



COMMENTARIES

Measuring the structure of dynamic visual signals

RICHARD A. PETERS*†, COLIN W. G. CLIFFORD‡§ & CHRISTOPHER S. EVANS*†

*Animal Behaviour Laboratory, Macquarie University

†Department of Psychology, Macquarie University

‡Macquarie Centre for Cognitive Science, Macquarie University

§Visual Perception Unit, School of Psychology, University of Sydney

(Received 4 May 2001; initial acceptance 27 June 2001;
final acceptance 3 December 2001; MS. number: SC-1207)

The evolutionary significance of perceptual processes is well documented (Endler 1991; Guilford & Dawkins 1991; Pagel 1993; Dawkins & Guilford 1996; Endler & Basolo 1998). Signals must be designed to stimulate the sense organs of intended receivers (e.g. social companions, opponents, or potential mates) and this effect must typically be achieved without attracting the attention of 'eavesdroppers' such as parasites and predators. Signals must also be memorable (Bernard & Remington 1991; Rowe & Guilford 1996; Speed 2000), so that nuances of structure are learned quickly amongst a welter of competing stimuli impinging upon the receiver.

Analysis of structure is an essential prerequisite for any exploration of signal design. The long history of successful work on acoustic communication can be traced to the development of the sound spectrograph in the 1950s (for a review see Hopp et al. 1998). Similarly, rapid advances have been made in recent years since spectral analysis (Endler 1990) has been applied in studies of static visual signals, such as ornaments and colour patterns.

In contrast, much less is known about signals that are defined by movement, such as courtship and aggressive displays. Such dynamic visual signals are ubiquitous: they are used in contexts as diverse as opponent assessment (Ord et al. 2001), female mate choice (where they may act synergistically with morphology; Rosenthal et al. 1996), pursuit deterrence (e.g. Hasson 1991; Caro 1995), alarm signalling (Hennessy et al. 1981) and even camouflage (e.g. Fleishman 1985). Motion is also an important component of many multimodal signals (Partan & Marler 1999), often in combination with sound (Evans & Marler 1994). The design of these motor patterns is a classic problem in evolutionary biology (Darwin 1871) and

understanding them is likely to be important for identifying general principles of signal evolution, but traditional techniques have been inadequate for analysing the movement sequences of which they are composed.

The conventional approach for analysing dynamic visual signals is to define a display action pattern (DAP), which depicts changes in the position of a chosen feature over time (e.g. Carpenter et al. 1970; Jenssen 1977; Fleishman 1988b; Martins & Lamont 1998). However, to understand design fully, such traditional ethological analysis is not enough. Pioneering work by Fleishman (1988b) illustrates the importance of more detailed description of motor patterns. Fourier analysis reveals that the territorial display of *Anolis auratus* is characterized by abrupt changes, which contrast with the more continuous sinusoidal pattern of background vegetation movement. Results of this kind provide insights about signal design by defining the pattern of sensory stimulation that the receiver's nervous system must recognize and discriminate from irrelevant events.

Motion perception is functionally important because this type of sensory processing often underpins prey detection, avoidance of predators and communication with conspecifics. Visual systems must not, however, respond equally to all types of movement in the environment. In many vertebrates, a 'visual grasp reflex' ensures that relevant moving stimuli detected in the periphery are brought on to the central fovea. Such responses are typically somewhat selective, reflecting the activity of sensory units that respond maximally to certain types of motion and thereby contributing to filtering of visual input. Motion-sensitive neurons that respond specifically to velocity and angular size are taxonomically widespread; for example, they have been identified in iguanas, *Iguana iguana* (Stein & Gaither 1983), flies, *Musca domestica* (Borst & Egelhaaf 1989), frogs, *Rana pipiens* (Lettvin et al. 1959), pigeons, *Columba livia* (Maturana & Frenk

Correspondence: R. Peters, Department of Psychology, Macquarie University, Sydney, NSW 2109, Australia (email: richard@galliform.psy.mq.edu.au). C. W. G. Clifford is at the School of Psychology, University of Sydney, Sydney, NSW 2006, Australia.

1963), rabbits, *Oryctolagus cuniculus* (Barlow et al. 1964), cats, *Felis catus* (Hubel & Weisel 1959) and monkeys, *Macaca mulatta* (Schiller et al. 1976).

The type of information that animal visual systems are likely to extract from a complex stimulus has been considered in other disciplines including physics, psychology and computer vision (e.g. Smith & Snowden 1994). According to most theoretical and empirical analyses of whole-scene motion, the critical input parameters are direction and speed. Capturing this information allows all other features to be extracted with subsequent higher-order processes (e.g. Grzywacz et al. 1994). We can think of motion in the three-dimensional world as projecting on to a two-dimensional retina (Marr 1982). This projection defines a spatial distribution of velocity vectors, referred to as optic flow fields (e.g. Gibson 1950; Regan 1986). Techniques used in computer vision and in the modelling of biological visual systems include efficient algorithms for image motion computation. These provide the basis for a new approach to the study of motion in animal signals by identifying the velocity characteristics of different image sequences. In the only published application of this kind, Zeil & Zanker (1997) have described the distribution of motion in the claw wave of fiddler crabs (genus *Uca*), concluding that recognition of this movement is likely to require spatial and temporal averaging mechanisms.

Here, we describe a new technique for analysing dynamic visual signals. First, we implement an algorithm for calculating velocity estimates. Results are presented as both optic flow plots and velocity histograms ('velocity signatures'). Optic flow plots identify where in an image motion occurs, while velocity signatures represent the direction and speed of movement. Qualitative differences are immediately apparent through visual inspection of these two types of graphical output. Further analysis yields summary data that helps to quantify more subtle differences between motor patterns and identifies those aspects of each movement that are distinctive. We introduce a novel approach that involves constructing a population of biologically plausible artificial sensory units sensitive to both direction and speed; these response characteristics are based upon those described in neurophysiological studies of real motion-sensitive neurons (e.g. Stein & Gauthier 1983). The combined output of these sensory units is suitable for quantitative analysis and hypothesis testing.

Our motivation for seeking an alternative method for the study of dynamic visual signals lies with the difficulty of using existing techniques to describe a complex display such as that of the Jacky dragon, *Amphibolurus muricatus*, an agamid lizard native to the east coast of Australia. Aggressive displays are composed of a rapid series of stereotyped motor patterns including tail flicking, arm waving and push-ups (Carpenter et al. 1970). As an example application, we consider the problem of generating a signal that is sufficiently conspicuous. The probability of detection by receivers can be increased by signalling at a time of day when environmental conditions enhance perceived signal intensity (Endler 1991), when receivers' sensory systems are most sensitive (Aho

et al. 1988), or when the signals of other species are absent (Greenfield 1988). Each of these temporal strategies increases the signal-to-noise ratio (Endler 1992). Structural characteristics that contrast with the visual environment similarly increase the likelihood of detection (Endler 1991; Fleishman 1992). We compare the introductory tail-flick with the irrelevant background movement against which it must be perceived. This analysis extends earlier work (Fleishman 1988b) by directly measuring motion variables that are likely to be processed by the visual system of a receiver.

Methods

Apparatus

For all video recordings we used a fixed Canon XL1 3-CCD digital video camcorder (optical resolution 625 lines) mounted on a Manfrotto tripod. We used a shutter speed of 1/250 s, an aperture of F8 and Sony DVM60EX digital tape (550 lines recorded resolution). The distance between the camera and the subject was constant.

Video sequences

We compared estimates of velocity produced by a computational motion analysis algorithm with a range of input image sequences. This included an aggressive display by a male Jacky dragon, control (nondisplay) sequences of the same lizard, and several sequences of wind-blown vegetation typical of the field site where these animals were originally captured. The lizard video footage used was from an archival collection recorded during studies reported elsewhere (Ord et al., 2002), while video sequences of wind-blown vegetation were filmed specifically for this analysis (see below). Figure 1 shows representative frames from each sequence.

Aggressive display of a Jacky dragon. Aggressive displays by male Jacky dragons are composed of a stereotyped series of discrete motor patterns, delivered in an obligatory sequence (Fig. 1a). We analysed each of the five successive components separately: tail-flick (231 frames), backward arm-wave (10 frames), forward arm-wave (7 frames), push-up (11 frames) and body-rock (10 frames). All 231 frames of the tail-flick were used to generate optic flow plots and velocity signatures. To obtain a matched-duration sample, we used a randomly selected window of 25 consecutive frames for comparison with vegetation sequences.

Control sequences. Nondisplay movement sequences depicted a lizard climbing on to the perch (101 frames) and the same animal turning and walking across the perch (152 frames; Fig. 1b). We also used sequences of a stationary lizard on a horizontal perch (25 frames) and of an empty perch (25 frames) to test the response of the algorithm to no movement (Fig. 1c). This analysis provides an indication of the effects of noise in the recording and analysis procedure.

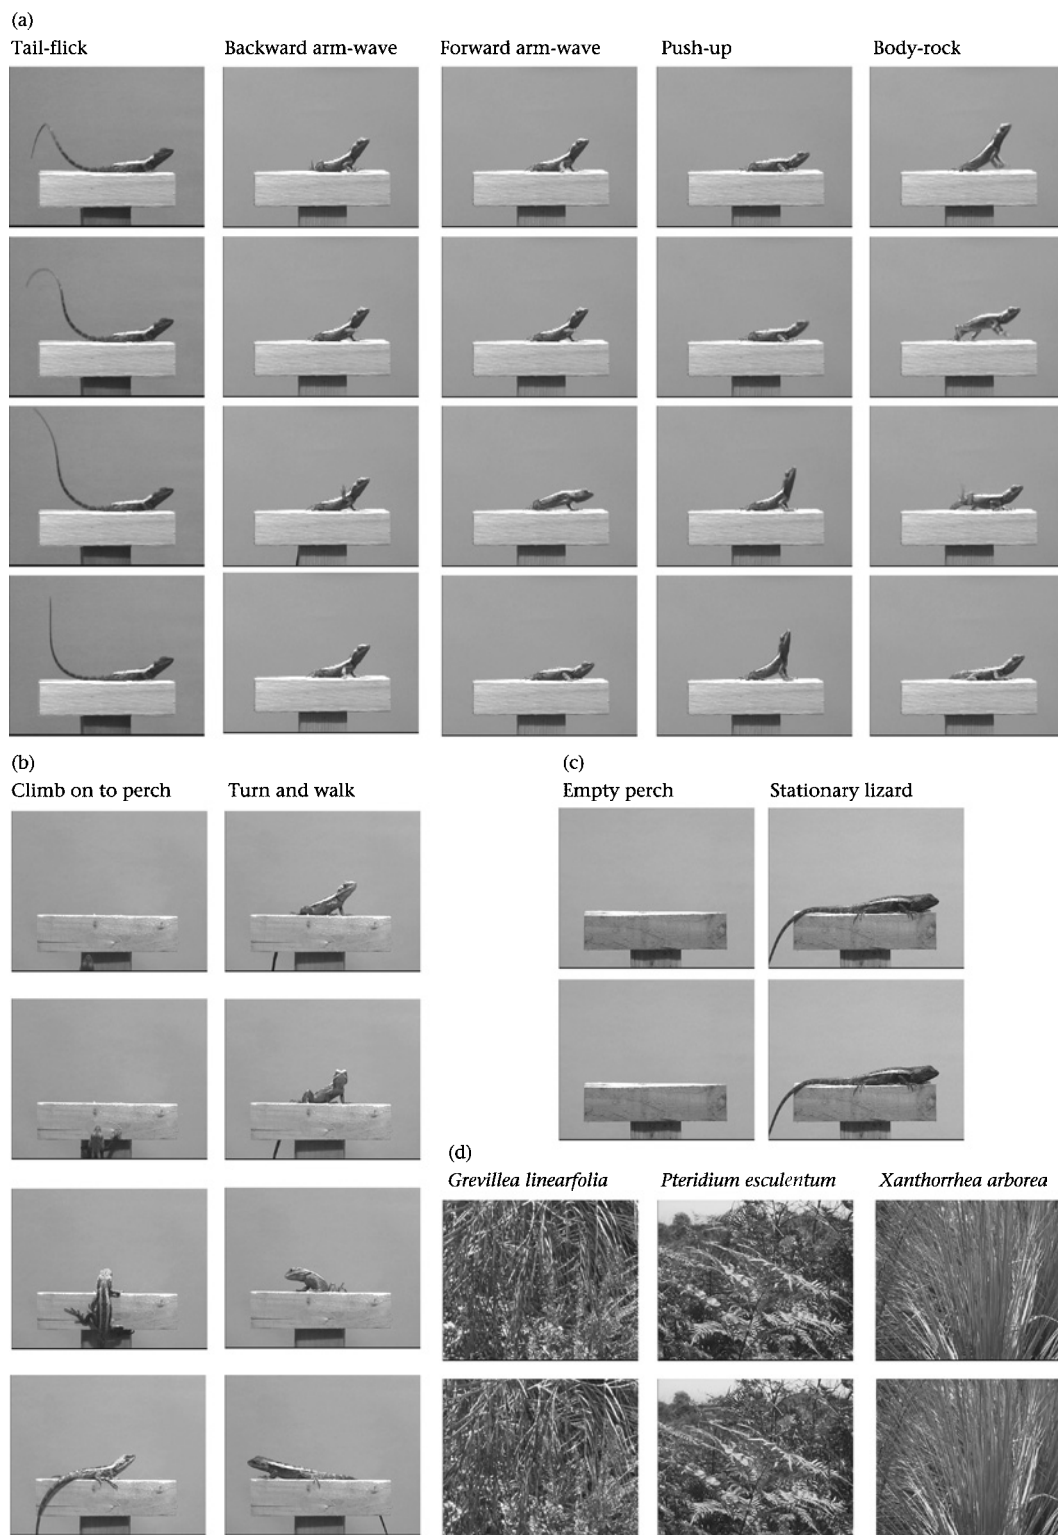


Figure 1. Representative frames from each of the image sequences used. (a) The aggressive display of *A. muricatus* separated into its five successive components: (from left) tail-flick (231 frames), backward arm-wave (10 frames), forward arm-wave (7 frames), push-up (11 frames), body-rock (10 frames). (b) Nondisplay lizard movement: a lizard climbing on to the perch (101 frames) and a lizard turning and walking across the perch (152 frames). (c) Control sequences: an empty perch and a stationary lizard on the perch (25 frames each). (d) Slow and fast wind speed sequences for three species of plants typical of the habitat of *A. muricatus* (25 frames each). Sequences (a), (b) and (c) were filmed indoors with additional lighting. Vegetation sequences (d) were filmed outdoors at various locations where *A. muricatus* is abundant.

Vegetation. A primary function of the visual system of any animal is to distinguish relevant signals from irrelevant background noise such as vegetation movement. As vegetation movement is likely to vary according to prevailing wind speed (Fleishman 1988a), we used two sequences, recorded at slow and fast wind speeds, for three plant species selected to represent the range of foliage types in the habitat of *A. muricatus* (Fig. 1d). We recorded video footage of the following plant species in the field, while simultaneously measuring wind speed with a hand-held anemometer (Dick Smith Model Q1411): *Grevillea linearifolia* (0.73 m/s, 1.75 m/s), *Pteridium esculentum* (0.73 m/s, 1.76 m/s), and *Xanthorrhoea arborea* (0.90 m/s, 1.67 m/s). All vegetation sequences were edited to a standard duration of 25 frames (Fig. 1d).

Conversion of video sequences to still images

Raw footage was transferred digitally from the camcorder to a DraCo nonlinear video editing workstation (MS MacroSystem Computer GmbH), using an IEEE 1394 'firewire' interface. We then used MovieShop v 5.2 software to convert the video sequences to individual still images. To avoid field jitter (an artefact caused by differences between the two interlaced fields of which video frames are composed), we exported the video sequences as individual fields (i.e. alternate lines of a single frame) and discarded the field containing the odd-numbered scan lines. The Nyquist theorem, which applies to any process in which a continuous function is represented by discrete samples (Hopp et al. 1998), defines the maximum frequency adequately reproduced in the final 25 fields/s sequence (12.5 Hz). It is perhaps easier to consider this constraint in terms of the corresponding time interval (80 ms). There is no evidence that complete movements as brief as this occurred in our video footage.

Graphic Converter v 3.9.1 software was used to rescale the images to restore the original frame aspect ratio (576 × 702 pixels) and convert them to 8-bit greyscale (i.e. the colour information in each frame was discarded, retaining only the intensity values recorded). To reduce computation time, we then downsampled images to 144 × 176 pixels and converted them to ASCII files for motion analysis (see Appendix).

Computational motion analysis

There are several possible methods for encoding image velocity (Barron et al. 1994) and choosing the most appropriate depends upon the task being carried out. We were interested in identifying aspects of the display that differ in structure from background vegetation movement and in analysing sequences likely to have independently moving objects. The best algorithms for tasks of separation like this are intensity-based schemes, in which motion is calculated from local changes in intensity. We have used the gradient, or differential, method because of its demonstrated reliability in accurately estimating the velocity field of a variety of real and synthetic video sequences (Barron et al. 1994). This approach is based on the assumption of locally constant image structure and

calculates the optic flow field from temporal and spatial derivatives of filtered versions of image intensity (Fleet & Langley 1995; see Appendix).

An initial stage of spatial filtering operates to compute spatial variation in intensity locally within the image. The responses of the filters centred at each spatial location are passed through recursive (or infinite impulse response) filters that store recent stimulus history implicitly in a small number of variables (Clifford & Langley 2000). These are updated through a series of equations, which keeps storage requirements down so that the algorithm can be implemented on inexpensive microcomputers (for further details about recursive filters see Fleet & Langley 1995; Clifford & Langley 2000). Calculating image velocity from independent local measurements makes no assumptions about the spatial structure of the velocity field. However, in natural image sequences, motion at neighbouring locations tends not to be independent, giving rise to a smoothly varying velocity field. We can make use of this smoothness constraint to reduce the effects of image noise on the computed velocity field by taking a weighted average over neighbouring filter outputs. The size of the kernel used in the spatial weighted average is consistent with other work (Fleet & Langley 1995) and reflects an optimal trade-off between reducing noise and maintaining high spatial resolution. Further details on the optic flow procedure used in the present paper are provided in the Appendix.

Implementation of the optic flow algorithm. The motion computation algorithm was implemented with Matlab v 5.2.1 for Macintosh and is available from R.A.P. The motion analysis algorithm uses greyscale values ranging from 0 to 255 to represent the intensity value of each of the 25 344 (144 × 176) image points (Fig. 2a). The algorithm calculates estimates of motion for each image point over a two-frame window and represents velocity as separate X and Y component vectors. Vector addition of these components determines the direction and magnitude of velocity at a given image point. These raw data can be represented in optic flow plots that maintain relative spatial location by plotting the vectors at each image point (Fig. 2b), using colour to label the direction of motion in 90° sectors. The vector appears on the graph as a line connecting the starting X and Y coordinate to the new location. Significant texture in the absence of motion is identified by the algorithm and can also be represented in the optic flow plot (depicted in black in Fig. 2b).

Velocity signatures are an alternative way to summarize motion, only the direction and speed of movement being considered (Fig. 2c). The location of each point in the X/Y plane identifies the direction of motion, while the distance from the origin defines velocity magnitude (speed). These plots provide a more concise description of the type of motion present, at the expense of relative spatial distribution. Optic flow plots and velocity signatures can be presented for individual frames to identify the motion sequence (Fig. 2), or superimposed to obtain collapsed plots summarizing movement occurring over a period of

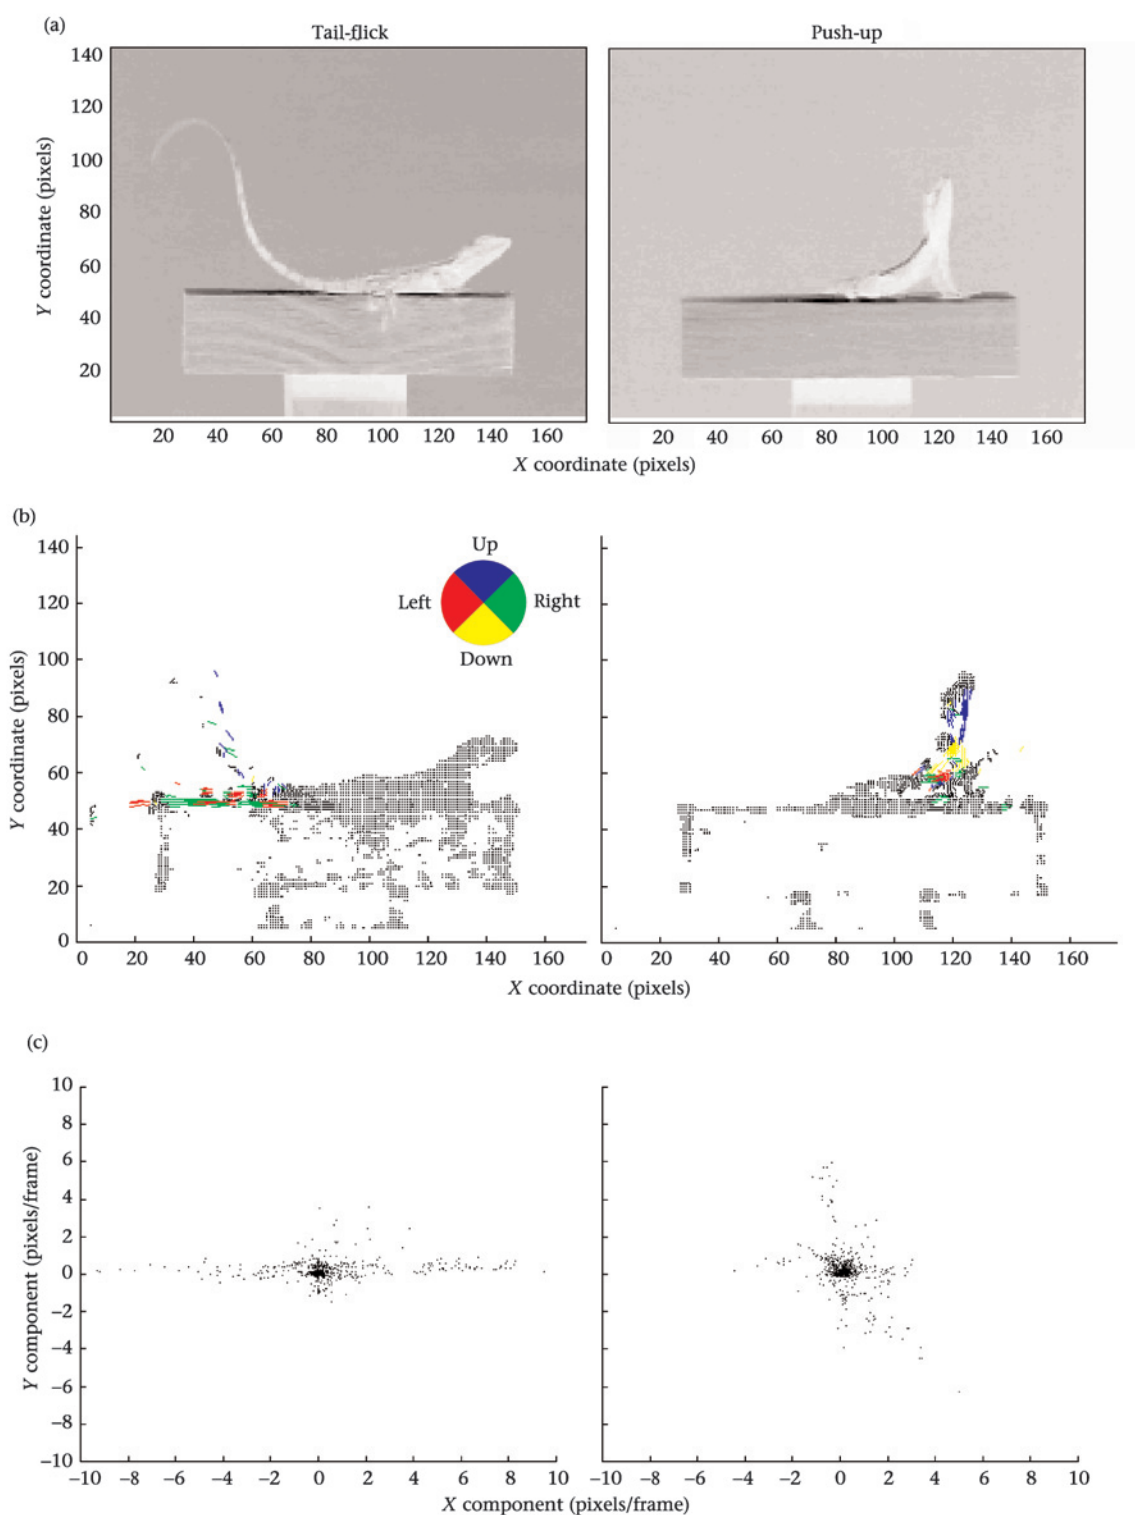


Figure 2. Computational motion analysis. Two separate motor patterns are represented: the tail-flick and push-up of the Jacky dragon. (a) Representative input frames of each type. The optic flow algorithm uses a matrix (X coordinate by Y coordinate) of greyscale values (0–255) to represent intensity. Output from the algorithm is calculated frame-by-frame, such that velocity estimates for a given frame represent movement over the preceding two frames (not shown). (b) Optic flow plots with colour coding in 90° sectors for straightforward identification of the general direction of motion. (c) A more concise description of the direction and magnitude of velocity is provided by velocity signatures. The location of data points in the two-dimensional plot defines direction (angle subtended from the origin) as well as speed (the distance from the origin).

time, which reduces graphical output. Henceforth, we report collapsed plots to describe the movement in each of the sequences analysed.

Summarizing motion information

We present three different approaches for obtaining a summary of the velocity information extracted from an image sequence: calculation of a mean vector and standard ellipse depicting the distribution of velocity estimates; derivation of a speed–time profile; and output from a population of artificial sensory units selectively tuned to particular velocities. For simplicity, we restrict this analysis to a comparison of the tail-flick and wind-blown vegetation.

Calculating a mean vector and standard ellipse.

Velocity is a vector quantity defined by the direction and speed of motion. Circular statistics are not appropriate for velocity data, as they are designed for analysis of direction only and treat vectors as having unit length; variation in the speed dimension would hence be ignored. This problem can be overcome with bivariate statistical methods (Batschelet 1981), which involve calculating two means (collectively referred to as the mean vector), two standard deviations and a correlation coefficient. Derivation of a standard ellipse then allows us to graph the orientation of movement in the X/Y plane; this summarizes mean direction, together with the X and Y components of variance.

Speed–time plots. It is likely that rapid changes in velocity enhance the conspicuousness of some display components and, conversely, that such movement characteristics are not expressed by animals seeking to remain cryptic (e.g. Fleishman 1985). The speed component of each local velocity estimate is simply the length of the vector. It is hence straightforward to calculate the average speed of movement for each frame in a sequence. A plot of these values over time will reveal episodes of sudden acceleration or deceleration.

Artificial sensory units. A final approach to summarizing the data involves processing with a population of artificial sensory units tuned to specific velocities (Fig. 3). Differences between image sequences are clearly apparent in the patterns of unit activity that they produce. This approach allows us to quantify differences in the direction and speed of movement, as well as interactions between these two parameters.

We determined the range of velocities present in Jacky dragon displays in preliminary analyses. We used these values to define the sensitivity characteristics of a population of 48 sensory units, according to a Gaussian function (see Appendix); direction tuning in biological neurons is believed to show similar response properties (Rosenberg & Ariel 1998). By design, a maximum response of 1 is returned to a pair of X and Y components that exactly match the preferred velocity; Fig. 3a shows the reduction in relative response to values around this

preference. Note that sensory unit bandwidth increases with preferred speed. Figure 3b shows the direction and speed sensitivity characteristics of the whole population of sensory units. This comprises a matrix of 16 directions (0, 22.5, 45, 67.5, 90, 112.5, 135, 157.5, 180, 202.5, 225, 247.5, 270, 292.5, 315 and 337.5°) by three speeds (3, 6 and 9 pixels/frame). Movements of less than 1 pixel/frame do not stimulate these sensory units. Pilot work showed that incorporating sensitivity to such very slow movement was not useful for our purposes; by using this minimum value we effectively included a high-pass filter in our design. Each sensory unit responds to all X and Y component pairs in an image, with the final response of each unit being the sum of unit responses to all image pairs in the frame. A 25-frame window was randomly selected for each of the image sequences and the responses to each frame were summed.

Our family of sensory units was set up based upon the observation that many animals have neurons that show directional selectivity and speed preferences during visual information processing (Stein & Gauthier 1981). By constructing sensory units tuned to specific directions at particular speeds we afford ourselves a degree of biological plausibility. However, our goal was to illustrate how estimates of velocity produced by optic flow algorithms can be reduced to a form suitable for statistical analysis, rather than to model the visual system of *A. muricatus*. Directional selectivity is a complicated process, which is influenced by a variety of factors such as the spatial and temporal context in which a stimulus appears (e.g. Worgotter & Eysel 2000). The simplified approach we have adopted is best considered as a descriptive tool.

Results

Output from the optic flow algorithm

Aggressive display of the Jacky dragon. Figure 4 shows optic flow plots and velocity signatures for the five components of the Jacky dragon aggressive display. The motion analysis algorithm captured the motor patterns characteristic of each component. The tail-flick is clearly distinct, both in spatial location and in the predominance of motion in the horizontal plane (Fig. 4a). Frame-by-frame inspection of video footage of the backward and forward arm-waves gives the impression that these movements are symmetrical, but computational motion analysis suggests that they are somewhat different. While there is limited motion in the backward arm-wave (Fig. 4b), velocity signatures reveal that the movement in the forward arm-wave also involves the upper region of the body (Fig. 4c). The vertical component of the push-up is clearly apparent in both the optic flow plot and the velocity signature (Fig. 4d). The body-rock features movement up and to the left as the lizard raises its body from the substrate, and down as it returns (Fig. 4e).

Control sequences. Figure 5 shows optic flow plots and velocity signatures for control sequences with no movement and for nondisplay lizard movements. Optic flow plots of representative frames from the empty perch and

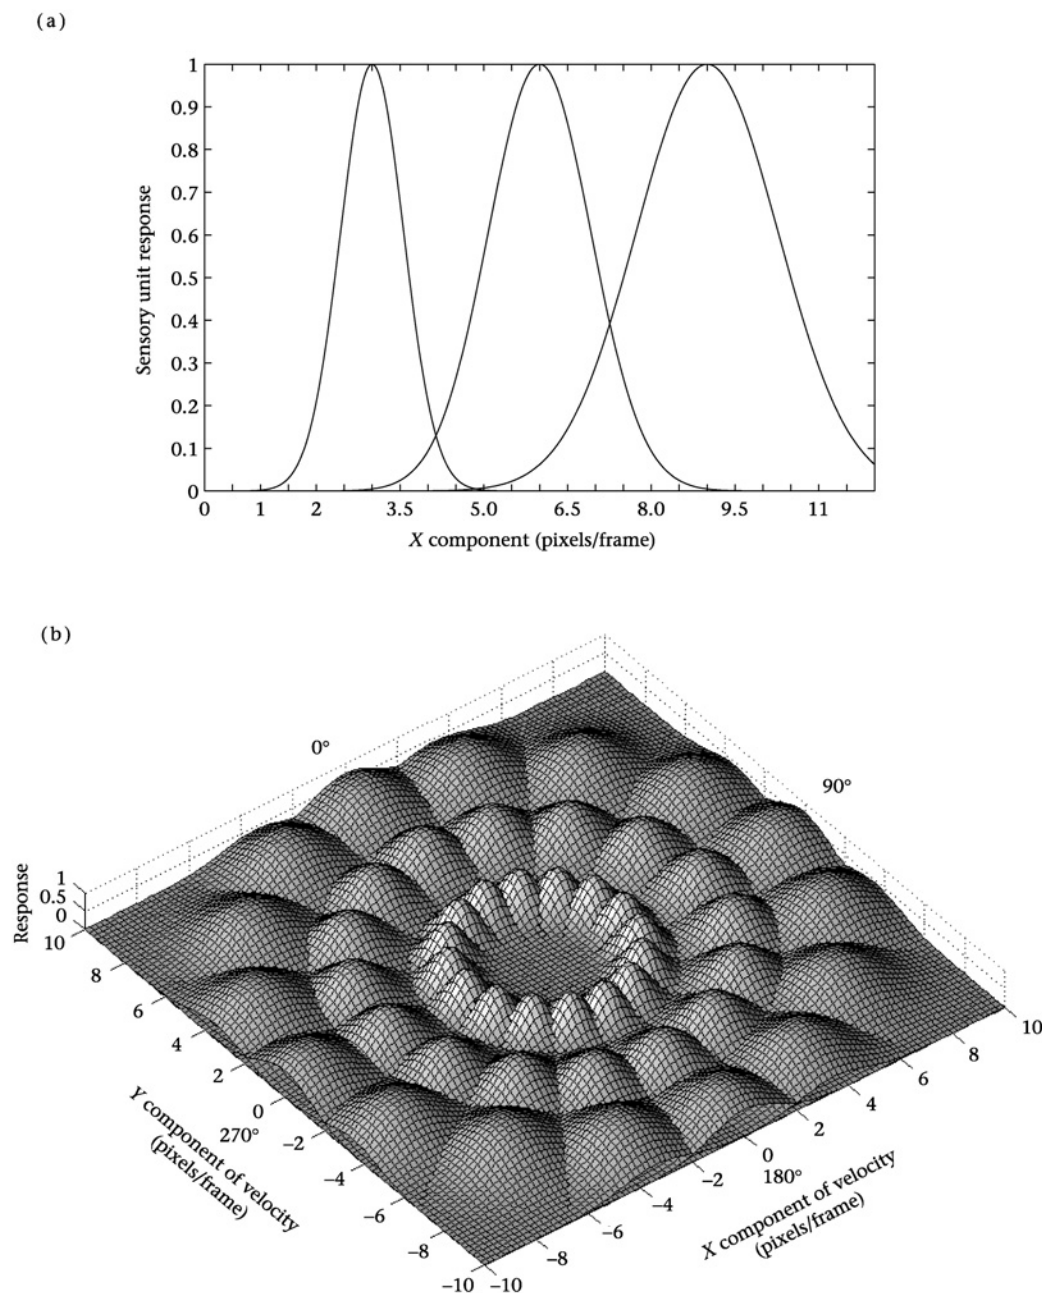


Figure 3. Description of sensory units tuned to specific angular velocities. A population of 48 units was defined which comprised 16 directions by three speeds (see text for details). (a) Plot of the response function for the subset of sensory units tuned to motion at 90° . (Preferred X components were 3, 6 and 9 pixels/s; the preferred Y component for each sensory unit was 0 pixels/s). The slope of the function was calculated such that sensory units tuned to low velocities had narrower bandwidths than those tuned to higher velocities (SDs of 0.8, 1.3 and 1.8 pixels/s, respectively). (b) Characteristics of the whole population of sensory units, illustrating the distribution of preferred X and Y components of velocity and response functions (Z axis).

stationary lizard sequences show texture but no movement (Fig. 5a). Collapsed optic flow plots of the two nondisplay movement sequences confirm the absence of motion, while the collapsed velocity signatures from these control sequences depict the small effect of texture values from the static image region. Figure 5b shows the nondisplay lizard movement sequences. As expected, the plots show vertical movement for the lizard climbing on

to the perch and movement in the horizontal plane for the lizard turning and walking.

Wind-blown vegetation. Figure 6 shows the optic flow plots and velocity signatures for vegetation movement at fast wind speeds. The velocity signatures, regardless of plant species, are circular or elliptical. At slower wind speeds (not shown) the velocity signatures contract

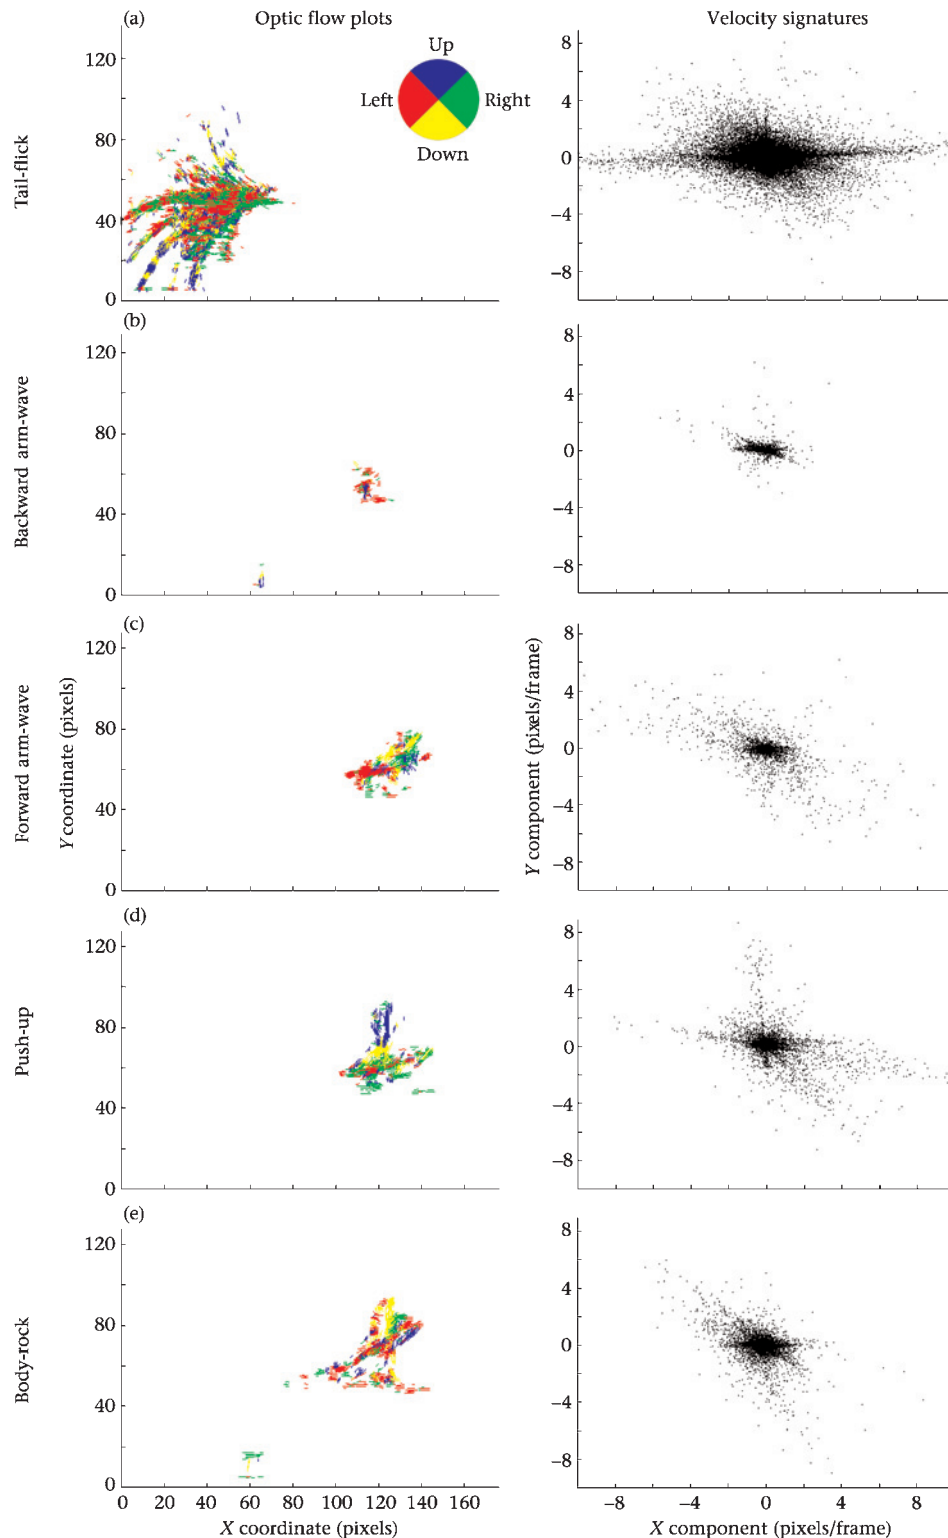


Figure 4. Collapsed optic flow plots and velocity signatures for each of the components in the aggressive display of *A. muricatus*. The tail-flick (a) appears at a different spatial location from the other components and spans a greater region of the image. Velocities to the left and right (red and green on the optic flow plot) with relatively large magnitudes (greater distance from the origin (0,0) in the histograms) are characteristic of the tail-flick. The backward arm-wave (b) has a restricted region of motion to the left, while the forward arm-wave (c) is composed of movement down (yellow) and to the right (green) as the lizard slaps its forelimb and upper body to the substrate. The optic flow plot of the push-up (d) shows preparatory movement down (yellow) and to the right (green), followed by the vertical (blue) component as the lizard lifts its head and thorax from the perch. The body-rock (e) involves movement to the left (red) and up (blue) as the lizard raises its body from the substrate, and then down (yellow) as it returns. All of these movements are also represented in the corresponding velocity signatures.

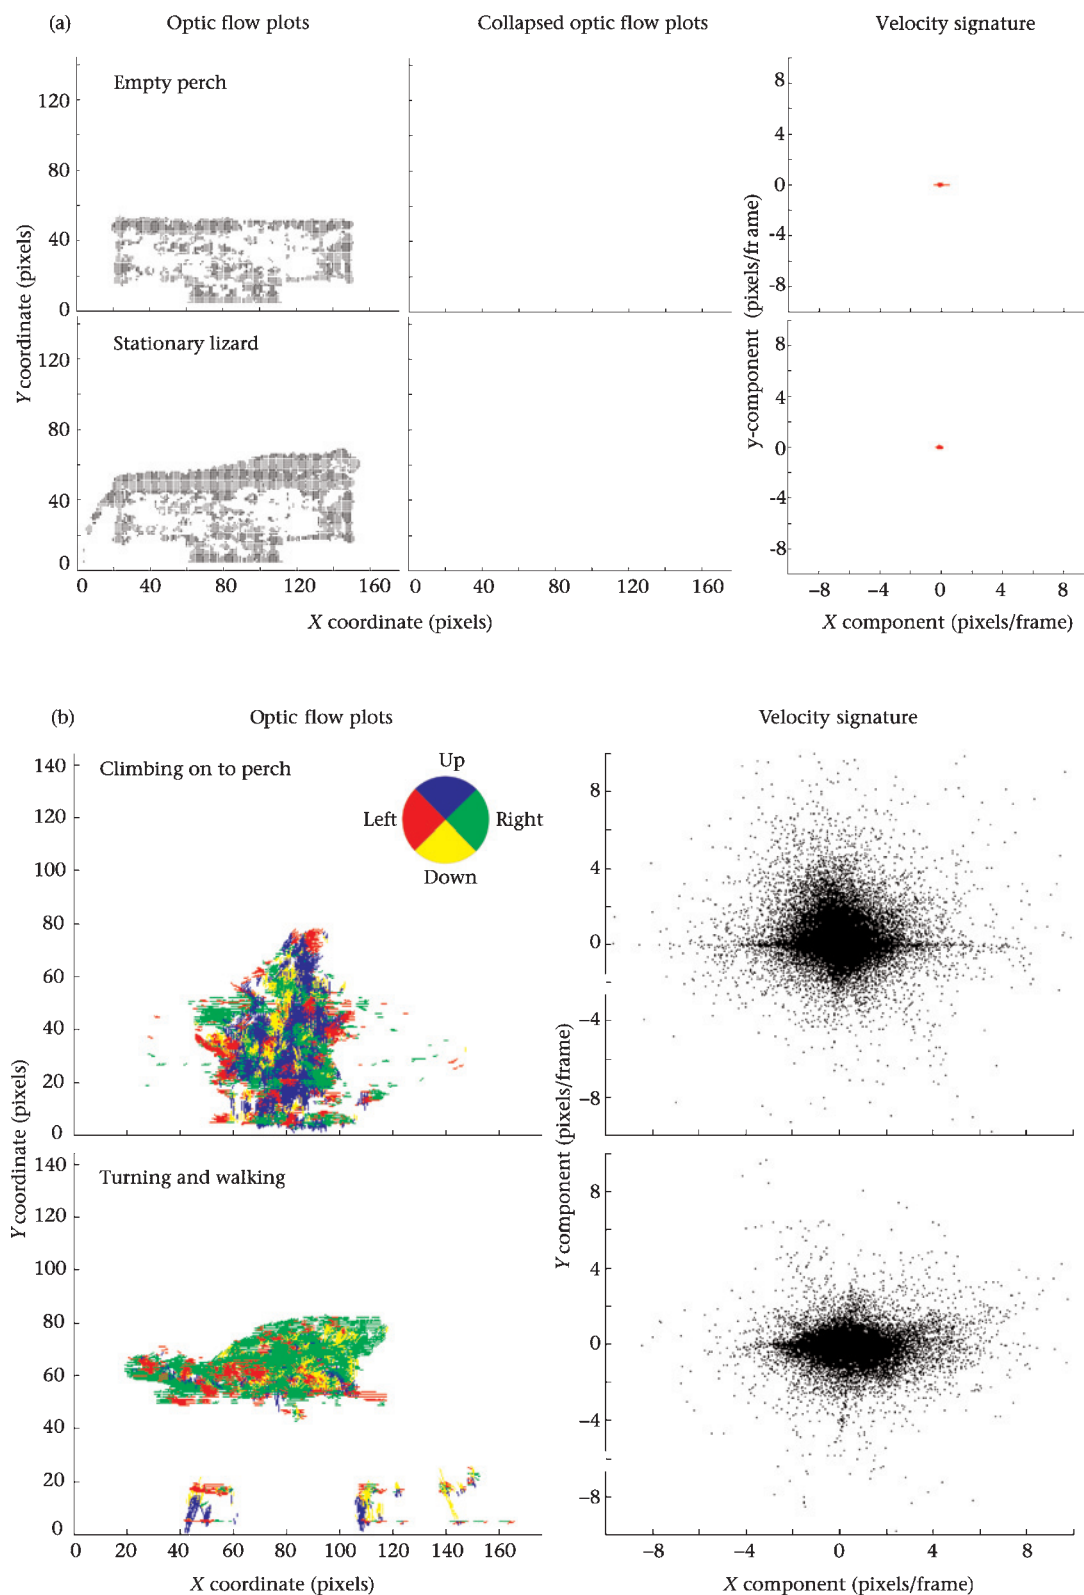


Figure 5. (a) Optic flow plots of a representative frame from each of the control sequences indicating the absence of motion, although the algorithm detected texture presented in the image (plotted in black). The central plots are optic flow plots for all frames collapsed to create a single summary plot for the sequence, which does not represent texture. Motion was not detected at any stage during these sequences. Velocity signatures for the entire sequence also show the absence of motion, with only the texture values providing data. (b) Control sequences of nondisplay lizard movement show a lizard climbing on to a perch and turning and walking across the perch. Optic flow plots and velocity signatures collapsed over the entire sequence summarize the type of motion present.

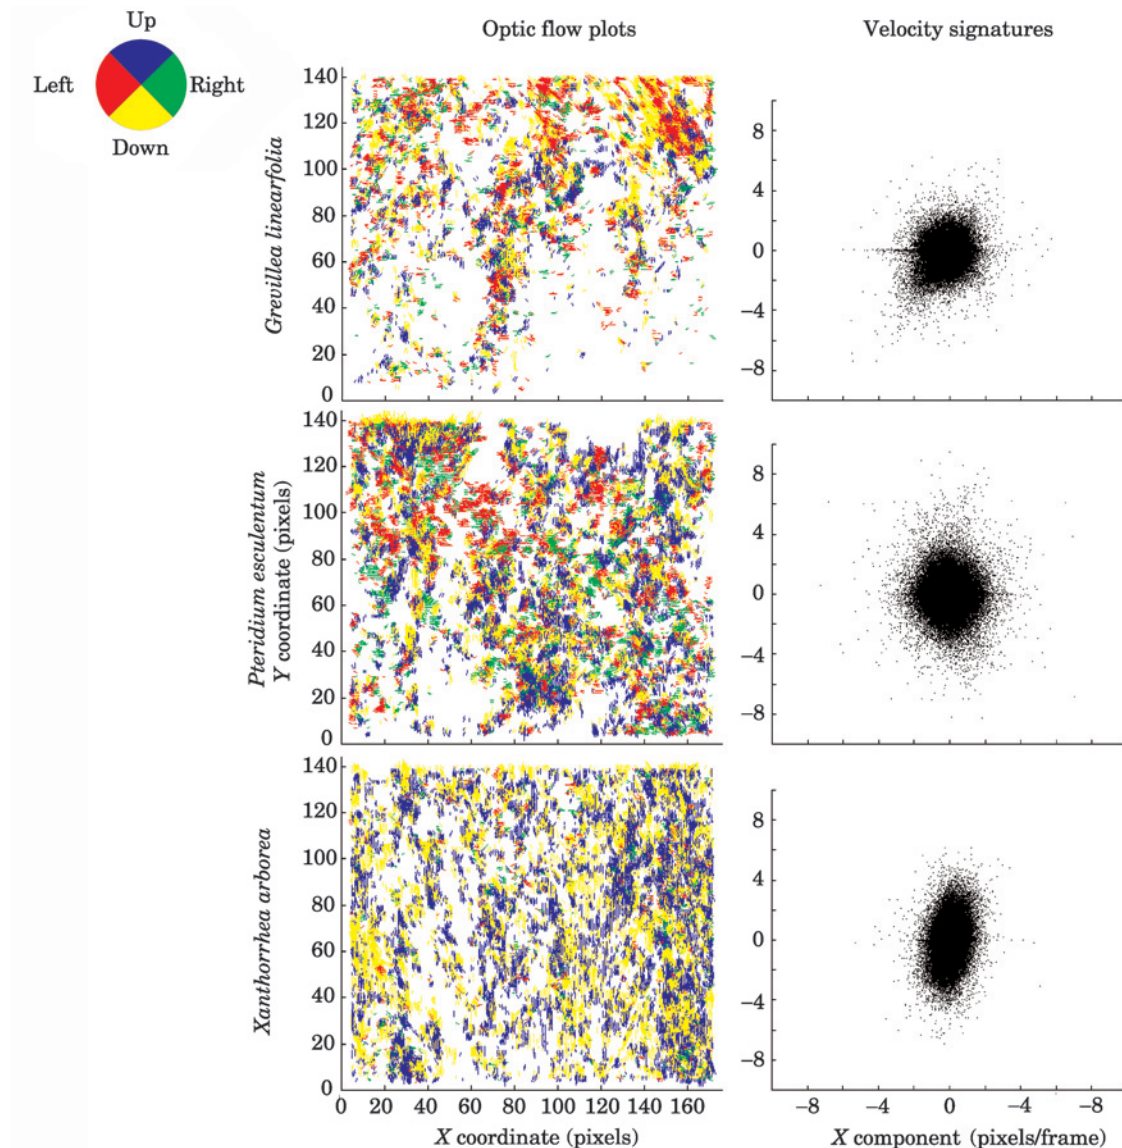


Figure 6. Collapsed optic flow plots and velocity signatures for the three species of plant typical of the habitat of *A. muricatus* at fast wind speeds. *Grevillea linearifolia* (1.75 m/s), *Pteridium esculentum* (1.76 m/s), *Xanthorrhoea arborea* (1.67 m/s).

towards the origin, while maintaining the same general shape.

Summarizing motion information

Figure 7 shows the optic flow plot and velocity signatures of the tail-flick (25-frame sample) overlaid on those generated by one of the moving vegetation controls. The two patterns are clearly distinct, but further analysis allows us to quantify the differences between them. Figure 8 shows the results obtained with each of the techniques that we are considering. The input in each case was velocity estimates from a single frame (10), which was selected to illustrate the contrast between the signal and background foliage movement.

Comparison of the velocity signatures for the tail-flick and vegetation (Fig. 8a) reveals that velocity estimates

for the signal component are distributed across the horizontal plane, while those of the vegetation are concentrated about the origin. Bivariate sample statistics provide five summary values which define a standard ellipse depicting the degree of dispersion and orientation of velocity estimates in the plane (Fig. 8b). The mean vector, $m = [\bar{x}/\bar{y}]$, provides the centre of the ellipse, while the standard deviations (s_x and s_y) and the correlation coefficient define the degree of dispersion (Batschelet 1981). The standard ellipse thus presents a graphical summary of the velocity signature (Fig. 8a) based on estimated population characteristics. The primary difference between the tail-flick and the vegetation sequences is in the standard deviation of the X component of velocity (s_x). Additional information about the distribution of velocity estimates, such as the mean direction, can also be calculated from these statistics.

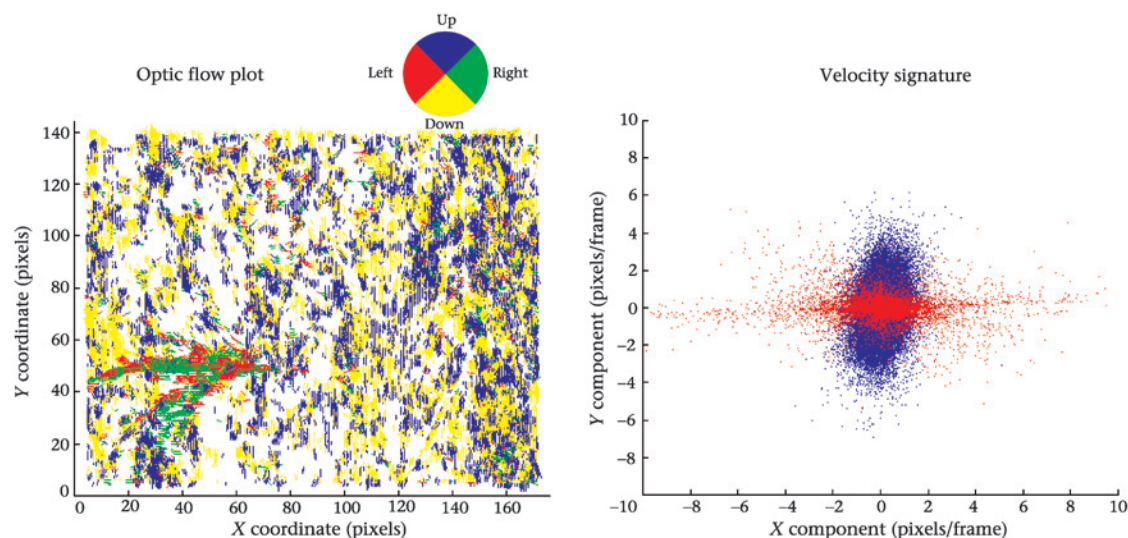


Figure 7. Velocity estimates obtained from a random 25-frame window of the tail-flick overlaid on to *X. arborea* foliage movement (fast wind speed). The tail-flick can be identified in the optic flow plot by the green and red vectors in the lower left portion of the graph. Similarly, the velocity signatures show that the tail-flick is characterized by fast movement in the horizontal plane (red), which contrasts with the relatively slower motion of the vegetation (blue).

Figure 8c shows the response of the family of sensory units to movement in frame 10. Maximum response to the tail-flick was obtained from units tuned to motion in the horizontal plane, either right (90°) or left (270°) at each of the three speeds. The vegetation sequences stimulated the sensory units tuned to slow speeds in an essentially uniform way but, as would be expected from the velocity signatures (Fig. 8a) and standard ellipses (Fig. 8b), they failed to evoke a response from sensory units tuned to faster speeds. We present a more detailed analysis of sensory unit responses over longer sequences below.

Our final statistical technique plots the average speed (velocity magnitude) for each frame of the display and vegetation sequences (Fig. 8d). In contrast to the vegetation sequences, the tail-flick is characterized by rapid changes in velocity over time.

Artificial sensory unit response to tail-flick and vegetation sequences. Figure 9 shows the summed response of each sensory unit to the entire 25-frame segments of the tail-flick and wind-blown vegetation sequences. As with the example in Fig. 8c, the units showing the maximum response to the tail-flick were tuned to motion in the horizontal plane at each of the preferred speeds. At slow wind speeds, the three species of plants produced maximal sensory unit response in the vertical plane, up (0°) and down (180°), at preferred speeds of both 3 and 6 pixels/frame. As would be expected from inspection of the velocity signatures (Fig. 6), vegetation movement failed to stimulate units tuned to faster speeds (9 pixels/frame). At faster wind speeds, the sensory unit responses to tail-flicks are masked by environmental noise at low velocity values (3 pixels/frame), become more conspicuous at intermediate velocities (6 pixels/frame), and account for almost all of the activity in the highest velocity class (9 pixels/frame; Fig. 9).

Discussion

We have presented a new approach for the study of complex dynamic visual signals and demonstrated its application using the aggressive display of the Jacky dragon. We used optic flow algorithms to determine the spatial distribution of velocity vectors, which characterizes the dynamic components of the signal (Fig. 4), other lizard movement (Fig. 5) and vegetation movement (Fig. 6), in terms of the direction and speed of motion in a two-dimensional plane. Depending on the question of interest, optic flow plots can be used to examine the spatial distribution of motion (e.g. to segment multiple moving objects in a scene), or to graph the direction and magnitude of motion as velocity signatures (Figs 4–7, 8a). Optic flow algorithms measure the speed and direction of motion directly and, since calculations of image velocity are spatially and temporally localized, it is possible to generate them in real time with sufficiently fast computers.

There are several options for more detailed analysis of these data. Calculation of mean vectors and standard ellipses (Fig. 8b), speed–time plots (Fig. 8d) and artificial sensory unit responses (Figs 8c, 9) provide both standard and novel means of quantifying movement. Collectively, these summaries provide additional insight about signal design by revealing the way in which the tail-flick is conspicuous against background vegetation movement (Figs 8 & 9).

Formal statistical comparisons could be conducted with the reduced data sets obtained with standard ellipses, which summarize the dispersal of velocity estimates and their orientation in the X/Y plane. For example, Hotelling's one- and two-sample tests could be used to assess whether the mean vector is significantly different from the origin (0,0), or from another sample,

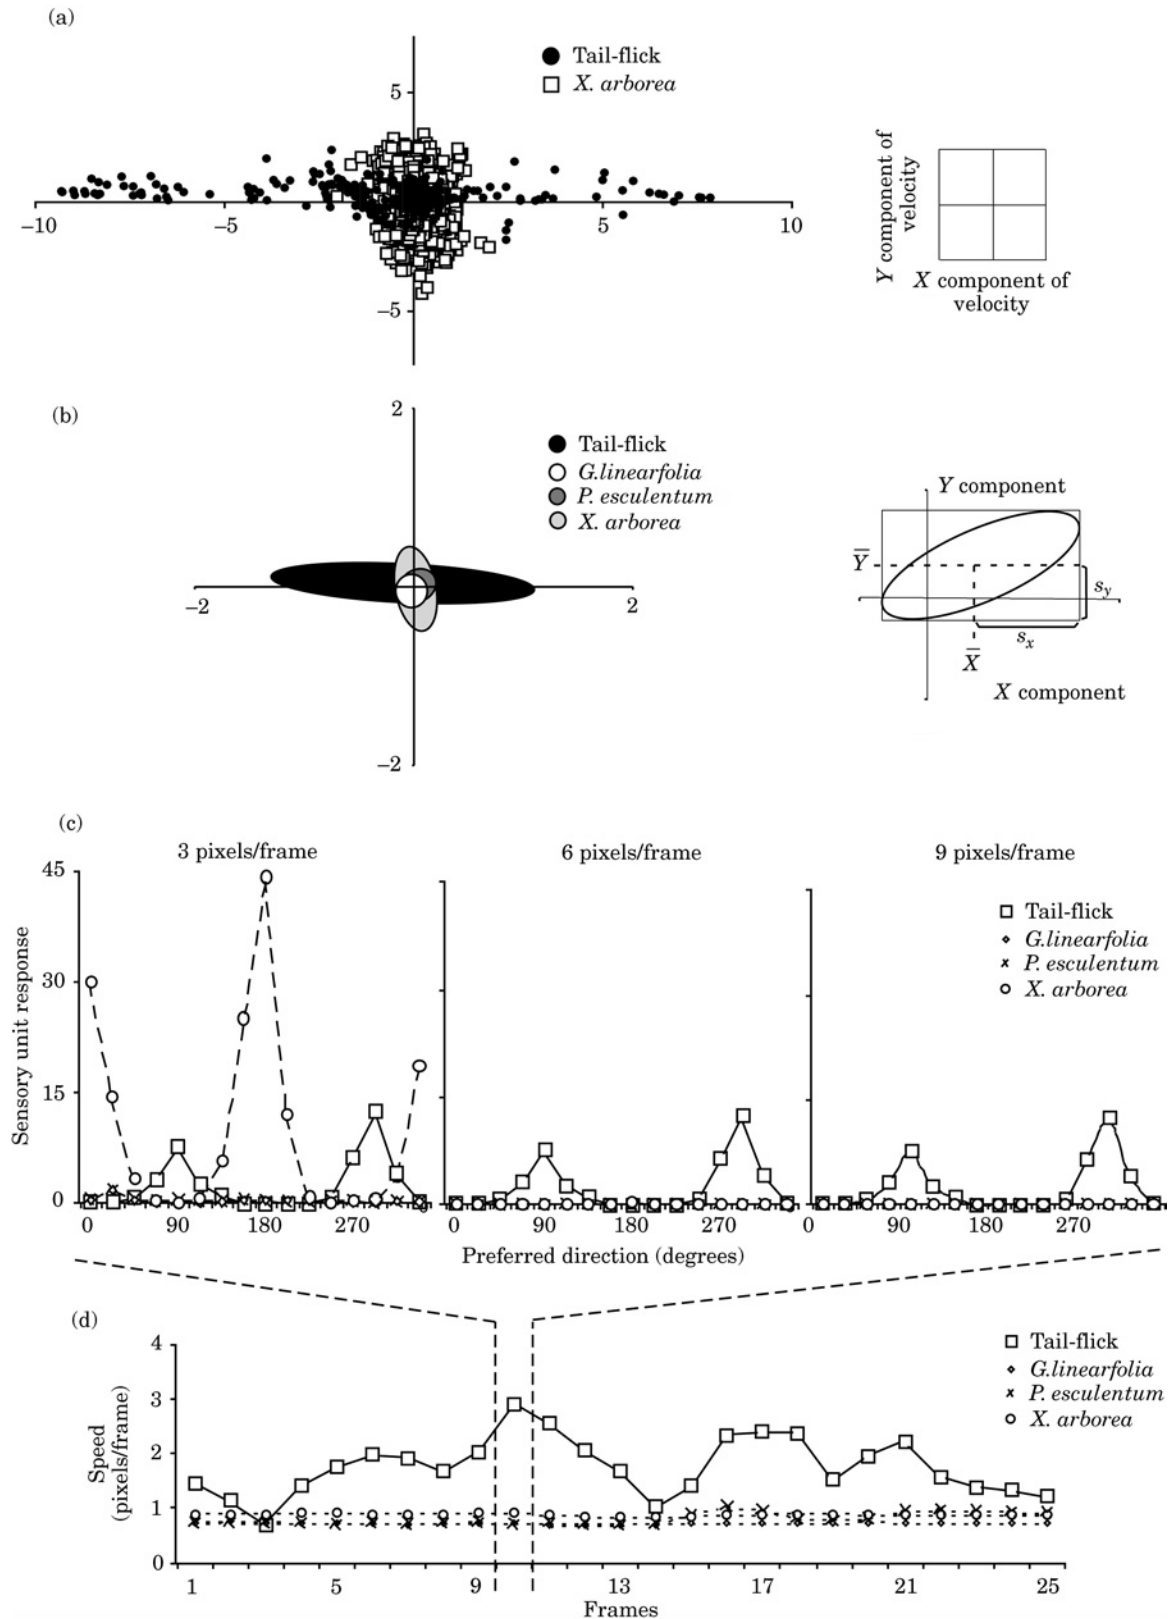


Figure 8. (a) Velocity signature of the tail-flick overlaid on a representative sequence of vegetation (*X. arborea*) for the 10th frame of each sequence. (b) The standard ellipse for the 10th frame for the tail-flick and each of the vegetation sequences at low wind speed. The centre of the ellipse is given by the mean vector (\bar{x}, \bar{y}) , while the standard deviations (s_x , s_y) and correlation coefficient (r) define the shape and orientation of the ellipse in the X/Y plane. (c) Sensory unit response to the tail-flick and three species of wind-blown plants. The preferred direction is plotted against response, separately for each speed. (d) Plot of average speed against time for the tail-flick and vegetation.

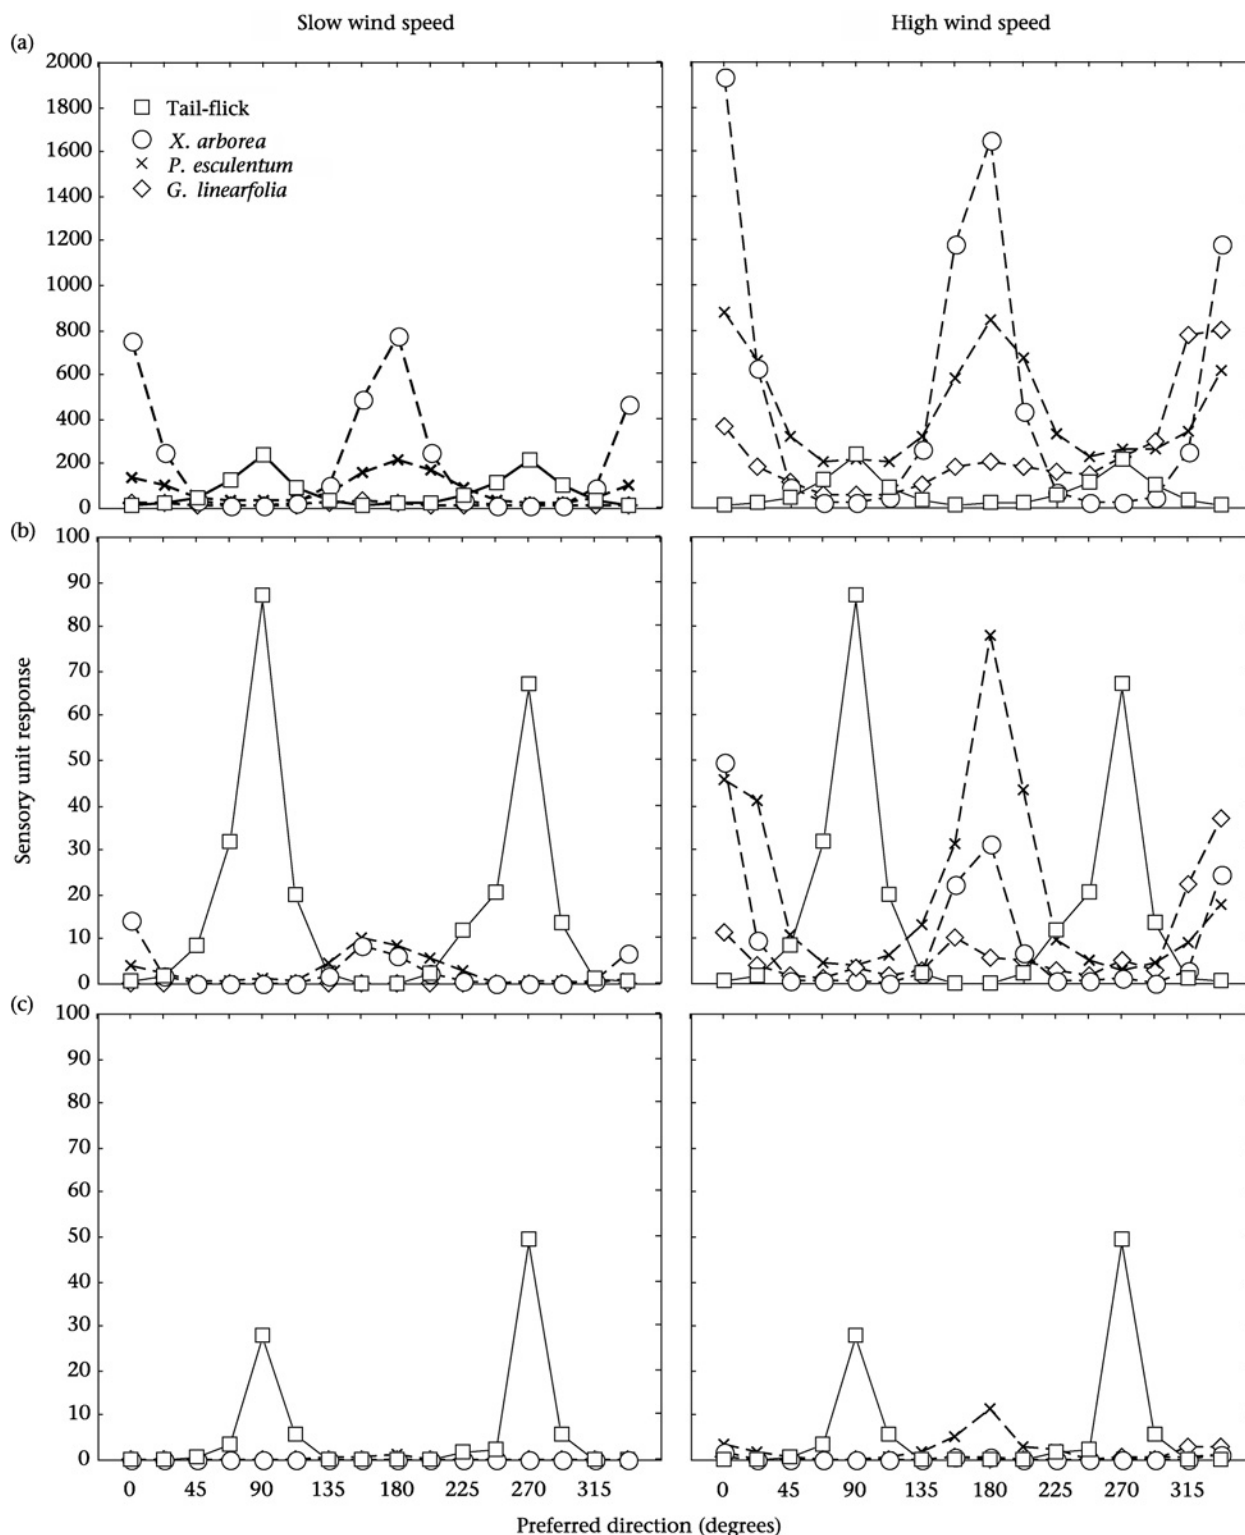


Figure 9. Response of each sensory unit to motion of the tail-flick and to three species of plants. The graphs are divided into slow wind speeds (0.73–0.90 m/s) and fast wind speeds (1.67–1.76 m/s): (a) 3 pixels/frame, (b) 6 pixels/frame and (c) 9 pixels/frame. The large number of low velocity estimates resulted in greater overall responses from sensory units tuned to this type of motion than those tuned to higher velocities. For this reason, a different scale is used for (a).

respectively. However, the likelihood of rejecting the null hypothesis would be high in both cases because of the large data sets upon which the means are based (tail-flick:

$N=3000$ velocity estimates; vegetation: $N=25\,000$ velocity estimates). Of greater interest is whether significant differences would consistently be detected across a larger

set of both movement types. One approach would be to use the summary statistics generated by bivariate analysis of individual sequences as the input data for a multivariate analysis.

Our system of artificial sensory units is a new technique for measuring structural differences between movements. Differences are apparent in the pattern of responses for the whole population, rather than in a single summary value. Conventional parametric statistics could be used to examine variation in these response profiles as a function of direction and speed. Artificial neural networks offer a promising alternative strategy for analysing sensory unit responses because of their ability to identify patterns in nonlinear input (e.g. Phelps & Ryan 1998). We have presented the cumulative response of units to entire movement sequences (Fig. 9), but it is also possible to look at the way in which responses change as a function of time. This approach is qualitatively similar to the speed-time plots described in the present study (Fig. 8d), although the four-variable (unit response \times direction \times speed \times time) output obtained is more challenging. Neural networks are especially suitable for the analysis of such data. For example, dynamic networks are designed to cope with sequential time-varying input (Elman 1990; Phelps & Ryan 1998).

Our computational approach complements earlier work by Fleishman (1992) and suggests the same conclusion in another system. We have shown that the introductory tail-flick in the aggressive display of *A. muricatus* is distinct from the background vegetation movement against which it is typically perceived. Acceleration is thought to be particularly important for triggering orienting responses (Fleishman 1992). Inspection of the speed-time plot (Fig. 8d) confirms that the tail-flick used in our analysis showed rapid changes in velocity, which reflect episodes of acceleration (or deceleration), whereas vegetation movement did not. Acceleration is simply the first derivative of speed over time. Like the velocity values extracted from optic flow plots, this is vector quantity with both magnitude and direction, so it is possible to calculate 'acceleration signatures' to obtain an analogous summary of this signal characteristic.

The goal in developing a new approach for the measurement of animal movement was to understand better the design of dynamic visual signals. However, quantitative analysis of motion using biologically relevant variables has wider applications in animal behaviour. Our particular use of computational motion analysis illustrates the more general point that this technique has attributes complementing those of traditional approaches. Systematic studies of the structure of motor patterns, which has hitherto been a recalcitrant problem, are now possible.

We thank two anonymous referees for their valuable comments on the manuscript. R.A.P. was supported by The Peter Rankin Fund for Herpetology (Australian Museum) and an Australian Postgraduate Award. C.S.E. is supported by an Australian Research Council grant. Research was conducted in partial fulfilment of the requirements for a doctoral thesis for R.A.P.

References

- Aho, A.-C., Donner, K., Hyden, C., Larsen, L. O. & Reuter, T. 1988. Low retinal noise in animals with low body temperature allows high visual sensitivity. *Nature*, **334**, 348–350.
- Batschelet, E. 1981. *Circular Statistics in Biology*. London: Academic Press.
- Barlow, H. B., Hill, R. M. & Levick, W. R. 1964. Retinal ganglion cells responding selectively to direction and speed of image motion in the rabbit. *Journal of Physiology*, **173**, 377–407.
- Barron, J. L., Fleet, D. J. & Beauchemin, S. S. 1994. Performance of optic flow techniques. *International Journal of Computer Vision*, **12**, 43–77.
- Bernard, G. D. & Remington, C. L. 1991. Color vision in *Lycaena* butterflies: spectral tuning of receptor arrays in relation to behavioral ecology. *Proceedings of the National Academy of Sciences, U.S.A.*, **88**, 2783–2787.
- Borst, A. & Egelhaaf, M. 1989. Principles of visual motion detection. *Trends in Neurosciences*, **12**, 297–306.
- Caro, T. M. 1995. Pursuit-deterrence revisited. *Trends in Ecology and Evolution*, **10**, 500–503.
- Carpenter, C. C., Badham, J. A. & Kimble, B. 1970. Behaviour patterns of three species of *Amphibolurus* (Agamidae). *Copeia*, **3**, 497–505.
- Clifford, C. W. & Langley, K. 2000. Recursive implementations of temporal filters for image motion computation. *Biological Cybernetics*, **82**, 383–390.
- Darwin, C. 1871. *The Descent of Man and Selection in Relation to Sex*. London: Murray.
- Dawkins, M. S. & Guilford, T. 1996. Sensory bias and the adaptiveness of female preference. *American Naturalist*, **148**, 937–942.
- Elman, J. L. 1990. Finding structure in time. *Cognitive Science*, **14**, 179–211.
- Endler, J. A. 1990. On the measurement and classification of colour in studies of animal colour vision. *Biological Journal of the Linnean Society*, **41**, 315–352.
- Endler, J. A. 1991. Variation in the appearance of guppy color patterns to guppies and their predators under different visual conditions. *Vision Research*, **31**, 587–608.
- Endler, J. A. 1992. Signals, signal condition and the direction of evolution. *American Naturalist* (Supplement), **139**, S125–S153.
- Endler, J. A. & Basolo, A. L. 1998. Sensory ecology, receiver biases and sexual selection. *Trends in Ecology and Evolution*, **13**, 415–420.
- Evans, C. S. & Marler, P. 1994. Food calling and audience effects in male chickens, *Gallus gallus*: their relationships to food availability, courtship and social facilitation. *Animal Behaviour*, **47**, 1159–1170.
- Fleet, D. J. & Langley, K. 1995. Recursive filters for optic flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **17**, 61–67.
- Fleishman, L. J. 1985. Cryptic movement in the vine snake *Oxybelis aeneus*. *Copeia*, **1985**, 242–245.
- Fleishman, L. J. 1988a. Sensory and environmental influences on display form in *Anolis auratus*, a grass anole of Panama. *Behavioural Ecology and Sociobiology*, **22**, 309–316.
- Fleishman, L. J. 1988b. Sensory influences on physical design of a visual display. *Animal Behaviour*, **36**, 1420–1424.
- Fleishman, L. J. 1992. The influence of the sensory system and the environment on motion patterns in the visual displays of anoline lizards and other vertebrates. *American Naturalist*, **139** (Supplement), S36–S61.
- Gibson, J. J. 1950. *The Perception of the Visual World*. Boston: Houghton Mifflin.
- Greenfield, M. D. 1988. Interspecific acoustic interactions among katydids *Neoconocephalus*: inhibition-induced shifts in diel periodicity. *Animal Behaviour*, **36**, 684–695.

- Grzywacz, N. M., Harris, J. M. & Amthor, F. R. 1994. Computational and neural constraints for the measurement of local visual motion. In: *Visual Detection of Motion* (Ed. by A. T. Smith & R. J. Snowden), pp. 19–50. London: Academic Press.
- Guilford, T. & Dawkins, M. S. 1991. Receiver psychology and the evolution of animal signals. *Animal Behaviour*, **42**, 1–14.
- Hasson, O. 1991. Pursuit-deterrent signals: communication between prey and predator. *Trends in Ecology and Evolution*, **6**, 325–329.
- Hennessy, D. F., Owings, D. H., Rowe, M. P., Coss, R. G. & Leger, D. W. 1981. The information afforded by a variable signal: constraints on snake-elicited tail flagging by California ground squirrels. *Behaviour*, **78**, 188–226.
- Hopp, S. L., Owren, M. J. & Evans, C. S. 1998. *Animal Acoustic Communication: Sound Analysis and Research Methods*. Berlin: Springer-Verlag.
- Hubel, D. H. & Weisel, T. N. 1959. Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology*, **148**, 574–591.
- Jenssen, T. A. 1977. Evolution of anoline lizard display behaviour. *American Zoologist*, **17**, 203–215.
- Lettvin, J. Y., Maturana, H. R., McCulloch, W. S. & Pitts, W. H. 1959. What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers*, **47**, 1940–1951.
- Marr, D. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman.
- Martins, E. P. & Lamont, J. 1998. Estimating ancestral states of a communicative display: a comparative study of *Cyclura* rock iguanas. *Animal Behaviour*, **55**, 1685–1706.
- Maturana, H. R. & Frenk, S. 1963. Directional movement and horizontal edge detectors in the pigeon eye. *Science*, **142**, 977–979.
- Ord, T. J., Blumstein, D. T. & Evans, C. S. 2001. Intrasexual selection predicts the evolution of signal complexity in lizards. *Proceedings of the Royal Society of London, Series B*, **268**, 737–744.
- Ord, T. J., Peters, R. A., Evans, C. S. & Taylor, A. J. 2002. Digital video playback and visual communication in lizards. *Animal Behaviour*, **63**, 879–890.
- Pagel, M. 1993. The design of animal signals. *Nature*, **361**, 18–20.
- Partan, S. & Marler, P. 1999. Communication goes multimodal. *Science*, **283**, 1272–1273.
- Phelps, S. M. & Ryan, M. J. 1998. Neural networks predict response biases of female tungara frogs. *Proceedings of the Royal Society of London, Series B*, **265**, 279–285.
- Regan, D. 1986. Visual processing of four kinds of relative motion. *Vision Research*, **26**, 127–145.
- Rosenberg, A. F. & Ariel, M. 1998. Analysis of direction-tuning curves of neurons in the turtle's accessory optic system. *Experimental Brain Research*, **121**, 361–370.
- Rosenthal, G. G., Evans, C. S. & Miller, W. L. 1996. Female preference for dynamic traits in the green swordtail, *Xiphophorus helleri*. *Animal Behaviour*, **51**, 811–820.
- Rowe, C. & Guilford, T. 1996. Hidden colour aversions in domestic chicks triggered by pyrazine odours of insect warning displays. *Nature*, **383**, 520–522.
- Schiller, P. H., Finlay, B. L. & Volman, S. F. 1976. Quantitative studies of single-cell properties in monkey striate cortex: I. Spatiotemporal organisation of receptive field. *Journal of Neurophysiology*, **39**, 1288–1374.
- Smith, A. T. & Snowden, R. J. 1994. Motion detection: an overview. In: *Visual Detection of Motion* (Ed. by A. T. Smith & R. J. Snowden), pp. 3–15. London: Academic Press.
- Speed, M. P. 2000. Warning signals, receiver psychology and predator memory. *Animal Behaviour*, **60**, 269–278.
- Stein, B. E. & Gaither, N. S. 1981. Sensory representation in reptilian optic tectum: some comparisons with mammals. *Journal of Comparative Neurology*, **202**, 69–87.
- Stein, B. E. & Gaither, N. S. 1983. Receptive-field properties in reptilian optic tectum: some comparisons with mammals. *Journal of Neurophysiology*, **50**, 102–124.
- Worgatter, F. & Eysel, U. T. 2000. Context, state and the receptive fields of striatal cortex cells. *Trends in Neurosciences*, **23**, 497–503.
- Zeil, J. & Zanker, J. M. 1997. A glimpse into crabworld. *Vision Research*, **37**, 3417–3426.

Appendix

Procedure used in calculating velocity estimates

Figure A1 summarizes the optic flow procedure used to calculate velocity estimates. (a) Video sequences are first exported as static images and converted to greyscale (see Methods). (b) Images are then subjected to spatial filtering to remove noise, followed by recursive (or infinite impulse response) temporal filtering. This results in filtered versions of the image with respect to the x and y directions of space (I_x and I_y , respectively) as well as time (I_t). (c) Filter outputs are then multiplied together. (d) Local spatial averages are then calculated with an exponential weighting function to smooth the velocity field. The resulting values (J_x^2 , $J_x J_y$, J_y^2 , $J_y J_t$, $J_x J_t$) are grouped in separate spatial and temporal matrices. (e) Finally, a velocity estimate (v) is obtained by combining these matrices using matrix algebra. The spatial weighted average of velocity estimates is incorporated to attenuate the effects of noise in local velocity estimates. There is an inherent trade-off between reducing noise and maintaining high spatial resolution. Here, a Gaussian window with a standard deviation of 0.85 pixels was implemented as a [0.25 0.5 0.25] convolution mask to provide a computationally efficient solution without parametric optimization.

Gaussian function used to define each sensory unit

This function is:

$$R = \exp - ((v_x - v_{x_{\text{pref}}})^2 / \sigma^2 + (v_y - v_{y_{\text{pref}}})^2 / \sigma^2)$$

where R is the response of the sensory unit to the velocity estimate; $v_{x_{\text{pref}}}$, $v_{y_{\text{pref}}}$ is the sensory unit's preferred x and y component of velocity, respectively; v_x and v_y are the observed x and y components of velocity, respectively. The slope of the function (or SD) is defined by σ . $R=1$, when v_x is equal to $v_{x_{\text{pref}}}$ and v_y is equal to $v_{y_{\text{pref}}}$.

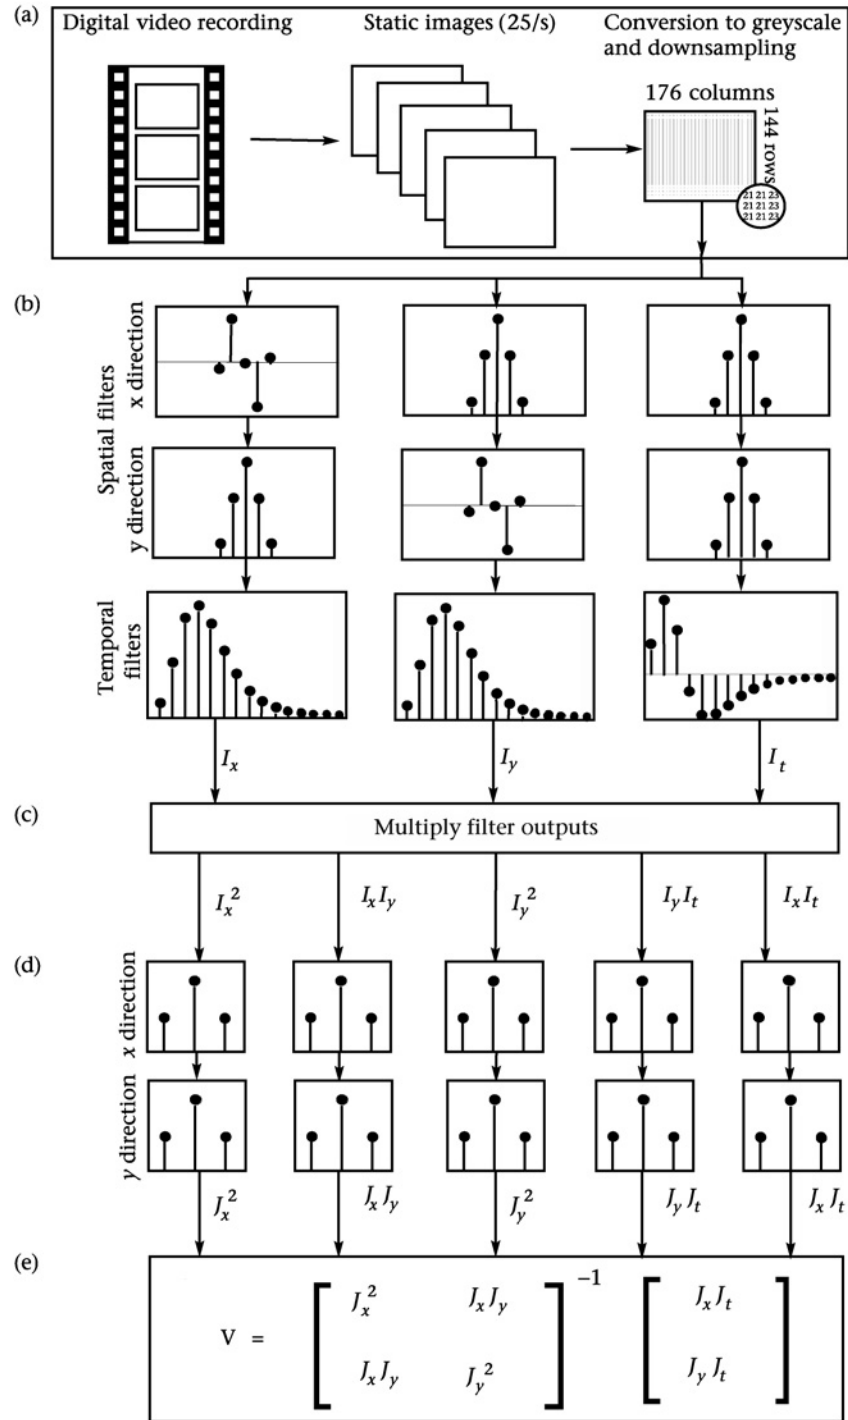


Figure A1. Flow chart depicting the computational motion analysis procedure used. (a) Digital video sequences are exported as static images (25/s), which are then converted to greyscale and reduced in size (144×176) to decrease computational requirements. (b) Spatial filters are then used to remove random variation in the intensity value of individual pixels (video noise), while recursive temporal filtering stores recent stimulus history. The filters result in partial derivatives, or impulse responses, of the image with respect to space (I_x and I_y) and time (I_t). We used a low-pass filter defined by a Gaussian function with SD of 1.5 pixels and its first-order derivative for spatial filtering. Similarly, filtering in the time domain was achieved with a low-pass temporal filter (third-order cascade of truncated exponentials; see [Clifford & Langley 2000](#) for details) and its derivative. (c) The filter outputs (I_x , I_y and I_t) were then multiplied, resulting in I_x^2 , $I_x I_y$, I_y^2 , $I_x I_t$, and $I_y I_t$. (d) To provide smooth estimates of velocity, each filter output was then replaced by an exponentially weighted local average of itself and its neighbours (I) using a masking filter (Gaussian with SD of 1.2 pixels). (e) These outputs were then grouped in separate spatial and temporal matrices, which were combined to yield a velocity estimate (v).