

## APPLIED MATHEMATICS

# The statistics of style

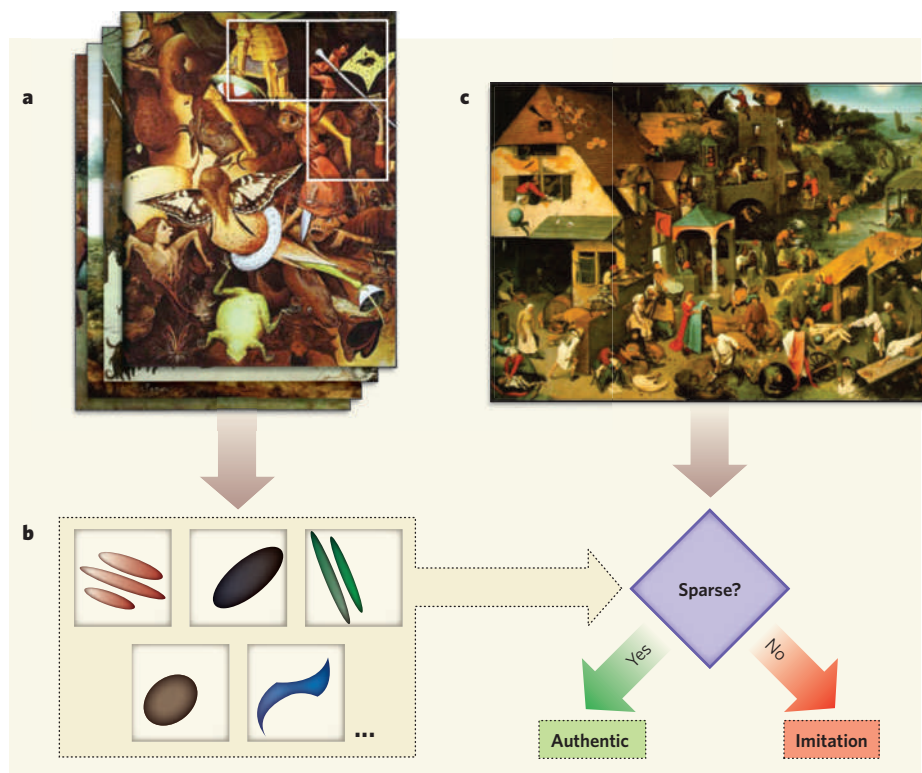
Bruno A. Olshausen and Michael R. DeWeese

**A mathematical method has been developed that distinguishes between the paintings of Pieter Bruegel the Elder and those of his imitators. But can the approach be used to spot imitations of works by any artist?**

What makes the style of an artist unique? For paintings or drawings, what comes to mind are specific brush or pen strokes, the manner in which objects are shaded, or how characters or landscapes are portrayed. Art historians are skilled at identifying such details through visual inspection, and art collectors and museums currently rely on this type of expert analysis to authenticate works of art. Might it be possible to automate this process to provide a more objective assessment? Is it possible to teach a computer to analyse art? In an article in *Proceedings of the National Academy of Sciences*, Hughes *et al.*<sup>1</sup> demonstrate that subtle stylistic differences between the paintings of Pieter Bruegel the Elder and those of his imitators, which were at one time misattributed by art historians, may be reliably detected by statistical methods.

Hughes and colleagues' work is the latest in a stream of research findings that have emerged over the past few decades in the field of 'image statistics'. The players in this field are an unlikely cadre of engineers, statisticians and neuroscientists who are seeking to characterize what makes images of the natural environment different from unstructured or random images (such as the 'static' on a computer monitor or television). Answering this question is central to the problem of coding and transmitting images over the airwaves and the Internet, and, it turns out, it is just as important for understanding how neurons encode and represent images in the brain.

The first image statisticians were television engineers, who, as early as the 1950s, were trying to exploit correlations in television signals to compress the signals into a more efficient format. Around the same time, pioneering psychologists and neuroscientists such as Fred Attneave and Horace Barlow were using ideas from information theory to work out how the particular structures contained in images shape the way that information is coded by neurons in the brain. Since then, others have succeeded in developing specific mathematical models of natural-image structure — showing, for example, that the two-dimensional power spectrum varies with spatial frequency,  $f$ , roughly as  $1/f^2$  (ref. 2), and that the distribution of contrast in



**Figure 1 | Sparse-coding analysis of artistic style.** Hughes and colleagues<sup>1</sup> show that small image patches taken from a collection of authentic works by Pieter Bruegel the Elder (a) can be used to generate a 'dictionary' of visual elements attuned to the statistics of his style (b). A test image (c) can then be authenticated by recreating it with a combination of dictionary elements. If recreation of the test image requires only a few dictionary elements, it is sparse, and labelled 'authentic', whereas if accurate encoding of the test image requires many dictionary elements, it is labelled as an 'imitation'.

local image regions is invariant across scale<sup>3-5</sup>.

Investigators also began applying these and related models to characterize the statistical structure of paintings by particular artists. It was shown, for example, that Jackson Pollock's drip paintings have fractal structure<sup>6</sup>, and that Bruegel's drawings could be distinguished from those of his imitators by the shape of the histogram of wavelet filter outputs, which represent how much spatial structure is present at different scales and orientations<sup>7</sup>. It is this latter work that formed the basis for Hughes and colleagues' study<sup>1</sup>. Instead of using standard wavelet filters, they apply a set of filters that are adapted to the statistics of Bruegel's drawings through a method known as sparse coding.

In a sparse-coding model, local regions of an image are encoded in terms of a 'dictionary' of spatial features; importantly, the dictionary is built up, or trained, from the statistics of an ensemble of images, so that only a few elements from the dictionary are needed to encode any given region. Essentially, sparsity forces the elements of the dictionary to match spatial patterns that tend to occur in the images with frequencies significantly higher than chance, thus providing a snapshot of structure contained in the data. Neuroscientists have shown that such dictionaries, when trained on a large ensemble of natural scenes, match the measured receptive-field characteristics of neurons in the primary visual cortex of mammals. These and

other empirical findings have lent support to the idea that sparse coding may be used by neurons for sensory representation in the cortex<sup>8</sup>.

Rather than attempting to form a generic code adapted to natural scenes, Hughes *et al.*<sup>1</sup> asked what sort of dictionary results from training on one specific class of image — the drawings of Pieter Bruegel the Elder. The dictionary that emerges, not surprisingly, differs from that adapted for natural scenes. In some sense, Hughes *et al.* have evolved an artificial visual system that is hyper-adapted to Bruegel's drawings. Such a visual system will be adept at representing other drawings from this class — that is, other authentic drawings by Bruegel — because they result in sparse encodings. However, it will not be so adept at representing images outside this class, such as drawings by other artists and even those attempting to imitate Bruegel, because they will result in denser encodings — more dictionary elements will be needed to describe each image region (Fig. 1). To put it another way, a picture may be worth a thousand words, but if it's an authentic Bruegel, it should take only a few Bruegel dictionary elements to represent it faithfully.

Can such an approach be used to authenticate works by any artist? And how robust can one expect it to be in practice? Key to the success of this study<sup>1</sup> is the fact that all of the analyses were performed on one particular type of artwork produced by Bruegel — drawings of landscapes. However, Bruegel worked in a variety of media, and his subject matter spanned a wide range of content. Moreover, an individual artist may use various styles. Developing algorithms capable of generalizing across these variations presents a much more challenging problem. Another concern is that it may be possible to defeat this method by generating images that are sparse for a wide range of dictionaries. For example, a geometrical abstract painting by Piet Mondrian would presumably yield a highly sparse representation using a dictionary trained on nearly any artist. Worse still, images randomly generated from the learned dictionary elements would also exhibit high sparsity but would look nothing like a real Bruegel. Thus, sparsity alone may be too fragile a measure for authentication.

One might question other technical choices made by the authors, such as the exclusive use of kurtosis (a statistical measure often used to quantify the degree of 'peakedness' of a probability distribution) to characterize the sparsity of filter outputs; and the analysis of statistical significance is at times puzzling. But Hughes and colleagues have taken a bold step. This is an exciting area of research that goes even beyond forgery detection. Indeed, it begs the question of whether it might be possible to fully capture the style of an artist using statistics. The field of natural-image statistics has advanced beyond the simple sparse-coding models used here, and it is now possible to characterize complex relationships among dictionary elements<sup>9,10</sup>. Intriguingly, all of these models are

generative — that is, they can be used to synthesize images matching the statistics captured by the model, as has already been done successfully with textures<sup>11</sup>. One exciting possibility is that computers could generate novel images that convincingly emulate the style of a particular artist. Perhaps someday the best Bruegel imitators will be computers. ■

Bruno A. Olshausen and Michael R. DeWeese are at the Redwood Center for Theoretical Neuroscience, Helen Wills Neuroscience Institute, School of Optometry (B.A.O.) and Department of Physics (M.R.D.), University of California, Berkeley, Berkeley, California 94720, USA.  
e-mails: baolshausen@berkeley.edu;  
deweese@berkeley.edu

1. Hughes, J. M., Graham, D. J. & Rockmore, D. N. *Proc. Natl Acad. Sci. USA* **107**, 1279–1283 (2010).
2. Field, D. J. *J. Opt. Soc. Am. A* **4**, 2379–2394 (1987).
3. Ruderman, D. L. *Network Comput. Neural Syst.* **5**, 517–548 (1994).
4. Geisler, W. S. *Annu. Rev. Psychol.* **59**, 167–192 (2008).
5. Hyvärinen, A., Hurri, J. & Hoyer, P. O. *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision* (Springer, 2009).
6. Taylor, R. P., Micolich, A. P. & Jonas, D. *Nature* **399**, 422 (1999).
7. Lyu, S., Rockmore, D. & Farid, H. *Proc. Natl Acad. Sci. USA* **101**, 17006–17010 (2004).
8. Olshausen, B. A. & Field, D. J. *Curr. Opin. Neurobiol.* **14**, 481–487 (2004).
9. Karklin, Y. & Lewicki, M. S. *Nature* **457**, 83–86 (2009).
10. Cadieu, C. F. & Olshausen, B. A. *Adv. Neural Inf. Process. Syst.* **21**, 209–216 (2009).
11. Portilla, J. & Simoncelli, E. P. *Int. J. Computer Vision* **40**, 49–71 (2000).

## CELL BIOLOGY

# A brake on lipid synthesis

Fikadu G. Tafesse and Joost C. M. Holthuis

**Although sphingolipids are vital cellular components, the path to their production is paved with toxic intermediates. Orm proteins allow cells to form these lipids without killing themselves in the process.**

Sphingolipids are an unusually versatile class of membrane lipid that, besides providing mechanical stability, have roles in molecular signalling and sorting as well as in cell recognition. Not surprisingly, therefore, a block in sphingolipid synthesis compromises cell growth and survival<sup>1</sup>. But sphingosine, ceramides and other intermediates of sphingolipid synthesis are also potent mediators of cellular stress pathways. When they accumulate in cells or are added experimentally, these compounds can trigger growth arrest and cell death<sup>2,3</sup>. So how does a cell handle the dilemma of generating sufficient amounts of sphingolipids without jeopardizing its viability? On page 1048 of this issue, Breslow *et al.*<sup>4</sup> identify Orm proteins as essential components of a sphingolipid rheostat that allows cells to fine-tune sphingolipid synthesis according to their needs.

Sphingolipid production starts in the intracellular organelle known as the endoplasmic reticulum with the condensation of serine and fatty acyl-CoA. This reaction is catalysed by the enzyme serine palmitoyltransferase (SPT) and yields the first of a series of long-chain bases (LCBs). The LCBs become acylated and are yet further modified to form ceramides — the backbone of all sphingolipids. Ceramides are transported to the Golgi complex, where they acquire a species-specific array of polar head groups to form the complex sphingolipids found primarily on the cell surface. In yeast, SPT consists of two related subunits, Lcb1 and Lcb2, which for maximal activity associate with a third subunit, Tsc3 (ref. 5). As the first and rate-limiting enzyme in sphingolipid synthesis, SPT is an attractive target

for homeostatic regulation of these lipids.

The ease of genetic manipulation in yeast has been invaluable for identifying the enzymes that mediate sphingolipid synthesis. The power of yeast genetics has gained further momentum with the advent of comprehensive strategies for mapping genetic interactions. Such interactions show how the effect of compromising one gene (such as impaired growth) is modulated by perturbing a second gene. An efficient way to study such functional relationships between genes is to create epistatic mini-array profiles, or E-MAPs<sup>6</sup>. This approach allows quantitative measurements to be made over the entire spectrum of possible genetic interactions — from those that aggravate the growth defect to those that alleviate it. E-MAPs also permit analysis of essential genes through the inclusion of hypomorphic alleles, which destabilize the genes' messenger RNAs.

Focusing on a collection of more than 1,400 yeast genes associated with the biology of the endoplasmic reticulum, Breslow *et al.*<sup>4</sup> noticed a striking inverse correlation in the interaction patterns of the *LCB1/2* genes with the gene for the endoplasmic-reticulum membrane protein Orm2. Increased expression of Orm2 — or its relative Orm1 — had the same effect as a reduction in *Lcb1/2* expression, indicating that *Lcb1/2* and Orm1/2 proteins have opposing roles. Moreover, cells lacking Orm1/2 accumulated LCBs and ceramides, whereas cells expressing higher than normal levels of Orm1/2 contained reduced levels of these intermediates. The authors also found that Orm1/2 proteins form a complex with *Lcb1/2* and Tsc3. Together, these findings indicate that Orm