Spatial Vision, Vol. 21, No. 1–2, pp. 149–164 (2007)

Also available online - www.brill.nl/sv

Statistical regularities of art images and natural scenes: Spectra, sparseness and nonlinearities

DANIEL J. GRAHAM * and DAVID J. FIELD

Department of Psychology, Uris Hall, Cornell University, Ithaca, NY 14853, USA

Received 12 June 2006; accepted 19 February 2007

Abstract—Paintings are the product of a process that begins with ordinary vision in the natural world and ends with manipulation of pigments on canvas. Because artists must produce images that can be seen by a visual system that is thought to take advantage of statistical regularities in natural scenes, artists are likely to replicate many of these regularities in their painted art. We have tested this notion by computing basic statistical properties and modeled cell response properties for a large set of digitized paintings and natural scenes. We find that both representational and non-representational (abstract) paintings from our sample (124 images) show basic similarities to a sample of natural scenes in terms of their spatial frequency amplitude spectra, but the paintings and natural scenes show significantly different mean amplitude spectrum slopes. We also find that the intensity distributions of paintings show a lower skewness and sparseness than natural scenes. We account for this by considering the range of luminances found in the environment compared to the range available in the medium of paint. A painting's range is limited by the reflective properties of its materials. We argue that artists do not simply scale the intensity range down but use a compressive nonlinearity. In our studies, modeled retinal and cortical filter responses to the images were less sparse for the paintings than for the natural scenes. But when a compressive nonlinearity was applied to the images, both the paintings' sparseness and the modeled responses to the paintings showed the same or greater sparseness compared to the natural scenes. This suggests that artists achieve some degree of nonlinear compression in their paintings. Because paintings have captivated humans for millennia, finding basic statistical regularities in paintings' spatial structure could grant insights into the range of spatial patterns that humans find compelling.

Keywords: Natural scenes; art; vision; sparseness; spatial frequency spectrum; paintings; power spectrum.

INTRODUCTION

From the standpoint of natural vision, painted artworks are interesting because they are willful, hand-made productions of a visual environment, and this is

*To whom correspondence should be addressed. E-mail: djg45@cornell.edu

true regardless of whether they seek to faithfully represent a part of the natural environment. There is now significant evidence that the early visual system takes advantage of the basic redundant statistical properties of natural scenes (e.g. Atick and Redlich, 1992; Field, 1987, 1994; Geisler *et al.*, 2001; Graham *et al.*, 2006; Hyvärinen and Hoyer, 2000). Therefore, it may be no surprise that artworks created to stimulate the visual system will also take advantage of these statistical regularities (see Note 1). For example, natural scenes show a spatial frequency amplitude spectrum that is well characterized by the function $1/f^k$, where f is spatial frequency and k is approximately 1 (Burton and Moorhead, 1987; Field, 1987, 1993; Ruderman and Bialek, 1994; Tolhurst *et al.*, 1992) (see Note 2). The scale invariance of spatial structure of images from the terrestrial world gives rise to this characteristic spectrum. Artworks that deviate strongly from this regularity may lack a key element of natural scene structure — one that helps identify it as such.

Artists create images that are intended to stimulate the human visual system. Since the visual system is matched to the statistics of the world, we might expect the images that artists create to match these statistics. Despite the freedom granted artists to put marks on canvas as they see fit, we argue that artists do not explore the full range of possible images. Instead, they choose to replicate the basic statistics of the world in their paintings, both for representational and nonrepresentational works. We will argue that even paintings such as those that belong to the abstract expressionist movement share statistical regularities with natural scenes, despite these images' apparent 'randomness'. Indeed, the work of few artists could be described as white noise which, from a statistics viewpoint, would be a truly random image (see Discussion).

A number of measures of image spatial structure besides the amplitude spectrum are in common use including the fractal dimension. The fractal dimension is related to the amplitude spectrum and has been shown to be regular for many natural forms (e.g. Mandelbrot, 1977) and for natural scene outlines (Keller *et al.*, 1987; Kube and Pentland, 1988; Pentland, 1984). Fractal dimension has been be related to perceptions of roughness and complexity (Cutting and Garvin, 1987; Pentland, 1984), discriminability (Geake and Landini, 1997; Knill *et al.*, 1990), and aesthetic preference (Aks and Sprott, 1996; Hagerhall *et al.*, 2004; Spehar *et al.*, 2003). For a 2D luminance surface, fractal dimension D_f is linearly related to the slope k of the spatial frequency amplitude spectrum as plotted on log–log coordinates:

$$k = 4 - D_f \tag{1}$$

(Knill *et al.*, 1990). This relation holds for all images whose amplitude spectra are well described by the function $1/f^k$. Our results are therefore relevant to studies of the regularity of fractal dimension in natural scenes and we give a prediction for a typical D_f to be found in paintings.

Along with these regularities, we examine ways that painted art varies compared to natural scenes, notably in terms of intensity distributions and modeled neural responses. There are likely to be differences in these statistics as a function of luminance because the possible range of luminances is much smaller for paintings than it is for natural scenes. Anecdotally, the typical intensity range for natural scenes may be 1000:1 or more; Jones and Condit (1949) proposed that the average is about 760:1, and this figure is viewed as a rough benchmark by computer graphics engineers who model real-world scenes. We will argue that whereas natural scene luminance values are a function of illumination and reflectance, paintings typically have only one illumination and therefore the reflected luminances are a function of their reflectances only. Since reflectances are rarely less than 3% or greater than 90%, paintings rarely produce an intensity ratio greater than 30:1 (see Dror *et al.*, 2004; Gilchrist, 1979). Typical illuminations are also much lower for art than for scenes since the former are generally displayed inside museums, though this has no effect on the *ratio* between the most intense and the least intense luminance.

It has been argued (e.g. Brady and Field, 2000; Mante *et al.*, 2005) that the compressive nonlinearity of the early visual system serves to produce an efficient representation of luminances and contrasts in natural scenes. It has also been noted that an early logarithmic transform serves to transform a differencing operation into a contrast measure (Field, 1994). In this study, we will explore the role that such a compressive nonlinearity has on both art and natural scenes. We argue that because of the large luminance range found in natural scenes, these early nonlinearities will have significant effects on both the skewness (the third statistical moment of a distribution) and the kurtosis (the fourth statistical moment) of luminant intensities of the image. However, we argue that because artists use a similar compressive nonlinearity in the production of their work (i.e. a kind of ‘artist’s gamma’), the nonlinearities have less impact on the statistics of the artworks.

In this study, we compare the distributions of intensities and contrasts of images before and after the application of a compressive nonlinearity. We also model the responses of visual neurons to these images.

We analyzed examples from two databases — one of calibrated natural scenes, one of paintings. Measurements were made of pixel statistics, Fourier spatial frequency amplitude spectra, and modeled center-surround and wavelet filter responses. Our findings are based on two large and diverse databases but it should be noted that these are biased image sets and therefore there are likely to be some variations in these statistics for different image sets (see Note 3).

MATERIALS

The images in this study were made up of 124 uncompressed TIFF images of works from the Herbert F. Johnson Museum of Art, Cornell University (Ithaca, NY), and 137 calibrated natural scene images (van Hateren and van der Schaaf, 1998). The art images were chosen randomly from the 1139 images classified by the museum as ‘paintings’: images were chosen randomly then included in the test set if they were acquired by the museum photographer between 1999 and 2000. Images acquired by this photographer during this time period all used the same lights (5500K fluorescent; Videssence, El Monte, California), camera (4×5; Sinar AG, Feuerthalen, Switzerland), scan back (Phase One Power Phase; Phase One Inc., Melville, NY) and copy stand (ZBE Inc., Carpenteria, California). The exposure of each image was set so that whites had pixel values around 250. The images were also white balanced and color corrected to visually match the painting. We converted the images to greyscale using the YIQ transform where $L = 0.299 * R + 0.587 * G + 0.114 * B$. The art images were rectangular with dimensions between 818 and 3072 pixels. The image size scaled roughly with the physical dimensions of the artwork. The paintings represent a range of time periods (approximately 12th century AD to the contemporary era), provenances (e.g. Europe, India, US, China — 42% were from Europe and America, 58% were from the Middle East and Asia), subject matter (e.g. still-life, abstract, landscape, portrait, scene painting) and artistic movements (e.g. Rajput miniatures, abstract expressionism, surrealism, Rococo). While some of the artists are well known (e.g. Dubuffet, Bouguereau, O'Keefe), others are not. Natural scene images were randomly selected from van Hateren's database (van Hateren and van der Schaaf, 1998) and scenes with significant blur or buildings were excluded (see Graham *et al.* 2006, for further description of the 137 natural scene images used).

STUDY 1: STATISTICAL REGULARITIES IN PAINTED ART

Methods

We extracted a patch of each image chosen at random from within the boundaries of the image. The patch size was 818×818 pixels (818 pixels is the smallest dimension in the set of paintings). We generated 2D amplitude spectra for each image and performed a rotational average to get the 1D spectra, then averaged these spectra and measured the slope of a least-squares fit to this spectrum on log–log coordinates. We also fitted each individual spectrum and averaged these slopes (see Note 4).

Results

The slope k of the best-fit line to the mean amplitude spectrum plotted on log–log coordinates was -1.23 ($R^2 = 0.97$) for the art and for the natural scenes it was -1.37 ($R^2 = 0.98$) where R is the correlation coefficient for the fit (see Fig. 1).

The mean of the slopes calculated for each individual image was -1.21 ± 0.017 (standard error used throughout) for the paintings and -1.40 ± 0.017 for the natural scenes. This value for natural scenes again reflects the fact that both databases we used are biased samples. The extrema of slopes for the art were -0.70 and -1.56 for *Spring Festival on the River* by Zhang and *Birth of the Virgin* by Giaquinto,

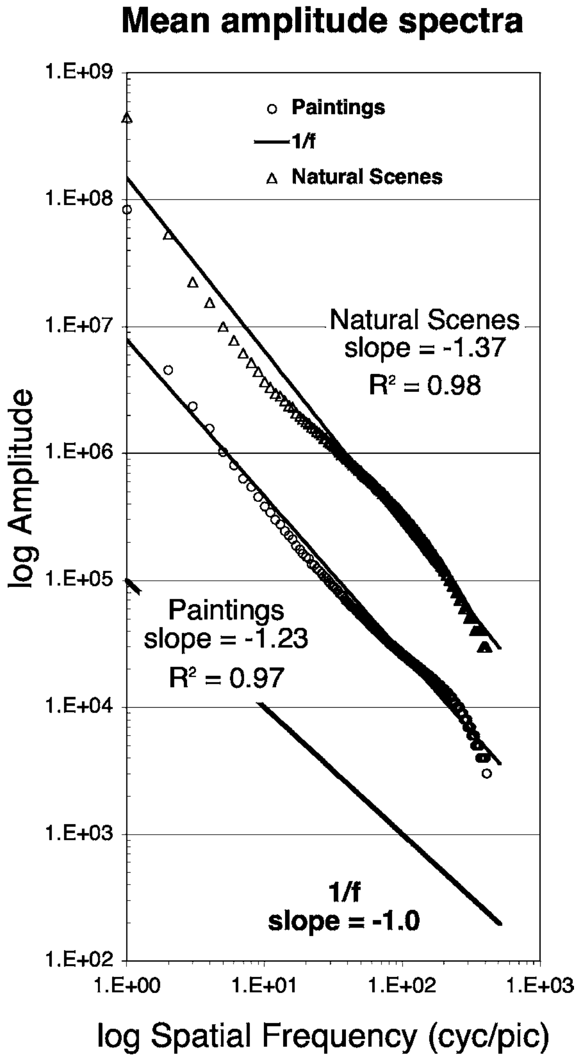

Figure 1. Log–log plot of spatial frequency (cycles/picture) versus amplitude (unitless) for the paintings and natural scenes, taken as an average across 1D spectra. A line with a slope of −1.0 in log–log coordinates is also shown.

respectively (see Fig. 8). A histogram of the distribution of amplitude spectrum slopes is shown in Fig. 2. It shows that the slopes for the paintings fall in a narrow range around the slopes of the natural scenes but that the two image classes show significantly different means ($p < 0.05$).

In general, paintings with abstract representations had slopes similar to the natural scenes, including *Bleeding Rain* by Bluhm, which had a slope of −1.14 (see Fig. 3).

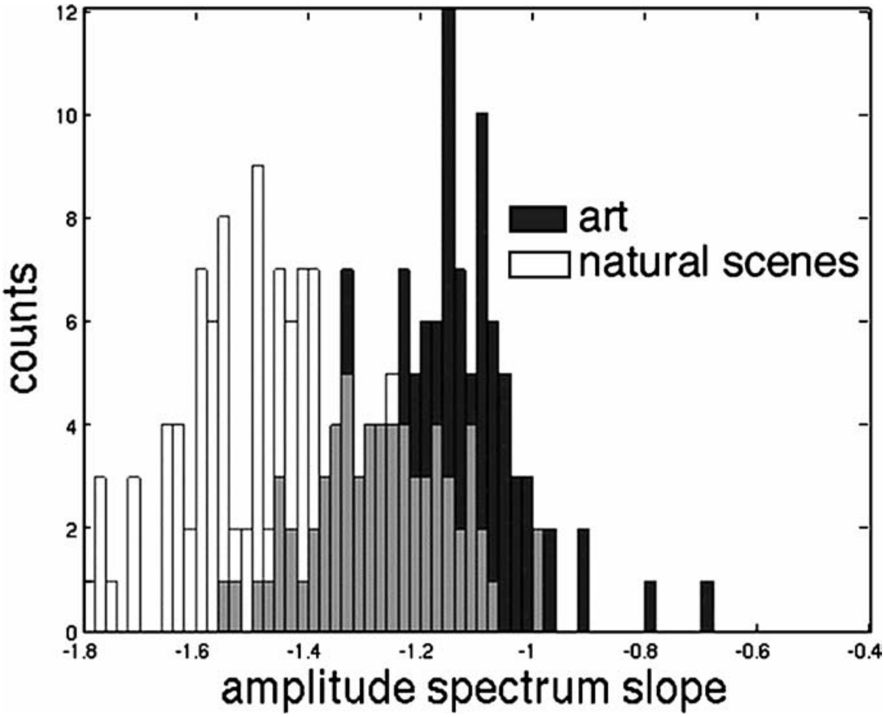

Figure 2. Histogram of the best-fit slopes for all art and natural scene images. Grey shaded areas indicate overlap between populations. Note that these slopes were calculated from fits before averaging the 1D spectra together, whereas the fit in Fig. 1 is for the mean spectrum.

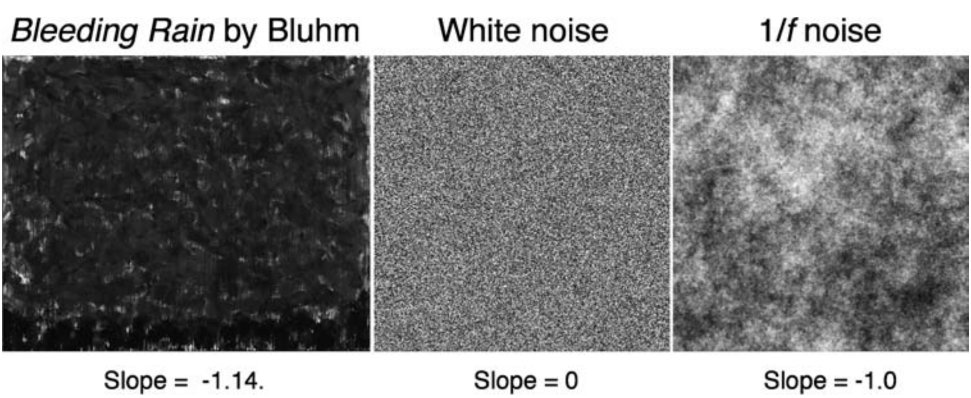

Figure 3. *Bleeding Rain* by Norman Bluhm (left) has an amplitude spectrum that is fitted by a line whose slope is -1.14 in log–log coordinates; white noise (center) and noise whose amplitude is distributed as $1/f$, where f is spatial frequency (right) have slopes of 0 and 1, respectively. Norman Bluhm, *Bleeding Rain*, 1956, Gift of Katharine Komaroff Goodman. Courtesy of the Herbert F. Johnson Museum of Art, Cornell University.

Discussion

Judging by our large but biased sample of paintings, these results suggest that a diverse set of painted art selected at random from the collection of a major university museum follows the amplitude spectrum statistical regularities of natural scenes. However, the paintings show a significantly different mean amplitude spectrum slope compared to the natural scenes. Slopes for the natural scenes are consistent with other studies: Tolhurst *et al.* (1992) put the value of k for natural scenes at -1.2.

Previous work by Taylor *et al.* (1999, 2002) has demonstrated that Jackson Pollock's drip paintings resemble the basic statistics of natural scenes: the fractal dimension of Pollock's paintings was found to have increased over a small range in a relatively orderly way during Pollock's experimentation with drip-painting techniques. Pollock paintings' fractal dimension value range spans a similar range as do 1D natural scene outlines (Spehar *et al.*, 2003). Our result implies that paintings as a group will display a fractal dimension D_f of approximately 2.8 (see equation (1)). This value is higher than the corresponding value for natural scenes ($\sim$2.6).

STUDY 2: SPARSENESS AND NONLINEARITIES

Methods

We used a convolution of a difference-of-Gaussians (DoG) filter with the images described above as a simulation of the ganglion cell response. We tested both the images and a transform of the images where each pixel value was scaled according to its logarithm. The log transform is a rough model of vertebrate cone photoreceptor luminance response (Naka and Rushton, 1966 and Baylor *et al.*, 1987, propose similar compressive functions) and the log recasts differences of intensities as ratios of same, a potentially useful property computationally (Field, 1994) (see Rodieck, 1965, for specifications for the DoG model). In our study, $\sigma_1 = 1.7$, $\sigma_2 = 10.2$ and frame size $=$ 129 pixels. We used 83 images of noise whose amplitude spectrum is distributed proportional to $1/f$, where f is spatial frequency, as a control. We also performed convolutions of the images with log Gabor filters at 3 scales and 4 orientations to model cortical responses (see Field, 1993). Filters uniformly span $\sim$1.5 octaves in frequency. The three scales correspond to filters of center wavelength 3 pixels, 6.3 pixels and 13.23 pixels, respectively. We measured the skewness and sparseness (kurtosis) of the responses for both types of filters to the image, using the linear (untransformed) images and the log transformed images.

Results

In the linear case (i.e. without the nonlinearity), the set of natural scenes gave a more sparse response than the paintings for the DoG filters, and for wavelet

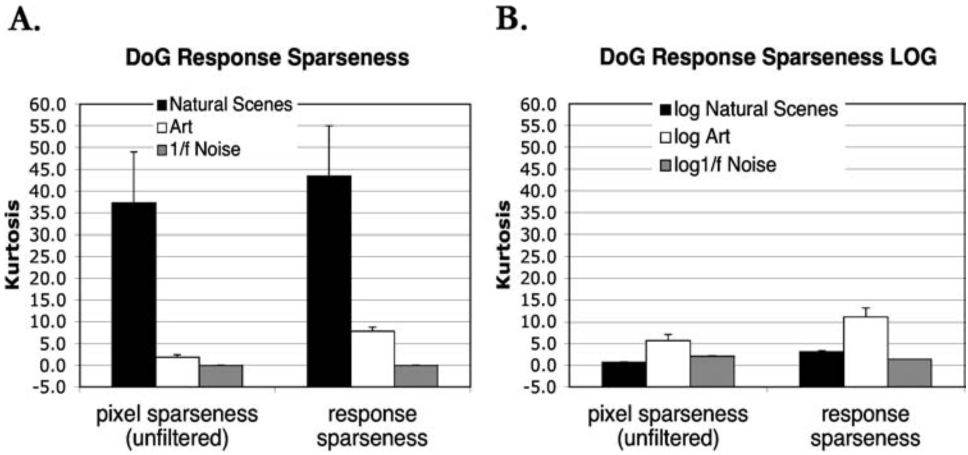

Figure 4. Sparseness (kurtosis) of the pixels and of the responses of a difference-of-Gaussians (DoG) filter to natural scene images, art images and $1/f$-distributed noise, for the linear images (A) and for the log-transformed images (B). Pixel sparseness is the pixel histogram kurtosis of the unfiltered image; response kurtosis is calculated on the convolved image, which is cropped to remove edge effects. Error bars show standard error. Natural scenes show greater sparseness compared to our set of paintings in terms of modeled retinal response (DoG convolution) in the linear case but the paintings gave more sparse responses in the nonlinear case.

filters at all orientations and scales (see Fig. 4). When a $\log(x)$ input nonlinearity was applied, the sparseness of the DoG and wavelet filters' responses to the log-transformed paintings was greater than or equal to the response sparseness of the log-transformed natural scenes (Fig. 5). The same pattern held for the skew (Fig. 5). Horizontal- and vertical-orientation filters generally showed a more sparse response to the natural scenes than diagonal orientations. Mean sparseness of the noise was approximately zero as predicted (-0.008 ± 0.006 before filtering and -0.003 ± 0.002 after filtering in the linear case).

Note that the high mean kurtosis of the linear natural scene images reflects the fact that a small number of these images (12) had extremely large kurtosis values (>70). These images were generally of sunlight filtered through dark foliage and their pixel intensity distributions showed very heavy tails. We plot the median kurtosis values for the linear images before and after DoG filtering in Fig. 6 for comparison.

Discussion

As would be expected, a modeled luminance nonlinearity has a greater compressive effect on the high-dynamic range natural scenes than it does on our set of paintings in terms of sparseness. Adding a luminance nonlinearity gives a more sparse response in modeled retinal and cortical cells for our paintings compared to our natural scenes; without the nonlinearity the natural scenes were more sparse. Cases of very high sparseness in the natural scenes' luminance distributions typically result from a few regions with relatively intense luminance. This produces both

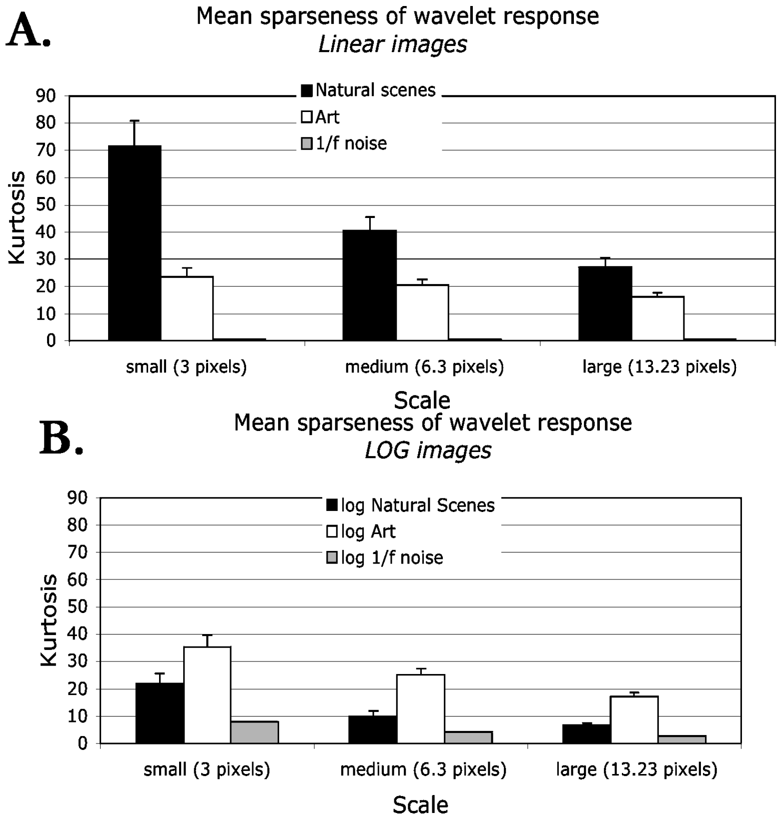

Figure 5. Plots showing response sparseness to wavelet filters across 3 scales (high, medium and low spatial frequency bands in the top, middle and bottom plots, respectively) averaged across orientation. (A) Shows results for linear images, (B) for log-transformed images. Note that whereas responses to the natural scenes were more sparse than the paintings in the linear case across all frequency bands and orientations, the paintings were more sparse when the log nonlinearity was applied before filtering. Plots showing response skewness to wavelet filters across 3 scales (high, medium and low spatial frequency bands in the left, middle and right plots, respectively) averaged across orientation.

high kurtosis (due to the heavy tail) and positive skew, since the tail is typically one-sided toward positive luminances. The compressive nonlinearity reduces this tail in the high luminances, reducing both skew and kurtosis, and it makes the roughly log-normal natural scene distributions more Gaussian. However, this nonlinearity has little effect on the paintings, whose luminance distributions already have a low skewness and a small dynamic range before compression.

Greater sparseness does not necessarily correlate with greater interpretability or aesthetic appeal. But these results suggest that artists apply a type of nonlinear luminance control that is manifest in their paintings. Indeed, if paintings simply took a linear scaling of the luminances in natural scenes, the resulting images would be dominated by low intensities and would appear to be very dark. Figure 7

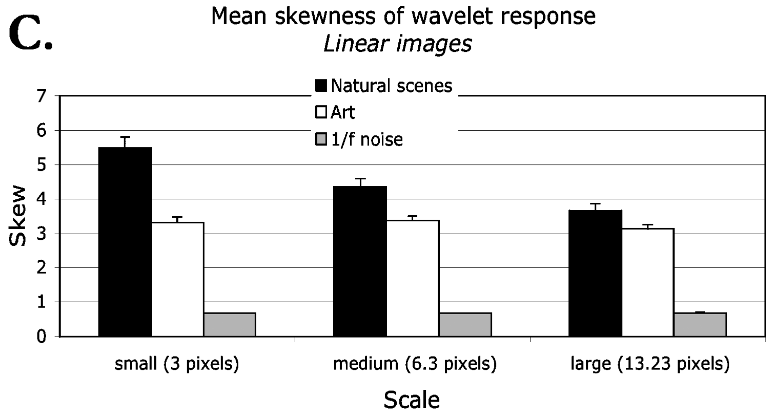

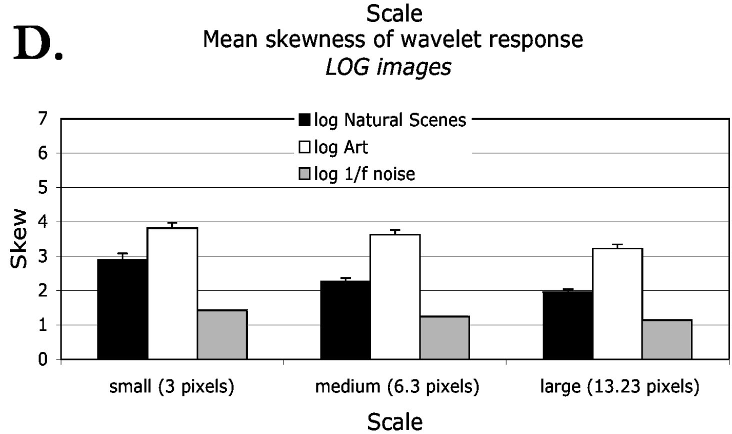

Figure 5. (Continued). (C) shows results for linear images, (D) for log-transformed images. Responses to the natural scenes were more skewed than the paintings in the linear case across all frequency bands and orientations, but the paintings were more skewed when the log nonlinearity was applied before filtering. High, medium and low frequency scales correspond to filters of center wavelength 3 pixels, 6.3 pixels and 13.23 pixels, respectively.

shows an example of a natural scene from our sample scaled linearly with respect to luminance, and scaled using a compressive nonlinearity (a log function).

GENERAL DISCUSSION

Our results suggest that a diverse set of paintings shares many of the basic spatial statistics of natural scenes but that the paintings show characteristic differences in terms of pixel intensity distribution statistics and modeled neural responses.

We used the class of art images labeled paintings in this study because this medium is typically concerned with replicating or interpreting a real or an abstract scene, whereas drawings — and especially caricatures — are more concerned with conveying the identity of objects or people (see Gombrich, 1961). In general, paintings are constructed by applying pigment over the entire surface of the canvas,

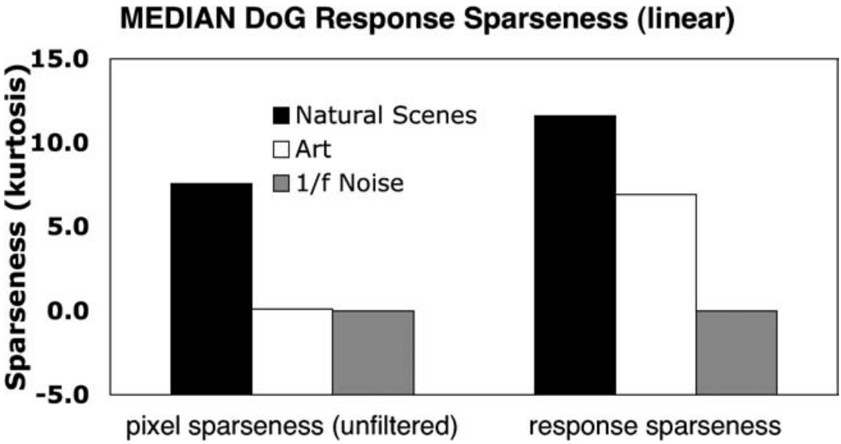

Figure 6. Plot of median sparseness (kurtosis) values for natural scenes, paintings and noise before filtering (pixel sparseness, left plot) and after DoG convolution (response sparseness, right plot). Data are for all 137 linear natural scenes, 124 linear paintings and 83 linear $1/f$ noise images. This plot suggests that the high mean kurtosis for natural scenes shown in Fig. 3 is the result of a handful of images ($\sim$12) with very high skewness and kurtosis. These images generally show sunlight filtered through a dark canopy of trees.

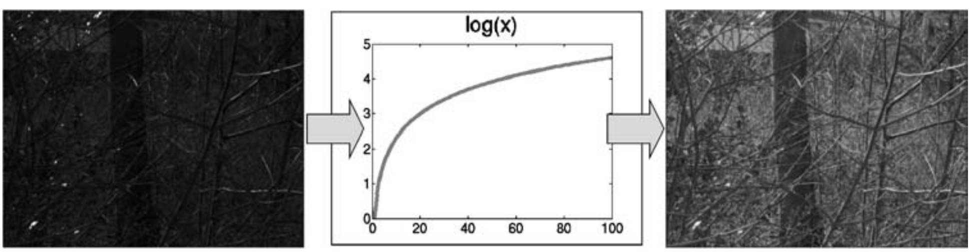

Figure 7. A natural scene from van Hateren's collection (image no. 619, van Hateren and van der Schaaf, 1998) displayed with linear scaling in luminance (left) and after application of a $\log(x)$ luminance nonlinearity. Note that the latter image appears less dominated by very dark and very bright regions.

thus representing material reflectances throughout the scene, while drawings are less involved with representing material reflectances. However, drawings as a class are likely to contain some of the same statistical regularities as paintings (drawings, to a lesser or greater degree, contain detail at many scales, which gives drawings roughly $1/f$ amplitude spectra). We are currently examining ways to analyze a number of classes of art, including drawings.

Why does the visual world have a $1/f$ structure, which is reflected in paintings? Field (1994) and Field and Brady (1997) have shown that a self-similar sum of functions produces global $1/f$ structure in synthetic images. Ruderman (1997) has shown a similar distribution of occluding objects will also produce such a spectrum. These results suggest that $1/f$ structure can be generated by simple sums of objects in the proper proportion of sizes and that a world of objects with fractal edges is not required to produce global 2D fractal structure.

A statistically random painting is simply white noise, i.e., it is an image in which the luminance at every point is randomly assigned from a Gaussian distribution, and the spectrum is flat. Results from our sample suggest that even abstract non-representational works do not approach this level of randomness. Instead, our sample showed similar amplitude spectrum characteristics across artworks.

Our results suggest that artists are applying a form of compressive nonlinearity in the process of utilizing the small range of reflectances available in paint. We are currently investigating ways of measuring this 'artist's gamma' and determining whether different artists use different gammas. We suspect that artists use similar strategies to achieve compression, but this has yet to be determined.

As with the reproduction of geometric form in painting, the above results may imply that we have fallen for a variation of the 'El Greco fallacy' (see Anstis, 2002). This fallacy is the argument that El Greco painted elongated figures because he suffered from astigmatism and saw the world elongated. However, the artist viewing his own canvas would still have the ability to see that his images did not match the world. In a similar way, if the artist was viewing the world through a compressive nonlinearity (as implied by both physiology and psychophysics) then we would not expect the artist to simply produce this nonlinearity on canvas since the image on the canvas must also pass through that nonlinearity. However, in the case of luminance, unlike geometry, the limitations of paint make it impossible to faithfully represent the world on the canvas. Therefore, we can ask the question: What is the most perceptually realistic transformation of luminances under the assumption that the luminance range must be significantly reduced? A linear transform might seem obvious. However, for images with high positive skew, a simple linear compression results in very dark images with a few bright regions (Fig. 7). The artist's gamma (i.e. the nonlinear transform of luminances from the world into paint) may be a general solution to this problem. Photography faces a similar problem of compressing the range of intensities onto the small range of intensities available in a print. Although photographic film also employs nonlinear luminance compression, the printing process allows and often employs local adjustments to exposure (burning and dodging). This can certainly improve the perceptual quality of the image but the method cannot be modeled as a simple function relating input and output. Indeed, artists likely also do something analogous to such local adjustments, but this is beyond the scope of the current study.

This study is not meant to prescribe a formula for acceptable art. It is intended primarily as a means of understanding the visual system through the analysis of artworks. The regularities we observe are of a low statistical order, meaning that they suggest only the most minimal constraints on art, given the infinite variety of possible images. Furthermore, as some have said, artists often view pronouncements about the typical characteristics of art as a challenge to create works that defy those characteristics and are yet undeniably art (Bloom, 2000). We suspect that the regularities we observe are similarly fallible.

Moreover, it is probable that many painters have already defied these regularities and still managed to produce works deemed art. For example, consider minimalist painter Agnes Martin's grid paintings from the 1960s, which are a group of six-foot square white canvases with perfect rectangular grids drawn in pencil. These images likely have a great deal of energy at the spatial frequency defined by the grid lines, and little in other frequency regimes, and thus their spectra could deviate strongly from a $1/f$ relationship. Also, any technique that results in a large degree of blur would serve to steepen the amplitude spectrum (i.e. make its slope more negative). Highly detailed works that show little low-frequency variation would show a shallower spectrum (slope nearer zero). Examples of the images from our sample that show extreme slope values give some idea of these effects (Fig. 8). Although the Giaquinto painting is not blurry, the details are smooth enough relative to the strong large-scale (low frequency) luminance changes to cause a steep tilt in the spectrum.

A

B

Figure 8. (A) Corrado Giaquinto, *The Birth of the Virgin*, 1751–1755. Herbert F. Johnson 'Class of 1922' Acquisition Fund. Courtesy of the Herbert F. Johnson Museum of Art, Cornell University. This is the painting with the steepest (most negative) amplitude spectrum slope in our sample, -1.56. (B) Zhang Ziduan, *Spring Festival on the River*, Late Qing (19th c.). Gift of Drs. Lee and Connie Koppelman. Courtesy of the Herbert F. Johnson Museum of Art, Cornell University. This painting had the shallowest slope (nearest zero), -0.70.

Artists have also challenged the characteristic dynamic range limitations of painting by using light sources as a pigment. Consider for example the work of American artists James Turrell or Dan Flavin (though one could argue these artists are not painters in the strictest sense).

One difficulty in determining the extent to which art as a class of images follows the statistical regularities of natural scenes is the establishment of a representative corpus. Is the art in museums or art books sufficiently representative? Some interesting work in this area attempts to define a canon of impressionist paintings by enumerating scholarly references to those images. Cutting (2006) has shown that even within impressionist art, the 'canon' is a relatively selective grouping and it is a function of the vagaries of early collectors.

Both the natural scene images and the art images we used represent biased samples of their respective image classes. For example, the natural scenes contain no large vistas, and the majority of the art images in our collection are of Asian provenance. We would expect some differences in these statistical relationships for other collections of images though we believe the deviations would be small. We are currently investigating statistical differences among different classes of art using an expanded database. Results suggest there are few low-level statistical differences among classes (unpublished data).

Whether because of taste, history, materials, the structure of the visual system, or the interaction of these factors, human art may tend to adhere to the basic statistical regularities of natural scenes even when the painting is highly abstract or non-representational. Our results suggest that nearly all painted art — including abstract painting — shares essential statistical regularities with natural scenes but that art is characteristically limited in terms of dynamic range in ways natural scenes are not.

The shared and divergent statistics of natural scenes and art may provide new tools for uncovering the coding strategies of the visual system. A full understanding of the similarities and differences between natural statistics and art statistics could lead to insights with regard to human visual coding of the natural environment.

Acknowledgements

D. J. Graham was supported by NEI Grant EY015393 and D. J. Field was supported by NGA Contract HM 1582-05-C-0007. We wish to acknowledge the many current and former staff of the H. F. Johnson Museum of Art in Ithaca, NY, who made this work possible: Frank Robinson (Director), Julie Magura (Chief Photographer), Whitney Tassie, Elizabeth Emrich (Rights and Reproductions), Renee Grinnell (Rights and Reproductions), Oya Rieger (Olin Library, Cornell), Rhea Garen (Olin Library, Cornell), Mike Bevans (Academic Imaging, Inc.) and Peter Hirtle.

NOTES

1. A similar notion has recently been proposed in relation to human writing systems and natural scenes (Changizi *et al.*, 2006).

2. Other studies have used spectra to classify natural scenes, e.g. Torralba and Oliva (2003).
3. Note that this study was carried out using greyscale images and thus it does not attempt to account for regularities in color statistics. Color plays a crucial role in the creation of art and a full theory of how regularities in art are related to the visual system would undoubtedly include an understanding of color.
4. Note that the spatial frequency amplitude spectrum is simply the square root of the power spectrum and therefore the slope of the amplitude spectrum is half that of the power spectrum.

REFERENCES

Aks, D. and Sprott, J. C. (1996). Quantifying aesthetic preference for chaotic patterns, *Empirical Studies of the Arts* **14**, 1–16.

Anstis, S. (2002). Was El Greco astigmatic? *Leonardo* **35**, 208.

Atick, J. J. and Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Comput.* **4**, 196–210.

Baylor, D. A., Nunn, B. J. and Schnapf, J. L. (1987). Spectral sensitivity of the cones of the monkey, *Macaca fascicularis*, *J. Physiol.* **390**, 145.

Bloom, P. (2000). *Descartes' Baby*. Basic Books, New York, USA.

Brady, N. and Field, D. J. (1995). What's constant in contrast constancy: the effects of scaling on the perceived contrast of bandpass patterns, *Vision Research* **35**, 739–756.

Brady, N. and Field, D. J. (2000). Local contrast in natural images: normalisation and coding efficiency, *Perception* **29**, 1041–1055.

Burton, G. J. and Moorhead, I. R. (1987). Color and spatial structure in natural scenes, *Appl. Optics* **26**, 157–170.

Changizi, M. A., Zhang, Q., Ye, H. and Shimojo, S. (2006). The structures of letters and symbols throughout human history are selected to match those found in objects in natural scenes, *American Naturalist* **167**, E117–E139.

Cutting, J. E. (2006). *Impressionism and Its Canon*. University Press of America, Lanham, MD, USA.

Cutting, J. E. and Garvin, J. J. (1987). Fractal curves and complexity, *Perception and Psychophysics* **42**, 365–370.

Dror, R. O., Willsky, A. S. and Adelson, E. H. (2004). Statistical characterization of real-world illumination, *J. Vision* **4**, 821–837.

Field, D. J. (1987). Relations between the statistics of natural images and the response profiles of cortical cells. *J. Opt. Soc. A.* **4**, 2379–2394.

Field, D. J. (1993). Scale-invariance and self-similar 'wavelet' transforms: an analysis of natural scenes and mammalian visual systems, in: *Wavelets, Fractals and Fourier Transforms: New Developments and New Applications*, Farge, M., Hunt, J. C. R. and Vassilicos, J. C. (Eds), pp. 151–193. Oxford University Press, Oxford, UK.

Field, D. (1994). What is the goal of sensory coding? *Neural Comput.* **6**, 559–601.

Field, D. J. and Brady, N. (1997). Wavelets, blur and the sources of variability in the amplitude spectra of natural scenes, *Vision Res.* **37**, 3367–3383.

Geake, J. and Landini, G. (1997). Individual differences in the perception of fractal curves, *Fractals* **5**, 129–143.

Geisler, W. S., Perry, J. S., Super, B. J. and Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance, *Vision Res.* **41**, 711–724.

Gilchrist, A. (1979). The perception of surface blacks and whites, *Sci. Amer.* **240**, 112–123.

Gombrich, E. H. (1961). *Art and Illusion*. Princeton University Press, Princeton, USA.

Graham, D. J., Chandler, D. M. and Field, D. J. (2006). Can the theory of 'whitening' explain the center-surround properties of retinal ganglion cell receptive fields? *Vision Res.* **46**, 2901–2913.

Hagerhall, C. M., Purcell, T. and Taylor, R. (2004). Fractal dimension of landscape silhouette outlines as a predictor of landscape preference, *J. Environ. Psych.* **24**, 247–255.

Hyvärinen, A. and Hoyer, P. O. (2000). Emergence of phase- and shift-invariant features by decomposition of natural images into independent feature subspaces, *Neural Comput.* **12**, 1705–1720.

Jones, L. A. and Condit, H. R. (1949). Sunlight and skylight as determinants of photographic exposure. II. Scene structure, directional index, photographic efficiency of daylight, safety factors, and evaluation of camera exposure, *J. Opt. Soc. Am.* **39**, 94–135.

Keller, J. M., Crownover, R. M. and Chen, R. Y. (1987). Characteristics of natural scenes related to the fractal dimension, *IEEE Trans. Pattern Anal. Machine Intell.* **9**, 621–627.

Knill, D. C., Field, D. and Kersten, D. (1990). Human discrimination of fractal images, *J. Opt. Soc. Amer.* **77**, 1113–1123.

Kube, P. and Pentland, A. (1988). On the imaging of fractal surfaces, *IEEE Trans. Pattern Anal. Machine Intell.* **10**, 704–707.

Mandelbrot, B. B. (1977). *The Fractal Geometry of Nature*. W. H. Freeman, New York, USA.

Mante, V., Frazor, R. A., Bonin, V., Geisler, W. S. and Carandini, M. (2005). Independence of luminance and contrast in natural scenes and in the early visual system, *Nature Neurosci.* **8**, 1690–1697.

Naka, K. I. and Rushton, W. A. (1966). S-potentials from colour units in the retina of fish (*Cyprinidae*), *J. Physiol.* **185**, 536–555.

Norman, R. A. and Werblin, F. S. (1974). Control of retinal sensitivity. I. Light and dark adaptation of vertebrate rods and cones, *J. Gen. Physiol.* **63**, 37–61.

Olshausen, B. A. and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images, *Nature* **381**, 607–609.

Pentland, A. (1984). Fractal-based description of natural scenes, *IEEE Trans Pattern Anal Machine Intell.* **6**, 661–674.

Rodieck, R. W. (1965). Quantitative analysis of cat retinal ganglion cell response to visual stimuli, *Vision Research* **5**, 583–601.

Ruderman, D. L. (1997). Origins of scaling in natural images, *Vision Research* **37**, 3385–3398.

Ruderman, D. L. and Bialek, W. (1994). Statistics of natural images: scaling in the woods, *Phys. Rev. Letts* **73**, 814–817.

Spehar, B., Clifford, C. W. G., Newell, B. R. and Taylor, R. P. (2003). Universal aesthetics of fractals, *Computers and Graphics* **27**, 813–820.

Taylor, R. P., Micolich, A. P. and Jonas, D. (1999). Fractal analysis of Pollock's drip paintings, *Nature* **399**, 422.

Taylor, R. P., Micolich, A. P. and Jonas, D. (2002). The construction of Jackson Pollock's fractal drip paintings, *Leonardo* **35**, 203–207.

Tolhurst, D. J., Tadmor, Y. and Tang, C. (1992). The amplitude spectra of natural images, *Ophthalmic Physiol. Opt.* **12**, 229–232.

Torralba, A. and Oliva, A. (1993). Statistics of natural image categories, *Network: Comput. Neural Syst.* **14**, 391–412.

van Hateren, J. H. and van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex, *Proc. Roy. Soc. B Biol. Sci.* **265**, 359–366.