown that symmetry perception in bandpass noise patterns is limited to a spatial integration region (IR) whose dimensions are inversely proportional to stimulus peak spatial frequency. This result is caricatured in Fig. 1 where the IR (dashed-line oval) is depicted for noise patterns filtered for three constant-octave spatial-frequency bands. The IR is scale invariant because its dimensions change with *retinal* spatial frequency in a way that includes a fixed number of *object* features (e.g. black and white texture elements). However, what re-

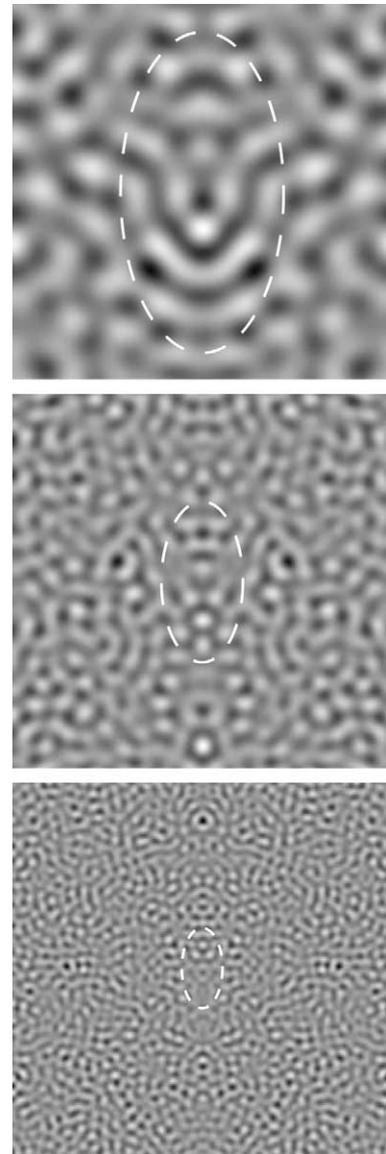


Fig. 1. The integration region for mirror symmetry in bandpass noise patterns. Vertically symmetric bandpass-filtered noise patterns of various scales similar to the ones used in Dakin and Herbert (1998). The approximate spatial IR measured for human observers is shown by dashed lines.

mains unclear in the Dakin and Herbert study is whether the changes in IR size are dependent on spatial frequency per se or on other stimulus properties that, like spatial frequency, covary with viewing distance.

In the natural environment, changes in viewing distance introduce simultaneous changes in several properties of the retinal image including spatial-frequency, size, numerosity, and density. Although the size of stimuli in Dakin and Herbert (1998) remained constant (see Fig. 1), changes in spatial frequency were accompanied by simultaneous changes in *numerosity* (e.g. the total number of black and white texture elements) and *density* (e.g. the total number of black and white texture

elements per unit area). In short, spatial frequency is confounded with numerosity and density. In the present study, we investigated which stimulus property—spatial frequency, size, numerosity, or density—determines the size of the IR for symmetry. Our data revealed that density, not spatial frequency, size, or numerosity, is the stimulus property the visual system uses to implement scale invariance.

From this point onward, we define numerosity as the square-root of the total number of texture elements within an image of equal width and height. In other words, numerosity refers to the number of elements per stimulus dimension rather than to the number of elements per image. For example, a  $4 \times 4$  deg image with 16 elements would have a numerosity of 4 elements. Defining numerosity in this way allows the number of elements along the width and height dimensions of an image to be considered separately as this will be more convenient when we report both the height and width of the IR instead of its area. Given our definition of numerosity, density is simply numerosity divided by stimulus size (i.e. width or height, depending which stimulus dimension we consider). In the example just given, the image has a density of 1.0 elem/deg.

Several psychophysical studies have investigated the dimensions of the IR for mirror symmetry (e.g. Barlow & Reeves, 1979; Bruce & Morgan, 1975; Dakin & Herbert, 1998; Dakin & Hess, 1997; Jenkins, 1982; Julesz, 1971; Rainville & Kingdom, 1999b; Rainville & Kingdom, 2000b; Tyler, Hardage, & Miller, 1995; Wenderoth, 1995) and while most report a preference for information proximal to the axis of symmetry, none have decoupled the effects of spatial frequency, size, numerosity, or density on IR size. However, two studies have explicitly addressed the issue of stimulus numerosity and/or density in the context of symmetry perception. Tyler and Hardage (1996) measured sensitivity to vertical even- and odd-symmetry in textures composed of black and white Gaussian blobs defined by either 10% or 100% pixel density. Whereas the envelope of low-density textures was heavily modulated across space, the envelope of high-density textures remained approximately flat. With these stimuli, the authors explored the differential contribution of Fourier (i.e. first-order) and non-Fourier (i.e. second-order) mechanisms to symmetry perception. Tyler and Hardage found that in the 10% density condition, performance remained approximately constant irrespective of whether symmetry was even or odd, thereby suggesting that non-Fourier mechanisms mediate symmetry perception in low-density stimuli. In the 100% density condition, in which only a small amount of non-Fourier structure was present, performance was higher for even symmetry than for odd symmetry, thereby suggesting that Fourier mechanisms mediate symmetry perception in high-density stimuli. While the results from Tyler and Hardage

do not distinguish between the effects of stimulus spatial frequency, numerosity, size, and density, they emphasize the contribution of non-Fourier mechanisms to symmetry perception.

Using a yes–no paradigm, Wenderoth (1996) asked observers to make judgments for random-dot patterns that were either perfectly symmetric or completely random. Percent correct was measured for several combinations of dot number and stimulus size. Data revealed no systematic relationship with numerosity, stimulus size, or density although performance generally improved as numerosity decreased. However, the fact that IR size was not measured makes results from Wenderoth (1996) difficult to interpret. Also, because stimuli were either perfectly symmetric or completely random, no psychophysical thresholds were computed.

### 1.3. Rationale

To decouple the effect of spatial frequency from those of stimulus numerosity and density, spatial frequency must be manipulated while holding the other stimulus properties constant. In Experiment 1 we measured the width and height of the IR using patterns consisting of a constant number of fixed-size bandpass microelements arranged in a mirror-symmetric but otherwise random fashion. Unlike bandpass noise patterns, the spatial frequency of our stimuli is specified by the carrier spatial frequency of the microelements and is therefore independent of microelement numerosity or density.

Decoupling numerosity and density is somewhat more difficult since the simple relationship  $\text{density} = \text{numerosity}/\text{size}$  illustrates that numerosity, size, and density are interdependent variables. Nonetheless, it is possible to control for one of these variables by holding one variable constant while covarying the other two (e.g. density is held fixed while numerosity and size covary). In Experiments 2, 3, and 4, we measured the width and height of the IR for stimuli of fixed numerosity, size, and density, respectively. The same methodology has been used by Dakin (2000) to study sampling efficiency in orientation-pooling mechanisms, and we compare his set of results with ours in Section 8. Part of the present research has been reported at conferences (Rainville & Kingdom, 1999a; Rainville & Kingdom, 2000a) and has been included in the first author's doctoral thesis (Rainville, 1999).

## 2. General method

### 2.1. Observers

The first author, SR, participated in all experiments. The second author, FK, participated in Experiments 2 and 3. A naive observer, LC, participated in

Experiments 1 and 4. All observers had normal or corrected-to-normal vision.

## 2.2. Hardware and calibration

Experiments were carried out using a Power Macintosh 7600/120 computer upgraded to a G3/180 main processor. Additional data were collected using a Power Macintosh G4/450. The two computers hosted standard 8-bit/gun color video cards driving a 17 in. Sony Multiscan monitor and a 21 in. Apple studio display monitor, respectively. Both monitors were set to a 68 Hz refresh rate. Luminance profiles were measured using a calibrated spot photometer for every fourth index of a 256-index grayscale lookup table (LUT) and modeled with best-fitting (least-squares) gamma functions. From the inverse gamma functions, we computed LUTs whose indexes corresponded to linear increments in luminance. After linearization, both display had an effective grayscale LUT depth of 7.1 bits and mean luminance was set to 33 cd/m<sup>2</sup>.

## 2.3. Stimuli

In all experiments, stimuli were scaled to mean luminance and 0.5 Michelson contrast. Patterns were composed of bandpass microelements positioned in a vertically symmetric but otherwise random fashion. If, by chance, elements were positioned such that part of their spatial profile fell outside the boundaries of the stimulus, left-over portions were neatly wrapped around to the other side of the image using the modulo operator. This ensured that no physical information was lost due to the occlusion of microelements by the stimulus aperture.

In Experiment 1, bullseye microelements consisted of radial sinusoids of variable spatial frequency windowed by a fixed-size spatial Gaussian function. The spatial profile of microelements,  $m$ , is given by

$$m(x, y) = \exp \left[ -\frac{(r - r_0)^2}{2\sigma^2} \right] \cos[2\pi(r - r_0)/\lambda] \quad (1)$$

where  $r$  is the radius given by  $\sqrt{x^2 + y^2}$  around an origin  $r_0$ , and where  $\sigma$  and  $\lambda$  are the envelope space constant and the carrier's spatial period, respectively.

In Experiments 2–4, the spatial profile of microelements,  $m$ , was defined in the Fourier domain by an isotropic bandpass log-Gaussian function  $M(u, v)$  where  $u$  and  $v$  are the dimensions of a two-dimensional Cartesian spatial-frequency coordinate system.  $M(u, v)$  is given by

$$M(u, v) = \exp \left[ -\frac{1}{2} \left( \frac{\ln(f/f_0)}{\ln(\sigma)} \right)^2 \right] \quad (2)$$

where  $f$  is defined as  $\sqrt{u^2 + v^2}$ . Center spatial-frequency  $f_0$  was fixed to 2.5 cdp and  $\sigma$  was set to 1.4 which corresponds to a full bandwidth at half-height of 1.2 octaves. The spatial profile  $m$  was obtained by computing the reverse Fourier transform  $m(x, y) = \text{Re}(\mathcal{F}^{-1}\{M(u, v)\})$ . Microelements were spatially localized by virtue of the fact that all Fourier components were in cosine phase.

To obtain meaningful measures of human performance, we degraded perfect symmetry by adding random jitter to the position of each microelement. Positional jitter was introduced independently in the  $x$  and  $y$  dimensions by randomly sampling from a uniform (i.e. flat) distribution of variable width—measured in degrees of arc—centered on zero. In all experiments, we fitted a two-parameter log  $x$  cumulative normal to the percent-correct vs. jitter data and computed the jitter level corresponding to 75% correct. We used this jitter level as our measure of jitter resistance, i.e., the maximum amount of spatial jitter observers can tolerate and still reach criterion performance. Geometric error bars were computed using a bootstrap technique (Efron & Tibshirani, 1993) in which we randomly resampled and refit the data 100 times, obtained a distribution of jitter-resistance values, and took the standard deviation of this distribution as our measure of error.

To measure the width and height of the IR, we used a procedure similar to the one in Dakin and Herbert (1998) and Rainville and Kingdom (2000b). For stimuli in the symmetric condition, microelements falling within a central window of variable width or height were assigned symmetric positions; elements outside the window were assigned random positions. From studies cited above, the expected effect of window size on performance is the following. For small window sizes, jitter resistance is poor since most of the stimulus is covered by non-symmetric elements. For larger window sizes, performance improves because more symmetric elements are made available to the observer. However, jitter resistance asymptotes once window size exceeds the dimensions of the IR since, by definition, observers are insensitive to mirror symmetry for elements falling outside the IR. We used the window size corresponding to asymptotic knee-point in performance as our measure of IR size.

## 2.4. Procedure

In all experiments, stimuli were computed in the MATLAB 5.2.1 environment and consisted of  $64 \times 64$ ,  $128 \times 128$ , or  $256 \times 256$  pixel matrices depending on stimulus size. Patterns were shown using high-level interfaces from PsychToolbox© (Brainard, 1997) calling lower-level routines from VideoToolbox© (Pelli, 1997). Viewing distance was set to 68 cm (17 in. monitor,  $640 \times 480$  pixels) or to 80 cm (21 in. monitor,  $640 \times 480$

pixels) such that one pixel subtended  $0.044 \times 0.044$  deg. This bracketed the spatial frequency content of the largest test patterns between 0.1 and 11.4 cycles per degree (cpd) although these limits were not attained because our stimuli were bandlimited.

In all experiments, observers discriminated between symmetric and non-symmetric images in a two-alternative forced-choice (2AFC) paradigm using a method of constant stimuli. The order of presentations was randomly interleaved across trials, and observers pressed one of two keys to report the interval that appeared more symmetric. Images were presented for 250 ms (or the equivalent of 17 screen refreshes at 68 Hz) and were separated by an inter-stimulus-interval (ISI) of 250 ms. A low-contrast fixation dot corresponding to the center of the stimulus was shown before every presentation to ensure that the axis of symmetry was foveated. Observers received auditory feedback on incorrect responses. Each run consisted of forty observations at one jitter level, and jitter level was randomized across runs. Although the total number of runs varied, a minimum of 80 observations were collected for every data point presented in the paper.

### 3. Experiment 1: the effect of spatial frequency

This experiment investigated the effects of spatial frequency on the IR without altering stimulus numerosity, size, or density. We used fixed-size stimuli (11.2 deg) with a numerosity of 5.7 elem and a density of 0.5 elem/deg. Microelements had one of three carrier spatial frequencies (1.9, 3.8, and 7.7 cpd) but their envelope size remained fixed ( $\sigma = 0.18$  deg). Due to their high contrast, patterns were easily visible even for highest spatial frequency (7.7 cpd). We measured jitter resistance as a function of the height or width of symmetric windows that systematically varied in equal log-steps from 0.20 deg (1.75% of stimulus) to 11.2 deg (100% of stimulus). Fig. 2 shows examples of stimuli for the three spatial frequencies used (rows). For illustration purposes, symmetric windows are shown here at half stimulus width (first column) and half stimulus height (second column) by dashed lines, and microelement position were not jittered.

Fig. 3 plots jitter resistance as a function of window width (top graphs) and window height (bottom graphs) for observers SR and LC. Filled triangles, squares, and circles correspond to spatial frequencies of 1.9, 3.8, and 7.7 cpd, respectively. Solid lines are the best fit of a four-parameter log–log cumulative normal to the data. Error bars show  $\pm 1$  mean and maximum standard deviation.

Fig. 3 shows that in all conditions, jitter resistance improves but tends towards an asymptote as window size increases. The key aspect is that data at different spatial frequencies are virtually superimposed and that

knee-points coincide. The agreement between data from different spatial-frequency conditions demonstrates that IR size is independent of spatial frequency. Our results also imply that, counter to the conclusions of Dakin and Herbert (1998), the IR size in bandpass noise patterns is determined not by spatial frequency but by some other stimulus property.

### 4. Parameters for experiments 2–4

The remaining three experiments focus on decoupling the effects of stimulus numerosity, size, and density on IR size and jitter resistance. As outlined in Section 1.3, each experiment requires that one stimulus property (e.g. numerosity) be held fixed while the other two properties (e.g. size and density) covary. Table 1 provides a summary of conditions in each experiment.

### 5. Experiment 2: fixed numerosity—size and density covary

This experiment controlled for stimulus numerosity and determined whether numerosity is, by itself, a good predictor of IR size. We fixed numerosity to 22.6 elem and allowed stimulus size to covary with density. We used three size–density pairings, namely {2.8 deg; 8.1 elem/deg}, {5.6 deg; 4.0 elem/deg}, and {11.2 deg; 2.0 elem/deg}. Because numerosity = size  $\times$  density, the product within each size–density pairing is necessarily 22.6 elem.

We measured jitter resistance as a function of the width and height of symmetric windows that systematically varied in equal log-steps from 0.20 deg until 100% of the stimulus was covered. Fig. 4 shows examples of stimuli for the three size–density pairings we used (rows). Symmetric windows are shown here at half width (first column) and half height (second column) by dashed lines, and the position of symmetric elements has not been jittered.

Fig. 5 plots jitter resistance as a function of window width (top graphs) and window height (bottom graphs) for observers SR and FK. Filled circles, squares, and triangles correspond to size–density pairings of {2.8 deg; 8.1 elem/deg}, {5.6 deg; 4.0 elem/deg}, and {11.2 deg; 2.0 elem/deg}, respectively. Solid lines are the best fits of four-parameter log–log cumulative normals for data from each of the size–density pairings. Error bars show  $\pm 1$  mean and maximum standard deviation.

Fig. 5 shows that in all conditions, jitter resistance improves but tends towards an asymptote as we increase window size. The key aspect is that data from different size–density pairings do not overlap. In particular, the size of the IR (defined by performance knee-points) in large low-density images is clearly larger than in small high-density images. Our finding that jitter resistance is

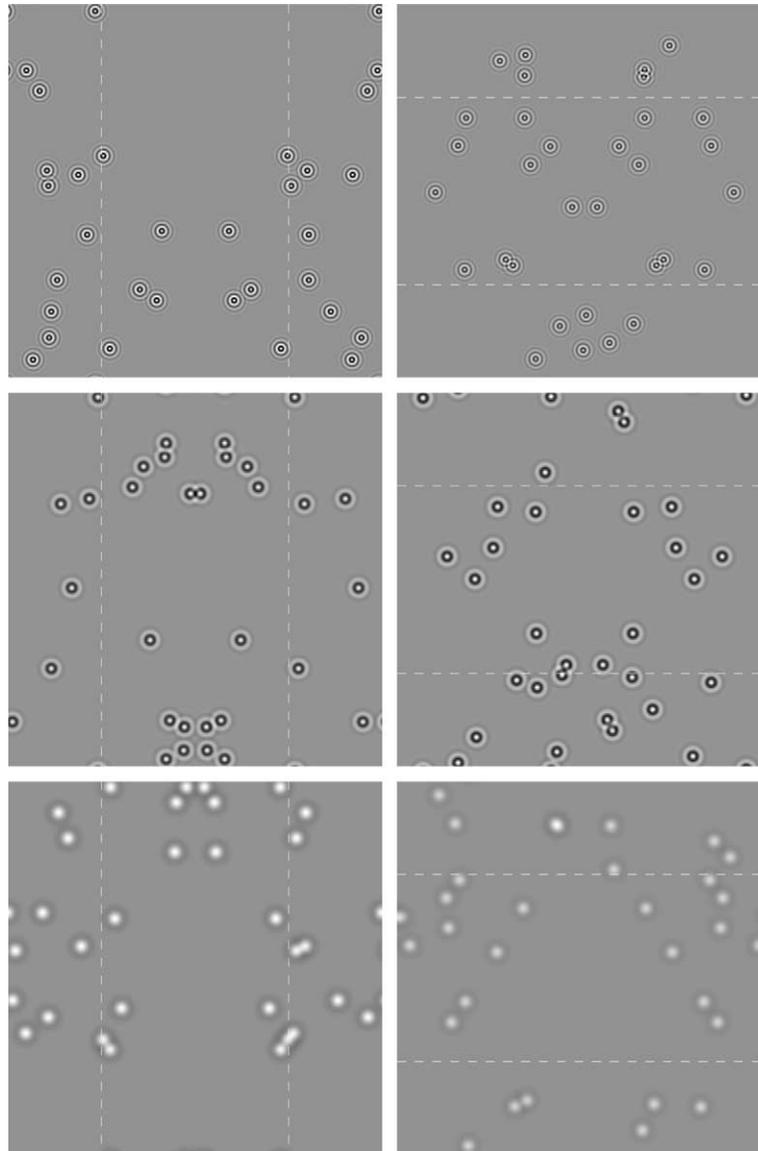


Fig. 2. Stimuli for Experiment 1. Stimuli consist of 32 bullseye microelements with one of three carrier spatial frequencies (1.9, 3.8, and 7.7 cpd) but fixed envelope size. Stimulus size, numerosity, and density also remained constant. Symmetric windows (dashed lines) are shown at half stimulus width (first column) and half stimulus height (second column).

not constant but instead depends on the size–density pairing is also of interest. We further discuss the implication of these findings in Section 8.

### 6. Experiment 3: fixed size—numerosity and density covary

This experiment controlled for stimulus size and determined whether stimulus size is, by itself, a good predictor of IR size. We fixed stimulus size to 11.2 deg and allowed numerosity and density to covary. We used three numerosity–density pairings, namely {5.7 elem; 0.5 elem/deg}, {22.6 elem; 2.0 elem/deg}, and {90.5 elem; 8.1 elem/deg}. Because size = numerosity/density, the

ratio within each numerosity–density pairing is necessarily 11.2 deg.

We measured jitter resistance as a function of the width and height of symmetric windows that systematically varied in equal log-steps from 0.20 to 11.2 deg. Fig. 6 shows examples of stimuli for the three size–density pairings we used (rows). For illustration purposes, symmetric windows are shown here for a fixed width (first column) and fixed height (second column) by dashed lines, and elements have not been jittered.

Fig. 7 plots jitter resistance as a function of window width (top graphs) and window height (bottom graphs) for observers SR and FK. Filled circles, squares, and triangles correspond to size–density pairings of {90.5 elem; 8.1 elem/deg}, {22.6 elem; 2.0 elem/deg}, and

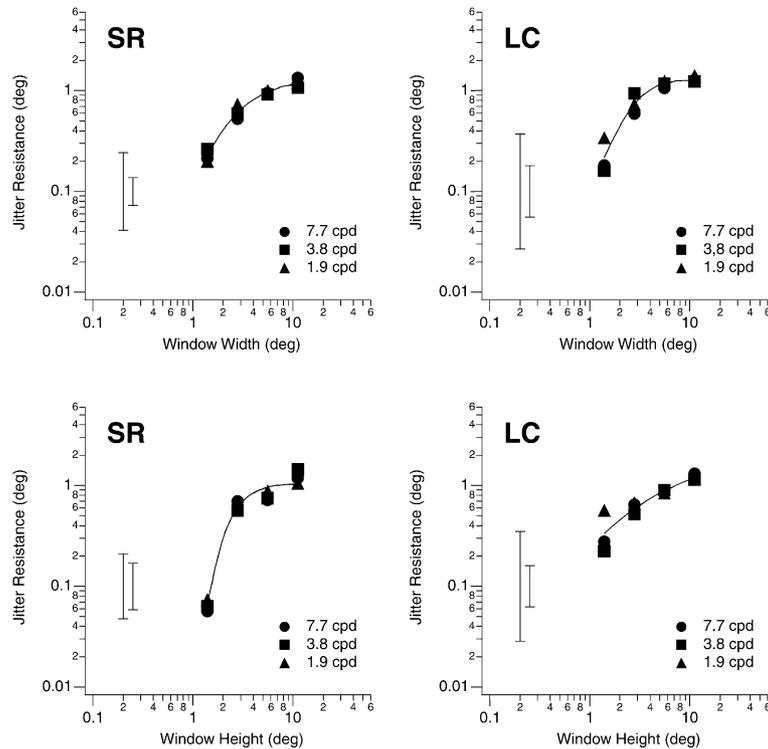


Fig. 3. Results of Experiment 1. Jitter resistance is plotted as a function of window width (top graphs) and window height (bottom graphs) for observers SR and LC. Filled triangles, squares, and circles correspond to spatial frequencies of 1.9, 3.8, and 7.7 cpd, respectively. Solid lines are the best fit of a four-parameter log-log cumulative normal to the data. Error bars show  $\pm 1$  mean and maximum standard deviation.

Table 1  
Parameters for Experiments 2–4

	Numerosity (elem)	Size (deg)	Density (elem/deg)
Experiment 2	22.6	2.8	8.1
(fixed numerosity)	22.6	5.6	4.0
	22.6	11.2	2.0
Experiment 3	5.7	11.2	0.5
(fixed size)	22.6	11.2	2.0
	90.5	11.2	8.1
Experiment 4	22.6	2.8	8.1
(fixed density)	45.3	5.6	8.1
	90.5	11.2	8.1

Columns list numerosity, size, and density conditions for experiments 2–4. Note that in each experiment, one stimulus property is held fixed while the other two covary.

{5.7 elem; 0.5 elem/deg}, respectively. Solid lines are the best fits of four-parameter log-log cumulative-normals for data from each of the size–density pairings. Error bars show  $\pm 1$  mean and maximum standard deviation.

Fig. 7 shows that in all conditions, jitter resistance improves but tends towards an asymptote as we increase window size. The key aspect is that data from different numerosity–density pairings do not overlap. In particular, the size of the IR (defined by performance knee-points) in low-numerosity low-density images is clearly larger than in high-numerosity high-density images. Our

finding that jitter resistance is not constant but instead depends on the numerosity–density pairing is also of interest. We further discuss the implication of these findings in Section 8.

## 7. Experiment 4: fixed density—numerosity and size covary

This experiment controlled for density and determined whether stimulus density is, by itself, a good predictor of IR size. We fixed stimulus density to 8.1 elem/deg and allowed numerosity and size to covary. We used three numerosity–size pairings, namely {22.6 elem; 2.8 deg}, {45.3 elem; 5.6 deg}, and {90.5 elem; 11.2 deg}. Because density = numerosity/size, the ratio of each numerosity–size pairing is necessarily 8.1 elem/deg.

We measured jitter resistance as a function of the width and height of symmetric windows that systematically varied in equal log-steps from 0.20 deg until 100% of the stimulus was covered. Fig. 8 shows examples of stimuli for the three numerosity–size pairings we used (rows). For illustration purposes, symmetric windows are shown at half width (first column) and half height (second column) by dashed lines, and element positions of were not jittered.

Fig. 9 plots jitter resistance as a function of window width (top graphs) and window height (bottom

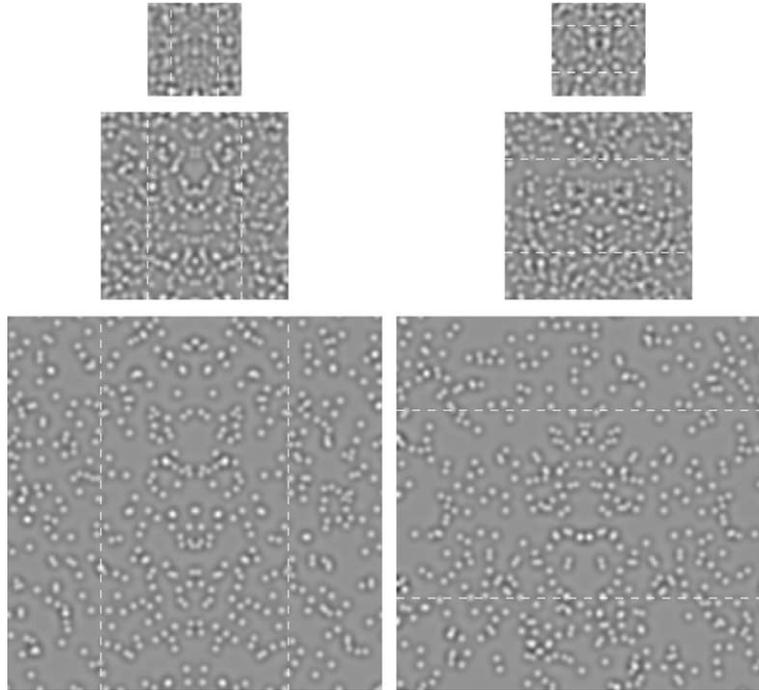


Fig. 4. Stimuli for Experiment 2. Fixed-numerosity stimuli with three size–density pairings: {2.8 deg; 8.1 elem/deg}, {5.6 deg; 4.0 elem/deg}, and {11.2 deg; 2.0 elem/deg} along the rows. Symmetric windows are shown here at half width (first column) and half height (second column) by dashed lines. The position of symmetric elements has not been jittered.

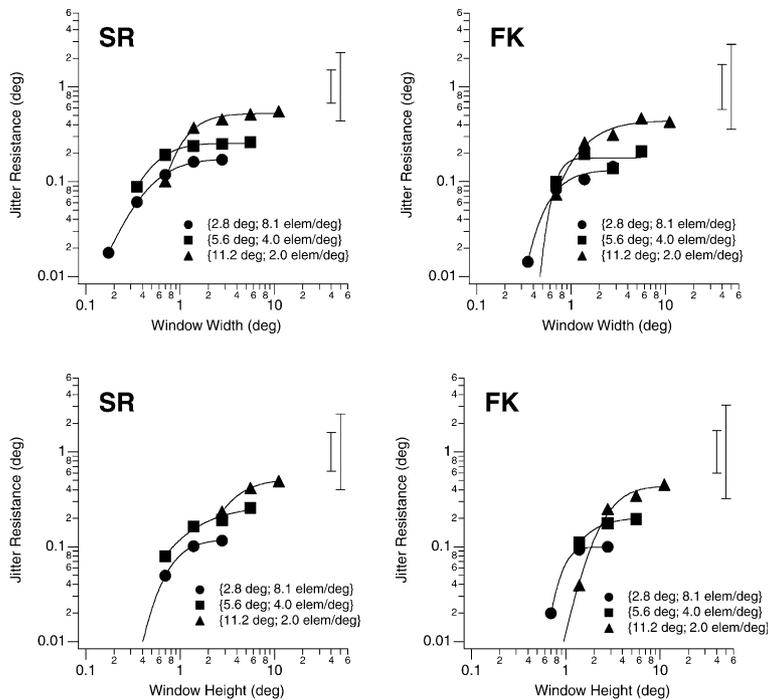


Fig. 5. Results of Experiment 2. Jitter resistance is plotted as a function of window width (top graphs) and window height (bottom graphs) for observers SR and FK. Filled circles, squares, and triangles correspond to size–density pairings of {2.8 deg; 8.1 elem/deg}, {5.6 deg; 4.0 elem/deg}, and {11.2 deg; 2.0 elem/deg}, respectively. Solid lines are the best fits of four-parameter log–log cumulative normals for data from each of the size–density pairings. Error bars show  $\pm 1$  mean and maximum standard deviation.

graphs) for observers SR and LC. Filled circles, squares, and triangles correspond to the numerosity–size pair-

ings of {22.6 elem; 2.8 deg}, {45.3 elem; 5.6 deg}, and {90.5 elem; 11.2 deg}, respectively. Solid lines are the

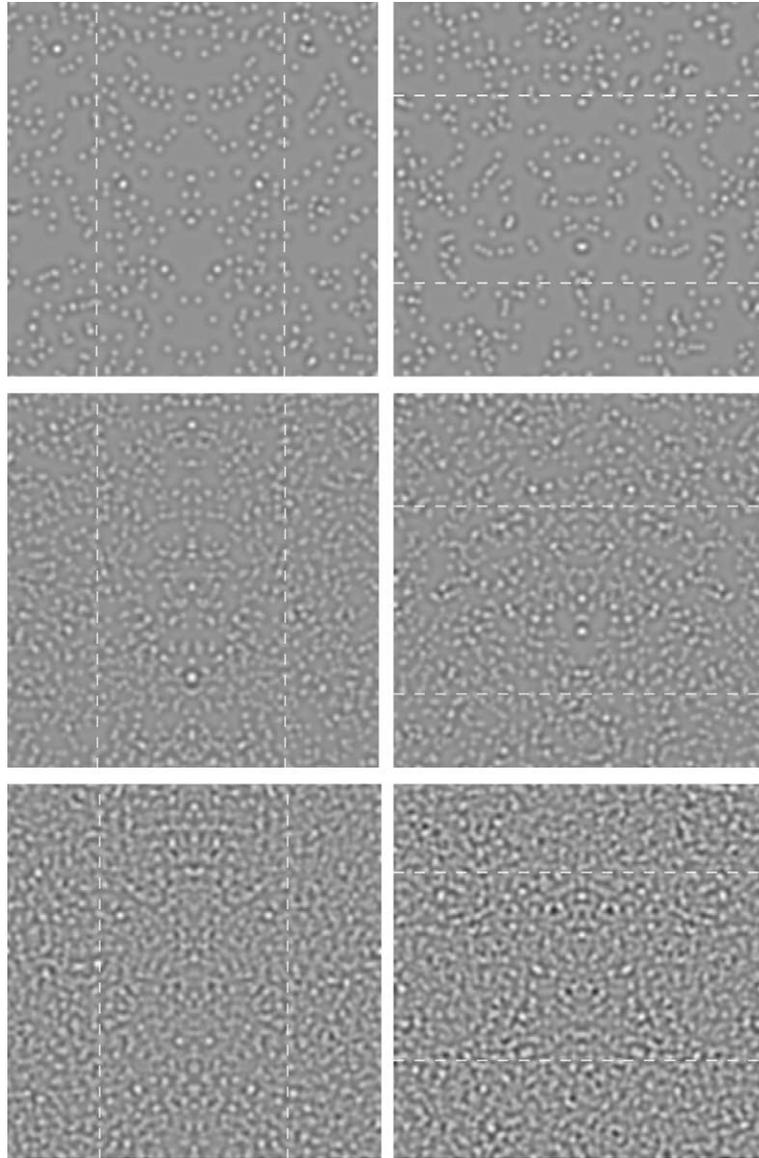


Fig. 6. Stimuli for Experiment 3. Fixed-size stimuli with three numerosity–density pairings: {5.7 elem; 0.5 elem/deg}, {22.6 elem; 2.0 elem/deg}, and {90.5 elem; 8.1 elem/deg} along the rows. Symmetric windows are shown here at half width (first column) and half height (second column) by dashed lines. The position of symmetric elements has not been jittered.

best fits of four-parameter log–log cumulative normals for data collapsed across all numerosity–size pairings. Error bars show  $\pm 1$  mean and maximum standard deviation.

Fig. 9 shows that in all conditions, jitter resistance improves but tends towards an asymptote as we increase window size. The key aspect is that data from different numerosity–size pairings are virtually superimposed and are well fit by a single function. In particular, the size of the IR (defined by performance knee-points) in small low-numerosity images is the same as in large high-numerosity images. Our results show that density, not spatial frequency, numerosity, or stimulus size, is the factor that predicts the size of the IR. Our finding that jitter resistance is constant for all numerosity–size pair-

ings is also of interest. We further discuss the implication of these findings in Section 8.

## 8. Discussion

### 8.1. Scale invariance is driven by stimulus density

Collectively, Experiments 1–4 demonstrate that stimulus density, not spatial frequency, numerosity, or size, determines jitter resistance and IR size. But how do IR size and jitter resistance change as a function of density? Data from Experiments 2 and 3 (see Figs. 5 and 7) are revealing in this respect since they show that jitter resistance and IR size *increase* as stimulus density

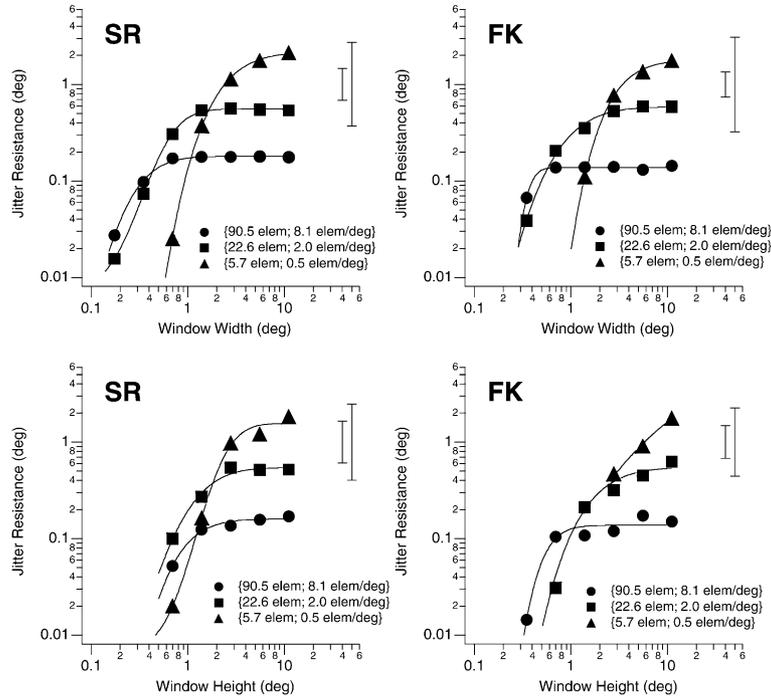


Fig. 7. Results of Experiment 3. Jitter resistance is plotted as a function of window width (top graphs) and window height (bottom graphs) for observers SR and FK. Filled circles, squares, and triangles correspond to size–density pairings of {90.5 elem; 8.1 elem/deg}, {22.6 elem; 2.0 elem/deg}, and {5.7 elem; 0.5 elem/deg}, respectively. Solid lines are the best fits of four-parameter log–log cumulative normals for data from each of the size–density pairings. Error bars show  $\pm 1$  mean and maximum standard deviation.

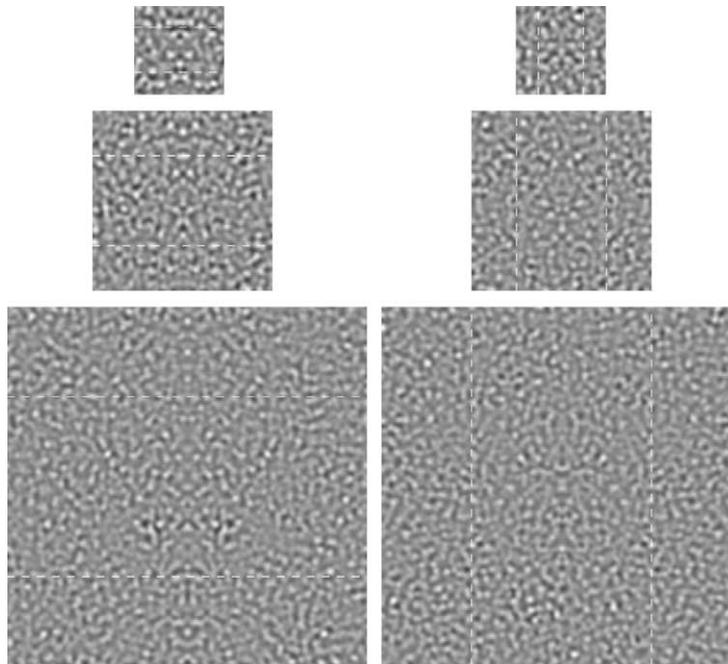


Fig. 8. Stimuli for Experiment 4. Fixed-density stimuli with three numerosity–size pairings: {22.6 elem; 2.8 deg}, {45.3 elem; 5.6 deg}, and {90.5 elem; 11.2 deg} along the rows. Symmetric windows are shown here at half width (first column) and half height (second column) by dashed lines. The position of symmetric elements has not been jittered.

decreases. If jitter resistance and IR size are inversely related to density ( $1/d$ ), then multiplying jitter and window size by  $d$  should remove the effect of density, in

which case all data from Experiments 1–4 should overlap. Elementary algebra reveals that multiplying window size  $w$  by density  $d$  leaves *elements* as the metric since

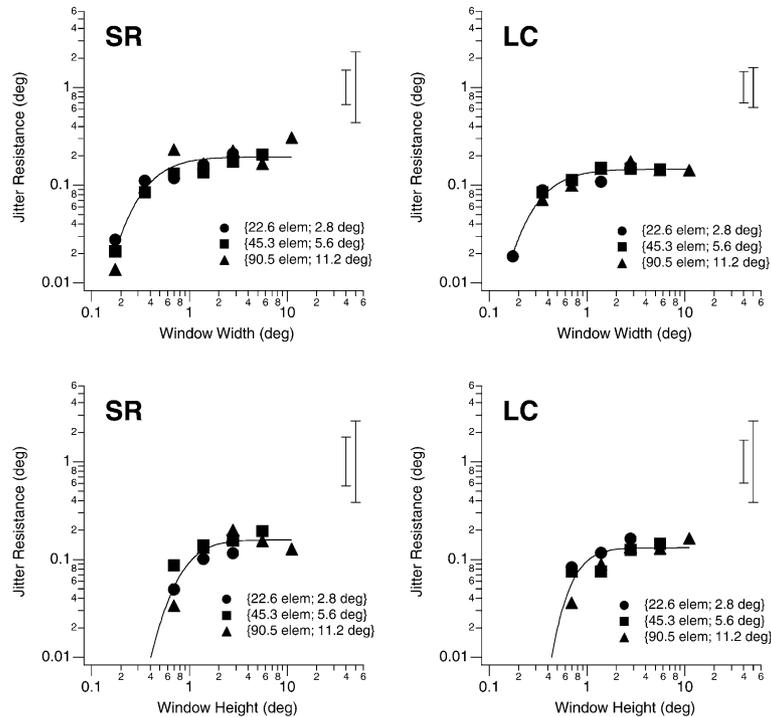


Fig. 9. Results of Experiment 4. Jitter resistance is plotted as a function of window width (top graphs) and window height (bottom graphs) for observers SR and LC. Filled circles, squares, and triangles correspond to the numerosity–size pairings of {22.6 elem; 2.8 deg}, {45.3 elem; 5.6 deg}, and {90.5 elem; 11.2 deg}, respectively. Solid lines are the best fits of four-parameter log–log cumulative normals for data collapsed across all numerosity–size pairings. Error bars show  $\pm 1$  mean and maximum standard deviation.

$$w \times d = \text{deg} \times \frac{\text{elem}}{\text{deg}} = \text{elem}.$$

Similarly, multiplying jitter  $j$  by density  $d$  also leaves elements as the metric. The consequence of multiplying jitter and window size by density is to transform those two variables from a *retino-centric* frame of reference (i.e. visual angle) into *object-centric* frame of reference (i.e. elements).

Fig. 10 replots jitter resistance for the three observers (rows) as a function of window width and height (columns) for all data gathered in Experiments 1–4. The only difference with previous data plots is that jitter resistance and window size are expressed in terms of elements (elem) rather than in visual angle (deg). Solid lines are the best fits of four-parameter log–log cumulative normals to the data in each graph. Arrows on the  $x$  axis indicate the window size corresponding to a factor-of-two drop in performance from maximal jitter resistance. We used this latter criterion as our measure of IR size.

Fig. 10 shows that, allowing for some intra-observer variability, data from all experiments collapse onto a single function when expressed in an object-centric rather than a retino-centric frame of reference. Averaged across observers, the IR width and height include  $\approx 3$  and 6 elements, respectively. The IR's height-to-width aspect ratio is nearly 2:1, the same ratio reported by Dakin and Herbert (1998) in bandpass-filtered random

noise patterns. Although the precise shape of the IR is unknown, the product of IR width and IR height suggests that the IR integrates  $\approx 18$  elements. Table 2 shows the breakdown of IR width, height, aspect ratio, and area across observers.

With respect to jitter resistance, Fig. 10 reveals that, on average, observers SR, FK, and LC tolerated a maximum of 1.4, 1.0, and 1.3 elements, respectively. Across observers, the geometric mean for jitter resistance was 1.2 elements. Jitter resistance, when expressed in elements rather than in visual angle, simply refers to the number of elements that fall within the region of space over which an element's position is randomized. The fact that jitter resistance is in neighborhood of one element means that positional jitter can be tolerated until it exceeds the average spacing between elements. If positional jitter exceeds average element spacing, then adjacent elements are confused for one another and positional information is lost to the observers.

Transforming retinal units into object units is key in achieving scale invariance. In the natural environment, changes in viewing distance introduce simultaneous changes in the spatial-frequency, numerosity, size, and density of the retinal image, and it is therefore ambiguous which of these four properties the visual system uses to achieve scale invariance. In the present study, we have shown that jitter resistance and IR size scale with the inverse of density ( $1/d$ ) in *retinal* units (i.e. visual

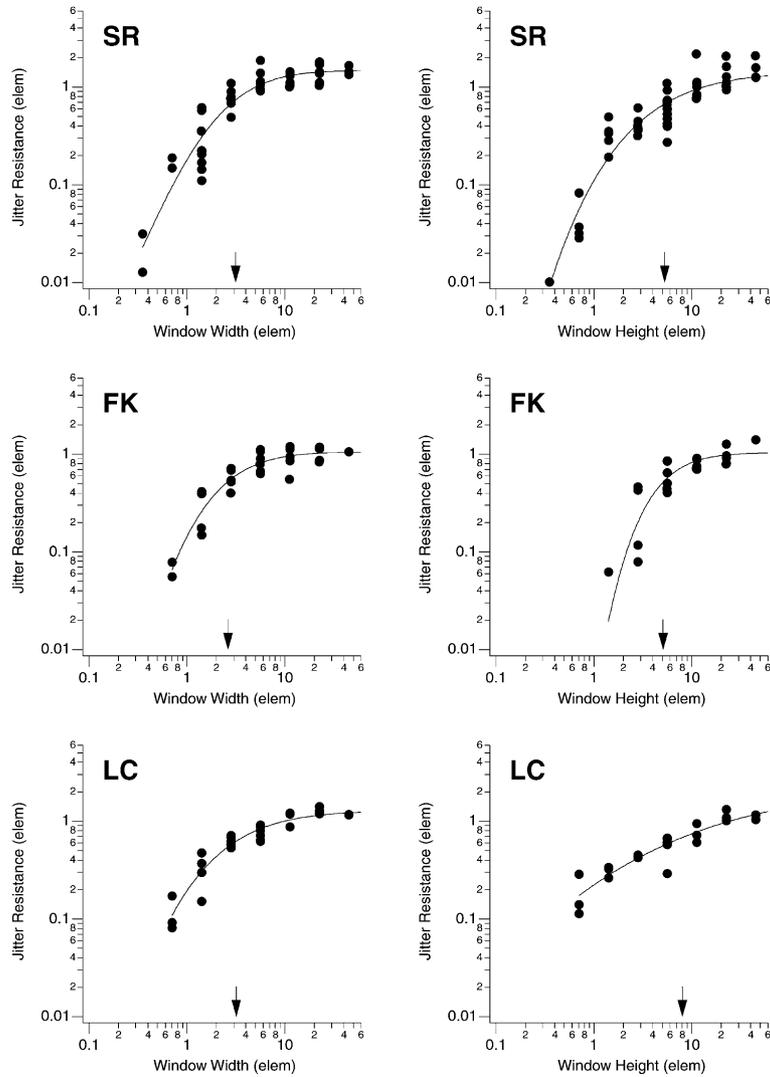


Fig. 10. Visual angle vs. number of elements. Jitter resistance is plotted for the three observers (rows) as a function of window width and height (columns) for all data gathered in Experiments 1–4. Unlike in previous plots, jitter resistance and window size are expressed in number of element (elem) rather than in visual angle (deg). Solid lines are the best fits of four-parameter log-log cumulative normals to the data in each graph. Arrows on the x axis indicate the window size corresponding to a factor-of-two drop in performance from maximal jitter resistance. We use this later criterion as our measure of IR size.

Table 2  
Width, height, and area of integration region (number of elements)

Observer	Width (elem)	Height (elem)	Width-to-height aspect ratio	Width–height product (elem <sup>2</sup> )
SR	3.14	5.31	1:1.69	16.69
FK	2.64	5.10	1:1.93	13.46
LC	3.19	8.02	1:2.51	25.61
Geometric mean	2.98	6.01	1:2.02	17.92

Estimates of IR width, height, aspect ratio, and area expressed for observers SR, FK, and LC. Measurements are expressed in terms of number of elements. Geometric means across observers are also given.

angle) but remain constant in *object* units (i.e. elements). Thus, our data reveal that for our stimuli, *scale invariance is driven by stimulus density, not by stimulus spatial frequency, numerosity, or size*. We should emphasize that our finding that the IR includes a constant number of

elements does not mean that stimulus numerosity is somehow a special variable. The number of elements integrated by the IR is a property of the visual system, and this quantity has no lawful relationship to stimulus numerosity (a variable we can control).

There are several reports of scale invariance in human vision (e.g. Joseph, Victor, & Optican, 1997; Landy & Bergen, 1991; Nothdurft, 1985; Parish & Sperling, 1991). Although their results are not incompatible with ours, they also do not distinguish between the effects of spatial frequency, numerosity, size, and density. However, two recent papers are directly relevant to our study. First, Dakin (2000) measured the effects of orientation variability on the ability of observers to estimate mean orientation in arrays of oriented microelements. Holding either texture numerosity, size, or density constant while allowing the other two variables to covary, the author revealed that performance is a power-law function of the number of orientation samples in the texture irrespective of microelement position. Data also revealed that internal noise is predominantly determined by local texture density, not texture numerosity or size. Results from Dakin (2000) agree with ours insofar as the spatial extent of pooling is not fixed (in terms of visual angle) but is determined instead by stimulus density. This similarity between our respective data sets is interesting given that symmetry detection and orientation pooling are fundamentally different tasks: unlike orientation pooling, symmetry detection necessarily requires that spatial position be encoded. Second, Kingdom and Keeble (1999) investigated mechanisms mediating scale invariance for orientation-modulated textures and found that scale invariance was more disrupted by spatial frequency than density. These results disagree with ours and raise the possibility that, in cases where both positional and local orientation information are important, density is not the only variable that governs scale invariance.

### 8.2. Fourier vs. non-Fourier mechanisms

Barlow and Reeves (1979) have demonstrated that the ideal observer for perfectly symmetric displays fails if dot positions are jittered even slightly because corresponding dots no longer fall in exact mirror locations. The same authors have also shown that to overcome the problem of positional jitter, the ideal observer must include a spatial filtering stage that precedes the computation of symmetry. Although this filtering reduces spatial resolution, it makes the ideal observer more robust by discarding positional information within the area where dot placement is randomized. Therefore, as positional jitter increases, the ideal observer must rely on increasingly coarse spatial filters.

Our study has revealed that jitter resistance is not fixed but rather varies with the inverse of stimulus density ( $1/d$ ). In accordance with Barlow and Reeves (1979), our results imply that symmetry detection is preceded by a filtering stage whose spatial scale is also a function of  $1/d$ . However, the fact that our stimuli were narrowband adds another complexity to this filtering

stage. As we show below, the ability of human observers to operate at spatial scales where no Fourier energy exists constitutes additional evidence that, as Tyler and Hardage (1996) first revealed (see Section 1.2), non-Fourier mechanisms are involved in symmetry detection.

Unlike their Fourier counterpart, non-Fourier mechanisms are non-linear by definition. Models such as the “backpocket” model of Chubb and Landy (1991) typically follow a filter–rectify–filter architecture where the output of bandpass Fourier mechanisms passes through a non-linearity (e.g. an energy operator) before undergoing further filtering. Human vision is sensitive to non-Fourier structures which includes spatial variations in contrast (e.g. Jamar & Koenderink, 1985; Sutter, Sperling, & Chubb, 1995), orientation (e.g. Kingdom, Keeble, & Moulden, 1995), or spatial frequency (e.g. Arsenault, Wilkinson, & Kingdom, 1999)—see Wilson (1999) for review, although see Kingdom and Hayes (2000) and Kingdom and Keeble (2000) for alternative viewpoints.

Fig. 11 shows a scale-space analysis of the Fourier and non-Fourier energy content of three bandpass luminance patterns (panel A) similar to the ones used in Experiment 3 (fixed size—numerosity and density covary). The three luminance profiles consist of 6, 22, and 90 bandpass elements and have densities of 0.5, 2.0, and 8.0 elem/deg, respectively. Profiles are mirror-symmetric about their midpoint, and element positions have not been jittered. Panel B depicts the distribution of *Fourier* energy (in white) over space ( $x$  axis) and spatial frequency ( $y$  axis) for each of the three luminance profiles. For any one given spatial frequency, Fourier energy was computed by convolving luminance profiles with bandpass quadrature-pair filters (cosine-and sine-phase) and summing their squared output. Spatial filters had constant-octave bandwidths and constant-area point-spread functions to ensure that filter outputs were not biased for any one particular scale (Brady & Field, 1995). Panel C shows the distribution of *non-Fourier* energy over space and frequency for each of the three luminance profiles. The only difference between Fourier and non-Fourier energy plots is that we applied our scale-space analysis to the *envelope* of the luminance profiles rather than on the luminance profiles directly. In all other aspects, Fourier and non-Fourier energy profiles were computed in the same way. Energy plots have not been individually normalized are therefore on the same intensity scale.

Panel B of Fig. 11 illustrates that, as in Experiments 2–4, Fourier energy lies in a 1.2-octave frequency band centered on 2.5 cpd irrespective of stimulus density. Yet, results from Experiment 3 (see Fig. 7) show that for a density of 0.5 elem/deg, human observers resisted jitter levels as high as 2 deg (i.e. 0.5 cpd), or more than two octaves away from the Fourier band’s center frequency.

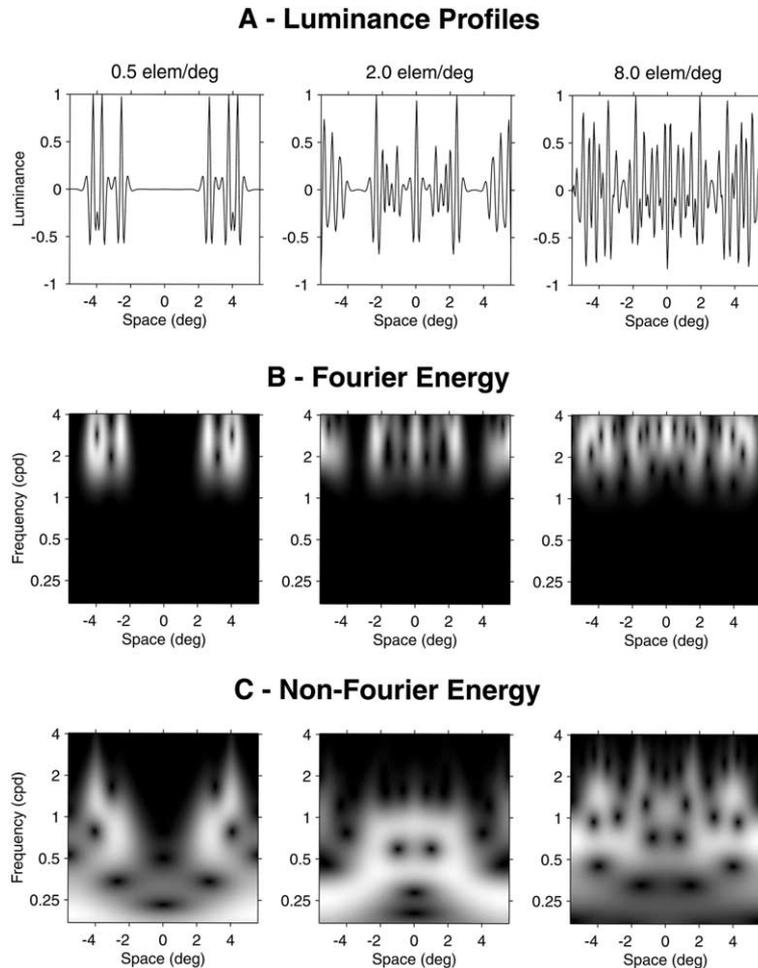


Fig. 11. Fourier and non-Fourier cross-scale image structure. (A) Three one-dimensional luminance profiles consisting of 6, 2, and 90 bandpass elements with densities of 0.5, 2.0, and 8.0 elem/deg, respectively. Profiles are mirror-symmetric about their midpoint, and element positions have not been jittered. (B) Distribution of Fourier energy (white) over space and spatial frequency for corresponding luminance profiles. Distributions of Fourier energy were obtained by summing the squared output from bandpass quadrature-pair filters of variable scale and position. (C) Distribution of non-Fourier energy over space and spatial frequency for corresponding luminance profiles. Non-Fourier energy were obtained exactly as for Fourier energy except that the envelope of the luminance profiles, not the luminance profiles directly, served as the input signal.

Although Fourier energy is absent at 0.5 cpd, panel C reveals the presence of significant non-Fourier energy which observers presumably can use to resist the effects of positional jitter.

### 8.3. Scale selection

The scale-space analyses of Fig. 11 show that non-Fourier energy is present at low spatial frequencies for the three densities used in our study. Although these coarse non-Fourier components can account for high jitter resistance in low-density stimuli, they also raise the problem of why jitter resistance is poor in high-density stimuli. Experiments 2–4 revealed that in the highest density condition (8.1 elem/deg), jitter resistance falls between 0.1 and 0.2 deg (i.e. between 5.0 and 10.0 cpd) yet, as shown in panel C of Fig. 11, non-Fourier energy is available at considerably lower spatial frequencies.

Additional evidence that non-Fourier energy is not tapped in high-density stimuli comes from patterns where mirror symmetry is defined by corresponding elements of opposite contrast polarities. In line with findings from Tyler and Hardage (1996), we have data from narrowband stimuli (Rainville, 1999) which show that contrast polarity has no effect on performance at low densities but hinders performance considerably at high densities. Unlike their Fourier counterpart, non-Fourier mechanisms are insensitive to contrast polarity and therefore cannot account for poorer performance in opposite-contrast high-density stimuli.

If, as panel C of Fig. 11 reveals, coarse non-Fourier components were available in our high-density stimuli, then why did human observers not use them to better resist the effect of positional jitter? One possibility is that the coarse non-Fourier content of our high-density stimuli contained little or no information about mirror

symmetry. The second possibility is that the coarse non-Fourier content contained useful information, but that observers discarded it and selected another spatial scale for symmetry detection. In this section, we report an ideal observer simulation demonstrating that poor jitter resistance for our high-density stimuli reflects a scale-selection property of the visual system rather than a lack of useful information in the non-Fourier content of our stimuli.

From an information-theory standpoint, the ideal way to measure the mirror-symmetry information content of an image is closely approximated by computing the cross-correlation between one half of the image and the mirror reflection of the other half (Rainville & Kingdom, 1999b; Tapiovaara, 1990). We implemented our ideal observer in the MATLAB environment, generated exemplars of the fixed-size stimuli used in Experiment 3 (fixed size—numerosity and density covary), and applied the cross-correlation model to several spatial scales of their non-Fourier energy content. To simplify computations, we gave the ideal observer prior knowledge of the location and orientation of the symmetry axis. Non-Fourier energy profiles were computed in the same way as in Fig. 11, i.e., we (i) filtered each stimulus with a pair of bandpass quadrature filters tuned to peak stimulus spatial frequency, (ii) computed the envelope as the square-root of the sum of squared odd and even filter outputs, (iii) filtered the envelope at 12 spatial scales using a bank of bandpass quadrature-pair filters, and finally (iv) computed non-Fourier energy profiles for each spatial scale as the sum of squared outputs of corresponding quadrature-pair filters. The only difference with Fig. 11 is that stimuli were two-dimensional ( $256 \times 256$  pixels).

In our simulation, the ideal observer made symmetry judgments on 500 symmetric/non-symmetric stimulus pairs for each cell of a  $3 \times 12 \times 5$  matrix containing three densities (0.5, 2.0, and 8.1 elem/deg), 12 peak spatial frequencies (equal log-steps between 0.21 and 4.04 cpd), and five levels of positional jitter (0.0 and four log-steps between 0.09 and 1.52 deg). For each image, the ideal observer returned a single value on a continuum bounded between 0.0 (image is random) to 1.0 (image is perfectly symmetric). In total, symmetry judgments were obtained for 188,000 images. Discriminability (expressed as  $d'$  values) was computed independently for each cell of the  $3 \times 12 \times 5$  matrix by comparing means and variances of the ideal observer's response to symmetric and non-symmetric stimuli (Green & Swets, 1988).

Results from the simulation are plotted in Fig. 12 where the top, middle, and bottom graphs correspond to stimulus densities of 0.5, 2.0, and 8.1 elem/deg, respectively. Each graph plots discriminability ( $d'$ ) between symmetric and non-symmetric stimuli as a function of the center spatial frequency of the non-Fourier energy

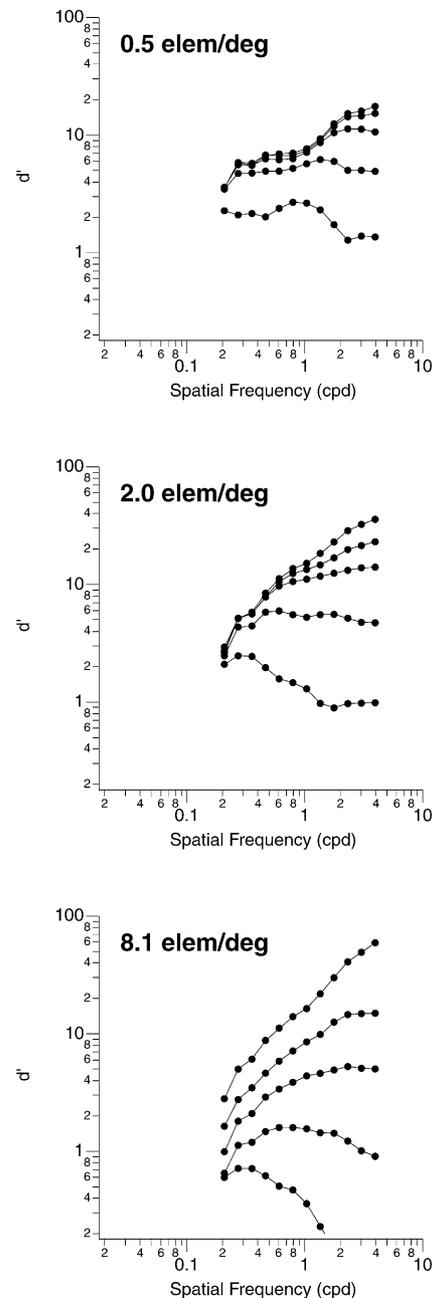


Fig. 12. Results of ideal-observer simulation. Graphs plot discriminability ( $d'$ ) between symmetric and non-symmetric stimuli as a function of peak spatial frequency of bandpass non-Fourier mechanisms. From top to bottom, graphs correspond to stimulus densities of 0.5, 2.0, and 8.1 elem/deg, respectively. Curves represent different levels of positional jitter.

band. Each curve corresponds to one of five levels of positional jitter. The effect of increased jitter is to decrease the overall  $d'$ .

Our simulation has produced two key results. First, in the absence of positional jitter, our high-density condition was more informative than the lower-density ones. This can be judged from the maximum  $d'$  reached for each density. Second, for all three density conditions,

the most informative spatial scale changes as a function of spatial jitter. This effect is particularly pronounced in the high-density condition (bottom graph) where the distribution of symmetry information is highpass at low jitter levels but lowpass at higher jitter levels. For the three density levels tested in this simulation, the optimal strategy is to use high frequencies in the absence of jitter but lower frequencies as positional jitter increases.

Data from our study show that in the high-density condition, human performance collapses even for small amounts of positional jitter. Yet, our ideal observer revealed that our high-density stimuli contained coarse non-Fourier components which, in principle, human observers could have used to overcome positional jitter. Therefore, poor jitter resistance in the high-density condition is a property of the visual system rather than a lack of symmetry information in the coarse non-Fourier content of our stimuli. Poor jitter resistance is also not attributable to a lack of sensitivity to coarse non-Fourier components since human observers were able to use them in low-density stimuli. We conclude from this that, instead of using the optimal strategy of adjusting spatial resolution to resist positional jitter, human vision appears to select its scale of analysis purely on the basis of stimulus density.

Our simulation shows that low spatial frequencies are informative for the particular stimuli we have used, but this does not hold true for all spatial-correspondence problems. For instance, Dakin (1997) has shown that in rotational Glass patterns, the spatial-correspondence information resides predominantly at one spatial scale for any one given stimulus density. If information in naturalistic stimuli is concentrated at a spatial scale that depends on local density, then a density-driven scale-selection mechanism would not only achieve scale invariance but would also highlight important features for further processing. Several computational models of scale selection have already been proposed (e.g. Canny, 1983; Elder & Zucker, 1996; Hayes, 1989; Morrone & Burr, 1997; Witkin, 1983) but to our knowledge, none of these models uses density as a scale-selection criterion. Models of perceived density have been also been proposed (e.g. Allik & Tuulmets, 1991; Durgin & Huk, 1997; see Beaudot & Mullen, 2000 for review) but these have yet to take scale selection into account.

## 9. Summary and conclusions

The present study has produced the following empirical findings:

- Jitter resistance and IR size are *not* determined by stimulus spatial frequency (Experiment 1).
- Jitter resistance and IR size are *not* determined by stimulus numerosity (Experiment 2).

- Jitter resistance and IR size are *not* determined by stimulus size (Experiment 3).
- Jitter resistance and IR size are *inversely proportional* to stimulus density (i.e.  $1/d$ ) (Experiment 4).
- Jitter resistance and IR size include a constant number of texture elements (all experiments).

In the present study, we have shown that jitter resistance and the IR for mirror symmetry are scale invariant because the visual system is able to discard the effects of viewing distance and transform information from a retino-centric to an object-centric frame of reference. The main contribution from this paper, however, is the demonstration that stimulus *density*, not stimulus spatial frequency, numerosity, or size, is used by the visual system to implement scale invariance. An analysis of the scale-space distribution of Fourier and non-Fourier energy, combined with an ideal-observer simulation, revealed that human performance is limited not by stimulus physics but by a neural scale-selection mechanism that uses density as a criterion.

## Acknowledgements

This research has been supported by a Post-Graduate Scholarship from the Canadian Natural Science and Engineering Research Council to SR and by grants from the Canadian Medical Research Council and the Canadian Natural Science and Engineering Research Council to FK. This research was also supported by US Public Health Service grants EY-4885 and EY-1319.

## References

- Allik, J., & Tuulmets, T. (1991). Occupancy model of perceived numerosity. *Perception and Psychophysics*, *49*, 303–314.
- Arsenault, A. S., Wilkinson, F., & Kingdom, F. A. A. (1999). Modulation frequency and orientation tuning of second-order texture mechanisms. *Journal of the Optical Society of America*, *16*(3), 427–437.
- Barlow, H. B. (1959). Sensory mechanisms, the reduction of redundancy, and intelligence. In *Proceedings of the National Physical Laboratory Symposium on the Mechanisation of Thought Processes* (pp. 537–559). London: H.M. Stationery Office.
- Barlow, H. B., & Reeves, B. C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research*, *19*(7), 783–793.
- Beaudot, W. H. A., & Mullen, K. T. (2000). Role of chromaticity, contrast, and local orientation cues in the perception of density. *Perception*, *29*, 581–600.
- Brady, N., & Field, D. J. (1995). What's constant in contrast constancy? The effects of scaling on the perceived contrast of bandpass patterns. *Vision Research*, *35*(6), 739–756.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 443–446.
- Bruce, V. G., & Morgan, M. J. (1975). Violations of symmetry and repetition in visual patterns. *Perception*, *4*(3), 239–249.
- Canny, J. F. (1983). *Finding edges and lines in images*. Boston, MA: MIT—AI Laboratory Technical Report.

- Chubb, C., & Landy, M. S. (1991). Orthogonal distribution analysis: a new approach to the study of texture perception. In M. S. Landy, & J. A. Movshon (Eds.), *Computational Models of Visual Processing*. Cambridge, MA: MIT Press.
- Dakin, S. (2000). An information limit on the spatial integration of orientation signals. *Journal of the Optical Society of America*, *41*(4), S805.
- Dakin, S. C. (1997). The detection of structure in glass patterns: psychophysics and computational models. *Vision Research*, *37*(16), 2227–2246.
- Dakin, S. C., & Herbert, A. M. (1998). The spatial region of integration for visual symmetry detection. *Proceedings of the Royal Society of London—Series B: Biological Sciences*, *265*(1397), 659–664.
- Dakin, S. C., & Hess, R. F. (1997). The spatial mechanisms mediating symmetry perception. *Vision Research*, *37*(20), 2915–2930.
- Durgin, F. H., & Huk, A. C. (1997). Texture density after effects in the perception of artificial and natural textures. *Vision Research*, *37*, 3273–3282.
- Efron, B., & Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. New York: Chapman and Hall.
- Elder, J. H., & Zucker, S. W. (1996). Local scale control for edge detection and blur discrimination. *Lecture Notes in Computer Science, Proceedings of the Fourth ECCV'96* (vol. 1064) (pp. 57–69).
- Enquist, M., & Arak, A. (1994). Symmetry, beauty and evolution. *Nature*, *372*(6502), 169–172.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A—Optics and Image Science*, *4*(12), 2379–2384.
- Green, D. M., & Swets, J. A. (1988). *Signal Detection Theory and Psychophysics*. Los Altos, CA: Peninsula.
- Hayes, A. (1989). Representation by images restricted in resolution and intensity range. *University of Western Australia*.
- Horridge, G. A. (1996). The honeybee (*apis mellifera*) detects bilateral symmetry and discriminates its axis. *Journal of Insect Physiology*, *42*, 755–764.
- Jamar, J. H., & Koenderink, J. J. (1985). Contrast detection and detection of contrast modulation for noise gratings. *Vision Research*, *25*(4), 511–521.
- Jenkins, B. (1982). Redundancy in the perception of bilateral symmetry in dot textures. *Perception and Psychophysics*, *32*(2), 171–177.
- Joseph, J. S., Victor, J. D., & Optican, L. M. (1997). Scaling effects in the perception of higher-order spatial correlations. *Vision Research*, *37*(22), 3097–3107.
- Julesz, B. (1971). *Foundations of Cyclopean Perception*. Chicago: University of Chicago Press.
- Kingdom, F. A. A., & Hayes, A. (2000). Mechanism independence for texture-property modulation. *Investigative Ophthalmology and Visual Science (Suppl.)*, *40*, 317.
- Kingdom, F. A. A., & Keeble, D. R. T. (1999). On the mechanism for scale invariance in orientation-defined textures. *Vision Research*, *39*, 1477–1489.
- Kingdom, F. A. A., & Keeble, D. R. T. (2000). Luminance spatial frequency differences facilitate the segmentation of superimposed textures. *Vision Research*, *40*, 1077–1087.
- Kingdom, F. A. A., Keeble, D. R. T., & Moulden, B. (1995). Sensitivity to orientation modulation in micropattern-based textures. *Vision Research*, *35*, 79–91.
- Landy, M. S., & Bergen, J. R. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Research*, *31*, 1073–1078.
- Møller, A. P. (1995). Bumblebee preference for symmetrical flowers. *Proceedings of the National Academy of Sciences of the United States of America*, *92*, 371–376.
- Morrone, M. C., & Burr, D. C. (1997). Capture and transparency in coarse quantized images. *Vision Research*, *37*, 2609–2629.
- Nothdurft, H. C. (1985). Texture segmentation and pop-out from orientation contrast. *Vision Research*, *25*, 1957–1968.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*(13), 607–609.
- Parish, D. H., & Sperling, G. (1991). Object spatial frequencies, retinal spatial frequencies, noise, and the efficiency of letter discrimination. *Vision Research*, *31*(7–8), 1399–1415.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Rainville, S. J. M. (1999). The spatial mechanisms mediating the perception of mirror symmetry in human vision. In *Unpublished Ph.D. Thesis*. Montréal, Que., Canada: McGill University.
- Rainville, S. J. M., & Kingdom, F. A. A. (1999a). The detection of mirror symmetry over a continuum of contrast polarities. *Investigative Ophthalmology and Visual Science*, *40*(4), S359.
- Rainville, S. J. M., & Kingdom, F. A. A. (1999b). Spatial-scale contribution to the detection of mirror symmetry in fractal noise. *Journal of the Optical Society of America A*, *16*(9), 2112–2123.
- Rainville, S. J. M., & Kingdom, F. A. A. (2000a). The detection of mirror symmetry in spatially-redundant stimuli. *Investigative Ophthalmology and Visual Science*, *41*(4), S218.
- Rainville, S. J. M., & Kingdom, F. A. A. (2000b). The functional role of oriented spatial filters in the perception of mirror symmetry—psychophysics and modeling. *Vision Research*, *40*(19), 2621–2644.
- Simoncelli, E. P., & Portilla, J. (1998). Texture characterization via joint statistics of wavelet coefficient magnitudes. *IEEE International Conference on Image Processing*, *1*, 62–66.
- Sutter, A., Sperling, G., & Chubb, C. (1995). Measuring the spatial frequency selectivity of second-order texture mechanisms. *Vision Research*, *35*(7), 915–924.
- Swaddle, J. P., & Cuthill, I. C. (1994). Preference for symmetric males by female zebra finches. *Nature*, *367*(6459), 165–166.
- Tapiovaara, M. (1990). Ideal observer and absolute efficiency in detecting mirror symmetry in random images. *Journal of the Optical Society of America*, *7*(12), 2245–2253.
- Tyler, C. W., & Hardage, L. (1996). Mirror symmetry detection: predominance of second-order pattern processing throughout the visual field. In C. W. Tyler (Ed.), *Human Symmetry Perception and its Computational Analysis*. Utrecht, The Netherlands: VSP.
- Tyler, C. W., Hardage, L., & Miller, R. T. (1995). Multiple mechanisms for the detection of mirror symmetry. The perception of symmetry. II. Empirical aspects (special issue). *Spatial Vision*, *9*(1), 79–100.
- Wenderoth, P. (1995). The role of pattern outline in bilateral symmetry detection with briefly flashed dot patterns. The perception of symmetry. II. Empirical aspects (special issue). *Spatial Vision*, *9*(1), 57–77.
- Wenderoth, P. (1996). The effects of dot pattern parameters and constraints on the relative salience of vertical bilateral symmetry. *Vision Research*, *36*(15), 2311–2320.
- Wilson, H. (1999). Non-Fourier cortical processes in texture, form, and motion perception. *Cerebral Cortex*, *13*, 445–477.
- Witkin, A. P. (1983). Scale-space filtering. In *Proceedings of the International Joint Conference on Artificial Intelligence* (pp. 1019–1021).