Sussex Research Online

Animal coloration patterns: linking spatial vision to quantitative analysis

Article (Published Version)

Synthesis

Animal Coloration Patterns: Linking Spatial Vision to Quantitative Analysis

Mary Caswell Stoddard[1],* and Daniel Osorio[2]

1. Department of Ecology and Evolutionary Biology, Princeton University, Princeton, New Jersey 08544;   2. School of Life Sciences, University of Sussex, Brighton BN1 9QG, United Kingdom

ABSTRACT: Animal coloration patterns, from zebra stripes to bird egg speckles, are remarkably varied. With research on the perception, function, and evolution of animal patterns growing rapidly, we require a convenient framework for quantifying their diversity, particularly in the contexts of camouflage, mimicry, mate choice, and individual recognition. Ideally, patterns should be defined by their locations in a low-dimensional pattern space that represents their appearance to their natural receivers, much as color is represented by color spaces. This synthesis explores the extent to which animal patterns, like colors, can be described by a few perceptual dimensions in a pattern space. We begin by reviewing biological spatial vision, focusing on early stages during which neurons act as spatial filters or detect simple features such as edges. We show how two methods from computational vision—spatial filtering and feature detection—offer qualitatively distinct measures of animal coloration patterns. Spatial filters provide a measure of the image statistics, captured by the spatial frequency power spectrum. Image statistics give a robust but incomplete representation of the appearance of patterns, whereas feature detectors are essential for sensing and recognizing physical objects, such as distinctive markings and animal bodies. Finally, we discuss how pattern space analyses can lead to new insights into signal design and macroevolution of animal phenotypes. Overall, pattern spaces open up new possibilities for exploring how receiver vision may shape the evolution of animal pattern signals.

Keywords: animal coloration patterns, animal spatial vision, Fourier transform, camouflage, communication, sensory ecology.

Introduction

Complex spatial patterns are often the most striking aspects of animal phenotypes. Indeed, many species' common names refer to pattern. Consider, for example, the variable checkerspot butterfly (*Euphydryas chalcedona*), the two-striped garter snake (*Thamnophis hammondii*), the speckled tanager (*Tangara guttata*), the side-blotched lizard (*Uta* spp.), and the chessboard blenny (*Starksia sluiteri*). The staggering diversity of patterned pelages, plumages, skins, and integuments is probably generated by a limited variety of developmental mechanisms (Maini 2004; Mills and Patterson 2009; Mallarino et al. 2016). Many selective forces shape this diversity, but where patterns function in communication and camouflage, they evolve in response to the vision of the signal receiver—which may be a conspecific, a predator, or a prey item.

Although animal patterns have long intrigued evolutionary biologists (Wallace 1878; Darwin 1888), a strong theoretical framework to describe them remains elusive. Visual stimuli are often considered to have three components: spectral (color), spatial (pattern), and temporal (motion). Quantifying two-dimensional spatial patterns is often more difficult than measuring spectral or temporal components (Rosenthal 2007; Osorio and Cuthill 2013) because spatial patterns typically reveal additional detail with increasing resolution, making "texture" challenging to define (Tuceryan and Jain 1993). Nonetheless, over the past decade, the widespread use of digital imaging to study animal phenotypes (Stevens et al. 2007; Troscianko and Stevens 2015) has led to a proliferation of quantitative methods for analyzing spatial patterns (Rosenthal 2007; Allen and Higham 2013; Troscianko et al. 2017; Pike 2018).

The uptick in research on animal coloration patterns presents challenges and opportunities. The range of approaches raises the questions of how these different methods work and which are the most suitable for characterizing patterns. In many contexts (e.g., camouflage, signaling), it is desirable to analyze patterns in a way that is appropriate for the animal viewer, as is now the norm in color research (Kemp et al. 2015; Renoult et al. 2017). A critical question emerges: to what extent can we specify a relatively small number of measures that describes the visual appearance of a coloration pattern to its natural receiver? These measures would define a low-dimensional perceptual space in

* Corresponding author; email: mstoddard@princeton.edu.
   ORCIDs: Stoddard, http://orcid.org/0000-0001-8264-3170; Osorio, http://orcid.org/0000-0002-5856-527X.

which coloration patterns could be represented; this is a *pattern space* (terms in italics appear in the glossary in the appendix; see also fig. 1), analogous to a color space (Kelber et al. 2003; Kemp et al. 2015; Renoult et al. 2017). A color space is a graphical representation of visual stimuli, informed by the color vision system of the animal viewer (Kelber et al. 2003; Renoult et al. 2017). A few parameters (e.g., the sensitivities of the color cone photoreceptors of the relevant species) specify simple color spaces that provide a powerful means of studying color traits. Color spaces have paved the way for major advances in animal communication research (Chittka and Brockmann 2005; Endler and Mielke 2005; Stoddard and Prum 2008; Schaefer and Ruxton 2015). A pattern space could enable similar progress, allowing visual ecologists to describe and compare animal patterns in a way that is relevant to sensory experience. Ultimately, pattern spaces could generate predictions about animal behavior and offer insights into signal design, development, and macroevolution.

But is a pattern space possible? Whereas color spaces are based on photoreceptor spectral sensitivities, we know much less about neural processes, which are essential to spatial vi-

sion. For spatial vision, the best-known systems are the cat and primate retinocortical pathways, where—as we will see—there has been a fruitful interaction among physiological, psychophysical, and computational approaches. However, even in these well-studied cases, there is much to learn. Nonetheless, the common observation that patterns directed at nonhuman receivers are often effective to our eyes (as camouflage or as conspicuous signals) implies that many features of animal spatial vision are conserved (see also Giurfa et al. 1997; Stevens 2007; Soto and Wasserman 2011; Cronin et al. 2014). Therefore, building a pattern space relevant to animal vision may be an achievable and worthwhile goal.

This synthesis explores four questions: (1) What is a pattern space, and why do we need one? (2) What principles of biological spatial vision could underpin a pattern space? (3) How can we construct a pattern space? (4) What new possibilities does a pattern space make possible? To begin, in the next section we introduce a phenotypic space of patterns (see "A Perceptual Space for Animal Coloration Patterns"). Because a pattern space ideally is informed by the basic principles of spatial vision, we review these concepts
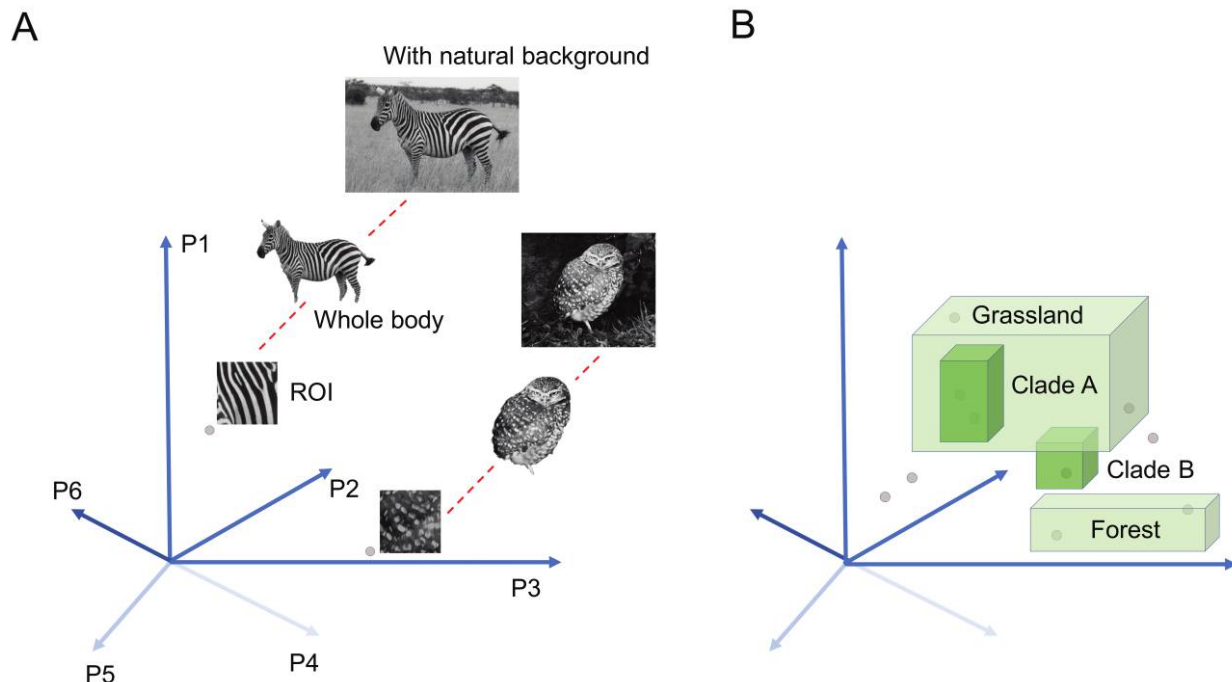
Figure 1: Hypothetical pattern space for animal coloration patterns. *A,* The axes, each representing a pattern parameter *P,* depend on the parameters of interest. A phenotypic pattern space can have *n* dimensions; six dimensions are represented here. A low-dimensional pattern space might include only metrics derived from first- and second-order image statistics, or additional metrics based on edge and feature detectors might be included. Animal patterns can be studied at multiple scales: regions of interest (ROIs), the entire body, or the body in the natural environment. *B,* Once animal patterns are mapped in a phenotypic space, the effects of environment, including microhabitat, and phylogeny on morphospace occupancy can be explored. Each dot represents an animal phenotype in the pattern morphospace. Photos: plains zebra (*Equus quagga*), credit: D. Rubenstein; burrowing owl (*Athene cunicularia*), credit: M. Stoddard.

(see "Biological Spatial Vision: A Brief Overview"). Next, we propose a multistage approach to building a quantitative pattern space (see "Building a Pattern Space Informed by Receiver Vision"). Finally, we examine the ways in which such a space can open new avenues of research (see "New Frontiers in Pattern Space").

A Perceptual Space for Animal Coloration Patterns

What is a pattern space? Ideally, we would like to represent patterns in a low-dimensional perceptual space (fig. 1) that specifies their visual appearance with a small number of measures that corresponds (or can be related) to visual mechanisms. This pattern space would allow quantitative description of patterns as seen by other animals. It would provide a comparison with the backgrounds against which animal patterns are viewed and an understanding of the adaptive landscape in which they evolve. Although this objective may seem abstract, color spaces are common in psychophysics (Kuehni 2001) and now in visual ecology (Renoult et al. 2017), and the same principles are applied to other aspects of perception (Zaidi et al. 2013). Encouragingly, complex patterns such as monkey faces have been described in morphospaces (Allen and Higham 2015), suggesting that pattern spaces can capture visually meaningful variation.

It is helpful to begin with a comparison to color vision (Kemp et al. 2015). One can measure reflectance spectra of objects of interest to an animal and then model photoreceptor excitations to them (Kelber et al. 2003). These excitations can be represented in a color space whose dimensionality is given by the number of spectral receptors (e.g., the color cones; Kelber et al. 2003; Endler and Mielke 2005; Stoddard and Prum 2008). Color spaces are useful for three main reasons: (1) they specify the spectral information that is encoded by the eye; (2) they are low dimensional, reducing spectral data to a small number of measures; and (3) because most reflectance spectra vary smoothly with wavelength, the color space of a tri- or tetrachromatic eye represents nearly all spectral information available in the physical stimulus (Maloney 1986; Vorobyev et al. 1997; Nascimento et al. 2002). Thus, a color space can be a physiologically meaningful, convenient, and nearly complete representation of spectral information.

The obstacles to finding a satisfactory pattern space are evident in light of the merits of color spaces: (1) The large number of photoreceptors in any eye, and hence pixels in the visual image, means that the receptor space has high dimensionality, so any conveniently low-dimensional measure of pattern will have to refer to its representation by neurons in the brain rather than the photoreceptors. However, unlike photoreceptor spectral sensitivities, the response properties of visual neurons are poorly known. (2) The number of these spatial mechanisms is unknown and may be large.

(3) Whereas reflectance spectra vary smoothly with wavelength, it is physically possible for any pixel in an image to have a value uncorrelated with its neighbors (as in spatial white noise; see "Building a Pattern Space Informed by Receiver Vision"). Hence, increasing image quality normally yields additional detail, which suggests that patterns cannot always be obviously reduced to a few dimensions (Tuceryan and Jain 1993).

Despite these three obstacles, the problem is not hopeless because natural images are not random but have a definite statistical structure (Ruderman and Bialek 1994; Olshausen and Field 1996; Ruderman 1997; Simoncelli and Olshausen 2001). Moreover, there is much interest in finding measures of the visual appearance of natural patterns, not only for biologists and psychologists but also for applications such as computer graphics and image data compression (Jain and Farrokhnia 1991; Portilla and Simoncelli 2000). Ideally, a pattern space should encode features that relate in some way to the viewer's perceptual experience. We will fall short of this ideal, since little is known about spatial vision in many species and we do not yet have an adequate understanding—in any visual system—of how the eyes and brains process patterns. That said, we should attempt to measure patterns in a way that is underpinned by the principles of biological vision. In the next section ("Biological Spatial Vision: A Brief Overview"), we provide an overview of our current understanding of how humans and other animals perceive spatial patterns.

Having defined a pattern space, how can it be used? What kind of new conceptual questions will a pattern space make possible? Imagine a clade of patterned animals: felid cats, coral snakes, leaf-mimic katydids, pheasants, poison dart frogs, or bird eggs, to name a few. Now imagine mapping the patterns in a simple pattern space (fig. 1), with relevance to the spatial vision of a chosen signal receiver—a mate, a competitor, a predator. The first questions you might ask are: What pattern information is available to that signal receiver, and what might this reveal about signal design? The second question is: Can the pattern space generate useful predictions about the viewer's behavior? The third question is: What are the limits of patterning in this clade, and what imposes those limits: developmental constraint or natural and sexual selection? The final question is: Comparing across diverse taxa, are there universal features of patterned signals—for camouflage, mate choice, mimicry, or individual recognition? We later expand on these questions (see "New Frontiers in Pattern Space").

As a hypothetical example, consider the remarkably diverse patterns (and colors) of different morphs of the poison dart frog *Dendrobates pumilio*, whose dorsal patterns are variably speckled, spotted, and solid (Siddiqi 2004; Wang and Shaffer 2008). First, plotting these patterns in a pattern space (fig. 1) with dimensions informed by frog or bird spatial vision could reveal differences in the information avail-

able to conspecifics and predators. Next, the pattern space could suggest that certain patterns are always found on red or orange morphs but never yellow or green morphs, leading to the hypothesis that the pattern has an aposematic function. This could be tested in the wild with experimental models, to determine whether avian predators more strongly avoid certain patterns. Then, the pattern space may show that the frogs have evolved speckles and spots but never squiggles or blotches, and exploring the reasons for this (e.g., natural selection, developmental constraint) could be fruitful. Finally, a broader analysis of all frog patterns may suggest that certain patterns (with certain values in the pattern space) are particularly effective for aposematism. Expanding the analysis to another clade—coral snakes (Kikuchi and Pfennig 2010; Davis Rabosky et al. 2016), for example—could reveal the extent to which aposematic patterns share common attributes.

It is easy to imagine how similar questions could be asked in other systems: in tiger moths, for example, a pattern space could be used to investigate trade-offs between aposematism, sexual displays, and camouflage (Nokelainen et al. 2011). Modeling the appearance of tiger moth color patterns from different viewing distances (resolutions) could, for example, reveal whether there are distance-dependent effects: perhaps some aspects of the pattern provide camouflage at a distance but are most effective for warning coloration and mating signals at close range, an idea that has received recent theoretical and experimental support using artificial prey (Barnett and Cuthill 2014).

Biological Spatial Vision: A Brief Overview

Since the 1940s, the study and interpretation of biological vision have been closely linked to progress in image processing, robotics, and artificial intelligence (Campbell and Robson 1968; Marr 1982; Wandell 1995; Yamins and DiCarlo 2016b). This link between physiology and engineering has the major advantage of producing well-specified models of visual processing that can be implemented with image-processing software but carries the risk that we then assume that the model is, in fact, the biological reality. The following account therefore first introduces biological vision and then explains how visual processing is modeled. This provides an understanding of current simulations of animal pattern perception, how they can be interpreted, and how they might be used to answer new questions.

Vision is commonly treated as a hierarchical process, where an optical image is sampled by an array of photoreceptors and then transformed by a series of neural stages. The initial stages are often modeled as *linear filtering* (or *convolution*) with tuned filters, and they are readily simulated using conventional image processing (box 1; fig. 2). The filter outputs provide well-defined measures that capture much visually relevant information about patterns (or visual textures; Bergen and Adelson 1988), but they do not directly identify objects in a scene. Subsequent (or parallel) neural processing is thought to use the filter outputs to locate local features, especially edges, which are then integrated across the image to locate objects. In general, these subsequent stages are much less easily defined as direct operations on the image; instead, they are characterized by their performance on a given task such as object recognition. As we will see later ("Building a Pattern Space Informed by Receiver Vision"), these two processes—spatial frequency analysis and edge and feature detection—can be approximated by two complementary forms of analysis involving (1) the power spectrum and *image statistics* and (2) algorithms for detecting features (including simple edges and complex objects). We now explore the stages of biological visual processing in more detail.

Visual Acuity and Contrast Sensitivity

The first job of the eye is to focus patterns of light in the environment (an image) on an array of photoreceptors in the retina (Land and Nilsson 2012). The spatial detail, or resolution, of this image is generally limited by the angular separation of the photoreceptors (e.g., rods and cones) in visual space (Land and Nilsson 2012). For example, humans and eagles have similar photoreceptor densities, but eagle eyes can resolve finer patterns because their deep foveas increase the effective focal length of the eye, thereby magnifying the retinal image (Reymond 1985). *Visual acuity* is a behavioral measure of spatial resolution, which can be defined as the finest pattern of equally spaced black and white lines—known as a grating—that is distinguishable from a uniform gray field. Acuity is often measured in cycles per degree (cpd), the number of pairs of black and white lines that fill a 1° angle of the visual space. For comparison, a thumb at arm's length subtends about 1°. Acuity is closely related to the quality of the retinal image (Uhlrich et al. 1981). Across animal species, visual acuity varies strikingly. Whereas humans can resolve 60–70 cpd, the Australian wedge-tailed eagle (*Aquila audax*) resolves about 140 cpd, cats resolve about 10 cpd, and bees resolve about 0.5 cpd (Reymond 1985; Land and Nilsson 2012). In all of these animals, and especially humans and eagles, acuity falls markedly from the area of highest acuity to the visual periphery.

Visual acuity measures the finest pattern that can be discerned, but it does not reveal how sensitive the visual system is to different spatial frequencies (see box 1). Contrast sensitivity provides this overall measure of spatial frequency tuning; it is defined as the lowest-contrast grating that can be detected at a particular spatial frequency (see box 1). The eye's overall response to contrast is commonly described by its *modulation transfer function* (MTF), which can be calculated

Box 1: Fourier transform and spatial filtering of images

Fourier analysis.   Fourier analysis (fig. 2) and related methods of spatial frequency analysis are widely used in image processing and provide a theoretical framework for much of visual psychophysics and neuroscience (Campbell and Robson 1968; Meese 2002). Fourier analysis breaks down any continuous signal, such as an image, into a set of sine waves, each with a specific amplitude and relative phase. Fourier power (or, equivalently, amplitude, the square root of the power) and phase can be represented as separate power and phase spectra. Local power spectra are relevant to the perception of pattern (especially preattentively; Bergen and Adelson 1988), and phase spectra are relevant to the location of local features such as edges (Morrone and Burr 1988; fig. 4E). In the power spectrum, each Fourier power component is mapped in the Fourier space as a function of its frequency (with low-frequency components at the center of the map) and orientation in the image. The Fourier power spectra of natural images usually vary smoothly with frequency and commonly follow a power law such that $p = 1/f^n$, where p is power, f is spatial frequency in cycles per unit distance (e.g., visual angle), and the exponent $n \simeq 2$ (Field 1987). It is therefore possible to describe the overall power spectrum by averaging the signal in a relatively small number of bands (typically one octave in width). This measure of the image is essentially equivalent to that made by taking the outputs of local spatial filters (known as wavelets), as used in image processing, and is thought to model important aspects of neural processing in animal visual systems. For 2-D images, one can calculate the average power within a series of circular bins drawn on the Fourier space (Field 1987), which is summarized in a graph of spatial frequency versus Fourier amplitude (fig. 4C). This spectrum represents the second-order image statistics, which account for the spatial relationships between pairs of pixels in the image. Complementary to the amplitude spectrum are spatial phase relations of the sinusoidal components, which are critical for specifying the locations of intensity changes in the image and the presence of local features such as edges and lines.

Spatial filtering.   Once an image is represented (via Fourier transform) in the frequency domain, it can be filtered. To achieve this, the Fourier-transformed input image is multiplied by the Fourier transform of the filter. A filter receives input and transmits some of the input as output, just as a sieve passes only particles below a certain size. Low-pass and high-pass filters transmit low- and high-frequency spatial frequency components, respectively. A band-pass filter transmits only Fourier components within a specified frequency range (or band). Filters may be insensitive to orientation (isotropic), while others are orientation specific. Finally, an inverse Fourier transform gives the output image in the spatial domain, as shown in figure 2. Multiplication in the frequency domain (equivalent to convolution) is easy with conventional image processing tools in MATLAB. This method can be used to simulate optical blurring in the eye, by convolution with a circular Gaussian function (fig. 3) of suitable dimensions. It can also be used to model the outputs of visual neurons with linear response functions, such as retinal neurons with center-surround receptive fields or cortical simple cells.

Global and local transforms.   The Fourier transform is global: it applies to an entire image. Alternative methods exist for breaking down an image into its frequency components. Common examples are circularly symmetrical difference-of-Gaussian (DoG) and the very similar Laplacian-of-Gaussian (LoG) filters (Marr and Hildreth 1980) or oriented wavelet functions, especially Gabor functions (i.e., a Gaussian multiplied by a sinusoid; Daugman 1988; Jain 1989; for simple Gaussian and Gabor filters, see fig. 3). All the foregoing methods produce a local spatial frequency analysis for a part of the image: they provide local information about space (2-D) and frequency, whereas the (global) Fourier power spectrum averages across the entire image (Graps 1995). Importantly, wavelets are thought to be a good model of receptive fields of simple cells in mammalian primary visual cortex, suggesting that similar computational principles might apply to biological vision. For spatial analysis, the specific choice of transform may depend on details such as whether orientation is important (as with striped patterns) or whether the analysis is to be done globally (via Fourier transform) or locally (i.e., piecewise on small areas of the image, via Gabor functions or wavelets). Global analysis, for example, might be relevant to describing the visual backgrounds against which camouflage or communication signals are most effective. In "Biological Spatial Vision: A Brief Overview," we describe how the visual system involves linear filtering. The Fourier, DoG, LoG, and wavelet transforms are types of linear filters.
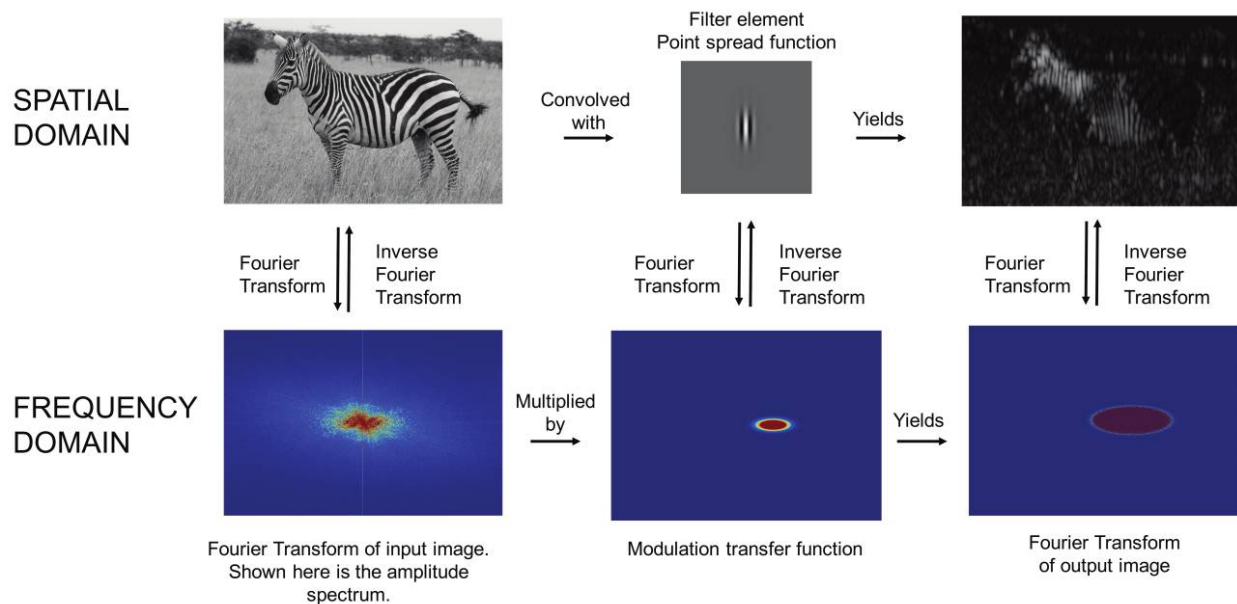
SPATIAL DOMAIN

Convolved with

Filter element
Point spread function

Yields

Fourier Transform / Inverse Fourier Transform

Fourier Transform / Inverse Fourier Transform

Fourier Transform / Inverse Fourier Transform

FREQUENCY DOMAIN

Multiplied by

Yields

Fourier Transform of input image. Shown here is the amplitude spectrum.

Modulation transfer function

Fourier Transform of output image

Figure 2: Filtering an image in the spatial and frequency domains. After Meese (2009). An image can be filtered in the spatial domain or the frequency domain (box 1). *Left*, Fourier analysis breaks down a complex waveform into its constituent sine waves, which can then be added back together (via an inverse Fourier transform) to give the original waveform. As a 2-D image, light intensity varies in two dimensions (*x*, *y*), and the Fourier transform represents information at any orientation in the frequency domain. The sinusoid amplitudes or power (amplitude squared) can be represented as a function of orientation (*lower left*), with low-frequency components at the center of the map. This plot can be averaged over all orientations to give a single amplitude or power spectrum (fig. 4C). *Middle*, now that an image is represented in the frequency domain via a Fourier transform, a filter can be applied. Much as a sieve passes particles below a certain size, a filter transmits some of the input signal and is selective for spatial frequency. Low-pass and high-pass filters preferentially transmit low- and high-frequency spatial frequencies, respectively, and a band-pass filter transmits some intermediate range. Some filters are insensitive to orientation, while others are orientation specific. The Fourier representation of a filter is called the filter's modulation transfer function (MTF). The MTF shows which frequency components the filter will transmit. To obtain a filtered image, the input signal (the Fourier-transformed input amplitude values; *lower left*) is multiplied by the MTF of the filter (*lower middle*), yielding the Fourier-transformed output image (*lower right*). *Right*, an inverse Fourier transform of the Fourier domain output (*lower right*) gives the output image in the spatial domain (*upper right*). Here, the filter (*middle*) is a Gabor filter with wavelength = 10 and orientation = 0 (vertical). For a low-pass filter, convolution is equivalent to blurring the image with a function that corresponds to the intensity distribution (expressed in pixels or angle) produced when a point source is transmitted through the filter, for example, a lens. This distribution is known as the point spread function (PSF) of the lens. More generally, the MTF (frequency domain) is identical to the Fourier transform of the PSF (spatial domain) of a filter. For sensory neurons that act as linear filters, the PSF corresponds to the cell's receptive field. Note that the scenario presented here is slightly oversimplified, since translation between the spatial and frequency domains requires some padding and shifting. Photo: plains zebra (*Equus quagga*), credit: D. Rubenstein.

in a variety of ways, based on the optics of the lens or the packing of retinal ganglion cells (Caves and Johnsen 2017). Related to the MTF is the *contrast sensitivity function* (CSF), which is a behavioral measure dependent on the optical (MTF) and subsequent neural contributions to contrast sensitivity (Michael et al. 2011). Both the MTF and the CSF vary across the visual field and with the light level; it becomes more challenging to resolve images as ambient light levels decline. In daylight conditions, humans have peak sensitivity to medium frequencies, with low- and high-frequency gratings requiring more contrast to be detected.

Receptive Fields and Lateral Inhibition

The early stages of visual processing beyond the photoreceptors are commonly modeled as linear filtering (Wan-

dell 1995), which is implemented by the operation known as convolution (see box 1; fig. 2). This filtering affects sensitivity to different spatial frequencies and, thus, the shape of the CSF. The relevant properties of visual neurons are described by the *receptive field*, defined as a region of the retina over which a cell responds to light (Hartline 1938). Within the receptive field, the response to a given change in light intensity depends on the cell's sign (positive or negative; i.e., whether the cell is excited by an increase or a decrease in light intensity) and sensitivity (i.e., the magnitude of the response to a given intensity change) to light at that particular location. The part of the receptive field that is activated by an increase in light intensity is known as an "on" region and the part activated by dimming as an "off" region. The cell's response is said to be linear if it can be modeled as

a simple summation of the intensity values at each location within the receptive field and nonlinear if this is not the case (Bruce et al. 2003). Simple cells in the primary visual cortex of mammals tend to exhibit roughly linear summation over their receptive fields, while complex cells show nonlinear responses (Hubel and Wiesel 1962, 1968; Bruce et al. 2003; see "Processing in the Mammalian Visual Cortex").

Lateral inhibition, which is a common first stage of visual processing, is implemented directly on photoreceptor outputs in both vertebrate and arthropod visual systems. In lateral inhibition, the stimulation at one location opposes stimulation at a neighboring location and vice versa. Computationally, lateral inhibition is modeled by subtracting from the light intensity signal at each point (or pixel) in the image an average of the responses in the neighboring points (Srinivasan et al. 1982; Bruce et al. 2003). This produces so-called center-surround receptive fields, which are circularly symmetric. For example, the retinal ganglion cells whose axons transmit visual signals up the optic nerve are commonly either center-on, being excited by light in the center of the receptive field and inhibited by light on the periphery (the surround), or center-off, being excited when light falls on the surround but not on the center. These responses can be modeled as a *difference-of-Gaussian* (DoG) function, with a sharp decline in sensitivity as distance from the center increases (Wandell 1995; Bruce et al. 2003). Lateral inhibition means that neurons act as spatial frequency band-pass filters, suppressing low frequencies but transmitting high frequencies (Srinivasan et al. 1982; Bruce et al. 2003). This gives the CSF of most animals a characteristic inverted U-shape, which for (light-adapted) humans peaks at about 3 cpd (Campbell and Robson 1968).

The above description of a receptive field applies well to cells such as photoreceptors or retinal neurons, but in the brain, neurons often respond selectively to complex patterns over wide areas. Here the concept of the receptive field can be generalized to refer to the cell's tuning to additional stimulus parameters. Thus, cells in the mammalian primary visual cortex (visual area 1 [V1]) are tuned to the orientation of lines and edges (Bruce et al. 2003), whereas at a later stage in the visual pathway, cells are tuned to some object- or face-relevant parameter, such as gaze direction, emotional state, or individual identity (Perrett et al. 1992; Quiroga et al. 2005).

Processing in the Mammalian Visual Cortex

In mammals, visual information is transmitted from the retina via the optic nerve and the lateral geniculate nucleus to the primary visual cortex (V1). Here receptive fields are mostly not circularly symmetrical but oriented (i.e., horizontal, vertical, diagonal; see box 1; figs. 2–4; Maffei and Fiorentini 1973; Meese 2002). V1 neurons are classified into two main types: simple cells and complex cells. Simple cells, which receive the most direct inputs from the retina, have approximately linear responses, summing center-surround inputs over the receptive field (Shapley 1997). In contrast, complex cells do not simply sum signal intensities and hence are nonlinear (fig. 4*E*; Shapley 1997). These cells are sensitive to edges and other local features, irrespective of their particular location within the receptive field or their contrast polarity (i.e., an edge with dark above light gives the same response as light above dark; fig. 4*E*; Riesenhuber and Poggio 1999). In other words, the nonlinear response of complex cells means that they respond to specific local features in a way that is more or less independent of their intensity and instead might be related to their relevance to object recognition.

Outside V1, visual information is further processed in many areas of the mammal cerebral cortex and, at least in primates, is thought to follow two visual streams. The ventral, or "what," stream includes areas V2 and V4 and then the inferior temporal (IT) cortex (Bruce et al. 2003) and processes information about the form, color, and recognition of objects. The dorsal, or "where," stream is concerned with spatial location and motion. In the ventral stream, neurons show selectivity to increasingly complex shapes and textures. Here, the spatial receptive fields become progressively more elaborate, with some cells sensitive to textures, borders, curves, and illusory contours (Krüger et al. 2013). In the IT cortex, some cells show selectivity to geometric shapes (Gallant et al. 1993) and even to faces (Perrett et al. 1992). Overall, the IT cortex extracts and integrates features of intermediate complexity from lower levels of the ventral stream, using them to build more sophisticated representations of objects (Krüger et al. 2013).

As a caveat, the general picture we have presented in this section is of bottom-up, feedforward, hierarchical visual processing. This perspective is helpful and has inspired most computer vision approaches (Medathati et al. 2016). However, top-down mechanisms, including attention (Beck and Kastner 2009; Carrasco 2011; Buschman and Kastner 2015), play critical roles mediated by extensive feedback between different parts of the visual pathway (Lee and Mumford 2003) and interactions between the two visual streams (Cloutman 2013; Medathati et al. 2016). For additional reading on spatial vision, we direct readers to textbooks on vision in humans (Wandell 1995; Bruce et al. 2003; Snowden et al. 2012) and animals (Land and Nilsson 2012; Cronin et al. 2014).

Species Differences in Spatial Vision

It is well known that species differences in photoreceptor spectral sensitivities lead to differences in color vision, with consequences for the evolution of color signals. Given the diversity of animal eyes and brains and the likelihood that

they are adapted to different ecological niches, what is the evidence for comparable differences in how animals perceive patterns?

The most obvious difference is in visual acuity, which varies widely across different animal eyes (Caves et al. 2018). Behavioral tests also reveal significant differences in contrast sensitivity (Land and Nilsson 2012). Beyond this stage, little is known about species differences in spatial vision. Neither behavioral nor neurophysiological studies of pattern recognition predict with confidence major species differences. For example, despite some anatomical differences in mammalian and avian brains (Jarvis et al. 2013)—and some intriguing behavioral differences in pattern processing (Qadri and Cook 2015)—spatial vision in mammals and birds likely involves analogous processes (Medina and Reiner 2000; Soto and Wasserman 2011). More generally, the weight of evidence points to remarkable similarity in the way in which visual animals, whether vertebrates, arthropods, or cephalopods, recognize patterns and objects (Cronin et al. 2014). However, evidence that illusions are often perceived differently by various animals (Kelley and Kelley 2014) suggests that there may be intriguing differences we do not yet understand, and much more work is merited.

Building a Pattern Space Informed by Receiver Vision

One can build a pattern space informed by general computational accounts of spatial vision. Broadly speaking, these propose that the image is first filtered (box 1; fig. 2) to transform the pixel-based image data into a more tractable form (Daugman 1988; Jain 1989). This spatial frequency analysis usually begins by breaking down the image into its composite frequencies, achieved by applying a *Fourier transform* (globally, to the whole image) or something similar (such as a *wavelet transform*, for a more local spatial analysis). Once in the frequency domain, the image can be filtered by selectively weighting frequency components from the transformed image (fig. 2). Linear filtering is achieved by convolving the image with the filter (in the spatial domain), which is equivalently—and conveniently—implemented by multiplying the image by the filter representation in the frequency domain (box 1; fig. 2). Which method to use depends on the size and nature of the filter: usually, analysis in the frequency domain is more computationally efficient.

Models of image processing can then make two qualitatively different uses of the outputs of the spatial frequency analysis of the image: (1) the power spectrum directly determines the first- and second-order image statistics, which specify the statistical characteristics of the light intensity values of the image pixels, and (2) objects and other physical structures are identified within the image by means of feature detection. This distinction no doubt simplifies biological reality but is rooted in fundamental aspects of visual

physiology and perception (see "Biological Spatial Vision: A Brief Overview"; Marr 1982; Wandell 1995; Bruce et al. 2003; Snowden et al. 2012) and is reflected in most visual models that are applied to animal coloration patterns. The following sections introduce the steps for building a pattern space. We show how they relate to principles of biological vision and link them to current approaches in the scientific literature.

Step 1: Capturing an Image

Quantitative analysis of a coloration pattern begins with an image (or a video). The image can be of the animal in its natural environment or of a skin in a museum collection. The pattern can be studied in its entirety or as a region of interest (fig. 1). Depending on the question at hand, it may be important to account for the intended signal receiver's visual field and viewing distance (see "Step 2: Modeling Visual Acuity"). In other cases, it may be preferable to capture images in a convenient, repeatable way, without faithfully replicating the receiver's vantage point.

Step 2: Modeling Visual Acuity

The next step is to filter the image to account for the visual acuity of the signal receiver. Humans have higher acuity than most animals, so, for a given viewing distance, there is a risk of grossly overestimating the spatial detail that can be resolved by other animals. A recent review (Caves et al. 2018) highlights the need for work on how visual acuity shapes signal evolution and lays out some useful hypotheses: for example, signals directed at large animals, which generally have large eyes and high acuity, should have more complex patterns, and very fine patterns should be displayed to the highest-acuity parts of the signal receiver's visual field. There are exceptions to these rules: jumping spiders are small but have large simple eyes and excellent resolution (Land and Nilsson 2012). Moreover, for many animals, viewing distance will counteract the effects of eye size.

As we saw previously (see "Visual Acuity and Contrast Sensitivity"), the eye's response to contrast (as a function of spatial frequency) is described by the MTF. We can use the MTF of a given visual system to determine how much spatial information can be resolved from a behaviorally relevant distance. This is achieved by convolving a pattern with a circular Gaussian filter (fig. 3) whose dimensions are given by the eye's resolution, based on anatomical or behavioral data, and the assumed viewing distance of the object of interest (Vorobyev et al. 2001; Land and Nilsson 2012). The tool AcuityView implements this function (Caves and Johnsen 2017): it first converts the image to the frequency domain via Fourier transform, then multiplies the image by the
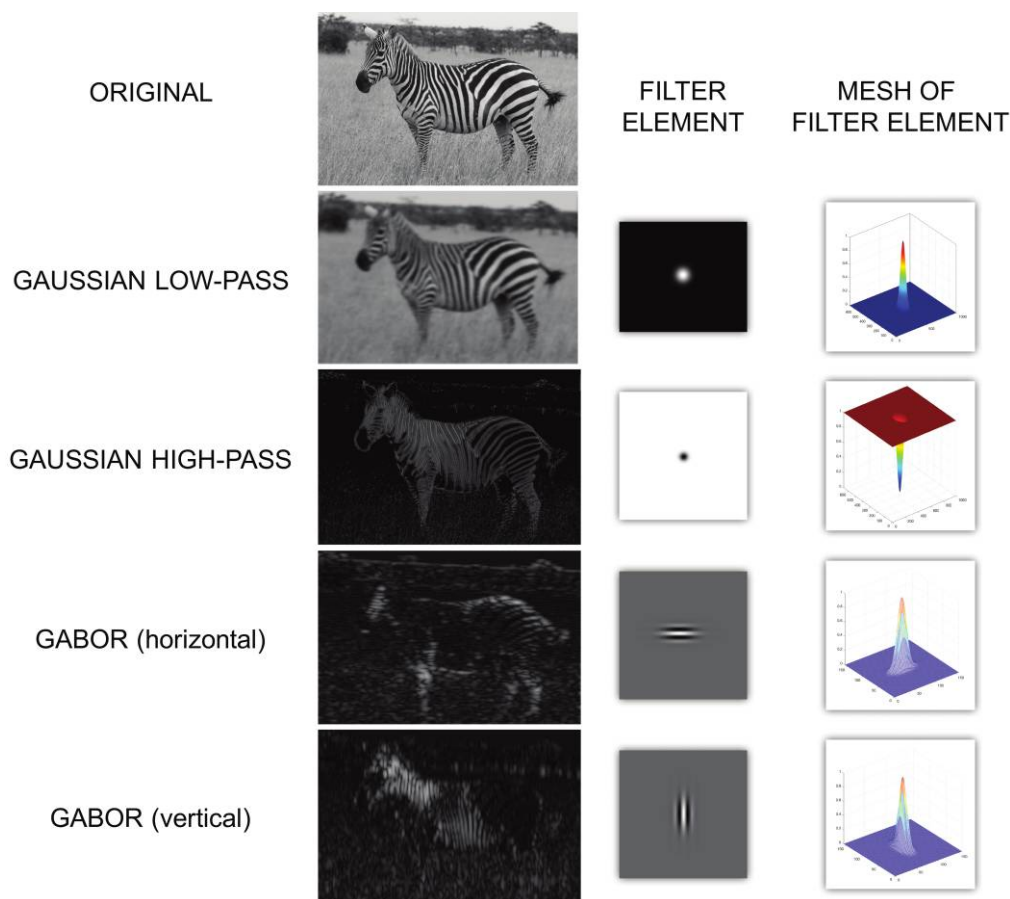
Figure 3: A variety of filters can be applied. Early stages of visual processing are often modeled as convolution with filters. These filters can correspond to optical blurring by the retina (Gaussian low pass), or the receptive fields of neurons such as the circularly symmetric center-surround cells of the retina, or the oriented simple cells of the primary visual cortex (V1). A Gaussian low-pass filter lets only low-frequency information through, blocking higher frequencies. A Gaussian high-pass filter lets only high-frequency information through, blocking lower frequencies. A horizontal Gabor filter accentuates horizontal information. A vertical Gabor filter accentuates vertical information. The filter element is shown next to its 3-D mesh representation, with red (intensity = 1) representing the frequencies transmitted and blue (intensity = 0) representing the frequencies blocked, with gradations in between. The 2-D and 3-D filter representations contain the same information, but the gradations are more apparent in the 3-D mesh versions. Images were generated using a modified version of Gaussian_image_filtering.m (S. Rao, MathWorks File Exchange). Photo: plains zebra (*Equus quagga*), credit: D. Rubenstein.

relevant MTF, and then returns the image so that only the resolvable spatial frequencies are retained.

The MTF, which is based on processes in the eye, can then be compared to the behavioral CSF (see "Visual Acuity and Contrast Sensitivity"), which corresponds to the MTF modified by postretinal neural processes (Michael et al. 2011). The CSF is tested by recording responses to grating patterns and is known for a wide range of animals; typically, the CSF varies with the light level (Land and Nilsson 2012). To simulate how an image would be modified by a given CSF—combining the effects of image blurring (by the eye's optics and photoreceptors) and lateral inhibition—an image can be filtered with a series of circularly symmetrical filters, such as DoG functions or oriented filters (figs. 2, 3;

see "Visual Acuity and Contrast Sensitivity"). Melin and colleagues (2016) used estimates of CSFs to model how zebra stripes are perceived by zebra predators and conspecifics. They concluded that zebra stripes are probably not cryptic, since zebras are just as detectable to a lion as are other similarly sized prey animals.

Step 3: Calculating Power Spectra and Image Statistics

Once the image has been adjusted for a signal receiver's visual acuity, the next step is to measure its first- and second-order image statistics or, equivalently, the spatial frequency *power spectrum* (Field 1987; van der Schaaf and van Hateren 1996; box 1; fig. 2). What are image statistics? An im-

age comprises a 2-D array of pixels, each having a certain intensity value. The first-order image statistics characterize the overall distribution of these pixel values—typically, their mean and variance—while the second-order statistics characterize the correlations between pairs of pixels (Field 1987; Ruderman and Bialek 1994).[1] Consider two extreme and unlikely images: (1) a uniform array, where all pixels have the same value so that, given the value of one pixel, the intensities of all are known, and (2) spatial white noise, where the intensities of each pixel in an image are drawn at random from some distribution, so knowing the value of any given pixel carries no information about any other. Natural images fall between these extremes. This is because nearby points tend to belong to the same object and to fall under the same illumination. Consequently, the intensities of neighboring pixels tend to be positively correlated, and this correlation falls with distance from the point in question (Field 1987; Ruderman and Bialek 1994).

Spatial correlation of pixel values, and hence the second-order image statistics, is usually characterized by the spatial frequency power spectrum, derived from a global Fourier analysis or a spatially local approximation (Baddeley 1992; van der Schaaf and van Hateren 1996). The Fourier transform, which is common in image processing, converts the familiar spatial pixel-based representation to a representation in the frequency domain, based on sine waves of varying *amplitude* and relative *phase*, which can then be represented as separate power (or amplitude) and phase spectra (see box 1; figs. 2, 4). In the 2-D power spectrum, each power or (amplitude) Fourier component is mapped in the Fourier space as a function of its frequency (with low-frequency components at the center of the map) and orientation in the image. This can be simplified by averaging across orientations (fig. 4C). In this simplified graph, power is the average power within a series of circular bins drawn on the Fourier space (Field 1987). The power spectrum makes it easy to see whether neighboring pixel intensities are correlated in the image. For images in which the correlation between pixels declines with distance (in the image), the result is an averaged power (or amplitude) spectrum with a negative relationship between frequency and power (or amplitude). Over very large distances (low spatial frequency information), pixels are uncorrelated (high variance, high power), but over short distances (high spatial frequency information), pixels are highly correlated (low variance, low power). As we stated previously, natural scenes tend to have this property (fig. 4C). For the purposes of analyzing animal patterns, image power

can usually be specified for a relatively small number of separate spatial frequency bands, as approximated by methods such as granularity analysis (Barbosa et al. 2008; Troscianko and Stevens 2015).

Conveniently, the Fourier analysis (box 1; fig. 2) is functionally analogous to filtering by linear neurons (circularly symmetric or oriented) tuned to a particular range of spatial frequencies. In this way, we can think of performing a Fourier transform—and thus describing the power spectra—as simulating the linear filtering processes in biological vision (see "Biological Spatial Vision: A Brief Overview"). Here, we refer to the process of performing a Fourier transform (or a similar transformation, such as the wavelet transform) to determine its underlying statistical properties as *spatial frequency analysis*. Given that natural images or animal patterns can often be defined in terms of a relatively small number of spatial frequency bands (Meese 2002; and that second-order image statistics are not dependent on spatial phase), this step gives us a low-parameter and visually relevant perceptual space in which to start characterizing animal coloration patterns (fig. 1).

There is abundant evidence that human observers are sensitive to the spatial frequency power spectrum in images and hence to first- and second-order image statistics. It has long been argued that human vision produces some type of local spatial frequency analysis, decomposing the image into separate spatial frequency bands (Campbell and Robson 1968; Meese 2002), perhaps implemented by the simple cells of the primary visual cortex (Maffei and Fiorentini 1973; Marčelja 1980; Field 1987). Importantly, it has been found (especially in a brief preattentive view) that the appearance of visual textures for humans and other animals depends on the power in different spatial frequency components of the pattern rather than the phase relations between these components (Bergen and Adelson 1988; Sutter et al. 1989; Malik and Perona 1990; Meese 2002). Moreover, the receptive fields of visual neurons, especially simple cells of the mammalian primary visual cortex, resemble *Gabor filters* (see box 1).

Spurred on by a few classic studies, researchers in the past decade have adopted various objective methods for characterizing patterns; among these, power spectra (and closely related metrics) are common. Godfrey and colleagues (1987) used a Fourier transform to calculate the power spectra for an image of a zebra and an image of a tiger. They found that at high spatial frequencies, the tiger and its background are well matched, while the zebra differs from the background. Similarly, Kiltie and Laine (1992) used measures of second-order image statistics to analyze visual textures in animal camouflage. In a subsequent study, the authors used wavelet analysis to calculate the degree of difference (based on energy distribution patterns) between two textures, applying this to questions about camouflage in

---

1. In principle, one can also characterize images and patterns by higher-order statistical relationships—between triplets, quadruplets of points, and so forth (Julesz et al. 1978; Julesz 1981)—but, in practice, higher-order image statistics have not been useful for classifying natural images (Ruderman 1997) or predicting discrimination of natural textures by humans (Bergen and Adelson 1988; Malik and Perona 1990).

A
ORIGINAL

B
FOURIER TRANSFORM
OF ORIGINAL

C
AMPLITUDE SPECTRUM

D
WHITENED IMAGE
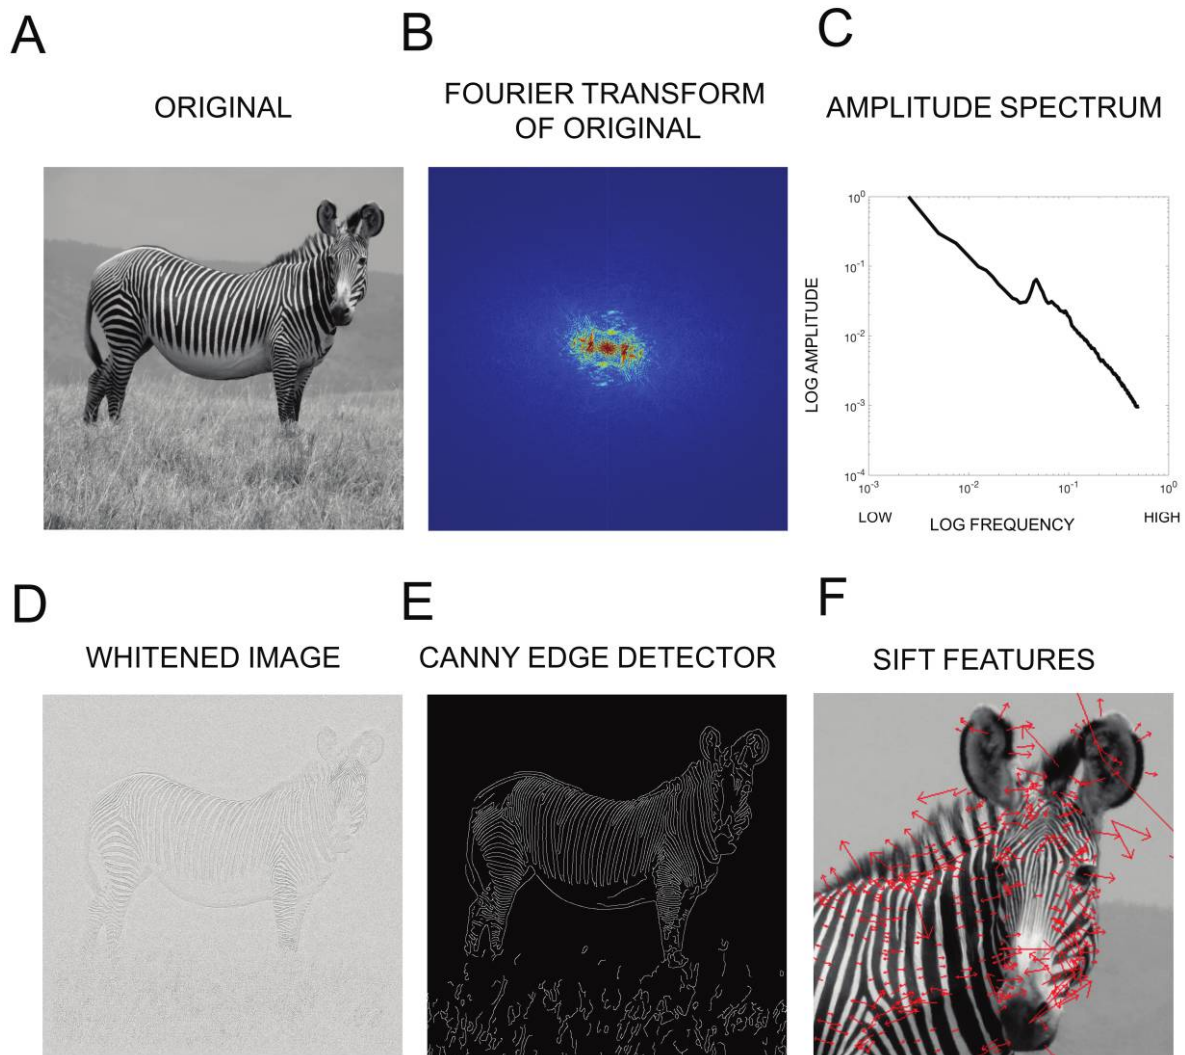
E
CANNY EDGE DETECTOR

F
SIFT FEATURES

Figure 4: Power spectra, edges, and features. The 2-D power spectrum (*B*) of the zebra image (*A*) shows a clear anisotropy (directionality), with greater power at high frequencies in the horizontal than the vertical directions, probably due to the zebra's vertical stripes. Averaging over all orientations (*C*), the power spectrum plotted on a log-log plot is approximately linear, with a slope of −1. Natural images typically obey such a power law (Field 1987). It may be that the noticeable deviation from the straight line is due to the zebra's stripes. In the whitened image (*D*), the power spectrum is altered to resemble white noise (zero spatial correlation). Power is equal at all frequencies, and when the spatial phase information is preserved, this transformation (which approximates high-pass filtering; fig. 3) does not render the image unrecognizable. There is no objective definition of a visual edge (*E*). Operators such as the Canny edge detector (Canny 1986) identify locations in the image that are likely to correspond to the borders of objects, thereby facilitating segregation of figure from ground. In this case, the edge signals created by the stripes might be incorrectly interpreted by the visual system as object borders, which is the principle thought to underlie disruptive camouflage (Osorio and Srinivasan 1991; Stevens and Merilaita 2009). *F*, Scale-invariant feature transform (SIFT) features identified using NaturePatternMatch (Stoddard et al. 2014). Image whitening was performed using the Advanced Digital Imaging Laboratory using MATLAB (L. Yaroslavsky). Photo: Grevy's zebra (*Equus grevyi*), credit: D. Rubenstein.

tigers, zebras, and peppered moths (Kiltie et al. 1995) and to egg crypsis (Westmoreland et al. 2007). Granularity analyses (Troscianko and Stevens 2015), which have been applied to the patterns of cuttlefish (Barbosa et al. 2008; Chiao et al. 2009) and the markings of bird eggs (Spottiswoode and Stevens 2010; Stoddard and Stevens 2010), use *octave*-wide *isotropic filters* (i.e., circularly symmetric) to measure image

power in a small number of frequency bands. Metrics derived from granularity analysis have sometimes been combined with color and luminance in a multidimensional space (Spottiswoode and Stevens 2010; Caves et al. 2017), yielding insights about which visual traits are used by parasitized birds in egg rejection decisions. This multidimensional space offers the benefit of merging aspects of color and pattern spaces.

How best to achieve this is certainly an area ripe for future work, especially in light of the fact that achromatic and chromatic processing in some animals may not be entirely independent, as once believed (Lind and Kelber 2011). The methods described so far in this subsection are all essentially based on spatial frequency analysis, power spectra, and image statistics.

Other pattern metrics may also be, in effect, measures of spatial frequency power spectra. For instance, fractal dimension (FD) has been used to describe the degree to which coloration patterns change when analyzing a pattern at different scales (Pérez-Rodríguez et al. 2017), where better body condition in male red-legged partridges appears to be linked to higher FD in the streaky black plumage on the bird's bib (Perez-Rodriguez et al. 2013). A similar approach has also been used to describe the distribution, size, and FD of speckles on bird eggs (Gómez and Liñán Cembrano 2017). It is likely (but would need to be confirmed) that fractal measures of animal coloration patterns are closely correlated with spatial frequency measures (Olshausen and Field 1996).

In summary, the first- and second-order image statistics characterize an animal pattern in a way that captures how animal visual systems process spatial information, and they are very useful for predicting the appearance of visual textures and patterns—at least to humans (Adelson and Bergen 1988; Meese 2002). From these statistics, one can choose parameters to map as axes in pattern space (fig. 1): for example, a small number of spatial frequency bands (as in granularity analysis), the dominant spatial frequency, overall image power, or variance in power across frequency bands (Stoddard and Stevens 2010). A pattern space specified by these parameters is relevant to many diverse questions in behavioral ecology, as evidenced by the studies discussed above. However, second-order image statistics (which describe the power spectrum and do not include phase information) on their own do not give a complete characterization of an image and are indeed largely irrelevant to the recognition of objects within images, which depends on specialized feature detectors (fig. 4E; see "Step 4: Detecting Local Features and Objects" and "Step 5: Detecting Higher-Level Objects"). In practice, in almost all cases, phase information is crucial for edge and feature detection. For example, abolishing power/amplitude information has little effect on the appearance of edges (fig. 4D).

Step 4: Detecting Local Features and Objects

To continue building the pattern space, the next step is to consider local features relevant to overall shape, especially edges. Moving from image statistics to local feature detection comes with the twin caveats that we know less about how animals detect these features, and they are less amenable to parameterization than image statistics (fig. 1). None-theless, edge detection algorithms (fig. 4E) are easily applied to animal coloration patterns, and it may also be possible to identify more complex features such as peacock eyespots, jaguar rosettes, bird egg markings, and disruptive camouflage patterns, which rely on pictorial edge or depth cues for their effect (Cott 1940; Osorio and Srinivasan 1991; Cuthill et al. 2005). Edges and local features are primarily important because they contribute to (or create) object boundaries. The real world is made of objects, that is, discrete regions of 3-D space, which are usually opaque, bounded by convex surfaces, and composed of one type of material. Images are generated by the interaction of objects with illumination. Even when the objects are physically alike, such as leaves on a tree, lighting effects allow us to see them. Objects are biologically relevant: animals eat, court, hide under, and navigate by them. It is no surprise that animals are concerned with objects, but it is perhaps less obvious that the visual task of isolating objects as discrete physical entities in natural images—figure-ground segregation—is computationally difficult. For example, shadows can be confused with object boundaries, and objects often occlude one another. Coloration patterns are well known to exploit mechanisms concerned with object perception. For example, disruptive camouflage acts by impeding figure-ground segregation, and, conversely, some signaling patterns accentuate the outline of the body (Cott 1940; Osorio and Srinivasan 1991; Troscianko et al. 2017). At the same time, we know that edges and other local features such as corners, lines, and axes of symmetry relevant to object recognition cannot be defined in terms of second-order image statistics, so how might their contribution to the appearance of a pattern be specified?

Visual edges cannot be objectively defined: they are simply locations within the image that are likely to correspond to the physical boundaries of objects. Operationally, this means finding discontinuities in brightness (Marr 1982; Canny 1986). In machine vision, image processing typically starts with linear filtering (i.e., convolution) by circularly symmetrical filters or oriented wavelet functions (see "Step 3: Calculating Power Spectra and Image Statistics"; box 1). The filter outputs are then used for identification of local features, especially edges but also corners (or points of occlusion) and lines, which are integrated across the image to define objects. Various methods of edge detection are used in computer vision (Marr and Hildreth 1980; Canny 1986; Morrone and Burr 1988), but all are nonlinear in that the output is not a continuous function of the input; instead, they classify a given location as edge or nonedge following some nonlinear operation (e.g., multiplication, rectification, or thresholding). In general, whereas measurement of second-order image statistics depends on linear operations on the image (i.e., convolution), feature detectors are generally nonlinear and sensitive to spatial phase

information (box 1). Overall, metrics derived from edge detection algorithms could form additional parameters, or axes, of the pattern space (fig. 1).

Edge detection algorithms have been used in numerous studies of animal patterning, especially in the context of camouflage. To investigate camouflage in moths, Stevens and Cuthill (2006) applied a modified version of Marr and Hildreth's (1980) algorithm to photographs of experimental moth stimuli. Their analysis supported the hypothesis that disruptive coloration works because false edges are detected at the periphery, obscuring the animal's true outline. Similarly, in a study of egg camouflage, Lovell and colleagues (2013) used an edge detector to derive a number of camouflage metrics related to background matching and disruptive coloration. They compared edge densities of the egg and the substrate, predicting that the egg would be well camouflaged (via background matching) when these densities were similar. A number of subsequent papers, including work on moths (Kang et al. 2015) and shorebird eggs (Stoddard et al. 2016), applied similar methods.

There is increasing interest in understanding which edge and feature detectors work the best for identifying cryptic prey. In a comprehensive study, Troscianko and colleagues (2017) evaluated seven classes of models used to analyze animal camouflage, including spatial frequency power (image statistics), edge detectors, and a new metric, GabRat, which uses a Gabor filter to calculate the average ratio of false to salient edges in a prey's outline. Briefly, a series of Gabor filters was applied to an image of a prey item. The GabRat measure gave the (relative) strength of false or disruptive edges in the pattern by taking the ratio of the outputs of the filters (i.e., spatial power) oriented perpendicular to those oriented parallel to the animal's real body contour. Ultimately, the GabRat metric best predicted human detection of computer-generated prey on a monitor, but, nonetheless, it accounted for only 8% of the variance in capture times by human participants. The low predictive ability of the model (at least for human viewers) suggests that we are far from understanding the underlying perceptual mechanisms involved in detecting objects in complex natural scenes. Nevertheless, systematic tests of the effects of different visual features (simple image statistics, edges, objects) on performance in visual tasks provide the ground-truthing necessary to ensure that pattern spaces are capturing meaningful information. How would other primates, or cats, or birds, perform in a similar task? Comparative pattern perception—across diverse animal taxa and in ecologically relevant contexts—is an exciting prospect for future research.

Step 5: Detecting Higher-Level Objects

The final step in fleshing out the pattern space may be to consider higher-level structures, such as objects and faces.

So far we have looked at pattern metrics that can reasonably be related to animal visual mechanisms. For higher-level patterns, such as monkey faces (Allen and Higham 2015), the markings on birds eggs (Stoddard et al. 2014), or cuttlefish and fish coloration patterns (Ramachandaran et al. 1996; Kelman et al. 2008), quantification can be challenging because we know much less about the sensory and cognitive processes involved in higher-level spatial vision. However, it is often possible to take multiple measures of a set of patterns, from which a low-dimensional representation can be derived by a suitable method of factor analysis. One can then test the visual salience of these representations, using the pattern space as a starting point for depicting higher-level, feature-rich patterns.

What methods exist for detecting features more complex than edges? Of particular note are methods used in machine vision, such as the scale-invariant feature transform (SIFT; Lowe 2004) and hierarchical model and X (HMAX; Riesenhuber and Poggio 1999; Serre et al. 2007), which resemble biological systems in that they start by convolving the image with filters that correspond to the circularly symmetrical DoG receptive fields of retinal neurons (SIFT) or Gabor function receptive fields of cortical simple cells (HMAX). SIFT computes the differences between Gaussian convolutions, at different spatial scales, and identifies a large number of statistical regularities in these DoG signals that potentially correspond to features of interest such as edges, corners, patches, and blobs (fig. 4F). The local image parameters derived from these convolutions are then used to identify objects in novel images; the SIFT algorithm is invariant to size and orientation, so a familiar object can be identified in a new scene even it appears to be, for example, smaller, rotated, or tilted. We note that individual SIFT features should be considered low- or mid-level "local features" (see "Step 4: Detecting Local Features and Objects"); we include SIFT in this section about "higher-level objects" because although SIFT finds local features, the algorithm is considerably more complex than conventional edge detection algorithms and may well correspond to higher-level neural processes in biological visual systems (Lowe 2000), though more research is needed. Overall, SIFT is highly effective in detecting high-contrast local features such as speckles on a bird's egg—or, indeed, the unique characteristics of human handwriting (Zhang et al. 2009)—that are otherwise difficult to define (fig. 4F). HMAX applies a series of Gabor filters to an image, the inputs of which are combined and represented as features in a feedforward, hierarchical manner (Riesenhuber and Poggio 1999). A detailed comparison of SIFT and HMAX is provided by Moreno et al. (2007).

Successors to SIFT and HMAX feature detectors in machine vision include hierarchical convolutional neural networks (HCNNs), which are based on neural network architectures that are compared to those of the cerebral

cortex (Yamins and DiCarlo 2016*a*). Convolutional neural networks (CNNs) and deep learning have become extremely popular in computer vision (LeCun et al. 2015). A CNN is made up of layers, each of which learns to detect different features. Filters are applied at different spatial scales, and the output of each convolved image becomes the input to the next layer. CNNs are then trained, usually on a very large set of images, so that the number and weightings of the different filters and layers can be optimized for the recognition task at hand. As phenotypic data sets become larger and more accessible (Beaman and Cellinese 2012), it may be possible to harness the power of CNNs to test hypotheses about animal visual signals; however, a real challenge will be to make sense of the CNN weightings, which can be difficult to interpret. In other words, a neural network is a black box in that it might identify useful features or objects, but deciphering how it does this—its underlying structure—is often unwieldy. Nevertheless, though the analogy with biological sensory processing invoked by machine learning techniques here may seem superficial, such methods can identify image parameters whose visual salience can be tested experimentally. Overall, in the pattern space, metrics derived from SIFT, HMAX, HCNNs, or related methods could form additional axes (fig. 1). This would be very useful because studies on object recognition in nonhuman animals remain rare. Most tests of object recognition are concerned with whether an animal can discriminate members of some behaviorally relevant (or sometimes irrelevant) data set using poorly parameterized test stimuli, which leaves open the question of whether discrimination is based on the low-level parameters such as image statistics or higher-level features.

In this context, visual patterns that have evolved under selection for deception, such as camouflage or mimicry, are of particular interest; visual discrimination is tested to its limits. In some cases, there may even be a coevolutionary arms race between the signaler and the signal receiver. For example, in a recent study of bird egg patterns, Stoddard and colleagues (2014) developed NaturePatternMatch, an algorithm that uses SIFT (fig. 4*F*), to test the hypothesis that host species have evolved recognizable egg pattern signatures in response to egg mimicry by common cuckoos (*Cuculus canorus*). They calculated the likelihood of a given host egg being correctly matched to its own clutch and used this as an index of pattern recognizability. Overall, host species subjected to the most intense cuckoo mimicry appear to have evolved the most recognizable patterns on their own eggs in response. Unlike most SIFT-based algorithms, NaturePatternMatch does not search for exact objects (or pattern features) on other eggs but rather compares two pattern textures.

Face recognition is a special case of pattern recognition. In a study of the diverse faces of guenon monkeys, Allen and colleagues (2014) used eigenfaces to calculate the visual distinctiveness of the faces of 22 guenon species. The eigenface method itself is not biologically motivated (Hansen and Atkinson 2010), but the features extracted are thought to be similar to those used by the primate brain in face processing (Allen et al. 2014). Using this computational approach to face processing, the authors demonstrated that guenon species in sympatry have evolved faces that are more visually distinctive, perhaps to aid in species recognition.

For further reading on recognition algorithms, we direct readers to Bruce et al. (2003) for a historical perspective, to DiCarlo et al. (2012) for a neuroscience perspective, to Krüger et al. (2013) and Medathati et al. (2016) for a computer vision perspective, and to Soto and Wasserman (2014) for a comparative perspective.

Additional Methods

So far, our approach to building a pattern space has involved mainstream methods of image processing: linear filtering is followed by successive stages of feature detection, typically leading to object recognition. We have shown how this approach is relevant to biological systems and is likely to yield insights into how animals perceive visual patterns. There are, however, alternative ways in which to characterize animal coloration patterns. Here we outline several such methods.

Regular, repeating textures may be good candidates for wavelet- and fractal-based analyses, but what about patterns with a relatively small number of well-defined colored regions, such as the different-colored patches that make up a parrot's phenotype or the irregular markings on poison frogs? Endler and Mielke (2005) introduced the LSED-MRPP method, which accounts for relationships among colors in an overall color pattern. This method combines least sum of Euclidean distances (LSED) regression analyses with multi-response permutation procedures (MRPP). The method is a nonparametric multivariate statistical approach for testing the hypothesis that two color patterns (represented by their color space coordinates) are different. Building on this approach, Endler (2012) proposed using adjacency analysis to determine which color patches are located next to one another. Briefly, the analysis includes collecting color data in a grid, placing the data on a zone map, and conducting statistical analyses about relative color frequency and pattern regularity. Compared to the earlier LSED-MRPP method, the color adjacency method provides more information about pattern complexity and texture. The method has recently been applied to the study of color pattern variation in poison frogs (Rojas and Endler 2013; Rojas et al. 2014). Like the LSED-MRPP, the adjacency method is not inspired by biological vision but could nonetheless be a powerful tool for capturing texture variation. A conceptually similar tool, the distance

transform, has recently been proposed to determine the similarity of a pair of patterns (Taylor et al. 2013). Two binary images are compared, pixel by pixel, using a distance transform. The authors tested the method on hoverflies, wasps, and butterflies, revealing that the metric could detect subtle variation in two patterns. It would be interesting to compare the predictions about visual appearance (e.g., discriminability) made by these methods to those (discussed previously) based on the outputs of linear filters or the spatial frequency power spectrum.

In a departure from existing methods, Allen and colleagues (2011) produced synthetic reference patterns to investigate the ecology and evolution of felid coat patterns. Using developmentally inspired reaction-diffusion equations, they produced a range of natural-looking patterns including cheetah spots and leopard rosettes. Human observers then selected the synthetic patterns most closely resembling the biological patterns. The synthetic patterns were used to test the hypothesis that coat pattern generation mechanisms are related to ecology. A major advantage of this approach is that it can deal with complex patterns and is linked to developmental processes, such as coat pattern formation. A disadvantage is that patterns are quantified by humans (as opposed to an objective method or one informed by the relevant signal receiver), but given how many aspects of spatial vision appear to be conserved across vertebrates (Stevens 2007; Cronin et al. 2014), the cost is likely minor (but see Kelley and Kelley 2014). Overall, this method has the important merit of emulating developmental mechanisms that might generate a range of patterns under natural selection. It also highlights the difficulty in relating even such a simple developmental morphospace to visual appearance, since human observers were required to decide how well the synthetic patterns matched natural ones.

A little-explored tool for the analysis of pattern textures is the elliptical Fourier analysis (EFA). EFA is a method that might be useful for characterizing the shapes of irregular animal color signals, such as the moustaches of guenon monkeys (Allen and Higham 2013). In contrast to traditional landmark-based approaches, where landmarks in the image must be selected, EFA works by representing a shape in terms of harmonic ellipses, with the first, second, and third harmonics capturing aspects of circularity, ellipticity, and triangularity, respectively.

Finally, some of the most exciting breakthroughs on animal pattern recognition have come from the growing field of animal biometrics, where the goal is often to identify individual animals for conservation, management, and monitoring. Computer vision techniques developed for this purpose—for example, to recognize individual zebras (Lahiri et al. 2011; Parham et al. 2017), penguins (Sherley et al. 2010), and manta rays (Town et al. 2013) in the wild—could be applied to studies of animal coloration patterns in the con-

text of signaling and camouflage. For example, it is well known in biometrics that there is a trade-off between optimizing methods for species versus individual recognition. What might this trade-off reveal about how signals evolve in nature under selection for recognizability? Going under the hood to explore why biometrics methods work—and when they fail—could be very profitable. For an excellent review, see Kühl and Burghardt (2013).

Which Parameters to Choose?

Ultimately, the choice of which parameters to include in a pattern space must be made. Should all aspects of an animal pattern be measured, or should a priori decisions be made about which aspects of pattern might be important? Given the data, which aspects of pattern can be easily measured, and what is known about the spatial vision of the signal receiver—if anything? These are just some of the questions that might arise. In the end, some approximations and compromises will likely be made when constructing a pattern space. Perhaps the most important consideration will be whether to include low-level image statistics only or to add local edges, features, and higher-level structures. Evidence from human psychophysics suggests that differences in spatial frequency power spectra often predict the discriminability of two patterns well (Bergen and Adelson 1988). However, the question of how well image statistics predict the detectability of a pattern as camouflage or communication signals remains largely open (Zylinski et al. 2011; Troscianko et al. 2017). Given this, in the absence of suitable evidence to the contrary, there are good reasons for starting the analysis by evaluating power spectra and other low-level image statistics. The benefits are similar to those for using color spaces: the measures are easily quantified, they tend to be quite simple, and we know that humans and other animals are sensitive to them (see "Step 3: Calculating Power Spectra and Image Statistics"). However, for many applications, quantifying local features such as edges and more complex structures such as objects will be vital to building a more complete and meaningful pattern space.

Once patterns are represented in pattern space, a number of additional metrics can be calculated, as in color space. For example, the Euclidean distance between two points in n-dimensional space (Spottiswoode and Stevens 2010) or the volume occupied by a set of points in pattern space may be useful starting points. Ultimately, much more psychophysical work must be done to determine whether such measures are biologically meaningful. In color space, the distance between two stimuli is, in theory, correlated with the perceived difference between them. Will some pattern spaces have this property? The fact that differences in second-order image statistics (power spectra) predict texture discrimina-

tion well (Bergen and Adelson 1988) suggests that this might be true, at least in some cases.

New Frontiers in Pattern Space

A pattern space provides a relatively simple way in which to map, compare, and analyze patterns, informed by basic principles of animal spatial vision. Moving forward, what new insights might be revealed by pattern space analyses?

Information and Signal Design

If animal patterns (say, of poison dart frogs) are subject to selection by multiple signal receivers—a conspecific frog, a predatory bird—it will be useful to compare those patterns across different pattern spaces. Where do they fall in a frog pattern space, with the parameters (axes) set by frog vision? Where do they fall in a bird pattern space? The equivalent analysis in color spaces is common. For example, to examine the dual influence of sexual selection and natural selection on color in *Dendrobates pumillio* dart frogs, Siddiqi (2004) analyzed frog colors by using models of frog color vision and bird color vision, respectively. They showed that each color morph possessed at least one color that would appear as a highly conspicuous color to frogs and to birds, suggesting that frog colors are effective signals to conspecifics and predators alike. A pattern space would permit the same analysis, accounting for variation in spatial vision—acuity, contrast sensitivity, or even a presumed sensitivity to certain edges or features—across different animal viewers. Analyses such as these may ultimately allow us to determine the extent to which animal patterns may have coevolved with spatial vision across taxa (sensu Lind et al. 2017) and whether ecological and sensory constraints on spatial vision have influenced signal design (Rosenthal 2007). In swordtails, for example, whether males of a particular species possess vertical stripes used in mating displays is apparently related to the critical flicker fusion frequency of females (Ryan and Rosenthal 2001; Rosenthal 2007).

Behavior

In studies of animal coloration, behavior plays a critical role: it provides the ultimate test of color vision models (Kemp et al. 2015). We must hold models of pattern vision to the same high standard. A pattern space gives a framework for analyzing patterns at multiple levels of complexity, from simple image statistics to low-level edges to higher-level structures and objects. In this sense, a pattern space helps to clarify predictions about what information might be important to a signal receiver in the context of mate choice, predation, or recognition. Which pattern information matters to behaving animals? An explosion of empiri-

cal work on camouflage (Stevens and Merilaita 2009; Stevens 2016; Hanlon and Messenger 2018) in recent years has contributed greatly to our understanding of how animals interpret spatial patterns, with no signs of slowing down (Troscianko et al. 2017; Pike 2018). Careful experiments designed to test how animals prioritize low- and high-level pattern information when making behavioral decisions will be very valuable. Eventually, we can start to use behavioral information to validate pattern spaces, which can in turn be used to generate better predictions about how animals act in the natural world.

Evolution

A pattern space can be useful for understanding the limits of animal coloration patterns on a broad scale. By analogy, efforts to quantify avian plumage coloration in the avian tetrahedral color space (Vorobyev et al. 2001; Stoddard and Prum 2011) have revealed that plumage colors have evolved to occupy some but not all of the theoretically visible color space of birds. What regions of the pattern morphospace are occupied by extant animal patterns (fig. 1)? What does the occupied morphospace include, and what does it exclude? Are gaps in the morphospace a consequence of physiological constraint or of natural and sexual selection? We are now well equipped to start addressing these questions in a way that accounts, on some level, for the spatial perception of signal receivers. In fact, because many aspects of spatial vision appear to be highly conserved across animals (Cronin et al. 2014), comparisons among very different taxonomic groups may be possible.

Moreover, if a phylogeny for the taxonomic group in question is available, these quantitative traits can be mapped onto a tree and analyzed in a comparative framework, providing powerful insights into the evolution of animal patterns (Allen et al. 2011, 2013; Davis Rabosky et al. 2016). Several recent studies on camouflage and mimicry have investigated the evolution of animal coloration patterns by using quantitative approaches. In the context of camouflage, Allen et al. (2011), as described above, used a morphospace informed by developmental mechanisms to investigate the evolution of felid coat patterns. Using a recent phylogeny of cats, the researchers used comparative methods to examine evolutionary changes in pattern over time—and in the context of habitat, arboreality, and nocturnality. Overall, felid coat patterns were highly labile (lacking a strong phylogenetic signal) and were correlated with habitat, presumably for background matching. Plain and patterned cats tended to live in homogenous and complex natural environments, respectively. A similar approach was used to evaluate the evolution of diverse coloration patterns on snakes (Allen et al. 2013)—and on humans (Talas et al. 2017), in the context of military camouflage. Although the phylogenetic con-

siderations were different, a quantitative analysis of military uniforms showed how historical events and political ideologies shaped the cultural evolution of camouflaged apparel (Talas et al. 2017). In this study, the authors represented camouflaged patterns in a pattern space defined by principal components derived from texture analysis using a bank of Log-Gabor filters, which resembled the responses of visual neurons in the cortex.

Finally, to test a classic hypothesis about mimicry—that shifts to mimetic coloration in nonpoisonous snakes are related to co-occurrence with poisonous coral snakes—Davis Rabosky et al. (2016) synthesized biogeographic, phylogenetic, ecological, and phenotypic data for more than 1,000 New World snake species. Their analysis provided strong support for the idea that coral snake mimicry followed the evolution of coral snakes and also showed that mimicry, far from being a stable endpoint, is frequently gained and lost in snakes. To simplify analyses, the color patterns of snakes were classified based on categorical codes. In the future, it would be exciting to map snake patterns in an avian color and pattern space—using some of the quantitative approaches championed in this synthesis—to estimate the extent to which snakes might appear mimetic or conspicuous to avian predators.

Universal Principles

Have natural and sexual selection influenced animal patterns in similar ways across very diverse taxa? If so, what general principles apply? A pattern space could help reveal these similarities. In fact, biological camouflage offers good evidence for the limited dimensionality of natural visual textures. Cryptic patterns are well matched to their backgrounds, even though they are probably generated by a comparatively small number of developmental mechanisms (say, <10; Carroll et al. 2013). Similarly, Kelman et al. (2006) found that to conceal itself on a wide range of seafloor backgrounds, a flatfish—the plaice *Pleuronectes platessa*—varies only two independent components in its body pattern: high-contrast scattered spots and blurry bars (see also Ramachandran et al. 1996). This "basis set" of two patterns that has evolved to generate camouflage is probably not a priori obvious from principles of image coding, highlighting the need for an empirical understanding of natural images and animal vision. The vision scientist Bela Julesz (1981, 1984), who was fascinated by camouflage, proposed that humans do indeed have a small number of texture measures (or channels), which he compared to the three color cones of trichromatic color vision. These channels could be identified by their ideal stimuli, or textons. Although texton theory has limited support as an account of human vision (Bergen and Adelson 1988; Malik and Perona 1990), animals such as flatfish and cuttlefish (Hanlon 2007; Hanlon and Messenger 2018) hint that

regularities of patterns in nature will reveal a small number of major modes of variation that are exploited by spatial vision and, consequently, by coloration patterns. Mapping the patterns of flatfish and cuttlefish alongside those of another group with cryptic patterns (e.g., nightjars)—all in a pattern space of a generic vertebrate predator—will allow us to determine whether animal patterns fall into predictable categories.

Conclusion

Quantifying animal coloration patterns is a challenging endeavor. For those who are interested in investigating the perception, function, and evolution of animal coloration patterns, an important question will be which approach to use. No single pattern space can completely capture the diversity of coloration patterns as they are seen by animals—our phenotypic space (fig. 1). Moreover, just as most animal color spaces fail to account for poststimulus processing (such as opponency or color categorization), a pattern space will necessarily be a simplification: the goal is a convenient, pragmatic depiction of animal coloration patterns, represented in a way that is relevant to an animal's sensory experience.

As a first step, animal coloration patterns can be analyzed in terms of image statistics, using spatial filters to measure power in different spatial frequency bands and orientations. Such methods can more or less directly yield a low-parameter, visually relevant perceptual space that can easily be compared across patterns. This approach is parsimonious and, in the absence of evidence to the contrary, should generally take priority over metrics based on higher-level features, for example, in describing the strength of a signal or the difference between a camouflage pattern and its background. As a second step, local features such as edges and higher-level features such as faces can be characterized by feature detection algorithms. This step adds new and sometimes complex parameters to the phenotypic space, but it can provide information not captured by image statistics. A critical goal for future research will be to test the extent to which this second step, the incorporation of edge- and feature-based parameters, provides essential information beyond a simple description based on image statistics. For example, for a host bird faced with recognizing and detecting an odd cuckoo egg, which is a better predictor of egg rejection: egg pattern metrics based on higher-level SIFT features or low-level power spectra? Evidence for selection acting on such characters—on camouflage patterns or parasitic eggs, for example—along with direct behavioral tests will show which pattern metrics matter to animals and in what ecological contexts (Stoddard et al. 2019).

Another question that will arise is: Which methods are the most biologically realistic? Marr's (1982) book on vision remains an excellent account of how computational prin-

ciples and methods can be related to biology. In general, there is little doubt that Marr's account of vision, based on spatial filtering followed by feature detection, is biologically relevant. Beyond this, models based on specific algorithms or computational schemes range from the precise (e.g., in optical models of the eye) to the speculative and metaphorical (e.g., comparisons of cerebral cortex to neural networks). In the end, detailed behavioral and psychophysical experiments, across many animal groups, will be needed to test which computational methods might actually resemble biological spatial vision. In this sense, we can consider computational methods (and therefore pattern spaces) as models that make predictions about spatial perception, to be verified (or not) through experiments.

Once animal coloration patterns are represented in a pattern space, many opportunities for sophisticated analysis come into focus. In this synthesis, we highlighted ways in which pattern space analyses can provide a powerful conceptual framework for investigating the production and perception of animal signals. Over the next decade, we expect that advances in neuroscience (Yamins and DiCarlo 2016*a*), combined with better, more accessible computational tools for visual ecologists (Weinstein 2017), will propel the study of animal patterns and animal spatial vision to new heights. Of course, pattern is just one aspect of a visual signal, and the tall order for future researchers will be finding ways to integrate color, pattern, and motion in a manner that relates to receiver perception (Rosenthal 2007; Cuthill et al. 2017).

Acknowledgments

APPENDIX

Glossary

Amplitude: A Fourier transform specifies the frequency, amplitude, and phase of all spatial frequency components of the image. The amplitude is the magnitude of the Fourier components; it is often represented by grayscale intensity or color in the 2-D power spectrum.

Contrast sensitivity function (CSF): A specialized form of the eye's MTF; it includes information about both the optical (the eye's MTF) and neural contributions to contrast sensitivity (Michael et al. 2011).

Difference-of-Gaussian (DoG): A linear filter with responses similar to those of the receptive fields of some visual neurons. The shape of a concentric receptive field is often modeled as a DoG function, which is constructed by taking the difference between the bell-shaped sensitivity profiles of the center region and surrounding regions as a function of distance from the center (Bruce et al. 2003). The Laplacian-of-Gaussian filter, which uses the second spatial derivative of a Gaussian to estimate the shape of the receptive field, is, for our purposes, essentially equivalent to the DoG function.

Fourier transform (FT): The classic approach for analyzing linear systems. The FT breaks down a signal such as an optical image into its constituent sine waves. In a 2-D image, light intensity varies in two dimensions (x, y). A Fourier transform converts this information to the frequency domain. The power spectrum is a representation of the magnitudes of the sinusoidal components as a function of spatial frequency; the representation does not include information about their spatial phase relations. Power spectra (or, equivalently, the first- and second-order image statistics) capture visually relevant characteristics of natural images. Linear filtering (i.e., convolution) and other transformations such as whitening (fig. 4*D*) are implemented on the FT image, and the modified image is recovered via an inverse Fourier transform.

Gabor filter: An orientation-sensitive filter that resembles the response properties of V1 simple cells (Daugman 1985; Jones and Palmer 1987; Soto and Wasserman 2011). The Gabor function (fig. 3) is obtained by multiplying a Gaussian envelope (the 2-D Gaussian curve) by a sine wave (Bruce et al. 2003). Gabor filters have the important property of representing both the spatial frequency and the spatial location of a signal with minimal joint error (Marčelja 1980) and for this reason are widely used as wavelets in image processing (Daugman 1988; Jain 1989). Multiple Gabor filters, tuned to different frequencies and orientations, can be used to represent the image in a manner that is efficient (as in the JPEG image compression system) and computationally convenient for edge detection and higher-level analyses (Lee 1996).

Image statistics: In general refers to any statistical relationship between intensity values in an image (Victor and Conte 1991). Here we are concerned only with first-order image statistics, which correspond to the mean and variance of the values of individual pixels, and second-order statistics, which specify correlations between pairs of pixels. Typically, this correlation falls with distance in the image (figs. 2, 4; van der Schaaf and van Hateren 1996). Image statistics are described by the power spectrum.

Isotropic filter: A filter that is not sensitive to the orientation of visual stimuli.

Lateral inhibition: A process by which neurons reduce the responses of their neighbors; it serves to increase con-

trast in an image, acting as a high-pass filter. High-pass filtering enhances the signal at high relative to low spatial frequencies. As a result, animals are often more sensitive to abrupt changes in image intensity, such as those present at edges, than more gradual changes in intensity, which are due to the effects of illumination gradients and shading.

Linear filtering and convolution: Transformation of an input image made by replacing the intensity value at each pixel/point by the weighted sum of the values over a region of the image. For low-pass filtering (i.e., blurring) characteristic of optical systems such as eyes, the function is more or less a circular Gaussian. Other linear filters often used to model visual mechanisms include difference-of-Gaussian (DoG) functions and Gabor functions. This transformation is known as convolution, where the image is "convolved" with the filter element (fig. 2) in the spatial domain. Convolution is most conveniently performed in the frequency domain by multiplying the Fourier transforms of the image and the filter functions (fig. 2) and (if desired) inverse transforming the output to recover the filtered image. Such operations are easily implemented in programming environments such as MATLAB. See figure 2 for more details.

Modulation transfer function (MTF): A common way of describing the eye's overall response to contrast as a function of spatial frequency. It can be calculated in a variety of ways, based on the optics of the lens or the packing of retinal ganglion cells (Caves and Johnsen 2017). An equation commonly used to estimate the MTF of an animal eye is provided by Snyder (1977) and is based on the minimum angular resolution (MAR) as a function of spatial frequency. The MTF can be compared to the contrast sensitivity function, which is a psychophysical measure of performance.

Octave: A logarithmic measure of frequency range. One octave corresponds to a factor-of-two difference (e.g., 1–2, 2–4, 4–8, or 8–16 cycles/degree). Visual cortex cells have an average bandwidth of 1.5 octaves (Bruce et al. 2003). Bandwidth refers to the range of spatial frequencies over which a sensor such as a neuron will respond, measured in octaves.

Pattern space: A low-dimensional space (fig. 1) that encodes features of natural patterns in a way that relates to the viewer's perceptual experience.

Phase: Here the relative phase between spatial frequency components is of most relevance because it can be used to identify lines and edges (Morrone and Burr 1988).

Power spectrum: The result of the Fourier transform (fig. 2), in which each Fourier power component is mapped in the Fourier space as a function of its frequency (with low-frequency components at the center of the map) and orientation in the image. The amplitude of a Fourier component is the square root of its power.

Receptive field: Traditionally defined as a region of the retina over which a cell responds to light (Hartline 1938)

but can generally refer to tuning in any sensory parameter space, such as spatial frequency.

Spatial frequency analysis: See box 1.

Visual acuity: A measure of spatial resolution, the ability to discriminate fine spatial patterns. Typically measured in cycles per degree (cpd), the number of pairs of black and white lines that fill a 1° angle of visual space.

Wavelet transform: Wavelets are often Gabor (or similar) functions, with the set of functions chosen to optimize data compression with minimal loss of image quality and information. Wavelets effectively provide a local Fourier transform of the image, which offers important advantages in image compression and machine vision, and are more plausible as a model of biological systems (Marčelja 1980; Daugman 1988).

Literature Cited

Allen, W. L., R. Baddeley, N. E. Scott-Samuel, and I. C. Cuthill. 2013. The evolution and function of pattern diversity in snakes. Behavioral Ecology 24:1237–1250.

Allen, W. L., I. C. Cuthill, N. E. Scott-Samuel, and R. Baddeley. 2011. Why the leopard got its spots: relating pattern development to ecology in felids. Proceedings of the Royal Society B 278:1373–1380.

Allen, W. L., and J. P. Higham. 2013. Analyzing visual signals as visual scenes. American Journal of Primatology 75:664–682.

———. 2015. Assessing the potential information content of multi-component visual signals: a machine learning approach. Proceedings of the Royal Society B 282:20142284.

Allen, W. L., M. Stevens, and J. P. Higham. 2014. Character displacement of *Cercopithecini* primate visual signals. Nature Communications 5:4266.

Baddeley, A. J. 1992. An error metric for binary images. Pages 59–78 *in* W. Forstner and S. Ruwiedel, ed. Robust computer vision. Wichmann, Karlsruhe.

Barbosa, A., L. M. Mäthger, K. C. Buresch, J. Kelly, C. Chubb, C. C. Chiao, and R. T. Hanlon. 2008. Cuttlefish camouflage: the effects of substrate contrast and size in evoking uniform, mottle or disruptive body patterns. Vision Research 48:1242–1253.

Barnett, J. B., and I. C. Cuthill. 2014. Distance-dependent defensive coloration. Current Biology 24:R1157–R1158.

Beaman, R., and N. Cellinese. 2012. Mass digitization of scientific collections: new opportunities to transform the use of biological specimens and underwrite biodiversity science. ZooKeys 209:7–17.

Beck, D. M., and S. Kastner. 2009. Top-down and bottom-up mechanisms in basing competition in the human brain. Vision Research 49:1154–1165.

Bergen, J. R., and E. H. Adelson. 1988. Early vision and texture perception. Nature 333:363–364.

Bruce, V., P. R. Green, and M. A. Georgeson. 2003. Visual perception: physiology, psychology, and ecology. Psychology, London.

Buschman, T. J., and S. Kastner. 2015. From behavior to neural dynamics: an integrated theory of attention. Neuron 88:127–144.

Campbell, F. W., and J. G. Robson. 1968. Application of Fourier analysis to the visibility of gratings. Journal of Physiology 197:551–566.

Canny, J. 1986. A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 8:679–698.

Carrasco, M. 2011. Visual attention: the past 25 years. Vision Research 51:1484–1525.

Carroll, S. B., J. K. Grenier, and S. D. Weatherbee. 2013. From DNA to diversity: molecular genetics and the evolution of animal design. Blackwell, Malden, MA.

Caves, E. M., N. C. Brandley, and S. Johnsen. 2018. Visual acuity and the evolution of signals. Trends in Ecology and Evolution 33:358–372.

Caves, E. M., and S. Johnsen. 2017. AcuityView: an R package for portraying the effects of visual acuity on scenes observed by an animal. Methods in Ecology and Evolution 9:793–797.

Caves, E. M., M. Stevens, and C. N. Spottiswoode. 2017. Does coevolution with a shared parasite drive hosts to partition their defences among species? Proceedings of the Royal Society B 284:20170272.

Chiao, C. C., C. Chubb, K. Buresch, L. Siemann, and R. T. Hanlon. 2009. The scaling effects of substrate texture on camouflage patterning in cuttlefish. Vision Research 49:1647–1656.

Chittka, L., and A. Brockmann. 2005. Perception space—the final frontier. PLoS Biology 3:e137.

Cloutman, L. L. 2013. Interaction between dorsal and ventral processing streams: where, when and how? Brain and Language 127:251–263.

Cott, H. B. 1940. Adaptive coloration in animals. Methuen, London.

Cronin, T. W., S. Johnsen, N. J. Marshall, and E. J. Warrant. 2014. Visual ecology. Princeton University Press, Princeton, NJ.

Cuthill, I. C., W. L. Allen, K. Arbuckle, B. Caspers, G. Chaplin, M. E. Hauber, G. E. Hill, et al. 2017. The biology of color. Science 357: eaan0221.

Cuthill, I. C., M. Stevens, J. Sheppard, T. Maddocks, C. A. Párraga, and T. S. Troscianko. 2005. Disruptive coloration and background pattern matching. Nature 434:72–74.

Darwin, C. 1888. The descent of man and selection in relation to sex. J. Murray, London.

Daugman, J. G. 1985. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. Journal of the Optical Society of America A 2:1160–1169.

———. 1988. Complete discrete 2-D Gabor transforms by neural networks for image-analysis and compression. IEEE Transactions on Acoustics Speech and Signal Processing 36:1169–1179.

Davis Rabosky, A. R., C. L. Cox, D. L. Rabosky, P. O. Title, I. A. Holmes, A. Feldman, and J. A. McGuire. 2016. Coral snakes predict the evolution of mimicry across New World snakes. Nature Communications 7:11484.

DiCarlo, J. J., D. Zoccolan, and N. C. Rust. 2012. How does the brain solve visual object recognition? Neuron 73:415–434.

Endler, J. A. 2012. A framework for analysing colour pattern geometry: adjacent colours. Biological Journal of the Linnean Society 107:233–253.

Endler, J. A., and P. W. Mielke Jr. 2005. Comparing entire colour patterns as birds see them. Biological Journal of the Linnean Society 86:405–431.

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 4:2379–2394.

Gallant, J. L., J. Braun, and D. C. Vanessen. 1993. Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual-cortex. Science 259:100–103.

Giurfa, M., and R. Menzel. 1997. Insect visual perception: complex abilities of simple nervous systems. Current Opinion in Neurobiology 7:505–513.

Godfrey, D., J. N. Lythgoe, and D. A. Rumball. 1987. Zebra stripes and tiger stripes: the spatial frequency distribution of the pattern compared to that of the background is significant in display and crypsis. Biological Journal of the Linnean Society 32:427–433.

Gómez, J., and G. Liñán Cembrano. 2017. SpotEgg: an image-processing tool for automatised analysis of colouration and spottiness. Journal of Avian Biology 48:502–512.

Graps, A. 1995. An introduction to wavelets. IEEE Computational Science and Engineering 2:50–61.

Hanlon, R. T. 2007. Cephalopod dynamic camouflage. Current Biology 17:R400–R404.

Hanlon, R. T., and J. B. Messenger. 2018. Cephalopod behaviour. Cambridge University Press, Cambridge.

Hansen, M. F., and G. A. Atkinson. 2010. Biologically inspired 3D face recognition from surface normals. Procedia Computer Science 2:26–34.

Hartline, H. K. 1938. The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. American Journal of Physiology 121:400–415.

Hubel, D. H., and T. N. Wiesel. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. Journal of Physiology 160:106–154.

———. 1968. Receptive fields and functional architecture of monkey striate cortex. Journal of Physiology 195:215–243.

Jain, A. K. 1989. Fundamentals of digital image processing. Prentice-Hall, Englewood Cliffs, NJ.

Jain, A. K., and F. Farrokhnia. 1991. Unsupervised texture segmentation using Gabor filters. Pattern Recognition 24:1167–1186.

Jarvis, E. D., J. Yu, M. V. Rivas, H. Horita, G. Feenders, O. Whitney, S. C. Jarvis, et al. 2013. Global view of the functional molecular organization of the avian cerebrum: mirror images and functional columns. Journal of Comparative Neurology 521:3614–3665.

Jones, J. P., and L. A. Palmer. 1987. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. Journal of Neurophysiology 58:1233–1258.

Julesz, B. 1981. Textons, the elements of texture perception, and their interactions. Nature 290:91–97.

———. 1984. A brief outline of the texton theory of human vision. Trends in Neurosciences 7:41–45.

Julesz, B., E. N. Gilbert, and J. D. Victor. 1978. Visual discrimination of textures with identical third-order statistics. Biological Cybernetics 31:137–140.

Kang, C., M. Stevens, J. Y. Moon, S. I. Lee, and P. G. Jablonski. 2015. Camouflage through behavior in moths: the role of background matching and disruptive coloration. Behavioral Ecology 26:45–54.

Kelber, A., M. Vorobyev, and D. Osorio. 2003. Animal colour vision—behavioural tests and physiological concepts. Biological Reviews 78:81–118.

Kelley, L. A., and J. L. Kelley. 2014. Animal visual illusion and confusion: the importance of a perceptual perspective. Behavioral Ecology 25:450–463.

Kelman, E. J., D. Osorio, and R. J. Baddeley. 2008. A review of cuttlefish camouflage and object recognition and evidence for depth perception. Journal of Experimental Biology 211:1757–1763.

Kelman, E. J., P. Tiptus, and D. Osorio. 2006. Juvenile plaice (*Pleuronectes platessa*) produce camouflage by flexibly combining two separate patterns. Journal of Experimental Biology 209:3288–3292.

Kemp, D. J., M. E. Herberstein, L. J. Fleishman, J. A. Endler, A. T. D. Bennett, A. G. Dyer, N. S. Hart, J. Marshall, and M. J. Whiting. 2015. An integrative framework for the appraisal of coloration in nature. American Naturalist 185:705–724.

Kikuchi, D. W., and D. W. Pfennig. 2010. Predator cognition permits imperfect coral snake mimicry. American Naturalist 176:830–834.

Kiltie, R. A., J. Fan, and A. F. Laine. 1995. A wavelet-based metric for visual texture discrimination with applications in evolutionary ecology. Mathematical Biosciences 126:21–39.

Kiltie, R. A., and A. F. Laine. 1992. Visual textures, machine vision and animal camouflage. Trends in Ecology and Evolution 7:163–166.

Krüger, N., P. Janssen, S. Kalkan, M. Lappe, A. Leonardis, J. Piater, A. J. Rodríguez-Sánchez, and L. Wiskott. 2013. Deep hierarchies in the primate visual cortex: what can we learn for computer vision? IEEE Transactions on Pattern Analysis and Machine Intelligence 35:1847–1871.

Kuehni, R. G. 2001. Color space and its divisions. Color Research and Application 26:209–222.

Kühl, H. S., and T. Burghardt. 2013. Animal biometrics: quantifying and detecting phenotypic appearance. Trends in Ecology and Evolution 28:432–441.

Lahiri, M., C. Tantipathananandh, R. Warungu, D. I. Rubenstein, and T. Y. Berger-Wolf. 2011. Biometric animal databases from field photographs: identification of individual zebra in the wild. Page 6 in Proceedings of the 1st ACM International Conference on Multimedia Retrieval. ACM, New York.

Land, M. F., and D. E. Nilsson. 2012. Animal eyes. Oxford University Press, Oxford.

LeCun, Y., Y. Bengio, and G. Hinton. 2015. Deep learning. Nature 521:436–444.

Lee, T. S. 1996. Image representation using 2D Gabor wavelets. IEEE Transactions on Pattern Analysis and Machine Intelligence 18:959–971.

Lee, T. S., and D. Mumford. 2003. Hierarchical Bayesian inference in the visual cortex. Journal of the Optical Society of America A 20:1434–1448.

Lind, O., M. J. Henze, A. Kelber, and D. Osorio. 2017. Coevolution of coloration and colour vision? Philosophical Transactions of the Royal Society B 372:20160338.

Lind, O., and A. Kelber. 2011. The spatial tuning of achromatic and chromatic vision in budgerigars. Journal of Vision 11:2.

Lovell, P. G., G. D. Ruxton, K. V. Langridge, and K. A. Spencer. 2013. Egg-laying substrate selection for optimal camouflage by quail. Current Biology 23:260–264.

Lowe, D. G. 2000. Towards a computational model for object recognition in IT cortex. Pages 20–31 in S.-W. Lee, H. H. Bülthoff, and T. Poggio, eds. Proceedings of the International Workshop on Biologically Motivated Computer Vision. Springer, Berlin.

———. 2004. Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision 60:91–110.

Maffei, L., and A. Fiorentini. 1973. The visual cortex as a spatial frequency analyser. Vision Research 13:1255–1267.

Maini, P. K. 2004. Using mathematical models to help understand biological pattern formation. Comptes Rendus Biologies 327:225–234.

Malik, J., and P. Perona. 1990. Preattentive texture discrimination with early vision mechanisms. Journal of the Optical Society of America A 7:923–932.

Mallarino, R., C. Henegar, M. Mirasierra, M. Manceau, C. Schradin, M. Vallejo, S. Beronja, G. S. Barsh, and H. E. Hoekstra. 2016. Developmental mechanisms of stripe patterns in rodents. Nature 539:518–523.

Maloney, L. T. 1986. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. Journal of the Optical Society of America A 3:1673–1683.

Marčelja, S. 1980. Mathematical description of the responses of simple cortical cells. Journal of the Optical Society of America 70:1297–1300.

Marr, D. 1982. Vision: a computational approach. Freeman, New York.

Marr, D., and E. Hildreth. 1980. Theory of edge detection. Proceedings of the Royal Society B 207:187–217.

Medathati, N. V. K., H. Neumann, G. S. Masson, and P. Kornprobst. 2016. Bio-inspired computer vision: towards a synergistic approach of artificial and biological vision. Computer Vision and Image Understanding 150:1–30.

Medina, L., and A. Reiner. 2000. Do birds possess homologues of mammalian primary visual, somatosensory and motor cortices? Trends in Neurosciences 23:1–12.

Meese, T. 2002. Spatial vision. Pages 171–183 in D. Roberts, ed. Signals and perception: the fundamentals of human of human sensations. Palgrave Macmillan, New York.

———. 2009. A tutorial essay on spatial filtering and spatial vision. Aston University, Birmingham, UK.

Melin, A. D., D. W. Kline, C. Hiramatsu, and T. Caro. 2016. Zebra stripes through the eyes of their predators, zebras, and humans. PLoS ONE 11:e0145679.

Michael, R., O. Guevara, M. de la Paz, J. Alvarez de Toledo, and R. I. Barraquer. 2011. Neural contrast sensitivity calculated from measured total contrast sensitivity and modulation transfer function. Acta Ophthalmologica 89:278–283.

Mills, M. G., and L. B. Patterson. 2009. Not just black and white: pigment pattern development and evolution in vertebrates. Seminars in Cell and Developmental Biology 20:72–81.

Moreno, P., M. J. Marín-Jiménez, A. Bernardino, J. Santos-Victor, and N. P. de la Blanca. 2007. A comparative study of local descriptors for object category recognition: SIFT vs. HMAX. Pages 515–522 in Pattern recognition and image analysis. Lecture Notes in Computer Science 4477. Springer, Berlin.

Morrone, M. C., and D. C. Burr. 1988. Feature detection in human vision: a phase-dependent energy model. Proceedings of the Royal Society B 235:221–245.

Nascimento, S. M., F. P. Ferreira, and D. H. Foster. 2002. Statistics of spatial cone-excitation ratios in natural scenes. Journal of the Optical Society of America A 19:1484–1490.

Nokelainen, O., R. H. Hegna, J. H. Reudler, C. Lindstedt, and J. Mappes. 2011. Trade-off between warning signal efficacy and mating success in the wood tiger moth. Proceedings of the Royal Society B 279:257–265.

Olshausen, B. A., and D. J. Field. 1996. Natural image statistics and efficient coding. Network-Computation in Neural Systems 7:333–339.

Osorio, D., and I. C. Cuthill. 2013. Camouflage and perceptual organization in the animal kingdom. Pages 843–862 in J. Wagemans, ed. The Oxford handbook of perceptual organization. Oxford University Press, Oxford.

Osorio, D., and M. V. Srinivasan. 1991. Camouflage by edge enhancement in animal coloration patterns and its implications for visual mechanisms. Proceedings of the Royal Society B 244:81–85.

Parham, J., J. Crall, C. Stewart, T. Berger-Wolf, and D. Rubenstein. 2017. Animal population censusing at scale with citizen science

and photographic identification. *In* AAAI Spring Symposium Series 2017. https://aaai.org/ocs/index.php/SSS/SSS17/paper/view/15245/14512.

Pérez-Rodríguez, L., R. Jovani, and F. Mougeot. 2013. Fractal geometry of a complex plumage trait reveals bird's quality. Proceedings of the Royal Society B 280:20122783.

Pérez-Rodríguez, L., R. Jovani, and M. Stevens. 2017. Shape matters: animal colour patterns as signals of individual quality. Proceedings of the Royal Society B 284:20162446.

Perrett, D. I., J. K. Hietanen, M. W. Oram, P. J. Benson, and E. T. Rolls. 1992. Organization and functions of cells responsive to faces in the temporal cortex. Philosophical Transactions of the Royal Society B 335:23–30.

Pike, T. W. 2018. Quantifying camouflage and conspicuousness using visual salience. Methods in Ecology and Evolution 9:1883–1895.

Portilla, J., and E. P. Simoncelli. 2000. A parametric texture model based on joint statistics of complex wavelet coefficients. International Journal of Computer Vision 40:49–71.

Qadri, M., and R. G. Cook. 2015. Experimental divergences in the visual cognition of birds and mammals. Comparative Cognition and Behavior Reviews 10:73–105.

Quiroga, R. Q., L. Reddy, G. Kreiman, C. Koch, and I. Fried. 2005. Invariant visual representation by single neurons in the human brain. Nature 435:1102–1107.

Ramachandran, V. S., C. W. Tyler, R. L. Gregory, D. Rogers-Ramachandran, S. Duensing, C. Pillsbury, and C. Ramachandran. 1996. Rapid adaptive camouflage in tropical flounders. Nature 379:815–818.

Renoult, J. P., A. Kelber, and H. M. Schaefer. 2017. Colour spaces in ecology and evolutionary biology. Biological Reviews 92:292–315.

Reymond, L. 1985. Spatial visual acuity of the eagle *Aquila audax*: a behavioural, optical and anatomical investigation. Vision Research 25:1477–1491.

Riesenhuber, M., and T. Poggio. 1999. Hierarchical models of object recognition in cortex. Nature Neuroscience 2:1019–1025.

Rojas, B., J. Devillechabrolle, and J. A. Endler. 2014. Paradox lost: variable colour-pattern geometry is associated with differences in movement in aposematic frogs. Biology Letters 10:20140193.

Rojas, B., and J. A. Endler. 2013. Sexual dimorphism and intra-populational colour pattern variation in the aposematic frog *Dendrobates tinctorius*. Evolutionary Ecology 27:739–753.

Rosenthal, G. G. 2007. Spatiotemporal dimensions of visual signals in animal communication. Annual Review of Ecology, Evolution, and Systematics 38:155–178.

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

Ruderman, D. L., and W. Bialek. 1994. Statistics of natural images: scaling in the woods. Physical Review Letters 73:814–817.

Ryan, M. J., and G. G. Rosenthal. 2001. Variation and selection in swordtails. Pages 133–148 *in* L. Dugatkin, ed. Model systems in behavioral ecology. Princeton University Press, Princeton, NJ.

Schaefer, H. M., and G. D. Ruxton. 2015. Signal diversity, sexual selection, and speciation. Annual Review of Ecology, Evolution, and Systematics 46:573–592.

Serre, T., L. Wolf, S. Bileschi, M. Riesenhuber, and T. Poggio. 2007. Robust object recognition with cortex-like mechanisms. IEEE Transactions on Pattern Analysis and Machine Intelligence 29:411–426.

Sherley, R., T. Burghardt, P. Barham, N. Campbell, and I. Cuthill. 2010. Spotting the difference: towards fully-automated population monitoring of African penguins *Spheniscus demersus*. Endangered Species Research 11:101–111.

Siddiqi, A. 2004. Interspecific and intraspecific views of color signals in the strawberry poison frog *Dendrobates pumilio*. Journal of Experimental Biology 207:2471–2485.

Simoncelli, E. P., and B. A. Olshausen. 2001. Natural image statistics and neural representation. Annual Review of Neuroscience 24:1193–1216.

Snowden, R., R. J. Snowden, P. Thompson, and T. Troscianko. 2012. Basic vision: an introduction to visual perception. Oxford University Press, Oxford.

Snyder, A. W. 1977. Acuity of compound eyes: physical limitations and design. Journal of Comparative Physiology A 116:161–182.

Sompolinsky, H., and R. Shapley. 1997. New perspectives on the mechanisms for orientation selectivity. Current Opinion in Neurobiology 7:514–522.

Soto, F. A. 2014. Mechanisms of object recognition: what we have learned from pigeons. Frontiers in Neural Circuits 8:122.

Soto, F. A., and E. A. Wasserman. 2011. Visual object categorization in birds and primates: integrating behavioral, neurobiological, and computational evidence within a "general process" framework. Cognitive, Affective, and Behavioral Neuroscience 12:220–240.

Spottiswoode, C. N., and M. Stevens. 2010. Visual modeling shows that avian host parents use multiple visual cues in rejecting parasitic eggs. Proceedings of the National Academy of Sciences of the USA 107:8672–8676.

Srinivasan, M. V., S. B. Laughlin, and A. Dubs. 1982. Predictive coding: a fresh view of inhibition in the retina. Proceedings of the Royal Society B 216:427–459.

Stevens, M. 2007. Predator perception and the interrelation between different forms of protective coloration. Proceedings of the Royal Society B 274:1457–1464.

———. 2016. Cheats and deceits: how animals and plants exploit and mislead. Oxford University Press, Oxford.

Stevens, M., and I. C. Cuthill. 2006. Disruptive coloration, crypsis and edge detection in early visual processing. Proceedings of the Royal Society B 273:2141–2147.

Stevens, M., and S. Merilaita. 2009. Animal camouflage: current issues and new perspectives. Philosophical Transactions of the Royal Society B 364:423–427.

Stevens, M., C. A. Párraga, I. C. Cuthill, J. C. Partridge, and T. S. Troscianko. 2007. Using digital photography to study animal coloration. Biological Journal of the Linnean Society 90:211–237.

Stoddard, M. C., B. G. Hogan, M. Stevens, and C. N. Spottiswoode. 2019. Higher-level pattern features provide additional information to birds when recognizing and rejecting parasitic eggs. Philosophical Transactions of the Royal Society B:20180197 (forthcoming). doi:10.1098/rstb.2018.0197.

Stoddard, M. C., R. M. Kilner, and C. Town. 2014. Pattern recognition algorithm reveals how birds evolve individual egg pattern signatures. Nature Communications 5:4117.

Stoddard, M. C., K. Kupán, H. N. Eyster, W. Rojas-Abreu, M. Cruz-López, M. A. Serrano-Meneses, and C. Küpper. 2016. Camouflage and clutch survival in plovers and terns. Scientific Reports 6:32059.

Stoddard, M. C., and R. O. Prum. 2008. Evolution of avian plumage color in a tetrahedral color space: a phylogenetic analysis of new world buntings. American Naturalist 171:755–776.

———. 2011. How colorful are birds? evolution of the avian plumage color gamut. Behavioral Ecology 22:1042–1052.

Stoddard, M. C., and M. Stevens. 2010. Pattern mimicry of host eggs by the common cuckoo, as seen through a bird's eye. Proceedings of the Royal Society B 277:1387–1393.

Sutter, A., J. Beck, and N. Graham. 1989. Contrast and spatial variables in texture segregation: testing a simple spatial-frequency channels model. Perception and Psychophysics 46:312–332.

Talas, L., R. J. Baddeley, and I. C. Cuthill. 2017. Cultural evolution of military camouflage. Philosophical Transactions of the Royal Society B 372:20160351.

Taylor, C., F. Gilbert, and T. Reader. 2013. Distance transform: a tool for the study of animal colour patterns. Methods in Ecology and Evolution 4:771–781.

Town, C., A. Marshall, and N. Sethasathien. 2013. Manta matcher: automated photographic identification of manta rays using keypoint features. Ecology and Evolution 3:1902–1914.

Troscianko, J., J. Skelhorn, and M. Stevens. 2017. Quantifying camouflage: how to predict detectability from appearance. BMC Evolutionary Biology 17:7.

Troscianko, J., and M. Stevens. 2015. Image calibration and analysis toolbox—a free software suite for objectively measuring reflectance, colour and pattern. Methods in Ecology and Evolution 6:1320–1331.

Tuceryan, M., and A. K. Jain. 1993. Texture analysis. Pages 235–276 *in* C. H. Chen, L. F. Pau, and P. S. P. Wang, eds. Handbook or pattern recognition and computer vision. World Scientific, Singapore.

Uhlrich, D. J., E. A. Essock, and S. Lehmkuhle. 1981. Cross-species correspondence of spatial contrast sensitivity functions. Behavioral Brain Research 2:291–299.

Van der Schaaf, A., and J. H. van Hateren. 1996. Modelling the power spectra of natural images: statistics and information. Vision Research 36:2759–2770.

Victor, J. D., and M. M. Conte. 1991. Spatial-organization of nonlinear-interactions in form perception. Vision Research 31:1457–1488.

Vorobyev, M., A. Gumbert, J. Kunze, M. Giurfa, and R. Menzel. 1997. Flowers through insect eyes. Israel Journal of Plant Sciences 45:93–101.

Vorobyev, M., J. Marshall, D. Osorio, N. Hempel de Ibarra, and R. Menzel. 2001. Colourful objects through animal eyes. Color Research and Application 26:S214–S217.

Wallace, A. R. 1878. Tropical nature, and other essays. Macmillan, London.

Wandell, B. A. 1995. Foundations of vision. Sinauer, Sunderland, MA.

Wang, I. J., and H. B. Shaffer. 2008. Rapid color evolution in an aposematic species: a phylogenetic analysis of color variation in the strikingly polymorphic strawberry poison dart frog. Evolution 62:2742–2759.

Weinstein, B. G. 2017. A computer vision for animal ecology. Journal of Animal Ecology 87:533–545.

Westmoreland, D., M. Schmitz, and K. E. Burns. 2007. Egg color as an adaptation for thermoregulation. Journal of Field Ornithology 78:176–183.

Yamins, D. L., and J. J. DiCarlo. 2016a. Eight open questions in the computational modeling of higher sensory cortex. Current Opinion in Neurobiology 37:114–120.

———. 2016b. Using goal-driven deep learning models to understand sensory cortex. Nature Neuroscience 19:356–365.

Zaidi, Q., J. Victor, J. McDermott, M. Geffen, S. Bensmaia, and T. A. Cleland. 2013. Perceptual spaces: mathematical structures to neural mechanisms. Journal of Neuroscience 33:17597–17602.

Zhang, Z., L. Jin, K. Ding, and X. Gao. 2009. Character-SIFT: a novel feature for offline handwritten Chinese character recognition. Pages 763–767 *in* Proceedings of the 10th International Conference on Document Analysis and Recognition. IEEE Computer Society, Washington, DC.

Zylinski, S., M. J. How, D. Osorio, R. T. Hanlon, and N. J. Marshall. 2011. To be seen or to hide: visual characteristics of body patterns for camouflage and communication in the Australian giant cuttlefish *Sepia apama*. American Naturalist 177:681–690.

<div align="right">Associate Editor: Alex Jordan
Editor: Daniel I. Bolnick</div>