nature COMMUNICATIONS

ARTICLE

Received 26 Oct 2012 | Accepted 18 Apr 2013 | Published 28 May 2013

DOI: 10.1038/ncomms2908 OPEN

Leg-tracking and automated behavioural classification in *Drosophila*

Jamey Kain[1], Chris Stokes[1], Quentin Gaudry[2], Xiangzhi Song[1,3], James Foley[1], Rachel Wilson[2] & Benjamin de Bivort[1,4,5]

Much remains unknown about how the nervous system of an animal generates behaviour, and even less is known about the evolution of behaviour. How does evolution alter existing behaviours or invent novel ones? Progress in computational techniques and equipment will allow these broad, complex questions to be explored in great detail. Here we present a method for tracking each leg of a fruit fly behaving spontaneously upon a trackball, in real time. Legs were tracked with infrared-fluorescent dyes invisible to the fly, and compatible with two-photon microscopy and controlled visual stimuli. We developed machine-learning classifiers to identify instances of numerous behavioural features (for example, walking, turning and grooming), thus producing the highest-resolution ethological profiles for individual flies.

[1] Rowland Institute, Harvard University, Cambridge, Massachusetts 02142, USA. [2] Department of Neurobiology, Harvard Medical School, Boston, Massachusetts 02115, USA. [3] College of Chemistry and Chemical Engineering, Central South University, Changsha 410083, People's Republic of China. [4] Center for Brain Science, Harvard University, Cambridge, Massachusetts 02138, USA. [5] Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, Massachusetts 02138, USA. Correspondence and requests for materials should be addressed to B.d.B. (email: debivort@rowland.harvard.edu).

A major goal of biology is to elucidate the mechanisms underlying behaviour and how they have evolved. Because of its vast genetic toolkit, *Drosophila melanogaster* is an ideal model system for understanding the underpinnings of behaviour. However, analysing behaviour is not trivial and therefore most studies rely on simple, robust behaviours such as phototaxis[1] and olfactory chemotaxis[2]. Classic paradigms, performed on groups of flies, quantify these behaviours efficiently, but coarsely. In contrast, methods to efficiently characterize the behaviour of individual flies at high levels of detail are rare. With recent advances in computer vision, a new generation of automated and sophisticated assays are being developed that allow for richer characterizations of behaviours, and consequently the genes, circuits and evolution underlying them.

The first high-resolution analysis of *Drosophila* walking was a labour-intensive, frame-by-frame movie analysis[3]. Newer, sophisticated techniques can automatically annotate behaviour at the level of conspicuous behaviours of single flies. For example, the most sensitive methods to date can track individual animals and detect the behavioural motifs of walking, lunging and wing extension, but cannot resolve individual legs[4,5]. Given the tiny size of *Drosophila*, legs are exceptionally difficult to track, particularly when flies are imaged within a large arena. Visualizing the legs is vital for identifying behaviours such as grooming or for a deeper analysis of walking. Using a total internal reflection based optical method, Mendes *et al.*[6] were able to achieve high temporal and spatial resolution of a walking fly, including its legs, to explore gait parameters in great detail. The legs of large insects, like cockroaches, marked with paint on their joints, have been tracked with high-speed cameras, and detailed analyses of their locomotion have been performed[7]. However, the data acquisition requires human assistance. Unfortunately, both of these methods only capture several seconds of walking because they require that the insect walks within the narrow field of view of the high-speed camera[6,7]. Recently, recordings for $\sim$1 min of a freely behaving cricket with many paint markers on its body have been made using manual camera tracking and automated offline video analysis[8]. In contrast to such arena assays, tethered *Drosophila* walking behaviour assays (for example, Buchner[9], Seelig *et al.*[10], Clark *et al.*[11] and Gaudry *et al.*[12]) can capture long periods (several hours) of data, but to date have not included the tracking of individual legs.

Here we present a method that overcomes these limitations. It is the first technique that tracks in real time the six individual legs and the fictive movement of a tethered fruit fly behaving either spontaneously or in response to controlled visual stimuli, for hours. We also developed a machine-learning classifier that automatically detects and categorizes distinct behavioural features (for example, walking, turning and grooming), rapidly providing the most detailed behavioural recordings of single flies yet achieved. Notably, this setup was designed for compatibility with either electrophysiology or optophysiological two-photon microscopy. This approach exploits the near-infrared fluorescence of two oxazine dyes to visualize leg position with conventional (and economical) optical components and detectors, in a compact arrangement suitable as the stage of custom multi-photon microscopes. Using this device, we have uncovered new insights about *Drosophila* behaviour. Specifically, we discovered significant individual-to-individual variation in freely behaving animals, and we found that individual variation is amplified by breaking the loop between motor behaviour and sensory feedback.

Results

Recording individual fly behaviour at high resolution. We started with a traditional floating-ball treadmill rig[9–12] because tethering a behaving fly simplifies leg imaging and permits the future incorporation of simultaneous electrophysiology or optophysiology. We substituted a transparent ball to allow imaging of the legs from below using a custom imaging system (Fig. 1a, Supplementary Fig. S1; Supplementary Methods). The sphere was tracked by two infrared laser sensors normally found in computer mice. As each sensor only detects two dimensions of motion, two sensors were used to capture all three rotational components (pitch—the result of forward/backward running; yaw—the result of turning in place; and roll—the result of sidestepping/crabwalking) (Supplementary Fig. S1d–f; Supplementary Methods).

Small pieces ($\sim 100 \times 100 \times 50\,\mu m^3$) of the dyes [2.2.1]-oxazine (221ox)[13] and julolidine-oxazine (julox) dye[14] were glued in alternation to each of the six legs (Fig. 1b,c; Supplementary Methods). Flies were allowed to adapt to the dye spots for at least 24 h, and then mounted above the floating sphere via a wire tether glued to the thorax. A HeNe laser illuminated the dye on the legs, which were imaged from below, through the clear sphere, by two cameras (Fig. 1a). Each camera was equipped with a band-pass filter optimized for either 221ox or julox, allowing each leg to be uniquely detected, and adjacent legs distinguished (Fig. 1d,e). The optical wavelengths used here (between 630 and 850 nm) were chosen for their invisibility to the flies[15]—which exhibited no conspicuous responses to illumination—and compatibility with two-photon and traditional fluorescent microscopy (Fig. 1e).

Animals were typically recorded in the dark for 2 h, although we have observed flies to behave beyond 16 h. After data acquisition, the fly can easily be removed from the tether and saved for future

Figure 1 | The leg-tracker apparatus and its properties. (**a**) Schematic of the leg-tracker apparatus. A fly is mounted by wire tether to a clear sphere supported by flowing air. Fictive motion of the fly is recorded by tracking sensors. The dye spots on the leg fluoresce when excited by the HeNe laser and a second sphere re-collimates the image of the dye spots that are tracked from below by two cameras. (**b**) Photograph of a fly with bits of 221ox and julox dyes (arrowheads) glued alternately on each leg. (**c**) Cartoon illustrating the positioning of the dye spots. Dye was placed upon the femurs of the fore and hind limbs and upon the tibia of the middle legs. (**d**) Imaging of dye spots and detection of their positions. Top panel shows differential detection of alternating legs marked with the two dyes. Because the dyed segments on the front and hind legs do not cross, leg identity can be readily inferred in single frames (bottom panel). (**e**) Diagram of the absorption wavelengths of the *Drosophila* rhodopsins (top). Diagram of the absorption (dashed lines) and emission (solid lines) wavelengths of the 221ox and julox dyes and their chemical structures (middle). The dyes are compatible with green fluorescent protein (GFP). Emission wavelengths of the HeNe laser, vertical cavity surface emitting laser (VCSEL) of the tracking sensors and Ti:Sapphire lasers are depicted as vertical lines (bottom). Also shown are the transmission properties of the rig dichroic filters. (**f**) Representative data showing the 15 data vectors (3 for ball motion and 12 for leg positions) being recorded over 7,200 s (top). Magnification of a 120 s data slice (middle). Higher magnifications of two typical behaviours (bottom). Red, black and green traces represent the ball's pitch, yaw and roll components, respectively. Grey and blue traces are the *x* and *y* coordinates, respectively, of each leg. Plots are in arbitrary units, with ball motion vectors and leg-position components respectively comparable. Instrument output is calibrated. (**g**) Analysis of swing and stance phases during the second half of the running bout magnified in **f**. Blue lines correspond to the last six data vectors in **f**. Colour bar indicates the number of legs in swing phase at each frame, that is, the gait index—with Te indicating tetrapod gait (green) and Tr indicating tripod gait (orange).

use. Custom LabView software records 15 vectors in real time: the x and y coordinates of each of the six legs, and the three rotational components of the floating sphere. The HeNe laser warmed the fly and ball slightly (increasing temperature 0.6 °C; see Supplementary Methods). However, neither this temperature difference nor the presence of the dye spots had any discernible effect on behaviour, by eye, or as measured by the motion of the ball; (1) the distribution of pitch motion in the floating ball, (2) the intervals between bursts of pitch motion in the floating ball and (3) the auto-correlation properties of ball pitch motion, for flies that were unilluminated by the laser and unmarked with the dye, fell within the range seen for dyed and illuminated flies (and near their median; Fig. 2).

Representative data is shown in Fig. 1f (see Supplementary Fig. S1g and Supplementary Methods), including instances of grooming and running behaviours. From this data, we can reiterate previous gait analysis findings, such as identifying when the animal's gait includes two or three legs simultaneously in swing phase (defined as tetrapod and tripod gaits, respectively). We find that fast running comprises primarily tripod gait[3] with brief excursions into tetrapod and non-canonical gaits[6], the detection of which is largely a

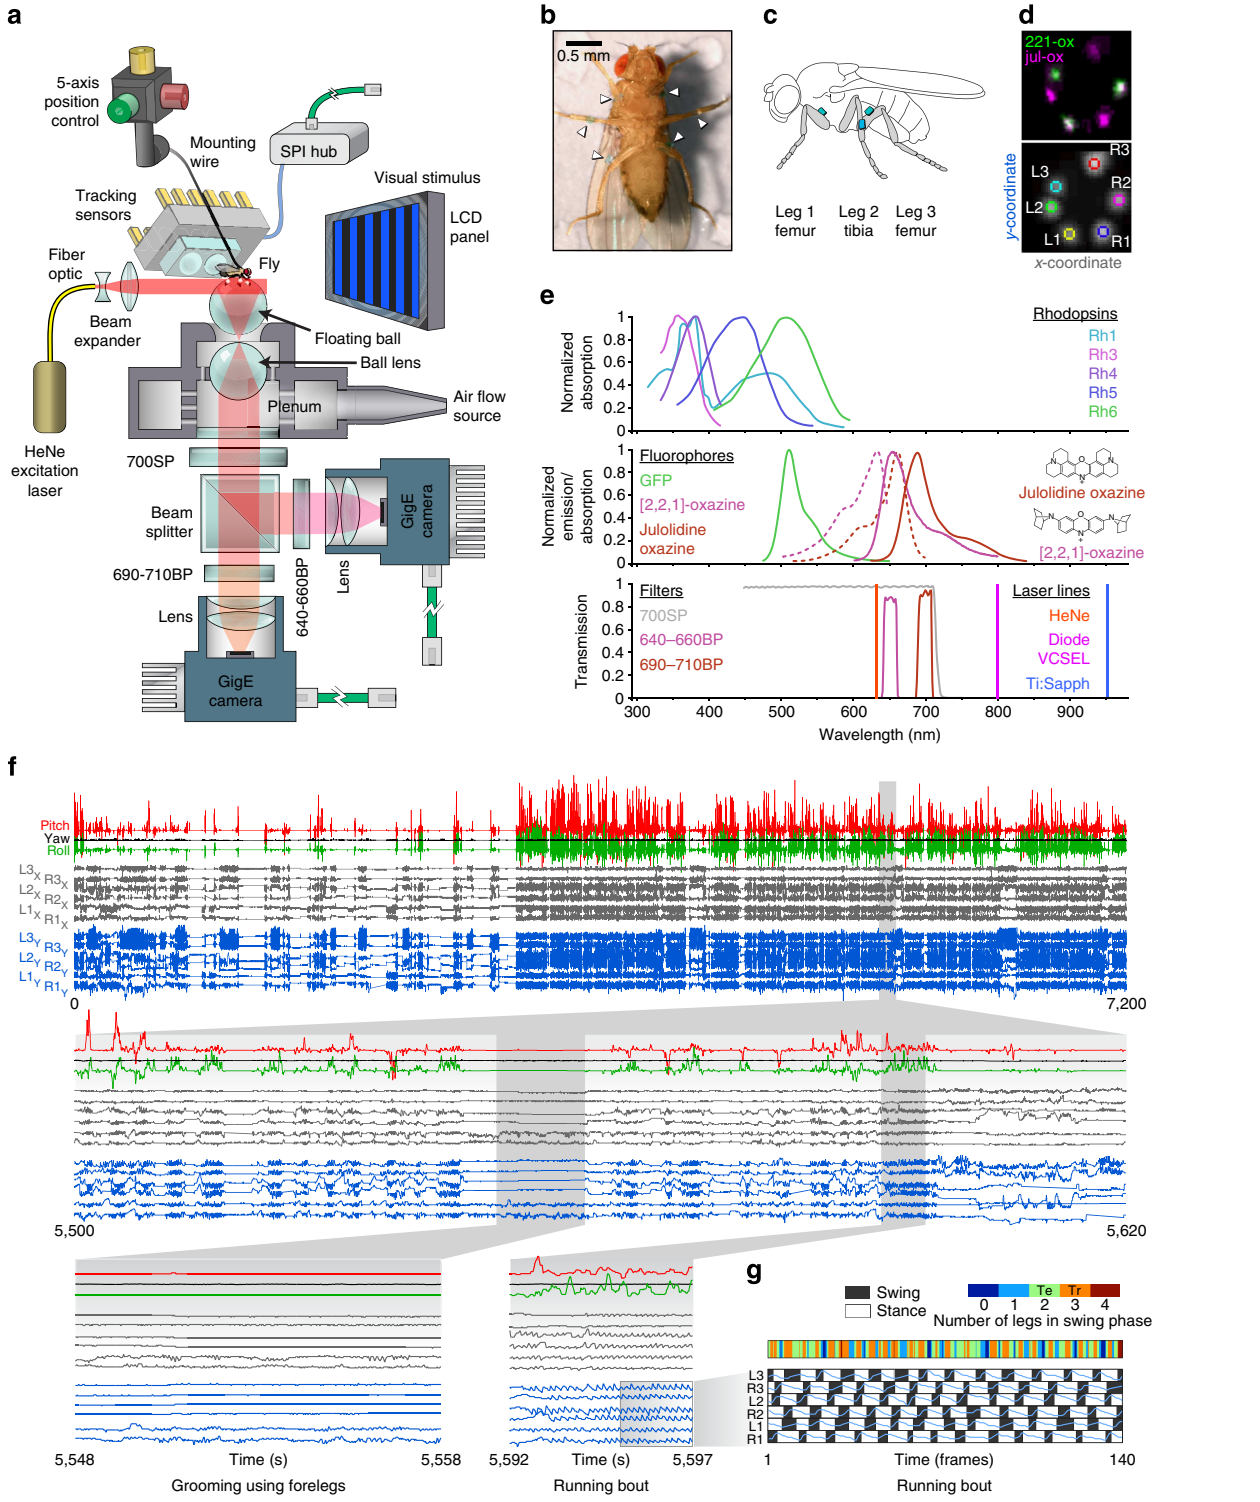

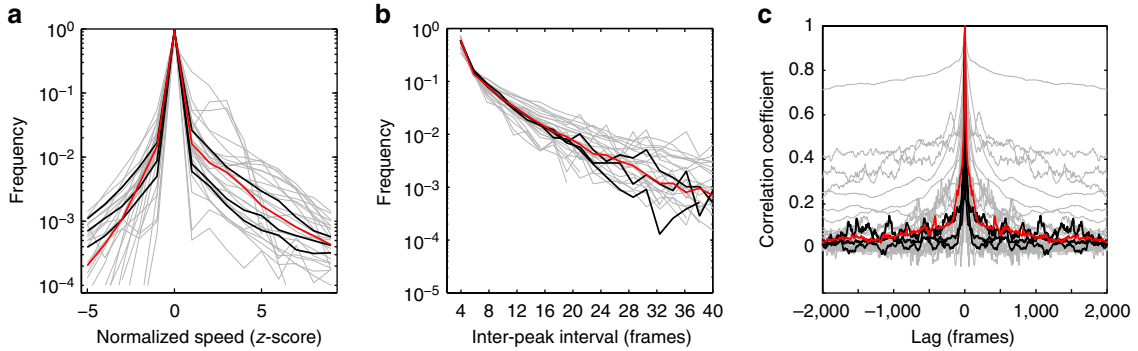

Figure 2 | Comparison of ball motion in dyed and undyed flies. (**a**) Histogram of normalized forward running speeds (ball pitch) from 2-h trials for normally dyed flies with the HeNe laser on (grey, $n=27$) and undyed flies with the HeNe laser off (black, $n=3$). Red indicates the median of dyed flies. (**b**) Histogram of intervals between peaks of forward running (ball pitch). Experiments and plot colouring as in **a**. (**c**) Auto-correlation plots for ball pitch vectors. Experiments and plot colouring as in **a**.

consequence of high temporal resolution (Fig. 1g). In addition to directly observing single leg strides via the dye and optics, the sphere itself was also responsive to single leg strides (Supplementary Fig. S1h).

Automated classification of behaviour. As our method records the position of legs at all times, we wanted to characterize behaviours other than running, such as grooming. Because of the volume of data collected (15 vectors per frame at 80 + Hz), we wanted to develop an automatic method for scoring behaviour. Non-linear classifiers have been quite successful at categorizing ethological data (see, for example, Dankert *et al.*[4], Branson *et al.*[5], Belongie *et al.*[16], Veeraraghavan *et al.*[17], Jhuang *et al.*[18] and Burgos-Artizzu *et al.*[19]). After exploring support vector machines and neural network classifiers, we chose to pursue *k*-nearest neighbours (KNN) analysis (Fig. 3a; Supplementary Methods)[20], which had the advantage of conceptual simplicity and performance comparable to the other methods. Briefly, KNN works by placing an unclassified data point into a space with all the data points whose classifications are known (the training set). The classification applied to the plurality of the *k* neighbours (from the training set) nearest to the unclassified point is taken as the label of that point. To train the classifier, two investigators independently scored, frame-by-frame, movies recorded concurrently with leg tracking (Fig. 3b). The training set consisted of data from three male and two female flies; across these and all other animals examined, we observed no conspicuous sex-dependent behavioural differences. Training data needs to be generated only once and can be used to score all future trials.

Hand scoring consisted of assigning to each frame of the movie one of the 12 possible behavioural labels (Fig. 3c). The definitions of these behavioural labels were not explicitly enumerated before scoring so that we could assess the degree of labelling discrepancies between investigators working independently—the appropriate benchmark for the automated training methods to follow. The two investigators achieved an unbalanced accuracy of 71%—that is, assigning the same score to 71% of frames, setting the goal for classifier performance (Fig. 3c). The greatest discrepancies between manual scorers appear between pairs of behaviours only distinguishable by arbitrary cutoffs (for example, complex motion versus forward/backward motion), as well as between abdominal and L3 grooming. The latter ambiguity reflects biological reality—we have observed flies to perform both of these behaviours simultaneously, stroking their abdomen as well as their hind legs. There was also a trend that rarer behaviours were scored less consistently between manual scorers. Classifiers trained on the original 15 vectors (the raw data) were mediocre (47% accuracy). However, augmenting the raw data with higher-order features, specifically the derivatives and local standard deviations of each of the 15 raw data vectors substantially improved performance (to 61% accuracy, see Supplementary Methods for calculations). Lastly, applying a low-pass filter to the classifier predictions further improved their accuracy. In the end, the classifier had an error rate only 5% greater than that seen between independent manual scoring (66% accuracy, Fig. 3d,f, see Supplementary Methods and Supplementary Movie 1). The accuracy was quite good (91%), if errors between behavioural categories that are only distinguishable by arbitrary cutoffs (the most common discrepancies in the manual scores) are ignored (see 'plausible accuracy' in Fig. 3d).

Fly-behavioural idiosyncrasy. With the classifier in hand, we categorized the behaviour of individual flies (for example, Supplementary Movie 2) across 2-h trials. Flies spent most of their time pausing or grooming and ran more at the start of trials (Fig. 4a,b). We built ethograms[21] to look at the direction and frequency of transitions between the 12 behavioural types (Supplementary Methods). We found that the most frequent transitions were between the two types of foreleg grooming and the two types of hind-leg grooming (Fig. 4d). As has been reported in higher-level insect behaviours—for example, escape response[22] or phototaxis[23], individual flies also showed variation in their behaviours. For example, one animal transitioned frequently from abdomen grooming to hind-leg grooming, whereas another animal did so more rarely. In contrast, the latter animal frequently alternated between abdomen grooming and pausing, whereas the former did so rarely (Fig. 4d).

This individual-to-individual variation could be an artifact of mounting bias. To examine this, we remounted the same individual flies on subsequent days. We found that the transition rates in the ethograms were significantly ($P=0.0022$ and 0.0011 as estimated by *t*-test and bootstrap resampling, respectively; see Supplementary Methods) more similar between re-mountings of the same fly than between flies, suggesting that much of the inter-trial variability comes from idiosyncrasies of the flies that persist for at least a few days (Fig. 4e). This was not observed (with statistical significance) in the simple percentage of time spent doing each behaviour (Fig. 4c), suggesting that there is more idiosyncrasy in the pairwise behavioural transition rates than in the overall off-rate of each behaviour. Moreover, the idiosyncratic nature of behavioural transitions is evident in

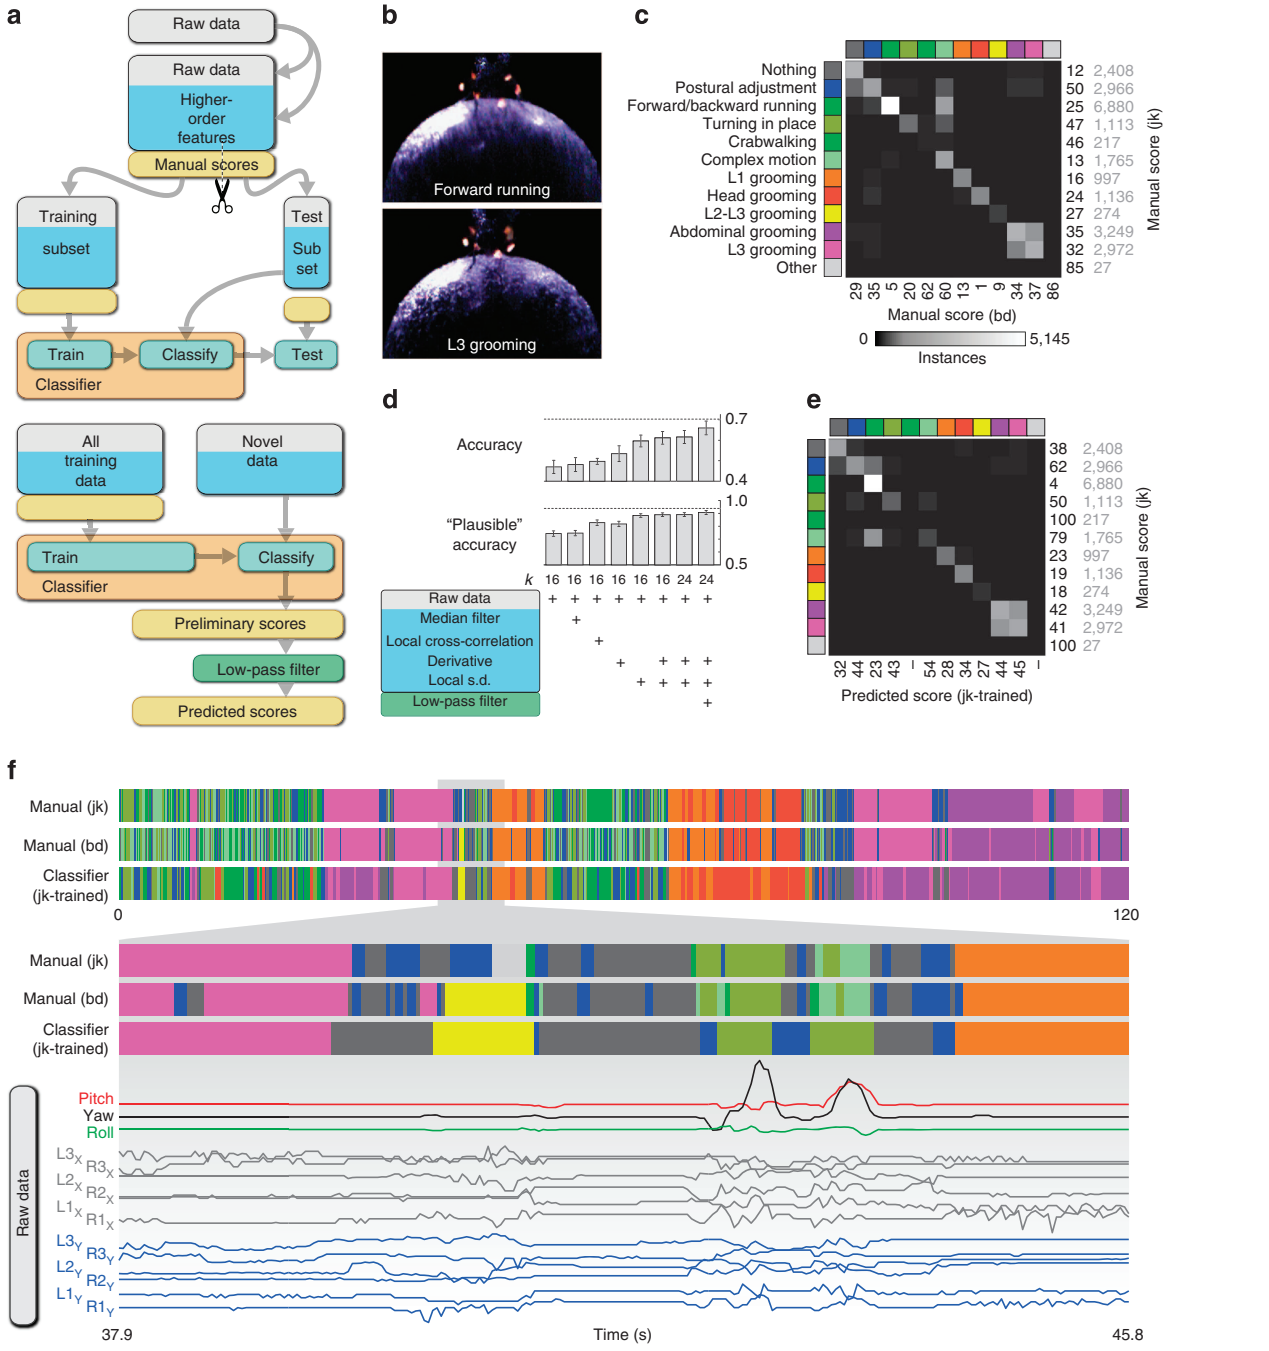

Figure 3 | KNN classifier for the automatic labelling of data. (**a**) Flowchart depicting development of the classifier (top) and its use to classify data (bottom) (see also Supplementary Methods). (**b**) Examples of movie frames used to hand-score behaviour for training the classifier. (**c**) Contingency matrix of the hand-scored movie-frame data sets from the two researchers. Shade of grey indicates the abundances of (dis)agreement between the two researchers. Colours here and in **e** and **f** indicate behaviour types. Black numbers at right and bottom, respectively, indicate the false-negative and false-positive rates (%) for each behavioural category with respect to the JK manual scoring. Grey numbers indicate the total number of frames manually annotated with each behaviour. (**d**) Including higher-order features with the raw data increases the classifier's accuracy (accuracy was measured as the total fraction of frames scored correctly, that is, the unbalanced accuracy; see Supplementary Methods). Behaviours such as forward/backward running and complex motion are qualitatively contiguous; the 'plausible' accuracy metric ignores errors between such pairs of behaviours (see text). Cross-validated error bars are $\pm$s.e.m. calculated across $n=5$ flies. (**e**) Contingency matrix of the JK hand-scored data set to the classifier trained on the JK-scored data set. Numbers as in **c**. (**f**) Sequences of behaviour scores from both JK and BD manual data sets and from the classifier for the same 120-s window (top). Magnification of a $\sim$8-sec subset along with the raw data used by the classifier (bottom).

fine-scale analysis of the mean positions of individual legs across said transitions (Fig. 4f). Thus, we conclude that certain ethological characterizations, namely the transitions between behaviours, are robust to any mounting artifacts, persist for days and vary across flies.

As the setup uses wavelengths invisible to the fly (Fig. 1e), we can present controlled visual stimuli to the animal to study behaviours such as phototaxis or optomotor responses[24]. As a demonstration, we recorded animals' responses to optic flow by placing them in front of a liquid–crystal display monitor

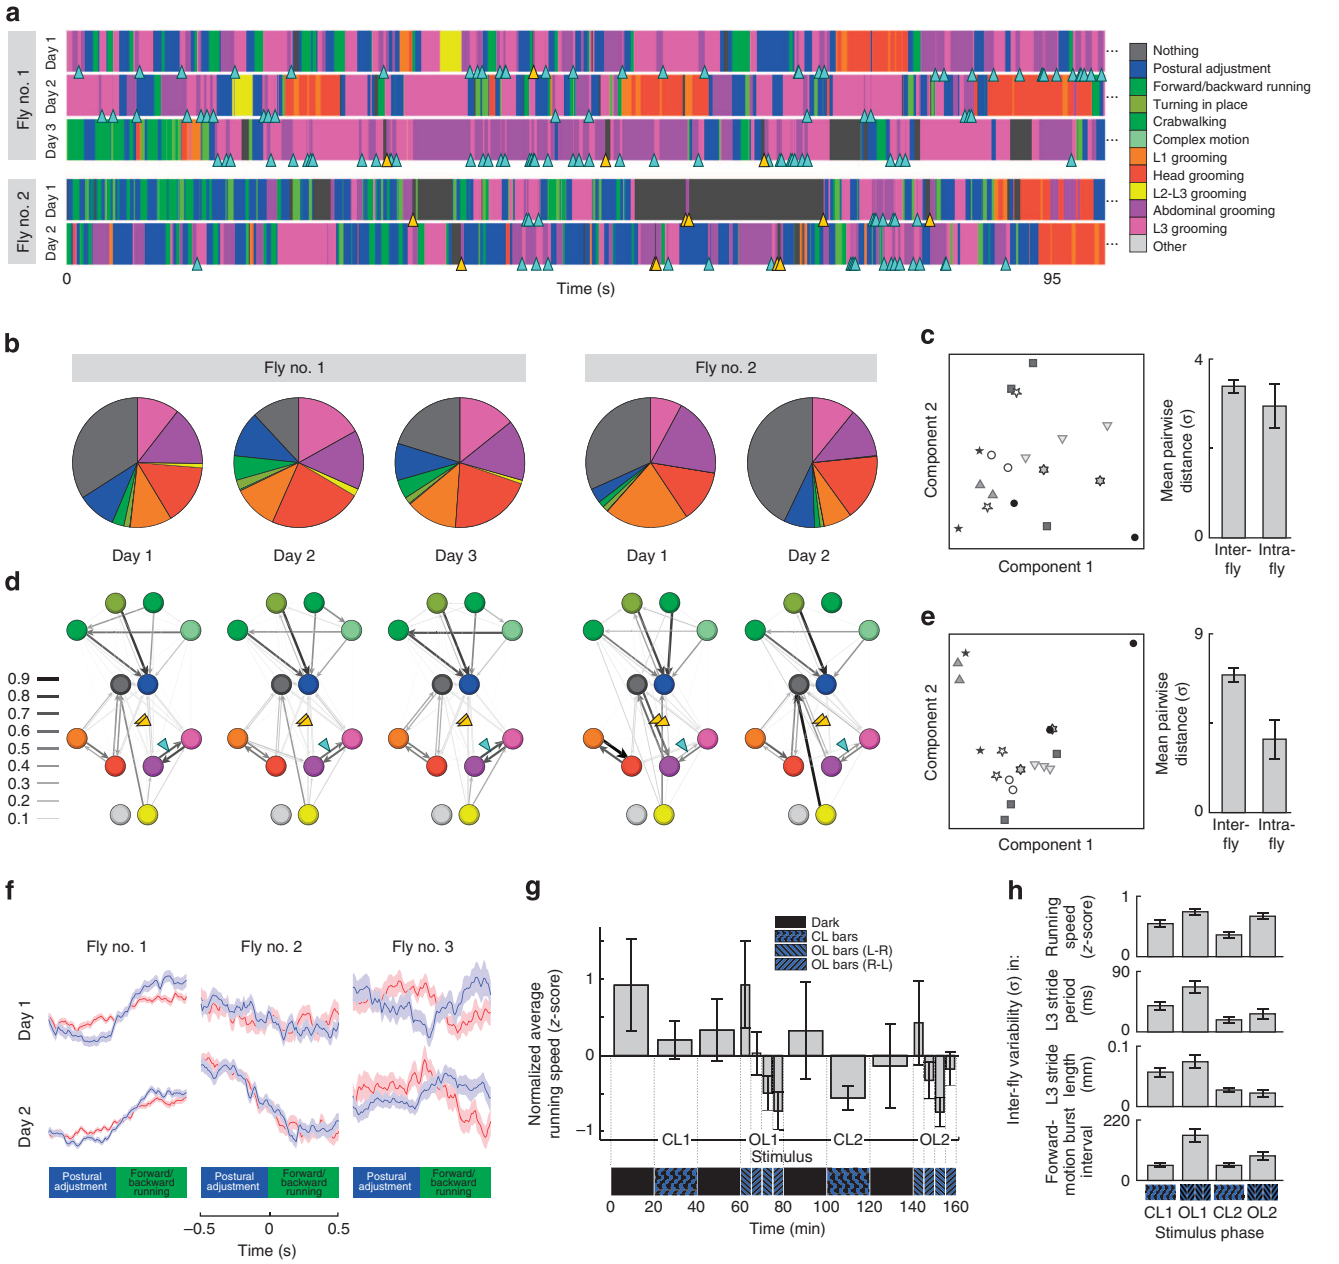

Figure 4 | The behavioural patterns of individual animals. (**a**) Multi-colour bars show snapshots (95 s) of classifier scores for multiple trials of two individual flies. Blue triangles indicate a behavioural transition more frequent in fly no. 1 and yellow triangles indicate a transition more frequent in fly no. 2. Colours here and in **b** and **d** indicate behaviour type. In **a**–**f**, distinct trials for each fly were recorded on separate days. (**b**) Pie charts depicting fraction of time spent performing each behaviour for distinct trials from two individual animals, based on 2-h recordings per experiment per animal. (**c**) Left: spatial visualization of principal components (PCs) 1 and 2 underlying the inter-trial variance in fraction of time spent per behaviour in **b**. Shapes indicate multiple trials of an individual fly. Right: mean pairwise distance between and within flies in PC1–PC2 space at left. Bars are ± s.e.m. across $n = 153$ inter-fly and $n = 12$ intra-fly pairwise distances. See Supplementary Methods. (**d**) Ethograms of the same trials from **b**. Arrows indicate direction of transition and thickness indicates frequency. Self-transitions are omitted for clarity. Triangles as in **a**. See Supplementary Methods. (**e**) Left: spatial visualization of principal components 1 and 2 underlying the variance in the ethograms from **d**. Shapes indicate multiple trials across days of an individual fly. Right: mean pairwise distance between and within flies in PC1–PC2 space at left. Bars are ± s.e.m. across $n = 153$ inter-fly and $n = 12$ intra-fly pairwise distances. See Supplementary Methods. (**f**) Individuality in the average y coordinates of the left (red) and right (blue) hind legs during the transition from postural adjustment to motion. Shaded regions indicate ± s.e.m., with n varying according to the number of times each fly performed a postural adjustment–forward/backward running transition in each trial ($48 \leq n \leq 282$). (**g**) Running speed (z-score normalized across all experimental phases) during optomotor experiments as a function of stimulus. Bars are ± s.d. across $n = 5$ flies. Blue/black patterns are representative top-to-bottom kymographs of the presented visual stimulus. CL, closed loop; OL, open loop. See Supplementary Methods. (**h**) Inter-fly variability (s.d.) of four different locomotor parameters, as a function of optomotor stimulus loop-closure status. Bar means in top graph correspond to error bars in **g**. The s.e. bars on the estimates of variability were determined by bootstrap resampling. L3 stride period and length were determined by inter-peak frequency analysis and are representative of all other legs. Forward-motion burst intervals were determined by the mode of the intervals between peaks of ball forward/backward velocity.

displaying vertical bars translating from a line of expansion under either open- or closed-loop feedback control. We found that flies showed less variance in their average running speed during closed-loop trials relative to open-loop trials ($P = 0.002$, χ^2-test of variance; Fig. 4g,h), an effect also seen in measures of flight symmetry in 'flight-simulator' assays (Michael Reiser, personal communication). This pattern was also present in the inter-fly variability in leg-stride period, the intervals between bursts of pitch motion in the floating ball, and to a lesser extent stride length (Fig. 3h). We also observed that inter-fly variability generally diminished as the experiment proceeded.

Discussion

This method provides unprecedented level of detail for the characterization of walking behaviour, revealing rich and diverse behavioural profiles of individual animals. Using two dyes, two cameras and some optics, we were able to track each leg of a spontaneously behaving fruit fly in real time. With this technique, we observed that individual flies from the lab wild-type strain behaved idiosyncratically. The behavioural personality of an individual fly was consistent from day-to-day, at least over the time frame we examined. We also found that under closed-loop conditions, the flies behave more similarly to each other; perhaps interactive stimuli engage circuitry that overrides their more idiosyncratic preferences. The automatic classification of behaviour described here represents a general approach to developing automatic ethological classifiers for the efficient collection of statistically powerful data sets.

This method will be a powerful tool for descriptive and comparative studies, for analysing subtle mutant phenotypes, and will reach its greatest potential upon integration with optophysiological recording, i.e. for probing the activity of neural circuits as they mediate individual behaviours. For example, the GAL4-targeted expression[25] of genetically encoded fluorescent physiological indicators[26] to neural subpopulations in the motor circuitry of the brain and ventral nerve cord, in simultaneous combination with leg and motion tracking, will be an effective technique for the identification of central pattern generators and decision-making circuitry.

Methods

Fly stocks. Flies (wild-type Canton S) were grown on standard cornmeal media (BuzzGro from Scientiis) in 25 °C incubators with a 12/12 h light–dark cycle. Behavioural experiments were recorded in an environmental room set at 23 °C, 40% humidity, with the lights off.

Apparatus design details and use. For the trackball, we used clear acrylic spheres (1/4″; McMaster-Carr) suspended on air. The spheres were scuffed with sandpaper for better tracking by the infrared sensors. The two infrared sensors (Avago chip no. ADNS-6090) were extracted from high-performance gaming mice. See Supplementary Methods for additional design details. The x and y coordinates acquired from the two sensors were transformed to generate the three motion vectors of the sphere using trigonometry. See Supplementary Methods for details.

For mounting, a fly was glued on the notum to a stainless steel no. 0 insect pin (Bioquip) using ultraviolet-cured adhesive (LED100 system from Electro-Lite Corporation and KOA 300 adhesive from Kemxert Corporation). The pin was glued to a flexible copper wire attached to a male audio jack. The male audio jack-wire-pin module was designed to be small and mobile so that it can be used at the dissecting scope. Once the fly was attached, the module press fits into a female audio jack that was attached to the micromanipulators (Siskiyou Inc) of the larger apparatus, which were used for fine adjustments to the animal's position upon the sphere. Animals were allowed to adapt to their positioning in the device for approximately 20–40 min, in the dark, before data collection.

The leg dyes were excited by a HeNe laser (632.8 nm; Thorlabs). The fluorescent light from the leg dye must pass through the clear trackball sphere to reach the imaging cameras (Prosilica GigE cameras; Allied Vision Technologies). Because this sphere will act as a lens, a second sphere is used to re-collimate the light before it reaches the cameras. Images were typically acquired at 80 Hz. Each camera used a band-pass filter (Edmund Optics) to optimize for the detection of one of the two leg dyes. In addition, a short-pass filter (NT49-829; Edmund Optics) was used to block stray light from the infrared tracking sensors and potential two-photon sources. See Supplementary Methods for additional details.

Laser illumination warmed the fly and sphere slightly (0.6 °C), an effect less than that of the GigE cameras, which warmed the stage by 3 °C. Temperature effects were measured using a digital thermometer with a metal probe placed across the floating-ball socket, before and after 60 min of laser illumination.

In real-time computational image processing, the 221ox image was masked out from the julox image within custom LabVIEW interfaces (National Instruments), so that the image from each camera now represented one specific tripod (that is, a set of three alternating legs). Leg identity was assigned by simple positional logic, for example, for the tripod consisting of the front-left, middle-right and back-left legs, the right-most image dot was assigned the identity of a middle leg, and the front-most of the remaining two dots a front-leg identity.

We used a digital microscope video recorder (ProScope; Bodelin Technologies) for trials where movies were being simultaneously recorded for hand scoring. We inserted a polyester film filter (Roscolux, medium-blue colour filter no. 83; Roscoe Laboratories) over the microscope to block light from the HeNe laser while still allowing visualization of the sphere, fly and fluorescent dye spots.

For the optomotor feedback experiments, we used a 15-inch liquid–crystal display monitor positioned 6 inches from the fly to display alternating blue- and black-coloured vertical bars. The trials comprised the conditions: 20 min in the dark; 20 min closed loop; 20 min dark; 20 min open loop; and repeat from the beginning once. Closed loop allowed the turns of the fly to control the position of the line of expansion for the translating vertical bars. See Supplementary Methods for additional details.

Dyeing the fly legs. The dye 221ox (phenoxazin-5-ium, 3,7-bis(7-azabicyclo[2.2.1]hept-7-yl)-chloride) was prepared according to the method described elsewhere[13]. The dye julox (1H,5H,11H,15H-diquinolizino[1,9-bc:1',9'-hi]phenoxazin-4-ium,2,3,6,7,12,13,16,17-octahydro-, chloride) was prepared according to the method described elsewhere[14]. Leg tracking would likely work using quantum dots[27] instead of the dyes, by selecting dots based on their diameters to match the emission spectra of the dyes, but we did not try this approach.

The dyes were dissolved in methyl chloride and then mixed with clear nail polish and spread out into a thin film upon a glass slide to dry. Small pieces were cut out of the dye film with a scalpel. A small dab of ultraviolet-cured glue was placed on each leg of the animal under CO_2 anesthetization (Fig. 1b,c). A piece of dye film was placed onto the glue drop and then exposed to several seconds of ultraviolet light, thus securing the dye to the leg. Dye pieces were placed on the femurs of the fore and hind legs and the tibias of the middle legs to better separate the dye signals in space for easier tracking. Marked flies were allowed to recover and acclimate overnight.

Raw data and analysis scripts. All data from this project, as well as all LabView and MATLAB scripts, can be downloaded at: http://lab.debivort.org/leg-tracking-and-automated-behavioral-classification-in-Drosophila/.

KNN classifier. After exploring several types of classifiers, we settled on the KNN approach. Two researchers, independently, designated each frame of a movie as falling into one of the 12 behavioural types. Five movies (representing five distinct flies) comprising either 4,000 or 8,000 frames were scored. These behavioural annotations, in combination with the corresponding 15 raw data vectors, were used to train the classifier. Supplying higher-order features (such as the instantaneous leg motion velocities, for example) in addition to the raw data improved the classifier performance. See Supplementary Methods for more details on the development of the classifier.

Data analysis. Ethograms were calculated by populating a 12×12 transition matrix at position $M_{a,b}$ with the number of instances for which a frame of behaviour a precedes a frame of behaviour b. Self-transitions were ignored to focus on inter-behavioural transitions, and because the pie chart representations predominantly reflect self-transitions. Rows of the transition matrix were normalized by their totals to yield transition probabilities, which were visualized in a network fashion.

Principle component analysis was performed on the (12) percentages of the pie charts and ($132 = 12 \times 11$) probabilities of the ethograms, with replicates consisting of the 11 experiments done across five different flies over 2 days (or 3 days in the case of two flies). In both cases, all input variables were z-score normalized before principle component analysis, and the first two principle components were found to explain a minority of the total variance. Statistical significance of the clustering of each fly's day-to-day replicates versus the overall distribution of the trials in Fig. 4c,e was determined in two ways: by t-test comparing the mean intra-fly distance in principal component (PC) space to the mean inter-fly distance; by a resampling approach in which the fly identity labels on the trial data were shuffled randomly and the intra-fly distances were re-calculated. The fraction of 50,000 trials in which the re-sampled mean intra-fly distance is less than or equal to the observed intra-fly distance is the P-value for this method, and was found to be in very close agreement with the t-test results. The intra-fly trials were not

significantly clustered when total percentage of time in each behaviour was considered ($P = 0.39$ and 0.11 by t-test and resampling, respectively; Fig. 4c), but were clustered when the transition rates between behaviours were considered ($P = 0.0022$ and 0.0011 by t-test and resampling, respectively; Fig. 4e).

The average position of each leg (Fig. 4f) was determined across behavioural transitions by identifying target behavioural transitions from the sequence of KNN-classified behaviours, and then averaging leg positions in a ± 50 frame window flanking the transitions, across all occurrences of the transition.

References

1. Benzer, S. Behavioral mutants of *Drosophila* isolated by countercurrent distribution. *Proc. Natl Acad. Sci. USA* **58,** 1112–1119 (1967).
2. Quinn, W. G., Harris, W. A. & Benzer, S. Conditioned behavior in *Drosophila melanogaster*. *Proc. Natl Acad. Sci. USA* **71,** 708–712 (1974).
3. Strauss, R. & Heisenberg, M. Coordination of legs during straight walking and turning in *Drosophila melanogaster*. *J. Comp. Physiol. A* **167,** 403–412 (1990).
4. Dankert, H. *et al.* Automated monitoring and analysis of social behavior in Drosophila. *Nat. Methods* **6,** 297–303 (2009).
5. Branson, K. *et al.* High-throughput ethomics in large groups of *Drosophila*. *Nat. Methods* **6,** 451–457 (2009).
6. Mendes, C. *et al.* Quantification of gait parameters in freely walking wild type and sensory deprived *Drosophila melanogaster*. *eLife* **2,** e00231 (2013).
7. Bender, J. *et al.* Computer-assisted 3D kinematic analysis of all leg joints in walking insects. *PLoS One* **5,** e13617 (2010).
8. Petrou, G. & Webb, B. Detailed tracking of body and leg movements of a freely walking female cricket during phonotaxis. *J. Neur. Methods* **203,** 56–68 (2012).
9. Buchner, E. Elementary movement detectors in an insect visual system. *Biol. Cybern.* **24,** 85–101 (1976).
10. Seelig, J. *et al.* Two-photon calcium imaging from head-fixed *Drosophila* during optomotor walking behavior. *Nat. Methods* **7,** 535–545 (2010).
11. Clark, D. A. *et al.* Defining the computational structure of the motion detector in *Drosophila*. *Neuron* **70,** 1165–1177 (2011).
12. Gaudry, Q. *et al.* Asymmetric neurotransmitter release enables rapid odour lateralization in *Drosophila*. *Nature* **493,** 424–428 (2013).
13. Song, X. & Foley, J. W. Dyes, compositions and related methods of use. *PCT Int. Appl.* (2010; WO 2010054183 A2 20100514).
14. Kanitz, A. & Hartmann, H. Preparation and characterization of bridged naphthoxazinium salts. *Eur. J. Org. Chem.* **4,** 923–930 (1999).
15. Yamaguchi, S. *et al.* Contribution of photoreceptor subtypes to spectral wavelength preference in *Drosophila*. *PNAS* **107,** 5634–5639 (2010).
16. Belongie, S. *et al.* Monitoring animal behavior in the smart vivarium. In Workshop on Measuring Behavior (Wageningen, The Netherlands, Citeseer, 2005).
17. Veeraraghavan, A., Chellappa, R. & Srinivasan, M. Shape-and-behavior encoded tracking of bee dances. *IEEE Trans. Pattern Anal. Mach. Intell.* **30,** 463–476 (2008).
18. Jhuang, H. *et al.* Automated home-cage behavioral phenotyping of mice. *Nat. Commun.* **1,** 1–9 (2010).
19. Burgos-Artizzu, X. *et al.* Social behavior recognition in continuous video. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 1322–1329 (2012).
20. Bishop, C. M. *Pattern Recognition and Machine Learning* (Springer, New York, 2007).
21. Chen, S. *et al.* Fighting fruit flies: a model system for the study of aggression. *Proc. Natl Acad. Sci. USA* **99,** 5664–5668 (2002).
22. Schuett, W. *et al.* Personality variation in a clonal insect: the pea aphid, *Acyrthosiphon pisum*. *Dev. Psychobiol.* **53,** 631–640 (2011).
23. Kain, J. *et al.* Phototactic personality in fruit flies and its suppression by serotonin and white. *Proc. Natl Acad. Sci. USA* **109,** 19834–19839 (2012).
24. Kalmus, H. Optomotor responses in *Drosophila* and Musca. *Physiol. Comp. Ocol. Int. J. Comp. Physiol. Ecol* **1,** 127–147 (1949).
25. Brand, A. & Perrimon, N. Targeted gene expression as a means of altering cell fates and generating dominant phenotypes. *Development* **118,** 401–415 (1993).
26. Nakai, J. *et al.* A high signal-to-noise Ca(2 +) probe composed of a single green fluorescent protein. *Nat. Biotechnol.* **19,** 137–141 (2001).
27. Ekimov, A. & Onushchenko, A. Quantum size effect in three-dimensional microscopic semiconductor crystals. *JETP Lett.* **34,** 345–349 (1981).

Acknowledgements

We thank D. Cox and N. Pinto for guidance with non-linear classifiers, and S. Buchanan for comments on the manuscript. J.K., C.S., X.S., J.F. and B.d.B. were supported by the Rowland Institute. Q.G. was funded by a research project grant from the National Institute on Deafness and Other Communication Disorders (R01 DC-008174). R.W. is an HHMI Early Career Scientist. We thank four anonymous referees for careful comments that significantly improved the manuscript.

Author contributions

J.K. and B.d.B. conceived the study, fabricated the apparatus and performed the experiments. Q.G. and R.W. helped with early versions of the spherical treadmill. C.S. helped fabricate the apparatus. X.S. and J.F. synthesized the dyes. All authors helped prepare the manuscript.

Additional information

Supplementary Information accompanies this paper at http://www.nature.com/naturecommunications

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at http://npg.nature.com/reprintsandpermissions/

How to cite this article: Kain, J. *et al.* Leg-tracking and automated behavioural classification in *Drosophila*. *Nat. Commun.* 4:1910 doi: 10.1038/ncomms2908 (2013).

Supplementary Information for:

Leg-tracking and automated behavioral classification in *Drosophila*

Jamey Kain[1], Chris Stokes[1], Quentin Gaudry[2], Xiangzhi Song[3], James Foley[1], Rachel Wilson[2], Benjamin de Bivort[1,4,5]

[1]The Rowland Institute at Harvard, Cambridge, Massachusetts, USA. [2]Department of Neurobiology, Harvard Medical School, Boston, Massachusetts, USA. [3]College of Chemistry & Chemical Engineering, Central South University, Changsha, P. R. China. [4]Center for Brain Science, Harvard University, Cambridge. Massachusetts, USA.

[5]Department of Organismic and Evolutionary Biology, Harvard University, Cambridge Massachusetts, USA.

Correspondence should be addressed to B.D. (debivort@rowland.harvard.edu).

CONTENTS

SUPPLEMENTARY FIGURE

Figure S1 (de Bivort et al.)

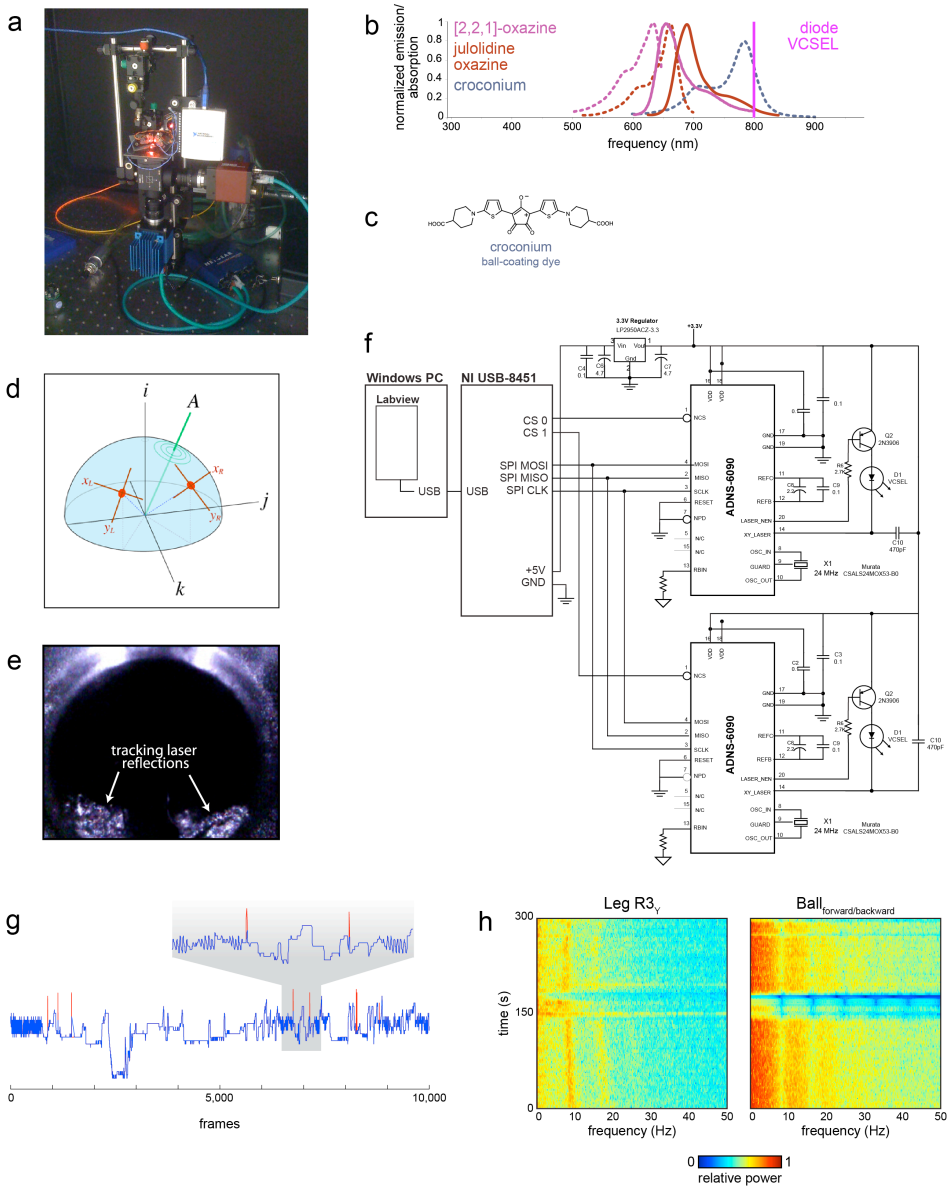

Supplementary Figure S1| Additional apparatus details.

(a) Photograph of the leg-tracker setup.

(b) Absorption and emission wavelengths as in Figure 1e, with the optional croconium ball dye (see Supplementary Methods) included.

(c) Chemical structure of the croconium dye.

(d) Motion of the ball around its axis of rotation (A) is described by a basis set of three component rotational axes (i, j, k), and inferred from the ball's motion at the tracking spots (red).

(e) The appearance of the floating ball in infrared as illuminated by the tracking sensor lasers.

(f) Wiring diagram allowing the Avago ADNS6090 tracking sensors to operate once extracted from their original application in computer mice, as well as communicate with the SPI hub.

(g) A representative vector data trace before (red) and after (blue) clean up using a median filter and removal and interpolated replacement of rare erroneous readings. See Supplementary Methods.

(h) Sphere motion is sufficiently sensitive to respond to single strides, as shown by time-frequency analysis of an individual leg vector (left) and the pitch component of ball motion (right). Color indicates relative power. The ~9 and ~18Hz leg-motion modes drive corresponding modes in the motion of the ball.

SUPPLEMENTARY METHODS

Apparatus design details and use

Previous reports of *Drosophila* spherical treadmills used styrofoam, high-density polyethylene, or polyurethane foam spheres [9, 10, 12, 28]. This was not an option for our setup because we needed an optically clear sphere for imaging the dye spots of the leg from below. At first we thought we would need a spherical treadmill that matched the inertial mass of the tiny fruit fly. This led us to try aerogel, a rigid, ultralight and translucent material produced through supercritical drying [29]. However, it was time-consuming and difficult to make uniform spheres of aerogel, the material absorbed infrared and was very hygroscopic.

Next, we tried clear, acrylic spheres (1/4'', McMaster-Carr), which were much heavier than a fruit fly, but considerably more spherical than our in-house manufactured aerogel spheres. We were able to drastically reduce the mass of the sphere (down to <7mg) by spliting them in half, hollowing them out using Dremel rotary tools, and then glueing them back together. However, we found that the spheres were too light, and that the fly had trouble running because each stride would displace the sphere instead of spinning it. Therefore, matching the inertial mass of the sphere to the fly was not the correct approach and ultimately, we found the clear, acrylic spheres without modification yielded behaviors that, to our eyes, were the most naturalistic.

We introduced small, scattering surface imperfections on the spheres by lightly rolling them over 400 grit sandpaper to allow the infrared sensors to efficiently track motion. The more scuffed the sphere, the better it was tracked, but at the cost of poorer transmission of the fluores-

cence from the leg dyes to the cameras below. Thus, we incrementally scuff a sphere and strike a balance between infrared tracking and fluorescence transmission.

Alternatively, the spheres can be coated in the near-infrared absorbing/reflecting croconium dye (Piperdinium, 4-carboxy-1-[5-[3-[5-(4-carboxy-1-piperdinyl0-2-thienyl]-2-hydroxy-4,5-dioxo-2-cyclopenten-1-ylidene]-2(5H)-thienylidene]-, inner salt) to make it visible to the sensors. The croconium dye was prepared using the procedure reported [30]. However, all of the data presented in this work was collected using the sandpaper method.

The housing of the treadmill sphere bearing consisted of the rim of a 150μL Eppendorf tube. During trials, the ball was floated on air from a cylinder of compressed air (Airgas). The air was first bubbled through water to prevent static charge from accumulating, then passed through an adjustable flow regulator, and finally into a plenum to evenly distribute the flow into the air bearing. We used an airflow rate of approximately 440mL per minute.

Infrared sensors (Avago chip #ADNS-6090) for tracking the sphere are available in high-performance gaming mice. The wiring diagram of the circuit that allows these sensors to operate once removed from their original circuit boards, as well as communicate with a Serial Peripheral Interface (SPI) hub (National Instruments) is shown in Supplementary Figure S1f. (See the section below for details on transforming the four vectors from the two sensors into the three motion vectors of the sphere.) Measurement of angular displacement in the three components of rotational motion of the ball was calibrated by manually turning a calibration sphere (rigidly attached to a long, stiff, wire handle) through π radians on the roll and yaw axes and $\pi/2$ radians on the pitch axis, numerous times and averaging the integrated sensor motion.

The dyes were excited by a HeNe laser (632.8 nm, Thorlabs) which passes through a beam expander to ensure the space surrounding the legs was fully illuminated. In preliminary

experiments, we determined that groups of agitated flies were attracted to red LED light sources centered at 625nm, but not the HeNe source. Since a sphere will act as a lens, a second sphere was positioned within the optics housing to re-collimate the light before reaching the cameras. In conjunction with a beam splitter, two Prosilica GigE cameras (Allied Vision Technologies) were used for imaging the dyed legs from through the spheres, and acquired images at 80Hz typically. Band-pass filters (Edmund Optics) were used to optimize each camera's detection for either 221ox (HQ650/20m band-pass filter, Chroma) or julox (HQ700/20m band-pass filter, Chroma). Additionally, a short-pass filter (NT49-829, Edmund Optics) was used to block stray light from both the infrared tracking sensors and potential two-photon sources.

For the optomotor feedback experiments we used a standard 15 inch LCD monitor to display alternating blue (RGB=0000FF, width=4.8cm) and black (RGB=000000, width=2.4cm) colored vertical bars. Optic flow consisted of these vertical bars translating out (10cm/s) from a line of expansion. The monitor was positioned six inches from the tethered fly. The optomotor feedback trial comprised the conditions: 20min in the dark, 20min closed-loop, 20min dark, 20min open-loop, repeat from the beginning once. Closed-loop allowed the turns of the fly to control the position of the line of expansion for the translating vertical bars based on the detected turning-in-place (i) vector. Open-loop disabled the feedback control from the fly and allowed the point of origin for the diverging vertical bars to drift (0.5cm/s) to the point that the entire screen was filled with bars moving either left-to-right or right-to-left. During the 20min open-loop trial, the drift would switch direction every 5min. The 5 animals recorded in the optomotor experiment were different than the 8 animals recorded for day-to-day idiosyncrasy experiments, and the 5 animals used to train the classifier, as well as several animals recorded for day-to-day analysis that yielded only one run of satisfactory data.

Transforming the tracking sensor data into motion vectors

The X, Y coordinates from the sensors, integrated over 10 frames in a sliding window fashion as a noise reducing measure, were transformed to generate the 3 motion vectors of the sphere as follows:

- The axis of rotation of the ball $\boldsymbol{A}$ has components whose relative proportions are the proportional extents to which the fly is running forward/backward, sidestepping, or turning in place (respectively, pitch, roll and yaw). Specifically, rotation around i indicates turning in place, j forward/backward walking, and k sidestepping. We need to determine a, b, c such that $\boldsymbol{A}=\{ai,bj,ck\}$ - these are the relative proportions of each type of motion.
- Data from the sensors comes as two vectors $\boldsymbol{L}$ and $\boldsymbol{R}$ each in its own coordinate system (basis): $\boldsymbol{L}=\{x_L, y_L,0\}$, $\boldsymbol{R}=\{x_R, y_R,0\}$. First, convert the basis of these vectors to the global ball basis ($B=\{\{i,0,0\},\{0,j,0\},\{0,0,k\}\}$) using a change of basis matrix $\boldsymbol{C}$ derived from the positions of the sensors. Call the converted vectors $\boldsymbol{L}_B$ and $\boldsymbol{R}_B$. These will always be latitude lines with respect to $\boldsymbol{A}$. Therefore $\boldsymbol{A}$ is parallel to a vector which is perpendicular to both $\boldsymbol{L}_B$ and $\boldsymbol{R}_B$. Therefore $\boldsymbol{A}$ is parallel to $\boldsymbol{A}'=\boldsymbol{L}_B \times \boldsymbol{R}_B$.
    - $\boldsymbol{A}'_i$ is the % yaw
    - $\boldsymbol{A}'_j$ is the % pitch
    - $\boldsymbol{A}'_k$ is the % roll.

- From $\boldsymbol{A}$, we don't know the speed at which the ball is turning. This is derived from the length of either $\boldsymbol{L}$ or $\boldsymbol{R}$ as a function of their distance from $\boldsymbol{A}$. We calculate it based on whichever sensor is farther from $\boldsymbol{A}$, as this will be a less noisy measurement. Assume, without loss of generality, that $|\boldsymbol{L}|>|\boldsymbol{R}|$, i.e. that $\boldsymbol{L}$ is farther from $\boldsymbol{A}$. The angle between $\boldsymbol{A}$ and $\boldsymbol{L}$ is given by $\cos^{-1}((\boldsymbol{A}\bullet\boldsymbol{L})/(|\boldsymbol{A}||\boldsymbol{L}|))$. The surface speed of a sphere moving with speed 1 at an angular distance θ from the axis of rotation is $\sin(\theta)$, thus the speed of rotation of the ball (and the magnitude of the vector whose components are the fly's behavioral components) is

$$M = |\boldsymbol{L}|/\sin(\cos^{-1}((\boldsymbol{A}\bullet\boldsymbol{L})/(|\boldsymbol{A}||\boldsymbol{L}|)))$$

- While $\boldsymbol{L}_B \times \boldsymbol{R}_B$ is mathematically parallel to $\boldsymbol{A}$, when $\boldsymbol{R}_B$ is nearly parallel to $\boldsymbol{L}_B$, $\boldsymbol{A}'$ will approach 0, and its orientation will become highly sensitive to noise in the measurement of $\boldsymbol{L}$ and $\boldsymbol{R}$. This occurs when $\boldsymbol{A}$ lies in the plane of the bearing housing, i.e. perpendicular to i. To the extent that $\boldsymbol{L}_B$ and $\boldsymbol{R}_B$ are parallel, we want to use a different formula than their cross product. Since $\boldsymbol{L}_B$ and $\boldsymbol{R}_B$ are parallel, we can replace their information with their mean, and we can let "the extent to which they are parallel" be the cosine of the angle between $\boldsymbol{L}_B$ and $\boldsymbol{R}_B$:

$$\beta = \cos(\cos^{-1}((\boldsymbol{L}_B\bullet\boldsymbol{R}_B)/(|\boldsymbol{L}_B||\boldsymbol{R}_B|))) = (\boldsymbol{L}_B\bullet\boldsymbol{R}_B)/(|\boldsymbol{L}_B||\boldsymbol{R}_B|)$$

- This is the weighting factor given to the term of our final equation when $\boldsymbol{L}_B$ and $\boldsymbol{R}_B$ are parallel. The term itself is simply the cross product of i with the mean of $\boldsymbol{L}_B$ and $\boldsymbol{R}_B$, which by assumption are both perpendicular to $\boldsymbol{A}$. Thus, the noise-resistant ball-basis axis of rotation is given by:

$$\boldsymbol{A}' = \beta*(\boldsymbol{L}_B+\boldsymbol{R}_B)/2 \times \{(|\boldsymbol{L}_B|+|\boldsymbol{R}_B|)/2, 0, 0\} + (1-\beta)*\boldsymbol{L}_B \times \boldsymbol{R}_B$$

Developing the *k*-nearest-neighbors classifier

We implemented feed-forward neural networks and support vector machines (SVM) as potential non-linear classifiers for our instrument data before settling on *k*-nearest neighbors analysis (KNN). Those alternative methods showed promise, but neural networks are hard to interpret and have path-dependent performance variability. Compared to KNN, SVM required more computationally intensive optimization of two parameters (in the case of the most generally useful Gaussian radial bias kernel). Moreover, the performance of SVM was not obviously better than that of the simpler KNN. Importantly, we intend to present our technique to building the classifier as a general approach, and the intuitive underpinnings, computational simplicity, and successful performance of KNN were compelling. The classifier was built as follows.

A training set was built by manually scoring 4000 or 8000 video frames of a fly behaving on the ball while simultaneously recording the 15 raw data vectors. This was done independently by two investigators to assess the variability in manual scoring. Five different flies were scored.

Data were recorded as fast as possible during experiments, resulting in varying inter-frame intervals, with an average recording rate of ~90Hz. Prior to any additional processing, raw data was linearly interpolated to a uniform rate of 100Hz. The videos acquired for training purposed were recorded at 30Hz, and therefore the frames used in conjunction with the manual scores were selected for their correspondence with the time stamps of the video frames. Illumination of the HeNe laser was used as a trigger to synchronize the video and instrument data.

Raw data contained rare frames with erroneous leg-track positions. These typically arose when the dye spots moved out of the HeNe laser excitation, e.g. if they were occluded by the ab-

domen. Any trial with more than ~1% error frames was discarded. Raw data for the KNN training, and all other experiments, was filtered to remove error positions as follows. A 3-frame median filter was applied to all raw data vectors eliminating single-frame errors. The first frames of errors persisting longer than a single frame were identified by their values deviating from adjacent frames by more than 5 standard deviations of the inter-frame differences. The final frames of these extended errors were identified as the next frame within 5 standard deviations of the value prior to the error, or within 1 standard deviation of the median position across all frames. Values thus flagged as erroneous were replaced with values interpolated linearly from flanking frames. This approach was conservative, but removed essentially all conspicuous errors (Fig. S1g).

To explore which higher-order features were required for sufficient KNN performance, we used cross-validation at the level of flies – i.e. setting aside all of one fly's data, training the KNN classifier with the raw values and higher-order features from the remaining four flies from the training set. Correlation coefficient was used as the KNN metric. Performance was assessed using the withheld fly's raw data and higher-order features. Data from different flies were made comparable by *z*-score normalization within the data from each fly (to eliminate any positioning differences between mountings), then concatenated and *z*-score normalized again (to assure that all training variables would contribute equally to the classifications). The data from the withheld fly was z-score normalized, and then re-normalized using the mean and standard deviation of concatenated training set data.

Performance of the KNN with respect to one set of manual scores was assessed using the unbalanced accuracy, i.e. the overall percent of classifications matching the manual score. Also evaluated for reference was a so-called "plausible accuracy" which tolerates errors in assignment

between behaviors that are only arbitrarily distinguishable (e.g. the precise distinction between standing still and making small postural adjustments is necessarily an arbitrary cutoff). Misclassifications tolerated in this analysis include the above, postural adjustment with the locomotory behaviors, the locomotory behaviors with complex motion, leg1- and head-grooming, and leg3- and abdomen-grooming. We have observed that flies will groom both their hind-legs and abdomen in the same stroke, so this reflects biological overlap in addition to definitional continuity. The plausible accuracy between the manual scoring data sets was 96%, and between the KNN classifier scores and manual (JK) scores 91%. Manual (BD) scores were slightly more difficult for KNN to classify, yielding an 88% plausible accuracy.

Higher-order features characterizing local temporal dynamics in the data were calculated with a sliding window of +/-5frames. Other window sizes were evaluated and did not improve performance. Vector derivatives were calculated as the difference of values from frame t-1 to t+1. Local cross-correlations were calculated in a pair-wise fashion between all 15 raw values in a sliding window of +/-5frames. Vectors of higher-order features were padded with 0s as necessary. The KNN k parameter was optimized to 16 using just the raw data, and re-optimized to 24 once the best training set (raw data + derivatives + local standard deviations) was identified. KNN performance was quite robust to the choice of $4<k<100$. Lastly we noticed that the model was prone to switch between behavioral types more frequently than the manual scorers. We explored two approaches to address this. First we implemented a sequential frame Bayesian hidden Markov Model with priors extracted from the training data set. This provided little improvement and thus was not kept. However, a simple low-pass filter that replaced each frame's classification with the most abundant classification in a +/-5frame sliding window did improve performance, albeit slightly (5%).

Raw data collected post-KNN-training was interpolated, filtered for errors, normalized, augmented with higher-order features, and low-pass filtered as above, with the exception that the training samples included data from all 5 training flies, i.e. none were withheld.

SUPPLEMENTARY REFERENCES

28. Gotz, KG. & Wenking, H. J. Visual control of locomotion in the walking fruitfly Drosophila. Comp. *Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* **85**, 235-266 (1973).

29. Kistler, SS. Coherent expanded aerogels and jellies. *Nature* **127**, 741 (1931).

30. Song, X. & Foley, JW. A new water-soluble near-infrared croconium dye. *Dyes and Pigments* **78**, 60-64 (2008).