

**Discriminating external and internal causes for heading changes
in freely flying *Drosophila***

Andrea Censi^{1#}, Andrew D. Straw^{2#}, Rosalyn W. Sayaman³, Richard M. Murray¹, Michael H. Dickinson^{4*}

1 Control & Dynamical Systems, California Institute of Technology, Pasadena, California, USA

2 Research Institute of Molecular Pathology (IMP), Vienna, Austria

3 Division of Biology, California Institute of Technology, Pasadena, California, USA

4 Department of Biology, University of Washington, Seattle, Washington, USA

These authors contributed equally to this work.

* E-mail: flyman@uw.edu

Copyright: © Censi *et al.* This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

As animals move through the world in search of resources, they change course in reaction to both external sensory cues and internally-generated programs. Elucidating the functional logic of complex search algorithms is challenging because the observable actions of the animal cannot be unambiguously assigned to externally- or internally-triggered events. We present a technique that addresses this challenge by assessing quantitatively the contribution of external stimuli and internal processes. We apply this technique to the analysis of rapid turns (“saccades”) of freely flying *Drosophila melanogaster*. We show that a single scalar feature computed from the visual stimulus experienced by the animal is sufficient to explain a majority (93%) of the turning decisions. We automatically estimate this scalar value from the observable trajectory, without any assumption regarding the sensory processing. A posteriori, we show that the estimated feature field is consistent with previous results measured in other experimental conditions. The remaining turning decisions, not explained by this feature of the visual input, may be attributed to a combination of deterministic processes based on unobservable internal states and purely stochastic behavior. We cannot distinguish these contributions using external observations alone, but we are able to provide a quantitative bound of their relative importance with respect to stimulus-triggered decisions. Our results suggest that comparatively few saccades in free-flying conditions are a result of an intrinsic spontaneous process, contrary to previous suggestions. We discuss how this technique could be generalized for use in other systems and employed as a tool for classifying effects into sensory, decision, and motor categories when used to analyze data from genetic behavioral screens.

Author summary

*Researchers have spent considerable effort studying how specific sensory stimuli elicit behavioral responses and how other behaviors may arise independent of external inputs in conditions of sensory deprivation. Yet an animal in its natural context, such as searching for food or mates, turns both in response to external stimuli and intrinsic, possibly stochastic, decisions. We show how to estimate the contribution of vision and internal causes on the observable behavior of freely flying *Drosophila*. We developed a dimensionality reduction scheme that finds a one-dimensional feature of the visual stimulus that best predicts turning decisions. This visual feature extraction is consistent with previous literature on visually elicited fly turning and predicts a large majority of turns in the tested environment. The rarity of stimulus-independent events suggests that fly behavior is more deterministic than previously suggested and that, more generally, animal search strategies may be dominated by responses to stimuli with only modest contributions from internal causes.*

36

37

Introduction

38 Active movement is one of the defining features of animals, and the use of locomotion to
39 search for resources within the environment is likely among the most ancient of behaviors.
40 Observations on motile organisms, ranging in scale from bacteria to whales, indicate that
41 search patterns are structured by a combination of internal processes and external cues [1,2].
42 Sensory systems enable organisms to detect favorable objects at a great distance [3–5] and
43 they use this ability to localize resources by either directed motion (taxis) or changes in
44 locomotor statistics (kinesis). Prior research suggests that, in the absence of external cues, the
45 animal behavior is generated by internal processes, and that the overall animal fitness is
46 sensitive to the exact characteristics of this internal process (e.g., Levy statistics) [6–18]; it has
47 also been questioned whether observed large-scale statistics can give any insight on an internal
48 process that generated the behavior, and whether the internal processes can dominate over
49 stimuli-elicited behavior [19–22]. As for the internal processes, these can be divided into truly
50 stochastic sources, and deterministic results of a deliberate, but unobservable, internal
51 mechanism based on internal metabolic/neural states. When observing an intact motile
52 organism, it is not easy to determine which components of its locomotion behavior are
53 triggered by internal processes versus external cues, yet such classification is essential for
54 deciphering the underlying logic of its movement and search behavior. The task is further
55 complicated by the fact that an external observer might not be able to distinguish between
56 truly stochastic processes and the deterministic results of a deliberate, but unobservable,
57 internal mechanism. For example, software pseudo-random number generators produce
58 strictly deterministic sequences, which appear to be random to an external observer who does
59 not have access to the internal state of the system [23]. A major goal of both cell biology and
60 neuroscience is explaining the molecular and cellular bases of these three qualitatively different
61 processes (sensory-driven, purely stochastic, and deterministically based on internal states).

62 If the salient features of the external world are known, it is possible to gain insight into sensory-
63 driven behaviors through the use of sensory-response correlation [24]. The analysis of the
64 internally-driven processes is much more challenging. Given uncertainty in measurement and
65 the inability to perfectly reproduce experimental conditions from trial to trial, variability in the
66 results of behavioral experiments has often been treated as a limit on our ability to measure
67 stimulus-driven behavior. In this view, variability in responses from trial to trial reflects
68 irrelevant components of behavior, which are averaged until the mean---interpreted as the
69 response the animal ideally would have produced---becomes clear [25]. From the opposite
70 perspective, many researchers have attempted to artificially remove all relevant sensory input
71 to an animal and measure behaviors in conditions of sensory deprivation to reveal intrinsic
72 properties, especially the statistical distributions of behaviors [26–28]. Although focusing in

isolation on either the stochastic [9,29,26] or the sensory components [30] of search behavior have provided key insights, neither of these extremes is sufficient to capture the full range of processes at play as an animal moves under natural conditions. Attempts to investigate the interaction of internal and external processes include studies of bacteria [31] and nematode worms [32,33], organisms for whom chemicals provide the most salient cues for food search. For larger animals with image-forming eyes, vision may provide another essential cue in search algorithms, because vision is the only sense which allows to perceive remote parts of the environment. Often vision cannot be considered separately from the mechanics of locomotion [34].

Flies are a model of computational efficiency and robustness, to date not equaled by artificial systems, which often seek to imitate nature [35,36]. Much is known about fly vision [37,38]. Since the pioneering work of Kennedy [39] and Mittelstadt [40], the behavioral responses of flies to experimenter-defined visual stimuli have been extensively investigated. Electrophysiological recordings have complemented and extended our knowledge of phenomena such as the neural basis of motion detection [41–44] and other key aspects of sensory processing, such as receptive field tuning [45]. However, there are many challenges in the identification of neural processing and how it produces complex behavior, especially as regards the characterization of “discrete” behaviors, such as the rapid turns (“saccades”) of *Drosophila*, which are the object of this study. In fact, many studies which offered complete characterization of the animal response are limited to “continuous” behavior, for which they provide linear (or “linearized”) models [46–49]; this allows using techniques such as linear system identification.

Identifying the neural causes for “discrete” behavior involves solving a different set of problems. Firstly, there are the problems of segmentation and classification of behaviors (including the definition of what “behavior” and “a behavior” are), for which it is often necessary the use of nonlinear machine learning methods [50]. Then, there is the problem of building models that can correlate the stimulus with the behavior(s). While it is possible to postulate models that also integrate well with our understanding of lower-level behavior [51,52], it is not clear how such methods can be identified from the data. On the practical side, it is evident that discrete decisions, such as turning decisions, are meant to guide exploration and therefore should be investigated in naturalistic situations. This poses practical problems of tracking the animal position in a large environment, and it also precludes (at the current level of technology) the uses of direct neural recording. In fact, comparatively few attempts have been made to correlate parameters of visual stimulus with behavioral responses in unrestrained conditions [53–56].

In this work, we present an analysis that can quantitatively discriminate the effect of visual stimulus as opposed to internal processes in the generation of saccades in the fruit fly. Our conclusions are that visual stimulus has a dominant role. One important message of our work is that it is very difficult to identify models of complex behavior that can explain everything, often because insufficient data can be collected. Therefore, it is important to “search for simplicity” [57], for example by framing the problem as dimensionality reduction, and to use models that a posteriori can justify their assumptions. While we describe this analysis for visual processing in *Drosophila*, our goal is to construct a general method that can be used for other sensory systems, other animal species, or in the context of genetic screens.

Methods

Fly care and experimental treatment

Flies from the laboratory stock derived from 200 wild-caught females were reared on a 16h:8h light dark cycle under standard laboratory conditions. Three day old adult female flies were anesthetized with cold and individually housed within centrifuge tubes containing a moist tissue paper. Flies were starved (but provided with water) in the tubes for four to six hours before being released into the flight arena. Most flies would immediately begin flying, and we terminated tracking after the fly landed. We then removed each fly with a wand attachment of a vacuum cleaner before introducing another fly. Thus, each recorded trajectory is derived from a fly’s initial experience exploring the novel environment.

The flight arena was a 2 meter diameter, 80cm high cylinder (see Figure 1A). 10cm x 10cm red and green gel filters (Roscolux) were attached to the arena in a regular checkerboard arrangement and provided a high contrast visual stimulus to flies near the wall. One meter from the wall (i.e., at the center of the arena), the angular wavelength of this pattern was $\sim 11^\circ$, and consequently would be twice the inter-ommatidial spacing of a $\sim 5.5^\circ$ in *Drosophila* [58]. The particular red and green filters were chosen to have similar infrared transmission to facilitate tracking using cameras outfitted with long (IR) pass filters. The arena was illuminated from outside with a circular array of eight 750W Fresnel stage lights pointing towards the arena center. These lights provided both visible and infrared light for fly visual responses and machine vision tracking, respectively.

Fly tracking

A detailed description of our tracking system may be found in [59]. Briefly, we used 11 cameras (6 monochrome Pt. Grey Firefly MV USB cameras and 5 monochrome Basler A602f cameras) with wide-angle lenses and infrared pass, visible cut filters (R72, Hoya Filters) to view the interior volume of the flight chamber. The cameras were positioned so that a fly within the tracking volume was viewed by 2 or more cameras at any given time, enabling a 3D estimate of

its position (Figure 1Bi). The cameras were first calibrated to compensate for image warping non-linearities (deviations from the pinhole model) and then the extrinsic and intrinsic parameters describing the pinhole model were found. Flies were tracked with an extended Kalman filter (EKF), in which the motion model was a linear constant velocity model, and fly maneuvering is captured by the stochastic component of the Kalman filter. Because tracking updates occurred at a high rate (60 fps) relative to fly maneuvering, we found this simplification to work well in practice. The 3D estimate of the fly position is recovered by triangulation from the 2D tracking data of each camera, and taking into account the relative uncertainty of each observations.

Saccade detection

Many species of flies, including *Drosophila*, exhibit rapid changes in heading as they fly, termed “saccades” [53]. Between saccades, flies tend to maintain an approximately straight course, and saccades account for at least 80% of the total net change in heading during flight [60]. There is little doubt that saccades can be triggered by visual stimuli, but the degree to which visual feedback plays a role in determining the velocity, duration, and amplitude of the resulting turn is unclear. Experiments using a magnetic tether, which permits free rotation about the yaw axis, suggest that flies do not respond to visual feedback during a saccade [61]. On the other hand, Stewart *et al.* [56] have observed a rebound effect after saccades in free flight, which they suggest is consistent with active optomotor feedback during the maneuver. This discrepancy is not of direct interest here, however, as we deal exclusively with the decision of *initiating* a saccade.

To analyze saccades within a flight trajectory, one should choose a detection algorithm that, given the trajectory data, returns a series of saccade events, possibly with other attributes such as direction, amplitude, velocity, etc. In the past, several detection algorithms have been proposed, each one implicitly using a slightly different definition of saccade, and each one able to compensate for different sources of noise. In practice, large saccades are such distinct events that all algorithms agree with respect to most classifications, but different algorithms may disagree on detection of small saccades. We make sure that our results are robust to the choice of the algorithm, by using two distinct algorithms based on different principles. The two algorithms are described in detail in Text S1 and their source code is available on line. Briefly, the Geometric Saccade Detector (GSD) detects saccades from the x-y planar trajectory. The Angular-Velocity based Saccade Detector (AVSD) works primarily by considering the smoothed angular heading rather than the planar position. Unless otherwise noted, the statistics shown through the paper are derived using GSD, which is *a posteriori* shown to be better suited for these particular experimental conditions and equipment. Alternative figures showing the same statistics obtained from the AVSD algorithm are available as part of Text S1.

Capturing behavior determinism and randomness using rate-variant Poisson processes

Figure 2A illustrates the conceptual approach of our analysis. We denote by $x(t)$ the animal's physical *spatial configuration* (its position and velocity in a fixed reference frame). The stimulus $y(t)$ is the set of all sensory cues perceived by the animal, and it is a function of both the spatial configuration $x(t)$ and the appearance of the world W . Whereas $x(t)$ is a concrete variable that we can possibly measure, the stimulus $y(t)$ and the world W are placeholders for things that, in general, are unknown. The actions $u(t)$ (e.g. saccades in our case) are the external manifestations of the internal neural processing, which depend both on the instantaneous stimulus as well as on $\xi(t)$, another placeholder variable that represents the animal's internal state (metabolic states, neural states, etc.), and which has dynamics of its own. We assume that it is possible to observe the spatial configuration $x(t)$ as well as infer the actions $u(t)$ from the observations, but that the internal state $\xi(t)$ is not observable.

We make a distinction between obtaining a functional model of an animal's behavior and identifying the underlying neural processes. Obtaining a functional description of behavior means obtaining a model that can predict the actions $u(t)$ given the spatial configuration $x(t)$ and a description of the world W . In principle, we can do this by observing an animal's behavior with enough samples of $x(t)$, W and $u(t)$. In general, however, there are a variety of neural models that could produce the same functional model. For example, many behaviors appear to be well-localized in time, suggesting an "action potential" neural model, but the underlying neural model can have very different properties [62] (in other words, the microscopic explanation might be quite different than what the macroscopic observations suggest). The model that we now describe and that we will identify should be interpreted as a purely functional model, which can inform the search for neural models, to make sure that they are compatible with the externally observable free flight behavior.

Figure 2B shows the particular model that we use in this paper. It is a particular form of the general model discussed above (Figure 2A). In this model, we propose that the animal's actions $u(t)$ can be summarized by the saccade events. We divide the saccade events in two classes: left and right saccades. In principle, one would want to consider additional attributes of the saccades, such as speed, duration, and amplitude. The analysis might also be expanded to consider other easily identifiable events [63]. However, limiting ourselves to a binary characterization of saccades allows us to model the behavior generation as Poisson processes, which offers relatively easy inference. We model saccade generation using rate-variant Poisson processes, i.e., we assume that, for each class of events, internal and external factors influence a time-varying event rate according to a quantitative relation that we will attempt to identify.

The most important assumption of our method (which can and will be verified *a posteriori*) is that, for the purpose of generating the behavior, the high-dimensional output $y(t)$ can be

compressed down to a low dimensional “feature” $z(t)$. This assumption is implicit in many other previous studies, and it is informed by the knowledge of the underlying neurobiology: the first level of sensory processing in flies and other animals consists in taking a very high-dimensional sensory stream and computing the few behaviorally-relevant features from it. Our only assumption is that this low-dimensional feature exists - we do not assume that we know this feature. However, we can attempt to automatically identify this feature from the observable data. It is important to note that we do not assume to know how this feature is computed from the stimulus. Indeed, the advantage of our method is that it allows identifying this feature based only on the observable behavior, without postulating anything on the sensory processing.

Figure 2B also shows explicitly that, in addition to the feature-dependent pathway in our model, other unmodeled processing influences the behavior. The effect of this unmodeled processing will be quantitatively estimated as well. The saccade events are assumed to be generated by a set of interacting Poisson process with variable rate $r_i(t)$, $i \in \{L, R\}$. The index i stands for either one of the two classes of events (L: left, R: right). The variable rate $r_i(t)$ is assumed to depend both on the stimulus $y(t)$ and the internal state $\xi(t)$, thus incorporating both random and deterministic effects. We write $r_i(t)$ as the sum of three factors:

$$r_i(t) = f_i(z(t)) + r_i^I(\xi(t)) + r_i^R, \quad (1)$$

where the term $f_i(z)$ is the contribution of the external stimulus through the feature z ; the term $r_i^I(\xi)$ is the contribution of the internal state ξ ; and the term r_i^R represents the contribution of a purely random stochastic process that does not depend either on an internal state or the stimulus. By omitting some of the terms in the equation above, one can recover many other simpler models. For example, purely random behavior is obtained by setting $r_i(t) = r_i^R$.

The Poisson processes interact by inhibition. If any process generates an event, then any event generated from that process or any other process for a period of length Δ is ignored. This is meant to model a feature of many fixed action patterns that, once initiated, must run to completion before a different motor program can be initiated.

Finally, Figure 2B shows another variable $c(t)$, which we call “reduced configuration”. We define $c(t)$ as the subset of the spatial configuration variables that actually influence the stimulus, for a particular class of environments W^* . In general, for a freely flying animal, $x(t)$ is at least a 12 dimensional quantity, including the 6 degrees of freedom for position/orientation and the corresponding 6 for velocities (additional degrees of freedom in the animal spatial configuration would be derived from the positions of body joints, such as the neck and wing

positions). For particular environments, however, the stimulus is only dependent on a subset of $x(t)$. For example, if the environment is distant enough, then the visual stimulus does not depend on the forward velocity. Therefore, even though the spatial configuration $x(t)$ is at least 12-dimensional, actually the stimulus depends on a smaller variable $c(t)$, i.e., the reduced configuration.

We will show that it is possible to identify all unknowns in this model. In particular, we will identify how the feature z depends on the reduced configuration c , and how the rates depend on the feature. Remarkably, it is possible to do this without assumptions on how the feature z is computed from the stimulus y or how the stimulus y depends on the reduced configuration c . We only assume to be able to observe the reduced configuration $c(t)$ and the generated saccade events. Before describing the method, we first discuss how this model based on rate-variant Poisson processes allows us to represent different functional models.

Predictions of different functional models

In **Figure 3** we illustrate the predictions of four qualitatively different functional models in terms of the observed statistics. On the left side we show the functional model, and on the right we show the expected observed event rates $f_i(z)$ as a function of the feature z . This exercise assumes that we know how to estimate the feature, which we will show later. Here we describe what we would expect to find, before embarking on the actual computation of z .

Figure 3A shows a “hard threshold” model, based on the computation of a single feature z , which is then thresholded to obtain the event rate. A Poisson process then generates the events based on this time variant rate. The “stochastic trigger” in the figure masks the fact that there are two processes generating two classes of events, and that these processes are interacting (see discussion above), which is not relevant to the present discussion. If the absolute value of the feature is below a threshold, no event is generated; otherwise, saccades to the left and right are generated at a fixed rate. A large fixed rate would mean that the model is practically deterministic, with a large stimulus feature z resulting in a behavioral event with only rare failures. On the right side of the figure, we show the observed event rates as a function of the feature z . For this simple model, the observed rates as a function of the feature are straight steps. We remark that our analysis does not assume necessarily that the feature exhibits a hard threshold as in this simple model. We choose this shape merely because it allows visualizing the effect of different sources of noise.

In particular, we are interested in understanding the implications of a noise source that acts on the computation of the feature (sensory noise) compared to noise that generates behavior in a parallel process independent of the stimulus-computed feature (decision-making or motor noise).

Figure 3B shows the effect of measurement noise on the hard threshold model. Random fluctuations in the feature turn the hard threshold into a *soft* threshold.

Figure 3C shows the effect of adding a spontaneous generation process in parallel to the feature pathway. This has the effect of raising the predicted event rate by a constant value, as the parallel process is independent of the feature. A parallel generation process that depended on an unobservable internal state would have the same expected statistics if the internal state is uncorrelated with the feature. This means that a constant baseline event rate that is independent of the feature must be interpreted as the joint contribution of a purely stochastic spontaneous event generation together with a deterministic response based on internal states.

It is also important to consider the effect of another unmodeled feature z' on the event rate statistics, if we only model the dependence of one feature z . This stems primarily from practical concerns, because the dimensionality of the feature that it is possible to identify depends primarily on the amount of data available. Therefore, once the dimensionality of the feature is fixed, we need a way to judge whether that dimension is sufficient to describe the behavior.

Figure 3D augments the model of **Figure 3A** with an additional pathway that uses a different feature z' . In such a case, if we plot the rates versus the feature, we will not find a clear functional dependency, indicating that the feature z is no longer sufficient to explain the event rates. Conversely, if we find a clear functional dependency, then we can say that the feature z is sufficient to capture the influence of the sensory stimulus on the behavior. This does not imply that z is the only behaviorally relevant feature of the stimulus, because there could be other features that are relevant for other behaviors not considered in the analysis.

Our identification algorithm, described in the next section, recovers the best one-dimensional feature z that explains the event rates. This permits constructing a function in which the experimental event rate is plotted against the feature curve. However, we anticipate that the experimental results, being dependent on experimental data, will have error bars both for dependent and independent variables. Strictly speaking, even if one finds a one-dimensional feature that uncovers a deterministic dependency between feature and rates compatible with the error bars, it is not possible to conclude that there is only one feature, because the effect of a second feature might be masked by the measurement noise. In this sense, our claims that one feature is sufficient is an application of parsimony.

In summary, we can identify the contributions of several qualitative factors by plotting the event generation rates as a function of z . Measurement noise will soften the curve (e.g., a hard threshold is turned into a soft threshold). A parallel purely stochastic event generation process has the same effect of a deterministic process based on an internal state uncorrelated with the feature, namely it raises the curve by a fixed baseline rate independent of z . If another

unmodeled feature z' influences the behavior, there is not a strict functional dependence between the rates and the feature z .

Identification of the feature z

We devised a procedure that obtains an estimate of the best one-dimensional feature of the input that predicts the observed event rates. We explain here the basic idea, and provide details in Text S1. Intuitively, the feature and event rates can be obtained from the spatial statistics of the observed behavioral output. With respect to the discussion so far, the main conceptual step consists in translating the problem from the time to the space domain. So far, we have written the feature $z(t)$ as a time-varying quantity. We have also assumed that $z(t)$ depends on the stimulus $y(t)$, and that the stimulus depends on the animal spatial configuration $x(t)$, or more precisely, on the reduced configuration $c(t)$. Therefore, we rewrite our model writing $z(c)$ instead of $z(t)$. The quantity $z(c)$ is a spatial field that we interpret as the feature computed from the typical stimulus experienced at the reduced configuration c . We will fit a model of the kind:

$$r_i(c) = f_i(z(c)) + r_i^0, \quad i \in \{L, R\} \quad (2)$$

where $r_i(c)$ is the average event rate for the i -th class (L : left, R : right) observed at the reduced configuration c ; $f_i(z)$ denotes the event generation rates for left and right saccades as a function of the feature, and r_i^0 is constant term that we call *baseline event rate*.

Note the differences with respect to the previous model (Eq. 1). First, we have written the rates as a function of the reduced spatial configuration instead of time. Moreover, we do not model explicitly the contribution of the internal state. As argued above, given that we cannot measure the unobservable internal states $\xi(t)$, we cannot distinguish between a purely stochastic contribution and the contribution of an internal state

Therefore, the constant term r_i^0 will be an estimate of the joint contribution of the two terms that we cannot distinguish:

$$r_i^0 = r_i^R + \mathbb{E} \{ r^I(\xi(t)) \}, \quad (3)$$

where $\mathbb{E} \{ \quad \}$ indicates the expected value taken over the whole trajectory.

We summarize here the three main phases for estimating $z(c)$ from the behavioral data, while leaving the details to Text S1. First, the reduced configuration space $c(t)$ is discretized into spatial cells with a resolution that depends on the amount of data available. For each of these

cells, basic statistics are computed, such as the average time spent in each cell, as well as the observed event rates in the cell. One advantage of the algorithm is that these spatial statistics, averaged over the whole trajectory, are intrinsically robust to measurement noise and uncertainty in the event detection algorithm. Next, the event *generation* rates $r_i(c)$ are computed from the *observed* rates. Because we assume that the Poisson processes interact with each other, and therefore the statistics of each process cannot be processed separately, and appropriate steps are required to take into account the interaction.

Once the average event generation rates $r_L(c)$, $r_R(c)$ are estimated, then we find the feature field $z(c)$ that explains both event generation rates, in the sense that there exist two functions f_L and f_R such that the constraint described by equation (2) holds. Writing the constraint explicitly for each cell c^k , and letting $z^k = z(c^k)$ the value of the feature to estimate, we can see that we have a system of constraints of the kind:

$$\begin{cases} r_L(c^k) = f_L(z^k) + r_L^0 \\ r_R(c^k) = f_R(z^k) + r_R^0 \end{cases} \quad (\text{for } k = 1, \dots, \#cells)$$

The generated event rates $r_L(c^k)$ and $r_R(c^k)$ on the left side have already been estimated, while both the feature value z^k and f_L and f_R have to be estimated. The constants r_L^0 and r_R^0 can be incorporated as part of f_L and f_R . Note that this can be interpreted as a dimensionality reduction problem, because we have to find one cause (the feature z) that explains two effects (left and right event rates) at the same time.

In our case, we solve a relatively simple instance of the problem in which z is assumed to be a scalar function. Therefore, the constraints can be algebraically manipulated to obtain a closed form solution, which also takes into account the uncertainty in all the data and provide error bars for the estimated feature. The details are given in Text S1. Our approach is very generic, and can be extended to scenarios with more than 2 behaviors and more than 1 feature.

The feature z should be considered a dimensionless quantity of arbitrary scale. In fact, the equations that define it have multiple solutions. For example, suppose that (z, f_L, f_R) is one solution of the system of equations given by (Eq. 2). If α is any invertible function, then one can verify that $(\alpha(z), \alpha^{-1}(f_L), \alpha^{-1}(f_R))$ is a solution as well. Therefore, once we have obtained a solution for z , we can rescale it using any function α that we find convenient. In the following, we choose the rescaling function such that z is uniformly distributed in the interval $[-1, +1]$.

Results

Event statistics and estimated feature

We tracked 88 flies for a total of 5130 seconds or approximately 1.4 hours. Of the total recorded time, we considered only the 4814 seconds of data in which the flight speed exceeded

388 5 centimeters per second. This threshold on the linear velocity allowed working on tracks for
389 which saccades were easier to detect. We detected a total of 6613 saccades with this criterion,
390 giving an average saccade rate of 1.37 saccades per second.

391 We chose a reduced configuration $c(t)$ that is two-dimensional. This follows from considering
392 only planar motion (which reduces the effective degrees of freedom to 3), and using the
393 symmetry of the circular arena (which reduces the degrees of freedom to 2). An implied
394 assumption (which can be verified a posteriori) is that the fly’s response is not dependent on
395 the variables not considered in the analysis; for example, even though it is known that flies [64]
396 and other insects [65] use gaze to stabilize vision, there is no gaze variable in our model. This is
397 because the resolution of our measurements is not enough to observe directly the relative pose
398 of head and body, in terms of pitch, roll, or yaw. All components of the spatial configuration
399 that are theoretically relevant for the stimulus, but cannot be measured, are “hidden” states
400 whose contribution is lumped into the constant term in (2).

401 The two-dimensional reduced configuration can be parameterized in different ways, the results
402 being independent of the particular parameterization. The primary parameterization that we
403 use for computation uses (d, φ) for coordinates: d is the distance to the wall and φ is the angle
404 that the fly heading forms with respect to the axis of the arena (Figure 1Ci). We chose this
405 parameterization because it corresponds to two behaviorally relevant variables. We preferred
406 the axis angle φ over other potentially valid representations for the heading (e.g., approach
407 angle) because the representation is not singular, as $\varphi \in [-180 \text{ deg}, +180 \text{ deg}]$ for any value
408 of d .

409 We compute all statistics in the (d, φ) space, but we also use another choice of coordinates to
410 visualize the same data. We rotate the original (x, y, θ) configuration of the fly around the
411 center of the arena, such that the new coordinates are $(x', y', +90\text{deg})$. These “fly-centric”
412 coordinates are displayed using a top-down view of the arena, in which the fly always points up
413 (Figure 1Cii).

414 The reduced configuration $c(t)$ was discretized in a grid with sides of 36 cells (for φ) and 20
415 cells (for d) (Figure 1Di). The angle φ was discretized in 36 cells of equal size 10 deg. The
416 distance d was discretized in 20 unequal intervals (note the unequal y axis in Figure 1Di).
417 Intervals for d are smaller at the center of the arena and larger near the border, in such a way
418 that each annulus of radius $1 - d$ and width Δd had the same area. To compensate for the
419 sparseness of the data, each cell extends 50% into the neighbor’s area. Although these choices
420 were somewhat arbitrary, we obtain qualitatively similar results if we vary the number of the
421 cells.

Figure 1D shows the distribution of time spent at each point of the arena, and Figure 1E shows the distribution of the detected saccades using the GSD algorithm (see Text S1 for figures using the saccades detected by the alternative AVSD algorithm). As clearly evident in Figure 1Eii, most of the detected saccades correspond to the fly avoiding the walls on the left or on the right. However, those are the configurations where the flies spent more time (Figure 1D). Therefore, we need to normalize this data to see the behavioral patterns.

Figure 4A shows the estimated saccade generation function $r_i(c)$ across the reduced configuration space. These rates are obtained by first computing the observed generation rates $m_i(c)$ by averaging the number of saccades (Figure 1E) by the time spent in each cell (Figure 1D). Then the rates $r_i(c)$ are obtained from $m_i(c)$ by correcting for an estimated inhibition interval $\Delta = 0.1$ s. Panels B and C show the data separately for left and right saccades (r_L and r_R). The most evident phenomenon is that the fly tends to turn left when the wall is on the right (and vice versa), however, there are many saccades of the opposite direction initiated, even when the turning would orient the fly towards the wall rather than away from it. This is the phenomenon that we want the feature $z(c)$ to explain: we want to find the best spatial scalar value $z(c)$ such that both $r_L(c)$ and $r_R(c)$ can be written as a function of z . Figure 5Ai-ii shows the estimated feature $z(c)$ as a function of the reduced configuration c . This is the unidimensional feature that best explains both the left and saccade rates. The estimated feature using the alternative saccade detector is qualitatively similar (Figure S1). We now have the spatial feature $z(c)$ as well as the rates $r(c)$ as a function of the reduced configuration c and can now plot r as a function of z (using c as an implicit variable). This is shown in Figure 5B, which shows, for each cell k , the value of $r_i(c^k)$ as a function of $z(c^k)$. Figure 5Bi shows the data as a scatter plot, while Figure 5Bii shows the error bars on the estimated rates $r_i(c)$ at the 95% significance level.

Predictive power of the estimated feature

The data in Figure 5B indicate the predictive power of the feature. If the feature was perfectly predictive of the event rates, then $r_i(c)$ would be a function of $z(c)$. In this case, taking into account the error bounds on the rates, it is possible to find two functions f_L, f_R that predict the event rates in approximately 93% of the environment. More specifically, given a generic cell corresponding to the spatial configuration c^k , we find that the predicted event rates $f_L(c^k)$ and $f_R(c^k)$ are compatible with the observed rates $r_L(c)$ and $r_R(c)$ at the 95% level of significance. In practice, this means that the data in Figure 5B can be explained by two smooth curves (f_L and f_R) that intersect 93% of the confidence intervals corresponding to each spatial configuration. In the remaining cells (51 of 720 cells), the rates cannot be predicted by this feature alone. Further inspection (data not shown) reveals that such points correspond to configurations with the fly pointing directly against the wall at a small distance (<0.3 m). Note

that being able to predict the rates from the feature does not mean that one is able to predict the direction of each single saccade event. For example, in the middle of the arena the probability of left and right saccade is 50%, and this percentage is perfectly predicted by the feature; however, it is impossible to predict the direction of the single saccade better than chance.

Bounds on the contributions of random process and internal states

As explained before, using only external observations of the animal spatial configuration, we cannot distinguish among the contribution of a purely random endogenous saccade generation process, a deterministic process based on an internal state, any unmodeled features computed from the stimulus, and any unobservable spatial configuration that we cannot observe due to the limited resolution of our instruments. These contributions are lumped together in a baseline saccade rate. By examining the curves in Figure 5Bi we can estimate a baseline event rate r_0^i of about ~ 0.4 saccades/sec. By comparing with a maximum estimated event rate of ~ 4 saccades/sec, we can estimate that roughly 90% of the saccades are stimulus-driven in the regions of maximum stimulus. This value depends on the geometry and texture of this particular arena (e.g., it would be different if the arena was larger or smaller). However, we predict that the baseline rate of ~ 0.4 saccades/sec that we measure at the center of the arena should be independent of the geometry, as the size and textures of this arena were chosen such that the fly cannot perceive significant visual contrast from the center.

We can make some informed guesses for the contribution of the various possible processes by considering circumstantial evidence from other experiments. In tethered flight experiments, deliberately performed in the absence of salient visual stimuli, spontaneous saccade rates are on the order of ~ 0.3 saccades/sec [66]. If we assume that these values obtained in tethered experiments are a good approximation of an assumed spontaneous generation process in free-flight, then we can account for approximately 75% of the unexplained 0.4 saccades/sec as the joint contribution of a random process and unobservable internal states. This leaves roughly 25% of unexplained data, which could possibly be explained by estimating an additional feature $z'(c)$, perhaps dependent on components of the spatial configuration that we cannot observe, such as the gaze direction. The contribution of a hypothetical feature z' is therefore very small with respect to the contribution of the estimated z , as z' could possibly explain about 0.1 saccades/sec versus the 4 saccades/sec explained by z .

We conclude that the saccade behavior of *Drosophila* that depends on external visual stimulus appears to depend for the most part on only a one-dimensional feature of the stimulus. These conclusions must be limited to the particular experimental condition, as we cannot exclude that more complex environments would elicit more complex responses that require a higher dimensional feature to be explained. However, even in our relatively simple flight environment,

our analysis implies that the vast majority of saccades we observed are stimulus-driven and are not due to an internal, stimulus-independent search algorithm (e.g. Levy flights), as has been suggested for *Drosophila* and many other species [6–18].

Approximating the feature field using known parametric structures for visual processing

We have been able to compute the feature field z from the observable fly trajectory, without any assumptions on the fly visual processing. Nevertheless, it is interesting to test whether this independently identified feature is compatible with existing models of the first stages of visual processing in flies. In particular, we test the hypothesis of whether the identified feature can be expressed as a linear function of the perceived optic flow.

We assume the following generative model for z :

$$z(t) = \alpha\left(\int \text{OF}(\theta, t)A(\theta)d\theta\right), \quad (4)$$

where $\text{OF}(\theta, t)$ is the optic flow, or angular velocity, at the retinal angle θ at time t , and $A(\theta)$ is a retinal input kernel. The value $\theta = 0$ corresponds to the animal's center front visual field.

The function α is an arbitrary nonlinear function that we include in the model, because the identification procedure allows us to know z only up to a monotone transformation (i.e., if $z(c)$ is a solution of the constraints system, then also $\alpha(z(c))$ is a valid solution). We can characterize the optimal $A(\theta)$ as the solution of an optimization problem:

$$A^*(\theta) = \arg \min_A E(A, \alpha), \quad (5)$$

where the error function E is given by:

$$E(A, \alpha) = \sum_k [z(c^k) - \alpha\left(\int \text{OF}(\theta, c^k)A(\theta)d\theta\right)]^2. \quad (6)$$

In this last expression, $\text{OF}(\theta, c^k)$ is the typical optic flow that the animal experiences at the reduced configuration c^k . By solving this optimization problem, we try to best approximate the estimated feature over the whole environment, assuming it can be expressed as a linear function of the optic flow.

Unfortunately, we found that this optimization problem is ill posed given our data. In particular, z is known only at a discrete set of values c^k (720 cells --- the density of these is constrained by the finite amount of data that we have), and it is quite noisy, whereas the unknowns (A, α) are of high dimension. Given that the resolution of the fly's visual system is around $\sim 5^\circ$, it makes

sense to use at least 70 numbers ($\sim 330/5$) for representing A . Furthermore, α can be any monotonic nonlinear function.

We tried to improve the results by penalizing large values and large spatial variations of A (measured either by the spatial derivatives $\partial A/\partial \theta$ or $\partial^2 A/\partial \theta^2$). The modified error function is:

$$E' = E(A, \alpha) + \beta \int \left| \frac{\partial^m}{\partial \theta^{(m)}} A(\theta) \right|^2 d\theta, \quad (7)$$

for $m = 0, 1, 2$ and different values of β . In general, by varying m and β , we found a multitude of solutions, all very different from each other, having approximately the same predictive power (Figure 6). We noticed that for increasing regularization values the estimated linear kernel tended to be shaped as an harmonic function, as illustrated by the kernel obtained by regularizing the second derivative ($m = 2$) and using a large value of β ($\beta = 10^5$), shown in Figure 6D. This kernel is still asymmetric. If we impose that the kernel must be symmetric, we find that the best approximation using one harmonic is:

$$A(\theta) \propto \cos(\theta) + 0.2. \quad (8)$$

This kernel and relative feature field is shown in Figure 6E, and it is a good approximation of the feature estimated from the data.

We conclude that the identified feature can be expressed as a simple function of the optic flow. However, while obtaining the behaviorally relevant feature z from the external observations alone is a well-posed mathematical problem, finding the function that maps the stimulus y to the feature z is an ill-posed problem, because the set of possible models is of very high dimension compared with the data that we have. Note that these issues are already evident when considering only linear functions of the optic flow, and would be even more pressing if we were to add other nonlinear components to the model that are known to exist in the neural circuits of the fly.

Most of the estimated kernels obtained using some form of regularization share a particular feature: $A(\theta)$ is never 0 for θ corresponding to the back of the animal, but has opposite sign in the front of the animal (Figure 6B,C,D,E). Further investigation shows that these non-zero values in the caudal region are responsible for the two small side lobes that appear in the feature field when plotted in fly-centric coordinates. If the kernel is set to zero in the back, these side lobes disappear. This is apparent by comparing the feature field in Figure 6F (corresponding to the kernel $\max\{\cos(\theta) + 0.2, 0\}$) with that in Figure 6E. These results suggest that the optic flow in the back of the animal influences the fly's turning decisions. This

response cannot be interpreted as pure obstacle avoidance, given that flies tend to fly forward and obstacles in the back are not expected to represent a threat for collision. For convex environments, the saccades initiated from this response would tend to align the fly's course in parallel to the environment boundaries and the overall result is to *follow* walls rather than completely avoid them (similar behavior has been observed in bees [67]). Such a visuo-motor system might provide a functional advantage with respect to the balance of collision avoidance and object search. An animal that balances attraction and obstacle avoidance would tend to remain relatively close to interesting visual features, whereas an animal whose primary reflex is to fly away from visual features would tend to find itself in large open areas, far from potential landmarks or food sources. The only way of quantitatively verifying this *attraction-deflection* hypothesis would be to obtain data from experiments within larger environments with more varied visual features. These results are also compatible with the observation in previous experiments on tethered flies that the optomotor response can be written as the function of a kernel in which the rear and front visual fields give opposite contributions [68], suggesting that a similar visual feature might be used for both behaviors.

Discussion

In this paper, we introduced a novel method to obtain an estimate of a low-dimensional feature of the stimulus that best predicts the observable behavioral event generation rates. The feature can be obtained from observable quantities, such as the recording of the trajectory of the animal, without any assumption on the nature of the stimulus and its underlying neural processing. Using this method, we have concluded that most of the saccade events generated by fruit flies exploring a structured laboratory environment are induced by visual stimuli, and that the instantaneous stimulus can be compressed down to a one-dimensional feature, while still being predictive of the event rates in $\sim 93\%$ of the environment. Using this method, it is not possible to distinguish between the contribution of an endogenous random process and a deterministic contribution dependent on an unobservable state. However, we can bound the contributions of these two terms in a baseline saccade rate that we estimate at 0.4 saccades/sec, roughly a tenth of the maximum rate. The strength of this method is that the feature z can be estimated working backwards from an animal's actions, rather than forward by postulating a model for the stimulus $y(t)$ and guessing what is the relevant feature. Once we know z , as a second step, it is possible to attempt to fit a parametric representation of neural processing to find the forward function from y to z , based on other assumptions about sensory processing, though this is not guaranteed to be a well posed problem, as one must optimize over all plausible models compatible with the animal's biology. In this particular case, we have shown that the feature z responsible for turning decisions in *Drosophila* can be written as a linear function of the optic flow, and that the particular linear kernel we obtain is compatible with that identified in tethered conditions, for a particular choice of regularization penalty to

601 make the problem well posed. Conversely, finding z from the behavioral data is a well posed
602 and intuitive problem, because it can be understood as a dimensionality reduction problem
603 (find the one feature that explains multiple behaviors).

604 The main advantage of this approach, compared with previous methods, is that it can be
605 applied to freely moving animals, and thus permits asking about responses to naturally
606 important stimuli. Moreover, it does not need any assumption of linearity between some
607 aspect of the stimulus and response, a precondition strongly needed in techniques such as
608 reverse correlation [48]. Even advanced reverse correlation techniques in single sensory
609 neurons [69] are not easy to generalize into models of network functionality that could be used
610 to predict behavior.

611 In the future, this method could be applied to different behaviors of the fruit fly and other
612 animals [70]. The formalization is quite generic, though some generalizations are possible. The
613 algorithm documented in Text S1 assumes that the feature is one-dimensional in order to
614 obtain a closed-form solution. To identify a feature of higher dimension, this must be
615 generalized, for example by using one of the various more computationally expensive
616 dimensionality reduction algorithms in machine learning (e.g., [71]). In any case, the rate-
617 variant interacting Poisson process model seems apt for modeling many other behaviors (e.g.,
618 landing, taking off) that can be reliably localized in time (i.e., they have a clear beginning and
619 end), and that can be caused by both external and internal causes.

620 Thinking in terms of the feature z as a proxy of the stimulus can potentially be useful in
621 understanding how different sensory modalities contribute to the same behavior. The feature is
622 independent of the sensory modality because it is just a function of the animal configuration,
623 and it is a proxy of the typical stimulus perceived at the location, so it could be used to study,
624 for example, the influence of olfaction instead of visual processing on turning behavior, or their
625 interaction, which has been the object of much research [72–75][56]. Note, however, that we
626 do have the strong assumption that the stimulus is a constant function of the configuration, so
627 the framework cannot be easily extended to time-varying stimuli.

628 This approach might also be useful to study different behaviors at the same time. *Drosophila*
629 has a large repertoire of behaviors/reflexes which are stimulus-triggered, such as landing, take-
630 off, chasing mates, and escaping from small targets. In this case, we focused on saccades, and
631 we found the feature z encoding the relevant function of the stimulus for saccade decisions. If
632 one repeated the analysis for a different behavior (e.g., landing), there would likely be another
633 feature z'' , that would be different from z . However, if this was repeated for all fly behaviors,
634 one would find that at some point the new identified features would be redundant; for
635 example, in the case of vision, the number of features is upper bounded by the number of

upstream signals towards the lobula. Ultimately, this exercise might provide a prediction of whether two behaviors are likely to share the same neural pathways.

Potentially, this technique could help in quantifying the behavioral differences of different genotypes. This model makes a distinction between the feature z and the event generation rate functions f_i . Whereas z is assumed to be correlated with computed percepts, f_i might be correlated more with the motor functions. This distinction could be used to obtain insight regarding the function of genetic manipulations such as a screen in which populations of neurons are “silenced” with a hyperpolarizing ion channel or synaptic release blockade. For example, if a modified animal gives the same feature z but modified rate functions f_i , it would be evidence that the silenced neurons are involved with motor generation rather than with stimulus processing. Consequently, with a large-scale screen [76,77], it might be possible to obtain a classification of phenotypes into sensory, decision making, and motor deficits. Similarly, we could use this feature to quantitatively compare the properties of different species.

Another interesting but more substantial extension of this work would be to expand the mathematical formalism to incorporate measurements of neuronal activity into the internal processing structure. This is now done in freely moving worms [78,79,32] and zebrafish [80]; in adult flies, most neural recording during behavior is being done on fixed flies [81–83,42].

Supporting Information

The following supplementary materials are available:

1. Text S1: Mathematical details of the identification technique and implementation details of the saccade detection algorithms (PDF)

Acknowledgements

We thank Kit Longden for insightful comments on a preliminary version of this work.

References

1. Berg HC, Brown DA (1972) Chemotaxis in *Escherichia coli* analysed by Three-dimensional Tracking. *Nature* 239: 500–504.
2. Tynan CT (1998) Ecological importance of the Southern Boundary of the Antarctic Circumpolar Current. *Nature* 392: 708–710. doi:10.1038/33675
3. Fuortes MG, S. Y (1964) Probability of occurrence of discrete potential waves in the eye of *Limulus*. *The Journal of general physiology* 47: 443–63.

- 668 4. Baylor DA, Lamb TD, Yau KW (1979) Responses of retinal rods to single photons. The
669 Journal of physiology 288: 613–34.
- 670 5. Menini A, Picco C, Firestein S (1995) Quantal-like current fluctuations induced by
671 odorants in olfactory receptor cells. Nature 373: 435–7. doi:10.1038/373435a0
- 672 6. Shimada I, Kawazoe Y, Hara H (1993) A temporal model of animal behavior based on a
673 fractality in the feeding of *Drosophila melanogaster*. Biological Cybernetics 68: 477–481.
- 674 7. Cole B, Cole J (1995) Fractal time in animal behaviour: the movement activity of
675 *Drosophila*. Animal Behaviour 50: 1317–1324. doi:10.1016/0003-3472(95)80047-6
- 676 8. Bartumeus F, Levin SA (2008) Fractal reorientation clocks: Linking animal behavior to
677 statistical patterns of search. Proceedings of the National Academy of Sciences 105:
678 19072.
- 679 9. Viswanathan GM, Afanasyev V, Buldyrev SV, Murphy EJ, Prince PA, et al. (1996) Lévy
680 flight search patterns of wandering albatrosses. Nature 381: 413–415.
681 doi:10.1038/381413a0
- 682 10. Bartumeus F, Peters F, Pueyo S, Marrasé C, Catalan J (2003) Helical Lévy walks:
683 adjusting searching statistics to resource availability in microzooplankton. Proceedings of
684 the National Academy of Sciences 100: 12771.
- 685 11. Humphries NE, Queiroz N, Dyer JRM, Pade NG, Musyl MK, et al. (2010) Environmental
686 context explains Lévy and Brownian movement patterns of marine predators. Nature
687 465: 1066–1069. doi:10.1038/nature09116
- 688 12. Maye A, Hsieh C-H, Sugihara G, Brembs B (2007) Order in spontaneous behavior. PloS
689 one 2: e443. doi:10.1371/journal.pone.0000443
- 690 13. Reynolds AM, Frye M a (2007) Free-flight odor tracking in *Drosophila* is consistent with
691 an optimal intermittent scale-free search. PloS one 2: e354.
692 doi:10.1371/journal.pone.0000354
- 693 14. Reynolds AM, Smith AD, Reynolds DR, Carreck NL, Osborne JL (2007) Honeybees perform
694 optimal scale-free searching flights when attempting to locate a food source. Journal of
695 Experimental Biology 4: 3763–3770. doi:10.1242/jeb.009563
- 696 15. Osborne JL, Andrew P, Carreck NL, Swain JL, Knight ME, et al. (2008) Bumblebee flight
697 distances in relation to the forage landscape. Journal of Animal Ecology 77: 406–415.
698 doi:10.1111/j.1365-2656.2007.01333.x

- 699 16. Nevitt G a, Losekoot M, Weimerskirch H (2008) Evidence for olfactory search in
700 wandering albatross, *Diomedea exulans*. Proceedings of the National Academy of
701 Sciences of the United States of America 105: 4576–4581. doi:10.1073/pnas.0709047105
- 702 17. Sims DW, Southall EJ, Humphries NE, Hays GC, Bradshaw CJA, et al. (2008) Scaling laws of
703 marine predator search behaviour. Nature 451: 1098–1102. doi:10.1038/nature06518
- 704 18. Nathan R (2008) An emerging movement ecology paradigm. Proceedings of the National
705 Academy of Sciences 105: 19050.
- 706 19. Reynolds A (2008) How many animals really do the Lévy walk? Comment. Ecology 89:
707 2347–2351.
- 708 20. Stumpf MPH, Porter MA (2012) Critical Truths About Power Laws. Science 335: 665–666.
709 doi:10.1126/science.1216142
- 710 21. Edwards AM, Phillips R a, Watkins NW, Freeman MP, Murphy EJ, et al. (2007) Revisiting
711 Lévy flight search patterns of wandering albatrosses, bumblebees and deer. Nature 449:
712 1044–8. doi:10.1038/nature06199
- 713 22. Bartumeus F (2007) Lévy processes in animal movement: an evolutionary hypothesis.
714 FRACTALS-LONDON- 15: 151.
- 715 23. Goldreich O (2001) Foundations of Cryptography: Basic Tools. Cambridge University
716 Press.
- 717 24. Hassenstein B, Reichardt W (1956) Systemtheoretische Analyse der Zeit-, Reihenfolgen-
718 und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkafers
719 *Chlorophanus*. Zeitschrift Für Naturforschung 11b: 513–524.
- 720 25. Srinivasan MV, Zhang SW, Lehrer M, Collett TS (1996) Honeybee navigation en route to
721 the goal: Visual flight control and odometry. Journal of Experimental Biology 199: 237–
722 244.
- 723 26. Reynolds AM, Frye MA (2007) Free-flight odor tracking in *Drosophila* is consistent with an
724 optimal intermittent scale-free search. PloS one 2: e354.
725 doi:10.1371/journal.pone.0000354
- 726 27. Maye A, Hsieh C-H, Sugihara G, Brembs B (2007) Order in spontaneous behavior. PloS
727 one 2: e443. doi:10.1371/journal.pone.0000443
- 728 28. Srivastava N, Clark DA, Samuel ADT (2009) Temporal Analysis of Stochastic Turning
729 Behavior of Swimming *C. elegans*. Journal of Neurophysiology: 1172–1179.
730 doi:10.1152/jn.90952.2008.

- 731 29. Edwards AM, Phillips RA, Watkins NW, Freeman MP, Murphy EJ, et al. (2007) Revisiting
732 Lévy flight search patterns of wandering albatrosses, bumblebees and deer. *Nature* 449:
733 1044–8. doi:10.1038/nature06199
- 734 30. Nevitt G a, Losekoot M, Weimerskirch H (2008) Evidence for olfactory search in
735 wandering albatross, *Diomedea exulans*. *Proceedings of the National Academy of*
736 *Sciences of the United States of America* 105: 4576–81. doi:10.1073/pnas.0709047105
- 737 31. Berg HC, Brown DA (1972) Chemotaxis in *Escherichia coli* analysed by Three-dimensional
738 Tracking. *Nature* 239: 500–504.
- 739 32. Piggott BJ, Liu J, Feng Z, Wescott SA, Xu XZS (2011) The Neural Circuits and Synaptic
740 Mechanisms Underlying Motor Initiation in *C. elegans*. *Cell* 147: 922–933.
741 doi:10.1016/j.cell.2011.08.053
- 742 33. Kawano T, Po MD, Gao S, Leung G, Ryu WS, et al. (2011) An Imbalancing Act: Gap
743 Junctions Reduce the Backward Motor Circuit Activity to Bias *C. elegans* for Forward
744 Locomotion. *Neuron* 72: 572–86. doi:10.1016/j.neuron.2011.09.005
- 745 34. Cowan NJ, Fortune ES (2007) The critical role of locomotion mechanics in decoding
746 sensory systems. *The Journal of neuroscience : the official journal of the Society for*
747 *Neuroscience* 27: 1123–8. doi:10.1523/JNEUROSCI.4198-06.2007
- 748 35. Beyeler A, Zufferey J-C, Floreano D (2009) Vision-based control of near-obstacle flight.
749 *Autonomous Robots* 27: 201–219. doi:10.1007/s10514-009-9139-6
- 750 36. Conroy J, Gremillion G, Ranganathan B, Humbert JS (n.d.) Implementation of wide-field
751 integration of optic flow for autonomous quadrotor navigation. *Autonomous Robots*: 1–
752 10.
- 753 37. Borst A, Euler T (2011) Seeing things in motion: models, circuits, and mechanisms.
754 *Neuron* 71: 974–94. doi:10.1016/j.neuron.2011.08.031
- 755 38. Borst A (2009) *Drosophila's* view on insect vision. *Current biology : CB* 19: R36–47.
756 doi:10.1016/j.cub.2008.11.001
- 757 39. Kennedy JS (1940) The visual responses of flying mosquitoes. *Proceedings of the*
758 *Zoological Society of London Series A* 109: 221–242. doi:10.1111/j.1096-
759 3642.1940.tb00831.x
- 760 40. Mittelstaedt H (1949) Telotaxis und Optomotorik von *Eristalis* bei Augeninversion.
761 *Naturwissenschaften* 36: 90–91.

- 762 41. Eichner H, Joesch M, Schnell B, Reiff DF, Borst A (2011) Internal structure of the fly
763 elementary motion detector. *Neuron* 70: 1155–64. doi:10.1016/j.neuron.2011.03.028
- 764 42. Clark D a, Bursztyn L, Horowitz M a, Schnitzer MJ, Clandinin TR (2011) Defining the
765 computational structure of the motion detector in *Drosophila*. *Neuron* 70: 1165–77.
766 doi:10.1016/j.neuron.2011.05.023
- 767 43. Reiff DF, Plett J, Mank M, Griesbeck O, Borst A (2010) Visualizing retinotopic half-wave
768 rectified input to the motion detection circuitry of *Drosophila*. *Nature neuroscience* 13:
769 973–978. doi:10.1038/nn.2595
- 770 44. Joesch M, Schnell B, Raghu SV, Reiff DF, Borst A (2010) ON and OFF pathways in
771 *Drosophila* motion vision. *Nature* 468: 300–304. doi:10.1038/nature09545
- 772 45. Krapp HG, Hengstenberg R (1996) Estimation of self-motion by optic flow processing in
773 single visual interneurons. *Nature* 384: 463–6. doi:10.1038/384463a0
- 774 46. Roth E, Reiser MB, Dickinson MH, Cowan NJ (n.d.) A Task-level Model for Optomotor Yaw
775 Regulation in *Drosophila Melanogaster* : A Frequency-Domain System Identification
776 Approach..
- 777 47. Aptekar JW, Shoemaker P a, Frye M a (2012) Figure Tracking by Flies Is Supported by
778 Parallel Visual Streams. *Current biology : CB*: 482–487. doi:10.1016/j.cub.2012.01.044
- 779 48. Theobald JC, Ringach DL, Frye MA (2010) Dynamics of optomotor responses in
780 *Drosophila* to perturbations in optic flow. *The Journal of experimental biology* 213:
781 1366–75. doi:10.1242/jeb.037945
- 782 49. Mronz M, Lehmann F-O (2008) The free-flight response of *Drosophila* to motion of the
783 visual environment. *The Journal of experimental biology* 211: 2026–45.
784 doi:10.1242/jeb.008268
- 785 50. Braun E, Geurten B, Egelhaaf M (2010) Identifying prototypical components in behaviour
786 using clustering algorithms. *PloS one* 5: e9361. doi:10.1371/journal.pone.0009361
- 787 51. Lindemann JP, Weiss H, Möller R, Egelhaaf M (2008) Saccadic flight strategy facilitates
788 collision avoidance: closed-loop performance of a cyberfly. *Biological Cybernetics* 98:
789 213–227. doi:10.1007/s00422-007-0205-x
- 790 52. Boeddeker N, Egelhaaf M (2005) A single control system for smooth and saccade-like
791 pursuit in blowflies. *4*: 1563–1572. doi:10.1242/jeb.01558
- 792 53. Land MF, Collett TS (1974) Chasing behavior of houseflies (*Fannia canicularis*): a
793 description and analysis. *J. Comp. Physiol.* 89: 331–357.

- 794 54. Wagner H (1986) Flight performance and visual control of the flight of the free-flying
795 housefly (*Musca domestica*). II. Pursuit of targets. *Philos. Trans. R. Soc. Long. B. Biol. Sci.*
796 312: 553–579.
- 797 55. Tammero LF, Dickinson MH (2002) The influence of visual landscape on the free flight
798 behavior of the fruit fly *Drosophila melanogaster*. *Journal of Experimental Biology* 205:
799 327–343.
- 800 56. Stewart FJ, Baker DA, Webb B (2010) A model of visual-olfactory integration for odour
801 localisation in free-flying fruit flies. *Journal of Experimental Biology* 213: 1886–1900.
802 doi:10.1242/jeb.026526
- 803 57. Stephens GJ, Osborne LC, Bialek W (2010) Searching for simplicity: Approaches to the
804 analysis of neurons and behavior. *Arxiv preprint arXiv:1012.3896*.
- 805 58. Heisenberg M, Wolf R (1984) *Vision in Drosophila: genetics of Microbehavior*. Springer-
806 Verlag.
- 807 59. Straw AD, Branson K, Neumann TR, Dickinson MH (2011) Multi-camera Realtime 3D
808 Tracking of Multiple Flying Animals. *J R Soc Interface* 8: 395–409.
809 doi:10.1098/rsif.2010.0230
- 810 60. Breugel F and van, Dickinson M. (2012) The visual control of landing and obstacle
811 avoidance in the fruit fly, *Drosophila*. *Journal of Experimental Biology* in press.
- 812 61. Bender JA, Dickinson MH (2006) A comparison of visual and haltere-mediated feedback
813 in the control of body saccades in *Drosophila melanogaster*. *Journal of Experimental*
814 *Biology* 209: 4597–4606. doi:10.1242/jeb.02583
- 815 62. Stephens GJ, Bueno de Mesquita M, Ryu WS, Bialek W (2011) Emergence of long
816 timescales and stereotyped behaviors in *Caenorhabditis elegans*. *Proceedings of the*
817 *National Academy of Sciences of the United States of America* 108: 7286–9.
818 doi:10.1073/pnas.1007868108
- 819 63. Geurten BRH, Kern R, Braun E, Egelhaaf M (2010) A syntax of hoverfly flight prototypes.
820 *Journal of Experimental Biology* 213: 2461–2475. doi:10.1242/jeb.036079
- 821 64. Rosner R, Egelhaaf M, Grewe J, Warzecha a. K (2009) Variability of blowfly head
822 optomotor responses. *Journal of Experimental Biology* 212: 1170–1184.
823 doi:10.1242/jeb.027060
- 824 65. Boeddeker N, Dittmar L, Stürzl W, Egelhaaf M (2010) The fine structure of honeybee
825 head and body yaw movements in a homing task. *Proceedings. Biological sciences / The*
826 *Royal Society* 277: 1899–906. doi:10.1098/rspb.2009.2326

- 827 66. Bender JA, Dickinson MH (2006) Visual stimulation of saccades in magnetically tethered
828 *Drosophila*. *The Journal of experimental biology* 209: 3170–82. doi:10.1242/jeb.02369
- 829 67. Serres JR, Masson GP, Ruffier F, Franceschini N (2008) A bee in the corridor: centering
830 and wall-following. *Naturwissenschaften* 95: 1181–1187. doi:10.1007/s00114-008-0440-
831 6
- 832 68. Tammero LF, Frye MA, Dickinson MH (2004) Spatial organization of visuomotor reflexes
833 in *Drosophila*. *J Exp Biol* 207: 113–122. doi:10.1242/jeb.00724
- 834 69. Pillow JW, Simoncelli EP (2006) Dimensionality reduction in neural models: an
835 information-theoretic generalization of spike-triggered average and covariance analysis.
836 *Journal of vision* 6: 414–28. doi:10.1167/6.4.9
- 837 70. Gomez-Marin A, Stephens GJ, Louis M (2011) Active sampling and decision making in
838 *Drosophila* chemotaxis. *Nature communications* 2: 441. doi:10.1038/ncomms1455
- 839 71. Tenenbaum JB, Silva VD, Langford JC (2000) A Global Geometric Framework for
840 Nonlinear Dimensionality Reduction. 290: 2319–2323.
- 841 72. Chow DM, Frye M a. (2009) The neuro-ecology of resource localization in *Drosophila*:
842 Behavioral components of perception and search. *Fly* 3: 50–61. doi:10.4161/fly.3.1.7775
- 843 73. Krishnan P, Duistermars BJ, Frye M a (2011) Odor identity influences tracking of
844 temporally patterned plumes in *Drosophila*. *BMC Neuroscience* 12: 62.
845 doi:10.1186/1471-2202-12-62
- 846 74. Frye M, Duistermars B (2009) Visually Mediated Odor Tracking During Flight in
847 *Drosophila*. *Journal of Visualized Experiments*.
- 848 75. Budick SA, Dickinson MH (2006) Free-flight responses of *Drosophila melanogaster* to
849 attractive odors. *Journal of Experimental Biology* 209: 3001–3017. doi:10.1242/jeb.02305
- 850 76. Iyengar BG, Chou CJ, Vandamme KM, Klose MK, Zhao X, et al. (2011) Silencing synaptic
851 communication between random interneurons during *Drosophila* larval locomotion.
852 *Genes, brain, and behavior* 10: 883–900. doi:10.1111/j.1601-183X.2011.00729.x
- 853 77. Katsov AY, Clandinin TR (2008) Motion processing streams in *Drosophila* are behaviorally
854 specialized. *Neuron* 59: 322–35. doi:10.1016/j.neuron.2008.05.022
- 855 78. Stirman JN, Crane MM, Husson SJ, Wabnig S, Schultheis C, et al. (2011) Real-time
856 multimodal optical control of neurons and muscles in freely behaving *Caenorhabditis*
857 *elegans*. *Nature methods* 8: 153–8. doi:10.1038/nmeth.1555

- 858 79. Leifer AM, Fang-Yen C, Gershow M, Alkema MJ, Samuel ADT (2011) Optogenetic
859 manipulation of neural activity in freely moving *Caenorhabditis elegans*. *Nature Methods*
860 advance on. doi:10.1038/nmeth.1554
- 861 80. Naumann EA, Kampff AR, Prober DA, Schier AF, Engert F (2010) Monitoring neural
862 activity with bioluminescence during natural behavior. *Nature neuroscience* 13: 513–20.
863 doi:10.1038/nn.2518
- 864 81. Maimon G, Straw AD, Dickinson MH (2010) Active flight increases the gain of visual
865 motion processing in *Drosophila*. *Nat Neurosci* 13: 393–399. doi:10.1038/nn.2492
- 866 82. Seelig JD, Chiappe ME, Lott GK, Dutta A, Osborne JE, et al. (2010) Two-photon calcium
867 imaging from head-fixed *Drosophila* during optomotor walking behavior. *Nature*
868 *methods* 7: 535–40. doi:10.1038/nmeth.1468
- 869 83. Chiappe ME, Seelig JD, Reiser MB, Jayaraman V (2010) Walking Modulates Speed
870 Sensitivity in *Drosophila* Motion Vision. *Current Biology* 20: 1470–1475.
871 doi:10.1016/j.cub.2010.06.072

872

873

Figure 1. Data collection, saccade detection, reduced coordinate space, time histogram, number of saccades histogram. Panel A shows the experimental setup: the fly is tracked in a circular arena of 1m radius. The retro-illuminated checkerboard pattern gives a uniform stimulus to the fly. Panel B-i shows some of the trajectories recorded in the arena. The trajectory can be interpreted as a mix of smooth turns and rapid turns, called "saccades", which are responsible for most of the total angular displacement of the animal. We wrote software to detect these saccades events, based on two different algorithms, documented in the Supplemental Materials. In this paper, we only consider these discrete saccade events (Panel B-ii). Panel C shows the two coordinate systems we use in this paper. We take advantage of the circular symmetry of the environment, along with a hypothesis of planarity, to reduce the degrees of freedom to 2. Panel C-i shows the choice of the axis-angle/distance from the wall coordinates. Panel C-ii shows the "fly-centric view". The fly configuration is reduced to 2 spatial coordinates by rotating the configuration so that the animal points "up" with respect to the diagram. We remark that all the results in this paper do not depend on the choice of coordinates. Panel D shows a density plot of T^k , which is the time spent at each configuration c^k . Panel E shows the number of saccades (both left and right) detected at each configuration.

Figure 2. General reference models for the animal behavior and decision making. Panel A illustrates the nomenclature that we use in this paper: $x(t)$ is the animal configuration (position/velocity), which ultimately depends on the past history of the animal decisions, the body dynamics, and environmental effects, here abstractly represented by the variable W . The stimulus $y(t)$ perceived by the animal is a function of the animal configuration $x(t)$ and the geometry/textures of the environment. In the most general terms, the actions of the animals, $u(t)$, are generated on the basis of the instantaneous stimulus $y(t)$ as well as the internal state $\xi(t)$, which includes, for our purposes, everything which is not observable, including metabolic and neural states. Panel B illustrates the specialization of the model that we postulate. The decisions of the animals are represented by series of observable events belonging to a fixed set of classes; in our case these are left and right body saccades. The events are assumed to be generated by a set of interacting rate-variant Poisson processes. The instantaneous rates $r_i(t)$ depend on several factors, including the unobservable states, and the external stimulus. The main hypothesis of this paper is that the contribution of the stimulus on the rate can be written as a function of a low-dimensional feature $z(t)$ computed from the stimulus. The inference problem in this paper consists in identifying the functions f_i that best explain the rates as a function of the stimulus ($r_i(t) = f_i(z(t))$). The diagram also shows the impact of other unmodeled neural processing based on internal states, acting as a disturbance in the model. We do not infer a functional description of this modeling, but we are able to bound its contribution and show that it is small with respect to the stimulus-induced contribution. The diagram also shows the *reduced configuration* $c(t)$, the subset of $x(t)$ on which the stimulus

actually depends. The reduced configuration depends on the particular experimental settings; in our case, we postulate that in a circular arena the stimulus is dependent on only two degrees of freedom. This is a hypothesis that can be verified a posteriori.

Figure 3. Simple models of decision making processes and relative experimental predictions.

This figure shows, on the left, several simplified saccade generation schemes, and their prediction in terms of the observed statistics. All models assume that the visual stimulus $y(t)$ is processed as to extract a one-dimensional feature $z(t)$ on which the animal decisions are based. The models presented are meant to represent a sample of qualitatively different functional models of behavior generations, and not necessarily biologically plausible models of neural computation. Panel A-i shows a “hard threshold” model: if the feature $z(t)$ is below a threshold, no event is generated, otherwise, the event is generated stochastically with a certain rate. Panel A-ii shows what would be the prediction of the model if we were to plot the saccade generation rate (an observable quantity) as a function of the feature $z(t)$, assuming we knew how to compute z . Panel B shows the same model, but with noise affecting the computation of the feature. The effect on the observed rate would be to transform the hard threshold in a soft threshold. Panel C shows a model in which there is a parallel saccade generation mechanism, which generates saccades randomly independently of the stimulus. The effect of this on the measured rate is to raise uniformly the curves. Also the contribution of some internal processing based on internal neural states which were not a function of the instantaneous stimulus would have the same effect on the rate statistics. Panel D shows the case where the behavior depends also on some other feature of the stimulus z' in addition to z . In this case, if we plotted the rates as a function of z , ignoring the dependency on z' , we would see that it is not possible for z to explain the rates by itself. Therefore, once we have identified the curves f_L , f_R , and the feature z , we are able to identify the contribution of a random generation process (or based on an internal state) as a uniform baseline saccade rate; and we can infer whether another feature is necessary to explain the behavior by the vertical spread of the rates.

Figure 4. Estimated saccade generation rates. This picture shows basic statistics of the processed data. Panel A shows the estimated saccade *rate* in polar coordinates (i) and fly-centered spatial coordinates (ii), indicated in the text as r^k . This density is obtained by taking the raw number of saccades in each cell n^k (Figure 1, Panel E), normalizing by the time spent in each cell T^k (Figure 1, Panel D), and then compensating for the interacting nature of the Poisson processes. Panels B-C show the rates for left and right saccades (r_L^k and r_R^k , respectively), which we plot in red in Panel B (left saccades) and in blue in panel C (right saccades). Note that the left and right saccade ratios appear roughly symmetric.

Figure 5. The estimated decision feature z . Panel A shows the estimated one dimensional feature z . This is the best one dimensional spatial feature that explains the left and right saccade rates. It is a dimensionless quantity, which we normalize in the interval $[-1, +1]$. Panel B-i shows, for each cell k , the rates r_i^k as a function of the estimated feature z^k ; Panel B-ii shows the same data, but with error bars corresponding to 95% confidence intervals (the bars are not symmetric because the posterior distribution of the estimated rates is not Gaussian; see Supplemental Materials for details). The single feature z is sufficient to predict the rate in $\sim 93\%$ of the environment, in the sense that $\sim 93\%$ of the rates can be considered (with the error bars) as lying on the same curve; these curves are the functions f_L and f_R discussed previously that allow predicting the rates from the feature. The remaining $\sim 7\%$ of data that this model cannot fit correspond to configurations with the fly pointing directly against the wall at a small distance (< 0.3 m).

Figure 6. Receptive fields of wide-field motion detection consistent with the feature z . These pictures show several receptive fields of wide field motion sensitive cells, the spatial feature that they compute, as well as a comparison with the feature z identified from the data. The pictures in first column show the kernel $A(\theta)$; the pictures in the second column show the corresponding feature field. The panels A through D show the kernels obtained as solutions of an optimization problem, respectively by solving a linear least-squares problem (panel A), and three different regularization problems: by penalizing the norm of the solution (panel B), by penalizing the norm of the spatial derivative (panel C), and by penalizing the curvature of the solution (panel D). Note all solutions are asymmetric due to the noise in the data. Panel E shows the kernel $\cos(\theta) + 0.2$, which is the closest harmonic function to the regularized solution in panel D. Panel F shows the result obtained by setting to zero this kernel in the back of the field of view. This shows that the contribution of the back of the field of view is necessary to recreate the small sidelobes of the estimated feature field.

Figure S1. The estimated decision feature z (results obtained with alternative detection algorithm). This figure shows the equivalent results of Figure 5 using the AVSD algorithm for saccade detection instead of GSD. Compared with Figure 5, the graphs are similar in many respects. Panel A shows the estimated one dimensional feature z as a function of the reduced configuration. Note that there are only two areas, corresponding to positive and negative z , instead of four, as in Figure 5A. Panel B shows the observed event rates as a function of the estimated feature. Compared to Figure 5B, in this case the estimated feature z^k is slightly less predictive of the rates r_i^k . This means that the events detected by the AVSD algorithm cannot be correlated with the stimulus as well as those detected by AVSD. This suggests that, for this

982 particular data, the GSD algorithm is able to better detect behaviorally relevant events in the
983 trajectory data.

984

Figure 1
[Click here to download Figure: Figure1.eps](#)

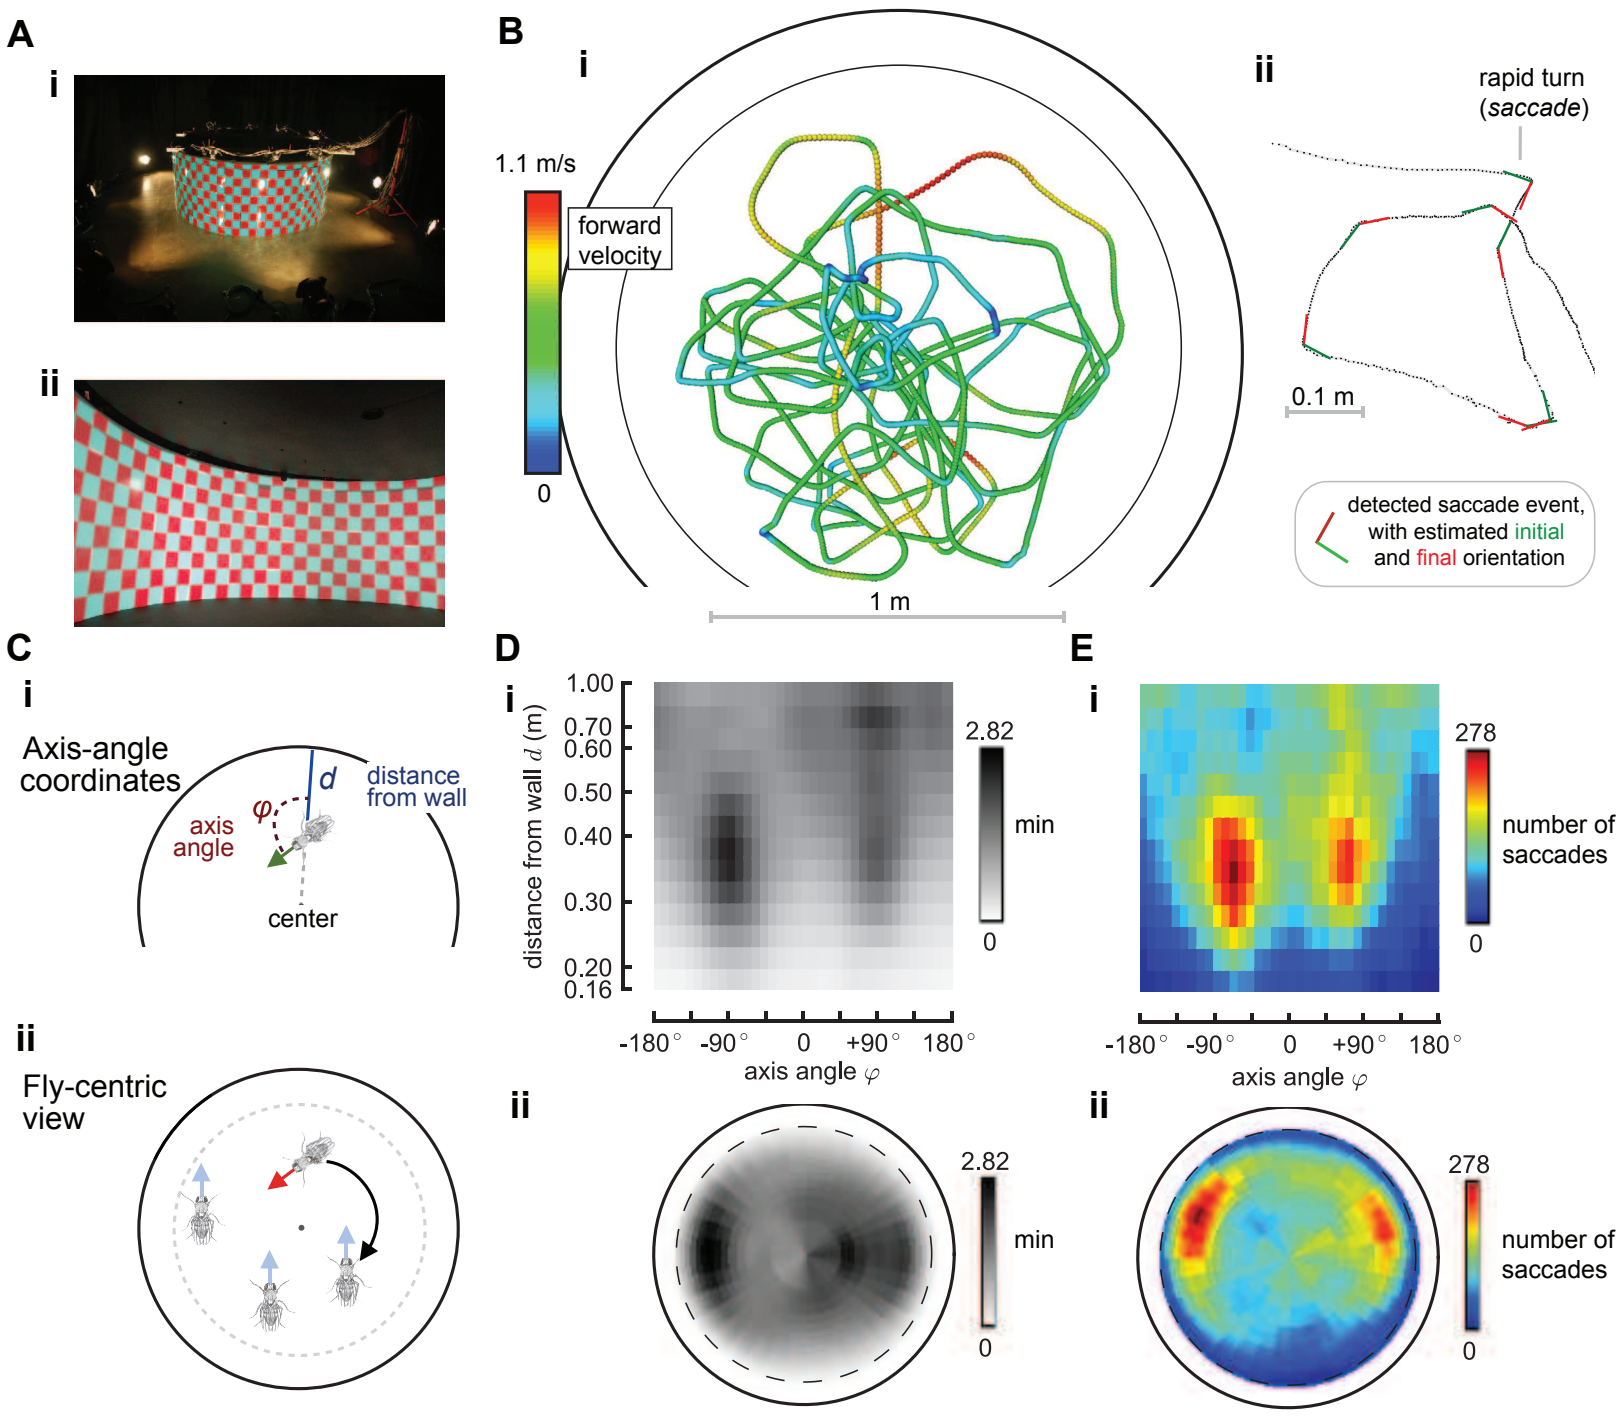
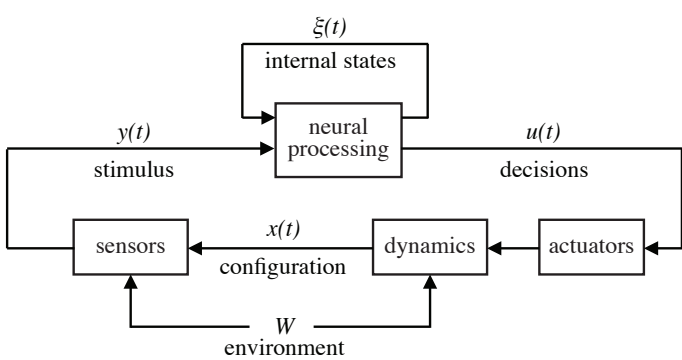


Figure 2
[Click here to download Figure: Figure2.eps](#)

A



B

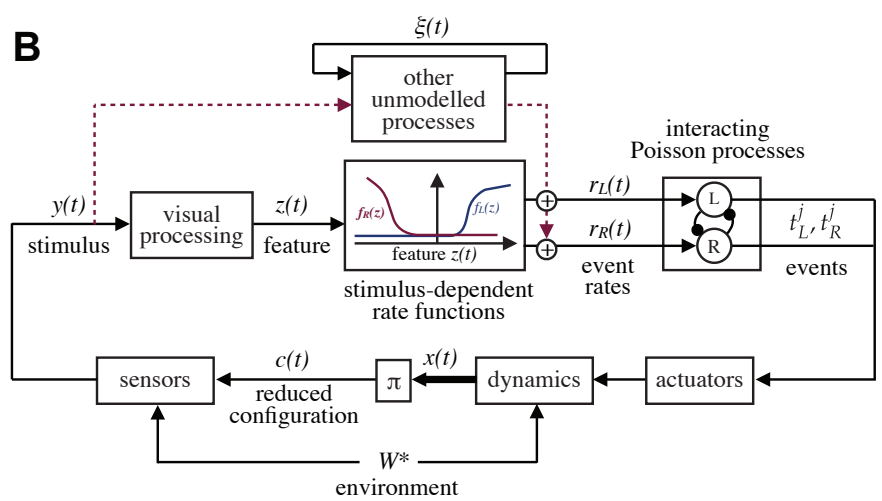


Figure 3
[Click here to download Figure: Figure3.eps](#)

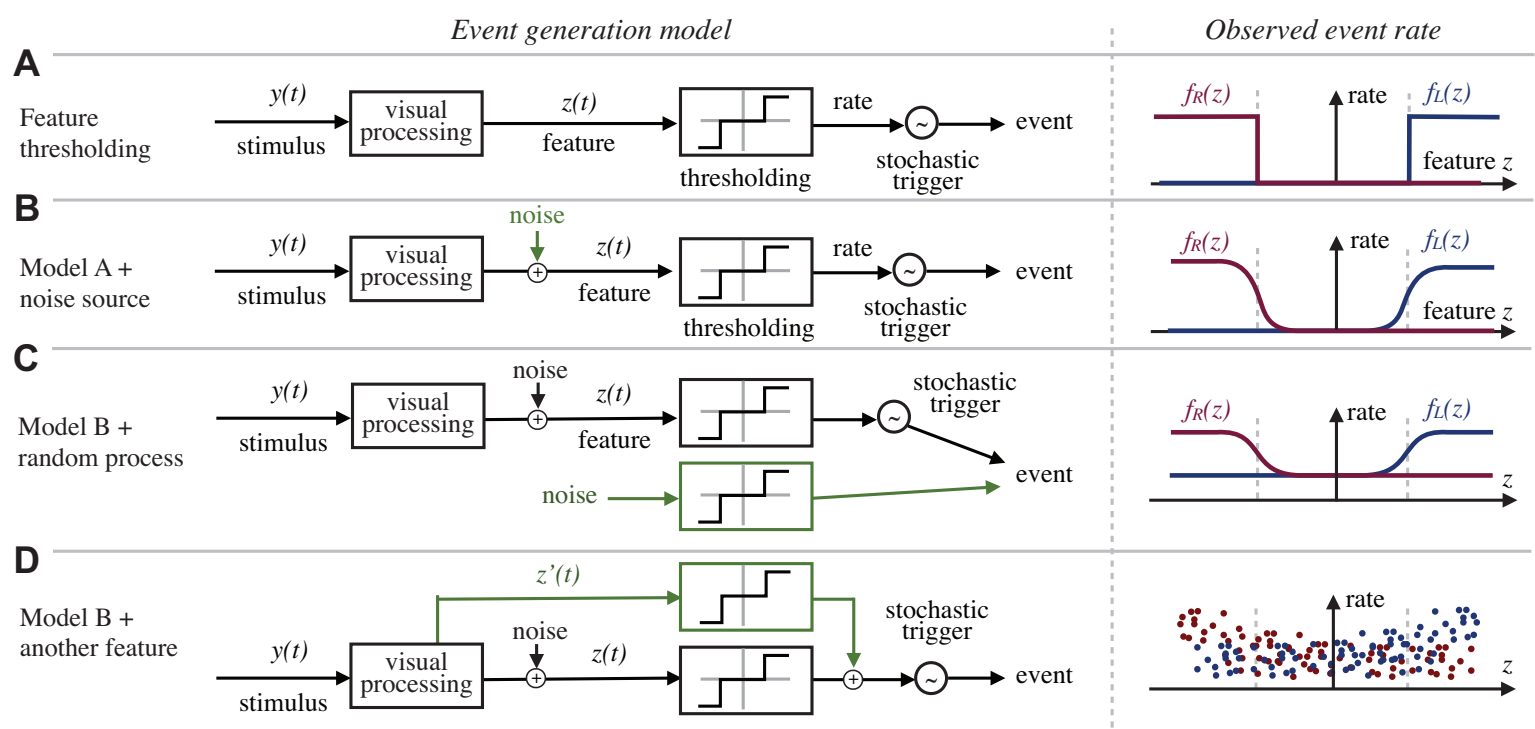


Figure 4
[Click here to download Figure: Figure4.eps](#)

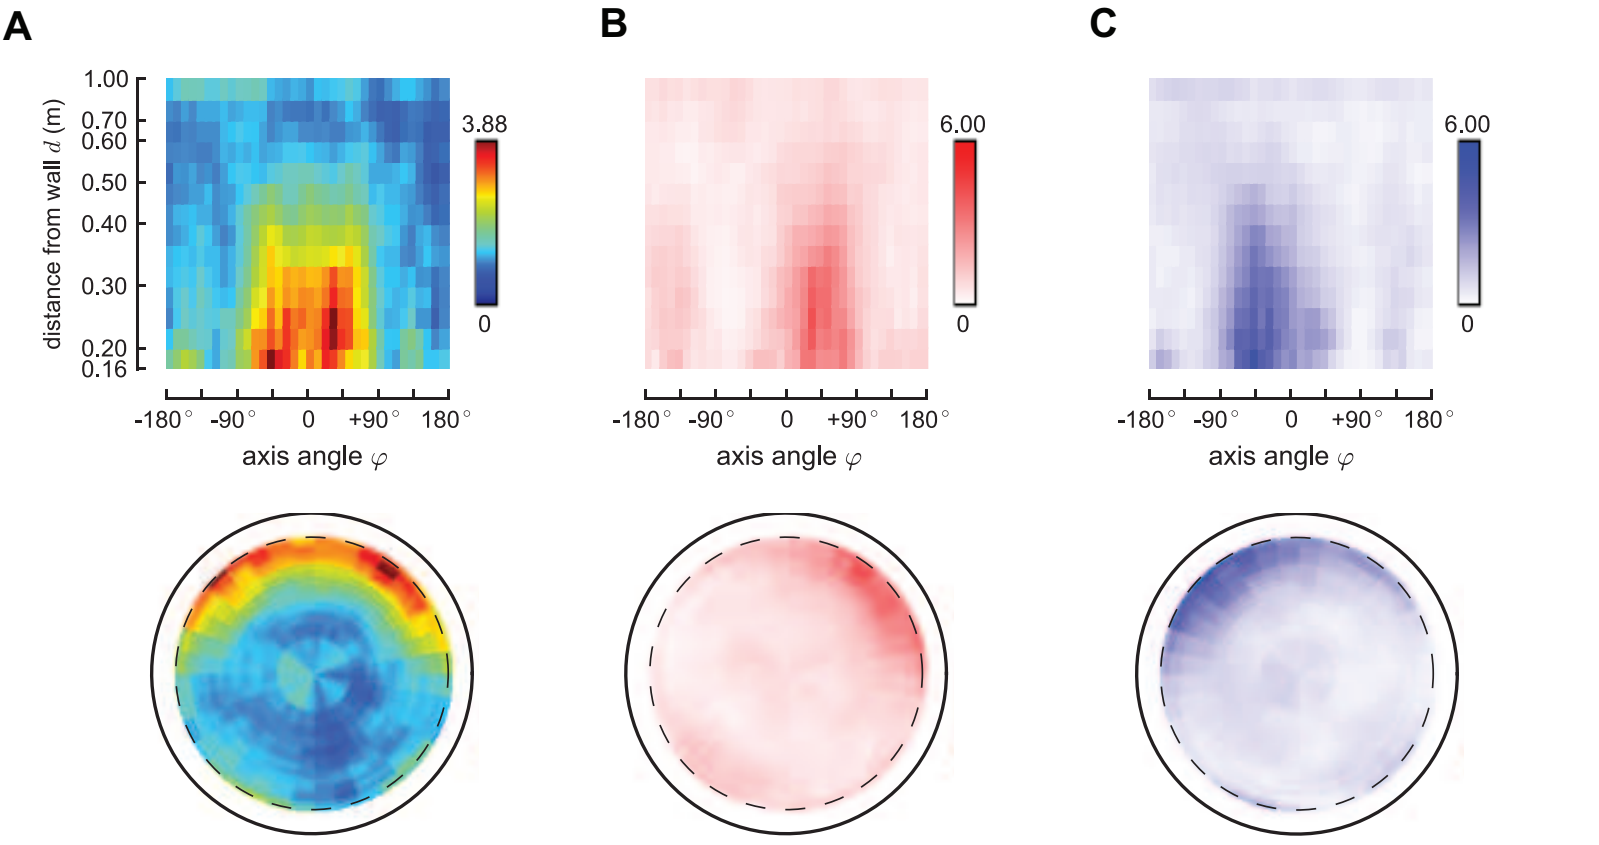


Figure 5
[Click here to download Figure: Figure5.eps](#)







Figure 6 Least squares solution
[Click here to download Figure: Figure6.eps](#)


A




B




C





D






E



F







0001
0002
0003
0004
0005
0006
0007
0008
0009
0010
0011
0012
0013
0014
0015
0016
0017
0018
0019
0020
0021
0022
0023
0024
0025
0026
0027
0028
0029
0030
0031
0032
0033
0034
0035
0036
0037
0038
0039
0040
0041
0042
0043
0044
0045
0046
0047
0048
0049
0050
0051
0052
0053
0054
0055
0056
0057
0058
0059
0060
0061
0062
0063
0064
0065
0066
0067
0068
0069
0070
0071

**DISCRIMINATING EXTERNAL AND INTERNAL CAUSES FOR HEADING
CHANGES IN FREELY FLYING *DROSOPHILA***

ANDREA CENSI, ANDREW D. STRAW, ROSALYN W. SAYAMAN, RICHARD M. MURRAY, MICHAEL H.
DICKINSON

Text S1

CONTENTS

1. Problem formalization	2
2. Modeling assumptions	5
3. Observability analysis	7
4. Estimation of event generation rates	8
5. Feature identification	10
6. Algorithm summary	16
7. Details of saccade detection algorithms	17
8. Guide to the experimental results	18
9. Complete plots (Geometric saccade detector)	22
10. Complete plots (Angular-velocity based detector)	28
References	33

1. PROBLEM FORMALIZATION

There are 5 kinds of objects in our analysis: *states*, *stimuli*, *features*, *events*, and *models*.

states States can be both *observable* and *unobservable*. This distinction is largely practical in nature, and dependent on the technology available at the moment.

Examples of observable states: position/velocity of the animal.

Examples of unobservable states: metabolic states, neural states.

stimuli These are all the sensory cues available to the animal for decision making. Stimuli are a function of the states. In this work, stimuli are abstract entities: they are assumed to exist, as a function of the states, but they are never computationally manipulated.

Examples of stimuli: perceived luminance at each photoreceptor, perceived odor traces, perceived acceleration.

features We define “features” as the behaviorally-relevant low-dimensional functions of the stimuli that are used for decision making. The existence of these features is postulated. The theoretical justification to investigate low-dimensional functions of the stimuli is that, while the stimuli are very high-dimensional, the decisions are usually very low-dimensional. Therefore, only a low-dimensional feature of the stimuli can possibly contribute to behaviors.

Example of feature: left/right optic flow imbalance.

(behavioral) events Behavioral events (or simply: events) are the external manifestations of decision making that we can observe. Our formulation applies to behaviors that can be clearly identified in time.

Examples of events: start of a saccades, landing, taking off.

(behavioral) models These are generative models that explain the observed events, as a function of the external stimuli and the internal states.




FIGURE 1.1.

1.1. Behavioral pathway of interest in this work. Naturally, a complete understanding of animal behavior can only be attained by considering all possible stimuli and all observable behaviors. However, the complexity of a model that can be reliably identified is bounded by the amount, diversity, and quality of the data that can be collected. Therefore, in practice, one can only consider a limited “behavioral pathway” that interests only a subset of the stimuli and a subset of the behaviors. We limit our analysis to the visual stimulus, which is a function of the animal position, and we take discrete body saccade events as the observable manifestation of behavior.

Whether the particular stimulus considered (visual stimulus) is sufficient to build a model for the chosen behavior is something that is not justified *a priori*, but rather will have to be confirmed *a posteriori* by the analysis.

1.2. Observable states. We call $x(t) \in X$ the observable animal *configuration*, which includes the position and linear/angular velocity in 3D space, and X the *configuration space*. The space X has dimension 12 (6 degrees of freedom for position/orientation, plus corresponding velocities).

An estimate of $x(t)$ is provided by our tracking software. While the observations are occasionally very noisy, we assume that $x(t)$ is observable, and we do not model the noise on x in this analysis. In the following, it will be clear that this does not impact the analysis much, as the method is very robust to bounded noise on $x(t)$.

The tracking system provides data at 60Hz. We ignore this time discretization and assume for convenience that $t \in \mathbb{R}$.

TABLE 1. Symbols used in this paper

<i>Symbol</i>	<i>meaning</i>
$t \in \mathbb{R}$	Time index.
$x(t) \in X$	Observable configuration (position and velocity).
$\xi(t)$	Unobservable states.
$y(t)$	Perceived stimuli.
C	Reduced configuration space
$c(t) \in C$	Reduced configuration at time t .
$r_i(t)$	Instantaneous event generation rates.
Δ	Inhibition interval.
$i \in \{L, R\}$	Index over event classes.
$t_i^1, t_i^2, \dots, t_i^j, \dots$	Detected events of the i -th class ($i \in \{L, R\}$).
j	Index over discrete events in time
$b_i(t)$	Detected events as a series of impulses in time.
f_i	Event generation rate function.
$\delta(t - \bar{t})$	Impulse function centered at \bar{t} .
$\{C^k\}_{k=1}^K$	Partition of C in K cells.
$k \in [1, K]$	Index over cells.
$c^k \in C$	Center of k -th cell.
n_i^k	Number of detected events in k -th cell.
n^k	Total number of detected events in k -th cell.
m_i^k	Measured event rate in the k -th cell.
r_i^k	Average event generation rate for i -th behavior in k -th cell.
$\mathbf{r}_i = \{r_i^k\}_{k=1}^K$	Vector containing all event generation rates over cells.
z^k	Feature associated to k -th cell.
$\mathbf{z} = \{z^k\}_{k=1}^K$	Vector containing all feature values over cells.
order	Function computing the order of each element in a vector.
$a \circ b$	Composition of the functions a and b .
$\text{Diff}(\mathbb{R})$	Diffeomorphisms of \mathbb{R} ; set of strictly monotone functions.
$\mathcal{N}(\mu, \sigma)$	Gaussian distribution with mean μ and standard deviation σ .
$\text{Unif}(a, b)$	Uniform distribution on the interval $[a, b]$.

1.3. Unobservable states. We call $\xi(t)$ the set of all other states that are behaviorally relevant but not observable:

$$\xi(t) \triangleq \text{all behaviorally relevant states} \setminus x(t).$$

These unobservable states include:

- unobservable kinematic properties (e.g., angle of fly neck);
- metabolic states that influence behavior (e.g., hunger);
- other unmodeled properties of the environment (residual odor traces, etc.).

1.4. Behavioral events. The observable behavioral events that we consider are the so-called “body saccades”, or just “saccades”. These are the moments where a fly turns rapidly in the horizontal plane.

There are two classes of events: left (L) and right (R) saccades. We limit our analysis to only the saccade direction; in principle, there are many observable properties of a saccade that would be interesting to analyze (top angular velocity, amplitude, etc.); those are left for future work.

The index “ i ” will be used to index over event classes: $i \in \{L, R\}$. Every time we write a formula with the generic index i , it is understood that $i \in \{L, R\}$.

The criterion for defining a saccade and the algorithms for detecting it are discussed in Section 7. Here, we assume that we have a procedure that, given a trajectory $x(t)$, returns a series of detected saccade events, represented by their initiation times $\{t_i^j\}$:

$$\text{detected left saccades: } t_L^1, t_L^2, \dots, t_L^j, \dots$$

$$\text{detected right saccades: } t_R^1, t_R^2, \dots, t_R^j, \dots$$

The index “ j ” is an index over the sequence of events.

We assume that these events are observable (i.e. we detect all saccades perfectly). In practice, we found that our saccade detection algorithms give false positives or false negatives on less than 1% of the data against a subset of the data manually annotated; this error can be neglected for our analysis, because all the statistics that we compute are essentially averages over spatial patches, therefore robust to random failures of the detection procedure.

We also introduce the variables $b_L(t), b_R(t)$. These variables are an alternative representation of the event sequences as trains of impulses:

$$b_i(t) = \sum_j \delta(t - t_i^j).$$

Here, $\delta(t - \bar{t})$ is the impulse centered at time \bar{t} . This notation will be convenient for further manipulations.

1.5. Stimulus and reduced configuration. We call $y(t)$ the set of all external stimuli perceived by the animal at time t . We assume that $y(t)$ is a function h of the animal configuration $x(t)$, corrupted by a noise process $v(t)$:

$$y(t) = h(x(t), v(t)).$$

In our case, $y(t)$ would be the luminance perceived on the retina, plus odor traces, and other sensory cues depending on position. The stimulus is also a function of the environment shape/textures/etc, which in the paper is indicated as “ W ”. However, because the environment is considered fixed, we omit it from the formulation.

We do not assume that $y(t)$ is observable, nor we are interested in estimating it; rather, we use this variable only as theoretical device to formalize our intuition of *environment symmetry*. We say that an environment has a symmetry if there exist two distinct points $x_1, x_2 \in X$ such that

$$h(x_1, v) = h(x_2, v).$$

If this is the case, then it makes sense to compress the state $x(t) \in X$, to a smaller, minimal representation $c(t) \in C$. We call $c(t)$ the *reduced configuration*, and C the *reduced configuration space*. We assume that C is given, along with a projection map $\pi : X \rightarrow C$, which maps the point $x(t)$ to $\pi(x(t)) = c(t) \in C$. In other words, the map π extracts from the whole state $x(t)$ the variables that are necessary to determine the stimulus. Moreover, we assume that $c(t)$ is a minimal representation, in the sense that, for all points x_1, x_2 :

$$\pi(x_1) = \pi(x_2) \quad \Leftrightarrow \quad h(x_1, v) = h(x_2, v).$$

In other words, $c(t)$ is sufficient to compute the stimulus (if one knew the function h), and it cannot be further reduced to an even smaller representation.

In our particular case, we have a circular symmetry, as the environment is a circular arena with uniform patterns on the walls. If one assumes X to describe position, orientation and linear/angular velocity of the animal, the circular symmetry allows to decrease by 1 the number of degrees of freedom of the data, by projecting the 12-D space X to a 11-D space. The reduced configuration stimulus depends on the details of each experiment and on the experimenter’s assumptions. For example, even if the arena was circular, the symmetry would not be valid if there were non symmetric patterns on the walls.

We consider the reduced configuration space and the projection map as other modeling choices whose validity must be justified *a posteriori*. Our assumptions regarding the reduced configuration space are explained in Section 2.5.

2. MODELING ASSUMPTIONS

This section introduces all modeling assumptions of our analysis and the implied approximations.

2.1. Behavioral events are generated by time-variant interacting Poisson processes. We model the generation of observable events as a set of interacting Poisson processes with time-variant rates $r_i(t)$ and inhibition interval Δ :

$$\{b_i(t)\}_{i \in \{L,R\}} \sim \text{InteractingPoisson}(\{r_i(t)\}_{i \in \{L,R\}}, \Delta).$$

This means that, once any process generates one event at time t , no other event can be generated in the interval $[t, t + \Delta]$, from the same or another process. A Poisson process with time-variant rate is sometimes called a Cox process [1].

2.1.1. Interpretation of the inhibition interval Δ . The inhibition time is meant to model the fact that we consider instantaneous events that correspond to the initiation of a motor program, and one motor program must complete before another can be initiated.

2.2. The event generation rates depend on the instantaneous stimulus, unobservable states, and endogenous random process. We assume that event generation rate $r_i(t)$ can be written as the sum of three terms r_i^S , r_i^E , r_i^R , which model the contributions of the stimulus, the unobservable states, and an endogenous random process:

$$(2.1) \quad r_i(t) = \underbrace{r_i^S(y(t))}_{\substack{\text{instantaneous} \\ \text{stimulus} \\ \text{contribution}}} + \underbrace{r_i^E(\xi(t))}_{\substack{\text{unobservable} \\ \text{states} \\ \text{contribution}}} + \underbrace{r_i^R}_{\substack{\text{endogeneous} \\ \text{random process} \\ \text{contribution}}}.$$

2.3. The high-dimensional stimulus y can be compressed down to a smaller feature z . The term $r_i^S(y)$ models the contribution of the stimulus to the event generation. We assume that this contribution can be written as a function of a low-dimensional feature $z \in Z$ of the stimulus y :

$$r_i^S(y) = f_i(z(y)).$$

We call the functions f_i the *event generation rates functions*.

2.4. The feature has monotone effect on the two behaviors. We introduce another constraint on the two functions f_R, f_L that allows to obtain a closed-form solution to the identification problem. This assumption is very specific to the particular problem studied. We assume that:

- The function f_L is monotonically increasing:

$$(2.2) \quad f_L(z_1) > f_L(z_2) \Leftrightarrow z_1 > z_2.$$

- The function f_R is monotonically decreasing:

$$(2.3) \quad f_R(z_1) > f_R(z_2) \Leftrightarrow z_1 < z_2.$$

These assumptions can be verified a posteriori after fitting the data.

2.5. Choice of the reduced configuration space C . Choosing the reduced configuration space C is a critical step of the analysis. There is a clear tradeoff: if the C space is too small, then we cannot represent the variability of behavior. If it is too large, then the data that we have will not be dense enough to have accurate statistics. This is the reason we do not use directly the values of $x(t)$.

In our case, we project the original 12 degrees of freedom space X to only two. The original degrees of freedom are, position (3), attitude (3), linear velocity (3), and angular velocity (3). We make the following assumptions:

- (1) The altitude can be ignored (removes 1 dof for the z position, and the corresponding 1 dof for velocity).
- (2) As for the attitude, only the yaw is relevant (removes 2 dof).

- (3) The angular velocity is ignored; the saccades start when the fly is flying approximately straight (removes 3 dof).
- (4) The fly has negligible sideways velocity (removes 1 dof).
- (5) The linear forward velocity is not relevant (removes 1 dof).

These assumptions bring the degrees of freedom from 12 to 3. The circular symmetry of the arena allows to further reduce the analysis to a 2D space. There is no particular algorithmic complication if the space has more than 2 dimensions. Given the amount of data we have, a 2D space allows the recorded trajectories to fill the space densely enough to compute the statistics necessary.

Note that the particular *parametrization* of the space C does not matter. In the paper we use two different parametrizations to display the results.

2.6. From time-variant to timeless spatial quantities. In the following, it will be convenient to write the rates r_i^S as a function of the configuration rather than as function of time. This is possible because we have assumed that $y(t)$ is a function of the reduced configuration $c(t) \in C$.

Consequently, also the feature z is a function of the reduced configuration. We call $\gamma : C \rightarrow Z$ the function that assigns a value of the feature to each reduced configuration:

$$z = \gamma(c).$$

2.7. Reduced model identified. While we assume (2.1) as the true model for the rate r_i , we will identify a reduced model that lumps together the contributions of the unobservable states and the endogenous random process. We cannot distinguish the contributions of r_i^E and r_i^R because we cannot access the unobservable states. However, we can show that we can still identify the stimulus-dependent term r_i^S by averaging over the trajectories.

We will identify the model:

$$(2.4) \quad r_i(c) = \underbrace{f_i(z(c))}_{\substack{\text{instantaneous} \\ \text{feature} \\ \text{contribution}}} + \underbrace{r_0}_{\substack{\text{unobservable states} \\ + \text{endogeneous process} \\ \text{average contribution}}},$$

where r_0 lumps together the contribution unobservable states and endogenous process.

To arrive at (2.4) from (2.1), we just need to average over time. Fix a point c_* and let C_* be a small area containing c_* . Define $\mathbb{E}_{c(t) \in C_*} \{ \cdots \}$ as the average of a quantity limited to those times in which $c(t) \in C_*$. Then, by averaging the observed instantaneous rate in C_* , we obtain

$$\mathbb{E}_{c(t) \in C_*} \{ r_i(t) \} = \mathbb{E}_{c(t) \in C_*} \{ f_i(z(c)) \} + \mathbb{E}_{c(t) \in C_*} \{ r_i^E(\zeta(t)) \} + r_i^R.$$

Define $r_i(c_*) = \mathbb{E}_{c(t) \in C_*} \{ r_i(t) \}$ to be the average rates around the point c_* .

If we assume that the feature is a continuous function of the configuration and that the neighborhood C_* is small enough, then we can approximate:

$$\mathbb{E}_{c(t) \in C_*} \{ f_i(z(c)) \} \simeq f_i(z(c_*)).$$

If we assume that the unobservable states are uncorrelated with the external configuration, then the value of the expectation $\mathbb{E}_{c(t) \in C_*} \{ r_i^E(\zeta(t)) \}$ does not depend on the configuration C_* :

$$\mathbb{E}_{c(t) \in C_*} \{ r_i^E(\zeta(t)) \} = \mathbb{E} \{ r_i^E(\zeta(t)) \} = \bar{r}_i^E$$

Therefore, by defining

$$r_0 = r_i^R + \bar{r}_i^E$$

we arrive at (2.4).

3. OBSERVABILITY ANALYSIS

Our first step towards identification of the model is an observability analysis. We first give some remarks on the dimensionality of the feature that we can identify; then we show that, fixed the dimensionality, the problem is still underconstrained because there are multiple solutions that satisfy the constraints.

3.1. Dimensionality of the feature. We can show that the dimension of the feature z such that the problem is well posed is bounded by the dimension of the reduced configuration space C and the number of event classes.

Proposition 1. *For the identification problem to be well-posed, we must have*

$$\dim(Z) < \min\{\dim(C), n_{\text{events}}\}.$$

In particular, for two event classes and a two-dimensional configuration space, we can only estimate a one-dimensional feature.

Proof. Our sets of constraints is

$$\begin{cases} z &= \gamma(c), \\ r_i^S(c) &= f_i(\gamma(c)), \quad i = 1, \dots, n_{\text{events}}. \end{cases}$$

Geometrically, we have that the function γ maps C to Z , and then the function $f = \{f_i\}$ maps Z to $\mathbb{R}^{n_{\text{events}}}$:

$$C \xrightarrow{\gamma} Z \xrightarrow{f} \mathbb{R}^{n_{\text{events}}}.$$

For the estimation problem to be well posed, we need to have the hourglass structure

$$\dim(C) > \dim(Z) < \dim(\mathbb{R}^{n_{\text{events}}}),$$

otherwise there are some trivial solutions.

(1) In the case

$$\dim(C) = \dim(Z) < \dim(\mathbb{R}^{n_{\text{events}}}),$$

we can choose $\gamma = \text{Identity}$ and $f_i(z) = r_i^S(c)$ as a trivial solution.

(2) In the case

$$\dim(C) > \dim(Z) = \dim(\mathbb{R}^{n_{\text{events}}}),$$

we can choose $f_i = \text{Identity}$, $z(c) = \{r_i^S(c)\}$ as a trivial solution.

□

3.2. Observability. The model is not fully observable, in the sense that we can find multiple solutions for the parameters that fit the data equally well. This is formalized in the following proposition.

Proposition 2. *Suppose that the configuration space is discretized in K cells, each with center $c^k \in C$. Let the bold vectors $\mathbf{r}_L, \mathbf{r}_R \in \mathbb{R}_+^K$ represent event generation rates associated to different animal configurations ($\mathbf{r}_i = \{r_i^k\}$, where k ranges over configuration), and similarly let $\mathbf{z} = \{z^k\}$ be the feature associated to the configurations $\mathbf{c} = \{c^k\}$. Assume that the model postulated so far holds exactly. Then we can write the constraints in vector form as:*

$$(3.1) \quad \mathbf{r}_L = \mathbf{f}_L(\mathbf{z}),$$

$$\mathbf{r}_R = \mathbf{f}_R(\mathbf{z}),$$

$$(3.2) \quad \mathbf{z} = \boldsymbol{\gamma}(\mathbf{c}).$$

Assume that the rates $\mathbf{r}_L, \mathbf{r}_R$ and the reduced configurations \mathbf{c} are *observable*, and that the functions $\mathbf{f}_L, \mathbf{f}_R, \boldsymbol{\gamma}$ are *unknown a priori*. Then it is possible to estimate \mathbf{z} only up to a monotone transformation, in the sense that it is not possible to distinguish between a solution \mathbf{z}_1 and a solution \mathbf{z}_2 , if $\mathbf{z}_1 = \alpha(\mathbf{z}_2)$ where $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly monotone function. Consequently, one cannot distinguish between the event generation rate functions $(\mathbf{f}_L, \mathbf{f}_R)$ and $(\mathbf{f}_L \circ \alpha^{-1}, \mathbf{f}_R \circ \alpha^{-1})$, where “ \circ ” denotes function compositions.

Proof. A “solution” of the system is a tuple (f_L, f_R, γ, z) for which the constraints (3.1)–(3.2) hold. If the system of constraints was completely observable, there could be only one solution. However, suppose $s = (f_L, f_R, \gamma, z)$ is a solution, and consider an invertible function α , and the solution tuple s^α defined as

$$\begin{aligned} s^\alpha &= (f_L^\alpha, f_R^\alpha, \gamma^\alpha, z^\alpha) \\ &= (f_L \circ \alpha^{-1}, f_R \circ \alpha^{-1}, \alpha \circ \gamma, \alpha(z)). \end{aligned}$$

One can verify that, assuming s is a solution, s^α is another solution of the system, because it respects the constraints (3.1)–(3.2):

$$\begin{aligned} r_L &= f_L(z) = f_L(\alpha^{-1}(\alpha(z))) = f_L^\alpha(z^\alpha), \\ r_R &= f_R(z) = f_R(\alpha^{-1}(\alpha(z))) = f_R^\alpha(z^\alpha), \\ z^\alpha &= \alpha(z) = \alpha(\gamma(c)) = \gamma^\alpha(c). \end{aligned}$$

□

This is the formal way to show that there is an ambiguity. A more intuitive way to see the same thing is by rewriting (3.1)–(3.2) in a slightly different way. Knowing γ is equivalent to knowing z , so we can write the constraints in terms of γ only:

$$(3.3) \quad \begin{aligned} r_L &= f_L(\gamma(c)), \\ r_R &= f_R(\gamma(c)). \end{aligned}$$

Intuitively, we have “2 equations for 3 variables”; because γ always appears composed together with f_L and f_R , it cannot be observed independently.

3.3. Interpretation of the unobservability. If the reduced feature is defined only up to a monotone transformation, then it should not be thought as a physical quantity, or as a measure of physiological activity. In fact, we cannot associate a meaningful measurement unit to it. Rather, the feature represents an *ordering* of the configurations (in fact, the order is what is conserved by any monotone function). All we can say is whether in a certain configuration the feature is weaker or stronger than in another.

Still, from the point of view of the analysis and the visualization, it is useful to consider z as a real-valued quantity associated to each configuration. From the point of view of the estimation, we cannot distinguish between z and $\alpha(z)$; therefore, we can choose any particular function α for visualizing the feature. We will choose a function α so that z varies between -1 and $+1$ across the environment.

4. ESTIMATION OF EVENT GENERATION RATES

In the first part of the algorithm, we estimate the event generation rates $\{r_i^k\}$, where $i \in \{L, R\}$, and k ranges over configurations. This operation is slightly more complicated than just dividing the number of detected events by time, because the events are generated by Poisson processes that interact with each other. For example, this means that the rate of observed left saccade events depends not only on r_L , but also the rate r_R : if r_R is very high, then we expect to see fewer left saccade events, as there is more inhibition.

4.1. Statistics collected from the data. Divide the reduced configuration space C in cells $\{C^1, C^2, \dots, C^k, \dots, C^K\}$, possibly overlapping, and each with center $c^k \in C$. It is assumed that the discretization is small enough to capture the variability of behavior, but large enough so that the samples are dense enough.

For each cell, we compute the following statistics from the data:

n_i^k	The number of events of the i -th class detected in the k -th cell.
T^k	The time spent in the k -th cell.

If one defines the variable $I^k(t)$ as

$$I^k(t) = \begin{cases} 1 & \text{if } c(t) \in C^k, \\ 0 & \text{otherwise.} \end{cases}, \quad \text{for } k \in [1, K].$$

then the two quantities n_i^k and T^k can be written as follows:

$$\begin{aligned} n_i^k &= \int I^k(t) b_i(t) dt, \\ T^k &= \int I^k(t) dt. \end{aligned}$$

We define also the following auxiliary statistics:

n^k The total number of events detected in the cell:

$$n^k = \sum_i n_i^k = n_L^k + n_R^k$$

m_i^k The *measured event rates*, given by

$$m_i^k = \frac{n_i^k}{T^k}.$$

m^k The total event rate m^k per cell:

$$m^k = \sum_i m_i^k = m_L^k + m_R^k.$$

4.2. Robustness of statistics to measurements noise. These are all the statistics that we need from the data. Note that the position $x(t)$ need to be only accurate enough so that it can be assigned to the correct cell. In particular, we do not need to compute higher derivatives of $x(t)$, therefore operations like smoothing are not necessary (some smoothing might be necessary to detect the saccade events). Also, misdetection of saccade events does not impact much the analysis; a 1% false positive/negative rate only changes the statistics m_i^k by 1%.

4.3. Measured saccade rates vs. saccade generation rates. Due to the inhibition period $\Delta > 0$, the measured event rates underestimate the event generation rates: $m_i^k < r_i^k$. This is true even in the case of 1 process, and the effect is more evident with multiple processes.

4.3.1. Event generation rate estimation for one process without inhibition. Consider first the case of only one Poisson process $b(t)$, with event generation rate r , and no inhibition period. Suppose we observe the process for T seconds, counting n events, obtaining the measured rate $m = n/T$. Then m is the maximum likelihood estimate of r .

4.3.2. Event generation rate estimation for one process with inhibition. However, if there is a non-zero inhibition period, Δ , then the measured rate m underestimates the true rate. For example, the observed rate m cannot be higher than $1/\Delta$, no matter how high the rate r is, because due to inhibition we can observe at most one event every Δ seconds.

A better estimate of the rate r can be found as follows:

$$\hat{r} = \frac{n}{T - \Delta n}.$$

The interpretation is easy: if there were n events, then the process was inhibited for Δn seconds. Therefore, $T' = T - \Delta n$ is the effective time in which the process was active and could generate events.

The same formula can be written in terms of the measured rate:

$$\hat{r} = \frac{n}{T(1 - \frac{\Delta n}{T})} = \frac{m}{1 - \Delta m}.$$

Note that $r \rightarrow m$ as $\Delta \rightarrow 0$: if there is no inhibition, the measured rate is the generation rate.

4.3.3. *Event generation rate estimation for multiple processes with inhibition.* Next, consider the case in which there are multiple processes $\{b_i(t)\}$ inhibiting each other. It is easy to see that not taking into account the inhibition can strongly skew the estimate. In fact, imagine that there is one process with very high rate. The rates of the other processes will be severely underestimated because the frequent process will often inhibit them. Fortunately, also in this case we can easily normalize the rates.

Suppose we observe a series of processes $\{b_i(t)\}$ over an interval T and we count n_i events for each process. Let $n = \sum_i n_i$ be the total number of events observed. Then we know that the effective time in which the processes were not inhibited is $T - n\Delta$ (note that according to this model, it does not matter which process inhibits which). Accordingly, the event generation rates can be estimated as

$$(4.1) \quad \hat{r}_i = \frac{n_i}{T - \Delta \sum_i n_i},$$

or, writing it as a function of the measured event rates:

$$(4.2) \quad \hat{r}_i = \frac{m_i}{1 - \Delta \sum_i m_i}.$$

Note that, as $\Delta \rightarrow 0$, $r_i \rightarrow m_i$ and the estimate of the rate of one process does not depend on the measured rates of the others. Also note that the correction factor necessary for one process depends on the cumulative intensity of all the others.

4.4. Estimation of confidence bounds for the event generation rates. The formulas (4.1)-(4.2) give the maximum likelihood estimators for the event generation rates. One can estimate confidence intervals using the method discussed in Guerriero *et al.* ed[2]. Upper and lower 95% confidence bounds are found as follows:

$$(4.3) \quad r \in [\underline{r}, \bar{r}] = \left[\frac{\left(1 - \frac{1.96}{\sqrt{n-1}}\right) n}{T - \Delta \sum_i n_i}, \frac{\left(1 + \frac{1.96}{\sqrt{n-1}}\right) n}{T - \Delta \sum_i n_i} \right].$$

5. FEATURE IDENTIFICATION

At this point, we assume to have an estimate of the event generation rates r_i^k . We write again our model with explicit dependence on the cell c^k :

$$\begin{aligned} r_L^k &= f_L(z^k), \\ r_R^k &= f_R(z^k), \\ z^k &= \gamma(c^k). \end{aligned}$$

Our goal is to find z . Once z is known, then the three functions f_L, f_R, γ can be recovered from these equations. For simplicity, we first assume that we know the event generation rates r_i^k precisely. Sections 5.1 and 5.4 show how to estimate z under this simplifying assumption. Then, Sections 5.3 and 5.4 show how to take into account the uncertainty in r_i^k .

5.1. The order(\cdot) function and its properties. We introduce the order(\cdot) function, sometimes also called “rank”.

Definition 3. Define the function $\text{order} : \mathbb{R}^K \rightarrow \text{Perm}(K)$, which takes a vector in \mathbb{R}^K and associates to it a permutation of length K , that gives the order of each element in the sequence.

For example, we would have

$$\text{order}([10, 20, 30]) = [0, 1, 2],$$

and

$$\text{order}([100, 3.14, 42]) = [2, 0, 1].$$

We will need three simple properties of this function.

Proposition 4. *Properties of the order(\cdot) function:*

- (1) *The order of the elements of a vector does not change if a strictly increasing function is applied to the vector.*

Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^K$, let $\beta : \mathbb{R} \rightarrow \mathbb{R}$ be a strictly increasing function, and let $b^k = \beta(a^k)$. Then

$$\text{order}(\mathbf{b}) = \text{order}(\beta(\mathbf{a})) = \text{order}(\mathbf{a}).$$

- (2) *Applying a strictly decreasing function inverts the order in the vector.*

Let $\mathbf{a}, \mathbf{w} \in \mathbb{R}^K$, let $\chi : \mathbb{R} \rightarrow \mathbb{R}$ be a strictly decreasing function, and let $w^k = \chi(a^k)$. Then

$$\text{order}(\mathbf{w}) = \text{order}(\chi(\mathbf{a})) = K - \text{order}(\mathbf{a}).$$

- (3) *Two vectors with elements in the same order are equivalent up to a increasing function.*

Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^K$ and suppose that $\text{order}(\mathbf{x}) = \text{order}(\mathbf{y})$. Then there exists a strictly increasing function $\psi : \mathbb{R} \rightarrow \mathbb{R}$ such that $y^k = \psi(x^k)$.

5.2. Estimating the reduced stimulus, assuming that there are no uncertainties. As a first propaedeutic step, assume that the values of r_K^i are known exactly without uncertainty. From the relation $\mathbf{r}_L = f_L(\mathbf{z})$, the assumption that f_L is strictly increasing, and property (1) of Proposition 4, one obtains that

$$(5.1) \quad \text{order}(\mathbf{r}_L) = \text{order}(\mathbf{z}).$$

By applying the order function on both sides of the equality, we were able to simplify f_L from the expressions, because applying a strictly increasing function does not change the order of the data.

Similarly, from the relation $\mathbf{r}_R = f_R(\mathbf{z})$, the assumption that f_R is strictly decreasing, and property (2) in Proposition 4, one obtains that

$$(5.2) \quad \text{order}(\mathbf{r}_R) = K - \text{order}(\mathbf{z}).$$

Here, because f_R is decreasing, the order of the elements is reversed.

At this point, we can use (5.1) and (5.2) to obtain an overdetermined system of equations for $\text{order}(\mathbf{z})$. The least square solution is obtained by simply averaging the two terms:

$$(5.3) \quad \begin{aligned} \hat{\text{order}}(\mathbf{z}) &= \text{estimate of order}(\mathbf{z}) \\ &= \frac{1}{2} \text{order}(\mathbf{r}_L) + \frac{1}{2} (K - \text{order}(\mathbf{r}_R)) \\ &= \frac{1}{2} \text{order}(\mathbf{r}_L) + \frac{1}{2} \text{order}(-\mathbf{r}_R). \end{aligned}$$

By property (3) of Proposition (4), we know that knowing $\text{order}(\mathbf{z})$ is equivalent to knowing \mathbf{z} up to a diffeomorphism:

$$\text{order}(\mathbf{z}) = \gamma(\mathbf{z}) \quad \text{for some } \gamma \in \text{Diff}(\mathbb{R}).$$

By the observability analysis of the problem (Proposition 2), we know that we can estimate \mathbf{z} only up to a diffeomorphism, therefore we are done and use as our estimate of the feature:

$$\hat{\mathbf{z}} = \hat{\text{order}}(\mathbf{z}).$$

Because $\hat{\mathbf{z}}$ is determined only up to a diffeomorphism, for purely esthetic reasons we can normalize it in the $[-1, +1]$ range, by setting

$$(5.4) \quad \hat{\mathbf{z}}' = \frac{\hat{\text{order}}(\mathbf{z}) - K/2}{K}.$$

Once we know an estimate of \mathbf{z} , the shape of the functions f_L and f_R can be obtained directly from the relations $\mathbf{r}_L = f_L(\mathbf{z})$, $\mathbf{r}_R = f_R(\mathbf{z})$.

This simplified procedure is valid only if the values \mathbf{r}_i are known without uncertainty. If uncertainty is present, then a slightly more complicated computation is needed, described in the next section. Most of the difficulty arises in understanding how uncertainty propagates through the order function.

5.3. Estimating the ranks of a collection of random variables. We now put uncertainty back in the picture. The values r_L, r_R are not known precisely; rather, we only have a posterior distribution estimated from the data. Here we want to show that the order function is very sensitive to noise, therefore the approximation (5.3) cannot be used directly.

We want to solve the following problem: given a collection of K random variables $\mathbf{X} = \{X^k\}_{k=1}^K$ with known probability distribution, estimate $\text{order}(\mathbf{X})$.

There are several unsatisfactory ways to solve this problem. One could just use the mean of the distributions to estimate the relative order:

$$(5.5) \quad \text{order}(\mathbf{X}) = \text{order}(\mathbb{E}\{\mathbf{X}\}).$$

However, this estimate is not satisfactory because it does not take into account the variance of the variables. An example of this situation is shown in the panels of Fig. 5.1. In this figure, we simulate a collection of random variables $\mathbf{X} = \{X^k\}$, where each random variable has a uniform distribution over an interval of length 0.1 with center $\exp(-0.02k)$:

$$X^k \sim \text{Uniform}(\exp(-0.02k) - 0.05, \exp(-0.02k) + 0.05)$$

The probability distribution of the variables is represented in Fig. 5.1a. The result of computing (5.5) is shown in Fig. 5.1b. The plot is a straight line because the means of the variables are already ordered by k .

Fig. 5.1c shows two realizations x_1, x_2 of the random variables \mathbf{X} , while Fig. 5.1d shows the order of the realizations $\text{order}(x_1)$ and $\text{order}(x_2)$. As one can see, the order of the variables changes dramatically, especially for large k , where the means of consecutive variables are very similar. This shows that applying the order operation to \mathbf{r} as in (5.3) is not a sensible way to estimate $\text{order}(\mathbf{r})$ if the data is noisy.

A reasonable estimate of $\text{order}(\mathbf{X})$, along with error bounds, can be obtained by simulation. Suppose that the distribution of \mathbf{X} is known:

$$p(X^k = x) = \Theta^k(x),$$

with Θ^k a known probability distribution. Then one can compute the distribution of $\text{order}(\mathbf{X})$ simply by drawing samples of \mathbf{X} from the known distribution and computing the observed order. More in detail, one computes a set of n samples x_1, x_2, \dots, x_n all with the same distribution: $x_i = \text{sample}(\Theta)$. Then the distribution of $\text{order}(\mathbf{X})$ can be approximated by the samples $\{\text{order}(x_i)\}$. In particular, we can derive mean and confidence bounds. This method is summarized as Algorithm 1.

Remark 5. This method is very simple and requires only the ability to draw samples from the distribution of \mathbf{X} . For completeness, we briefly mention the analytical difficulties to

Algorithm 1 OrderBySampling

Input:

$\Theta = \{\Theta^k\}_{k=1}^K$ The probability distributions of K random variables $\mathbf{X} = \{X^k\}_{k=1}^K$, such that $X^k \sim \Theta^k$.

Parameters:

N Number of simulations.

Output:

$\Psi = \{\Psi^k\}_{k=1}^K$: An estimate of the probability distribution of $\text{order}(\mathbf{X})$.

Algorithm:

function $\Psi = \text{OrderBySampling}(\Theta, N)$:

(1) For j in $1, \dots, N$:

(a) Draw the sample $x_j \sim \Theta$.

(b) Compute $o_j = \text{order}(x_j)$

(2) Compute the density Ψ^k as the observed distribution of $\{o_j^k\}_{j=1}^N$.

obtaining a more analytical solution. (Skipping this remark does not impact understanding of the rest of this document). The reader will have noticed that we did not give an analytical characterization of the distribution of $\text{order}(\mathbf{X})$. In theory, the distribution of $\text{order}(\mathbf{X})$ can be thought as a sum of binomial distributions. In fact, we have:

$$\begin{aligned} \text{order}(X^k) &= \text{number of variables in } \{X^j\}_{j=1}^K \text{ such that } X^k \geq X^j \\ &= \sum_{j=1}^K b_{jk}, \end{aligned}$$

where b_{jk} is a binomial variable defined as

$$b_{jk} = \begin{cases} 1 & \text{if } X^j \geq X^k, \\ 0 & \text{if } X^j < X^k. \end{cases}$$

The problem is that the variables b_{jk} are *not* independent. Therefore all the convenient theoretical results about sums of independent binomials cannot be used. Assuming that the distribution of \mathbf{X} is uni-modal, one can expect the distribution of $\text{order}(\mathbf{X})$ to be uni-modal as well, and if the number of variables K is large enough, a Gaussian approximation could be appropriate. For example, Fig. 5.1 shows the distribution of $\text{order}(X^{10})$, $\text{order}(X^{15})$, and $\text{order}(X^{40})$.

5.4. Estimating the reduced stimulus, taking into account the uncertainty in \mathbf{r} . We now refine the procedure in Section 5.4 taking into account the uncertainty in \mathbf{r} . Let Θ_L^k, Θ_R^k be the posterior distributions of r_L^k, r_R^k estimated from the data:

$$\begin{aligned} p(r_L^k | b_L) &= \Theta_L^k, \\ p(r_R^k | b_R) &= \Theta_R^k. \end{aligned}$$

Using Algorithm 1, we can estimate the distribution of $\text{order}(r_L^k), \text{order}(r_R^k)$; call the resulting distributions Φ_L^k and Φ_R^k .

$$\begin{aligned} p(\text{order}^k(\mathbf{r}_L) | b_L) &= \Phi_L^k = \text{OrderBySampling}(\Theta_L^k), \\ p(\text{order}^k(\mathbf{r}_R) | b_R) &= \Phi_R^k = \text{OrderBySampling}(\Theta_R^k). \end{aligned}$$

From (5.1) and (5.2) we can compute the corresponding probability distributions of $\text{order}(\mathbf{z})$, which we call Γ_L^k and Γ_R^k :

$$(5.6) \quad \Gamma_L^k(m) = p(\text{order}^k(\mathbf{z}) = m | b_L) = \Phi_L^k(m),$$

$$(5.7) \quad \Gamma_R^k(m) = p(\text{order}^k(\mathbf{z}) = m | b_R) = \Phi_R^k(K - m).$$

We use a Gaussian approximation for these densities:

$$(5.8) \quad \Gamma_L^k \simeq \mathcal{N}(\mu_L^k, \sigma_L^{2k}),$$

$$(5.9) \quad \Gamma_R^k \simeq \mathcal{N}(\mu_R^k, \sigma_R^{2k}).$$

Assuming a non-informative prior on $\text{order}^k(\mathbf{z})$:

$$p(\text{order}^k(\mathbf{z}) = m) = 1/K,$$

we can find the posterior distribution of $\text{order}(\mathbf{z})$ by fusing together (5.8)-(5.9), obtaining

$$\begin{aligned} (5.10) \quad \text{order}^k(\mathbf{z}) &\sim \mathcal{N}(\mu^k, \sigma^{2k}), \quad \text{with} \\ \mu^k &= \left(\frac{1}{\sigma_L^{2k}} + \frac{1}{\sigma_R^{2k}} \right)^{-1} \left(\frac{\mu_L^k}{\sigma_L^{2k}} + \frac{\mu_R^k}{\sigma_R^{2k}} \right), \\ \sigma^{2k} &= \left(\frac{1}{\sigma_L^{2k}} + \frac{1}{\sigma_R^{2k}} \right)^{-1}. \end{aligned}$$

Normalize the feature distribution in the $[-1, +1]$ range:

$$(5.11) \quad \begin{aligned} \bar{\mu}^k &= \frac{\mu^k - K/2}{K}. \\ \bar{\sigma}^k &= \frac{\mu^k}{K}. \end{aligned}$$

This gives the best estimate for the feature z^k .

At this point we can estimate the functions f_L, f_R directly using the relations

$$r_L^k = f_L(z^k) \quad \text{and} \quad r_R^k = f_R(z^k).$$

Note that the shape of f_L, f_R can be visualized directly by plotting the points (z^k, r_R^k) and (z^k, r_L^k) , along with confidence intervals (for both r_L^k, r_R^k and z^k), and it is not strictly necessary to impose some parametric form.

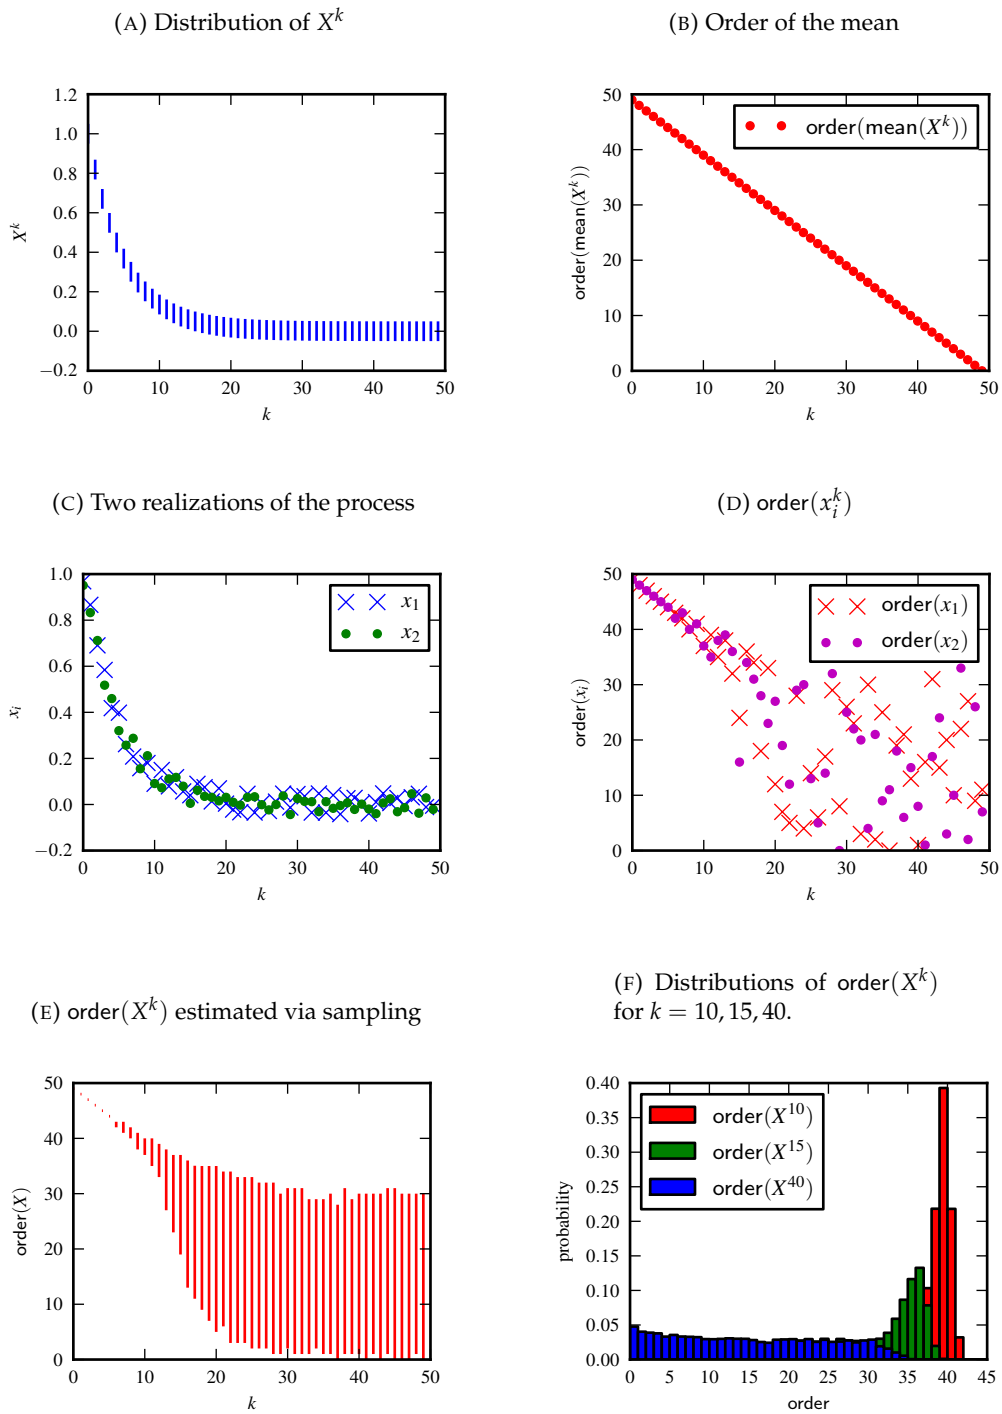


FIGURE 5.1. Synthetic data used to illustrate the properties of the order function and the order estimation procedure of Algorithm 1.

6. ALGORITHM SUMMARY

This is a summary of the identification algorithm.

Input:

- $x(t) \in X$ Recorded configuration.
 t_i^j The series of detected events, where $i \in \{L, R\}$ ranges over behaviors and j ranges over events.

Parameters:

- $\pi : X \rightarrow C$ Projection map from the configuration space to the reduced configuration space.
 $\{C^k\}_{k=1}^K$ Partition of the reduced configuration space C in cells: $C = \cup_{k=1}^K C^k$.

Output:

- $\{z^k\}_{k=1}^K$ Reconstructed feature over the reduced configuration space.
 f_L, f_R Reconstructed event generation rate functions.

Procedure:

- (1) Compute the reduced configuration $c(t)$ using the projection π :

$$c(t) = \pi(x(t)).$$

- (2) Define the variable $I^k(t)$ as 1 if the animal is in the k -th cell at time t :

$$I^k(t) = \begin{cases} 1 & \text{if } c(t) \in C^k, \\ 0 & \text{otherwise.} \end{cases}, \quad \text{for } k \in [1, K].$$

Compute the permanence time in each cell:

$$T^k = \int I^k(t) dt, \quad \text{for } k \in [1, \dots, K].$$

- (3) Define $b_i(t)$ as a series of impulses centered at the observed events t_i^j :

$$b_i(t) = \sum_j \delta(t - t_i^j), \quad i \in \{L, R\}.$$

Count the number of events observed in each cell:

$$n_i^k = \int I^k(t) b_i(t) dt, \quad \text{for } i \in \{L, R\}, k \in [1, K].$$

- (4) Compute the observed event rates m_L^k, m_R^k :

$$m_i^k = \frac{n_i^k}{T^k}, \quad \text{for } i \in \{L, R\}, k \in [1, K].$$

- (5) Estimate the inhibition interval Δ from the inter-event statistics.

- (6) Estimate the *event generation rates* using

$$r_i^k = \frac{m_i^k}{1 - \Delta \sum_i m_i^k}, \quad \text{for } i \in \{L, R\}, k \in [1, K].$$

Let the bold symbol \mathbf{r}_i indicate the set of values for all cells:

$$\mathbf{r}_i = \{r_i^k\}_{k=1}^K.$$

- (7) *Deterministic approximation:*

- (a) Compute an estimate of $\text{order}(z)$ using (5.3).
(b) Compute the normalized feature using (5.4).
(c) Having estimated z^k and r_L^k, r_R^k , fit the functions f_L, f_R directly from the relations $r_L^k = f_L(z^k)$; and $r_R^k = f_R(z^k)$.

- (8) *Method taking into account the uncertainty of the data:*

- (a) Compute 95% confidence intervals $[r_i^k, \bar{r}_i^k]$ for r_i^k using (4.3):

$$[r_i^k, \bar{r}_i^k] = \left[\frac{(1 - 1.96/\sqrt{n_i^k - 1})n_i^k}{T^k - \Delta \sum_i n_i^k}, \frac{(1 + 1.96/\sqrt{n_i^k - 1})n_i^k}{T^k - \Delta \sum_i n_i^k} \right].$$

Take $\Theta_i^k = \text{Unif}([r_i^k, \bar{r}_i^k])$ as an approximation of $p(r_i^k | b_i)$.

- (b) Estimate the probability distribution $\Phi_i^k = p(\text{order}^k(r_i)|b_i)$ using the method described as Algorithm 1:

$$\{\Phi_i^k\}_{k=1}^K = \text{OrderBySampling}(\{\Theta_i^k\}_{k=1}^K).$$

- (c) Compute the distributions $\Gamma_i^k = p(\text{order}^k(z))$ from Φ_i^k using (5.6)–(5.7).
 (d) Approximate Γ_i^k as a Normal distribution $\mathcal{N}(\mu_i^k, \sigma_i^{2k})$ using mean and variance.
 (e) Compute the posterior distribution $\mathcal{N}(\mu^k, \sigma^{2k})$ for z^k using (5.10).
 (f) Normalize the values of the feature in the $[-1, +1]$ range using (5.11) obtaining $\mathcal{N}(\bar{\mu}^k, \bar{\sigma}^k)$, which is our final estimate for z^k .
 (g) Having estimated z^k and r_L^k, r_R^k (both with appropriate confidence bounds), estimate f_L, f_R directly from the relations $r_L^k = f_L(z^k)$; and $r_R^k = f_R(z^k)$.

7. DETAILS OF SACCADDE DETECTION ALGORITHMS

The Python source code for both algorithms is available online at http://github.com/AndreaCensi/geometric_saccade_detector

The Kalman filter/smoothing implementation are available as part of Flydra.

7.1. Geometric saccade detector (GSD). The geometric saccade detector (GSD) algorithm works using x, y tracking data, rather than using angular velocity. This makes it most useful for noisy data, as it does not need to derive the data twice (once to obtain the linear velocity, and again to obtain the angular velocity). However, it cannot be used for tethered experiments, for which the x, y data is not available.

The algorithm can be summarized as follows:

- (1) Obtain the trajectory $p(k) = \langle x(k), y(k) \rangle$.
 We use the trajectory returned by Flydra which has been processed with a causal Kalman filter. The algorithm is robust enough to be used on noisy data; so we do not use smoothing.
- (2) Consider separately each instant \bar{k} .
 - (a) Translate the coordinate frame, such that $p(\bar{k}) = \langle 0, 0 \rangle$ becomes the origin.
 - (b) Fix an interval Δ and consider the samples in $[\bar{k} - \Delta, \bar{k} - 1] \cup [\bar{k} + 1, \bar{k} + \Delta]$.
 Δ is a parameter which, for our data, is set to $\Delta = 5$ time steps ($\simeq 0.07s$).
 - (c) Compute the polar coordinates of the samples with respect to the origin:

$$\alpha(k) = \arctan 2(y(k), x(k))$$

- (d) Compute the average orientations before and after:

$$\begin{aligned} \theta_{\text{before}}(\bar{k}) &= \frac{1}{\Delta} \sum_{k=\bar{k}-\Delta}^{\bar{k}+\Delta} \alpha(k) \\ \theta_{\text{after}}(\bar{k}) &= \frac{1}{\Delta} \sum_{k=\bar{k}-\Delta}^{\bar{k}+\Delta} \alpha(k) \end{aligned}$$

In these computations, we consider that angles are defined modulo 360° .

- (e) Similarly, compute the dispersion $\sigma_{\text{before}}(\bar{k}), \sigma_{\text{after}}(\bar{k})$.
- (f) Define the amplitude of a potential saccade as

$$A(\bar{k}) = \theta_{\text{after}}(\bar{k}) - \theta_{\text{before}}(\bar{k})$$

and the “score” of a saccade as

$$S(\bar{k}) = \sigma_{\text{before}}(\bar{k}) + \sigma_{\text{after}}(\bar{k}).$$

- (3) Mark potential saccades as the points where

$$S(\bar{k}) \geq S_{\min} = 20^\circ$$

and

$$A(\bar{k}) \geq A_{\min} = 20^\circ.$$

- (4) At this point, we have a sequence $S(\bar{k})$ that describes the likelihood that there is a saccade at time \bar{k} . To segment the data, examine each point \bar{k} in decreasing order of S , and mark the points in the interval $[\bar{k} - \Delta, \bar{k} + \Delta]$ as unavailable as well. Repeat until all points are marked unavailable.

7.2. Angular-velocity based saccade detector (AVSD). The AVSD algorithm operates using the angular velocity. The advantage of this algorithm is that it can be used also for tethered data, where only the animal heading is available. However, if one starts with x, y data, then one must derive the data twice to obtain the angular velocity.

The algorithm can be summarized as follows:

- (1) Obtain the angular velocity $\omega(k)$.
In our case, this is done using a Kalman smoother on the position data, then deriving once to obtain the translational velocity, obtain the angular heading as the planar direction of the velocity vector, then derive again to obtain the angular velocity.
- (2) Define saccades as the intervals where $|\omega(k)| > \omega_{\min}$.
In our case, set $\omega_{\min} = 300 \text{ deg/s}$.

8. GUIDE TO THE EXPERIMENTAL RESULTS

8.1. Configuration space. The Flydra tracking system tracks the position and velocity of flies in the cylindrical Mamarama arena. The arena has height 0.8m and radius 1m. The data is returned at a temporal resolution of 60Hz, with spatial noise on the order of 0.5cm. Much of the complexity of the saccade detection algorithm (explained in Section 7) is due to handling this limited resolution. However, for the sake of simplicity, we are going to ignore these issues in this section. We then consider the data to be a continuous signal.

We define the following quantities:

- $p(t) \in \mathbb{R}^3$ Position of the animal with respect to a fixed coordinate frame.
- $v(t) \in \mathbb{R}^3$ Linear velocity.
- $R(t) \in \text{SO}(3)$ Attitude (represented as a rotation matrix).
- $\omega(t) \in \mathbb{R}^3$ Angular velocity.

These quantities constitute the original observable configuration $x(t)$ of dimension 12:

$$x(t) = \langle p(t), v(t), R(t), \omega(t) \rangle.$$

8.2. Reduced configuration space. As explained in Section 2.5, we project down the data to a reduced configuration space C of dimension 2. This is done in two steps. In the first step, we only assume that the planar configuration of the animal is relevant. This reduced the configuration to $\langle p_1(t), p_2(t), \theta(t) \rangle$, where p_1 and p_2 are the planar components of the position, and θ is the planar orientation. This reduces the dimension from 12 to 3. By an arbitrary choice of reference frame, we let $p_1 = 0, p_2 = 0$ correspond to the center of the arena.

The second step consists in taking into account the symmetry of the environment. Due to the circular symmetry of the environment, we assume that the animal perceives approximately the same stimulus if its planar configuration is rotated around the center of the arena. Therefore, the two variables that contribute to the stimulus are the distance from the center and the animal orientation.

More formally, we choose as reduced configuration space the two variables d, φ defined as follows:

$$\begin{aligned} d &= 1 - \sqrt{p_1^2 + p_2^2}, \\ \varphi &= \theta - \arctan 2(p_2, p_1) \end{aligned}$$

The angle φ , which we call *axis angle*, is the angle that the animal heading forms with the axis that joins the animal position to the arena center. These are the two quantities that are invariant to a rotation around the center of the arena.

The reduced configuration is defined as $c = (d, \varphi) \in C$. The bounds of the domain C are as follows:

$$\begin{aligned} d &\in [0.15, 1\text{m}], \\ \varphi &\in [-180\text{deg}, 180\text{deg}]. \end{aligned}$$

We express angles in degrees. Note that all the operations on φ are to be executed modulo 360deg ($\varphi = 180$ and $\varphi = -180$ are the same point).

For the distance d , we have:

$$d = \begin{cases} 1 & \text{fly at the center of the arena} \\ 0.15 & \text{limit for reliable data} \\ 0 & \text{fly landed on the wall.} \end{cases}$$

We censor the data at around $d \geq 0.15$. Albeit the tracking system returns data in the whole domain of the arena, it is sometimes unreliable at $d \leq 0.15$, as sometimes the tracking cannot be obtained with full quality, notwithstanding the use of 11 cameras. In the interval $d \geq 0.15$, the data is always very reliable and of homogeneous quality.

For the axis angle, we have

$$\varphi = \begin{cases} -180^\circ & \text{fly pointing directly away from the wall} \\ -90^\circ & \text{closest point on the wall is at the **left**} \\ 0 & \text{fly pointing towards the closest point on the wall} \\ +90^\circ & \text{closest point on the wall is at the **right**} \\ +180^\circ & \text{fly pointing directly away from the wall} \end{cases}$$

Fig. 8.1a shows an example plot in these coordinate. Note that the discretization for d is not uniform in $[0, 1]$, but it is chosen such that each cell in the φ, d space corresponds to an equal area in the p_1, p_2, θ space.

8.3. Fly-centric coordinate space. The $c = (\varphi, d)$ space is the space where we collect statistics and do all of our computations. For visualizing the results in a more intuitive way, we use also another representation, which is just a change of coordinate, shown in Fig. 8.1b. In this representation we use two “aligned” spatial coordinates x^a, y^a . The animal always points “up”.

More formally, in these coordinates the dynamics of the animal is given by:

$$\begin{aligned} \frac{d}{dt}x^a(t) &= 0, \\ \frac{d}{dt}y^a(t) &> 0. \end{aligned}$$

The change of coordinates is given by:

$$\begin{bmatrix} x^a \\ y^a \end{bmatrix} = \begin{bmatrix} \cos(-\theta) & -\sin(-\theta) \\ \sin(-\theta) & \cos(-\theta) \end{bmatrix} \begin{bmatrix} p_1 \\ p_2 \end{bmatrix}.$$

This means that the original configuration $\langle p_1, p_2, \theta \rangle$ is rotated around the arena center to obtain the configuration $\langle x^a, y^a, 0 \rangle$.

8.4. On the choice of coordinates. When going from the configuration (φ, d) to (x^a, y^a) , there is a singularity at the center of the arena. Uniform cells in the (φ, d) domain appear as elongated “pizza slices” in the (x^a, y^a) domain. This appears as a slight artifact when the cells obtained in the (φ, d) domain are plotted in the (x^a, y^a) domain.

Note that all the operations done in the analysis are invariant to the choice of the coordinates. However, once chosen one or the other, one cannot avoid a singularity in the change of coordinates when converting between the two systems. Our rationale for choosing the


(φ, d) domain is that these variables are behaviorally relevant even for environments with different geometry (e.g., rectangular). This should make future comparisons with experiments with different environment easier.

8.5. Nuisances in the analysis. Finally, we summarize all the approximations/limitations of this analysis, which should be kept in mind when interpreting the results; Fig. 8.2 contains a diagrams highlighting some of the factors.

- (1) There might be unobservable states that influence behaviors. The contribution of these states appears as a baseline event rate not explained by the feature.
- (2) The reduced configuration space C might be too small to be a proxy for the true configuration. If this is true, then the estimated features cannot be predictive of all events.

For example, we ignore the altitude and the velocity of the fly.

- (3) The dimension of the feature identified is bounded by the number of event classes considered. To identify a feature of dimension n , one needs at least $n + 1$ events. Therefore, the stimulus might contain more behaviorally-relevant information than what is revealed just by the feature identified from the particular event classes considered.
- (4) There are measurements errors:



(B) Fly-centric view (x^a / y^a plot).




FIGURE 8.1. Explanation of the two kinds of plots used.

- The state $x(t)$ is noisily observed, and the noise is not negligible with respect to the partition of C .
- The events are not exactly detected.

The analysis is generally robust to this kind of noises, but still they are unmodeled phenomena.

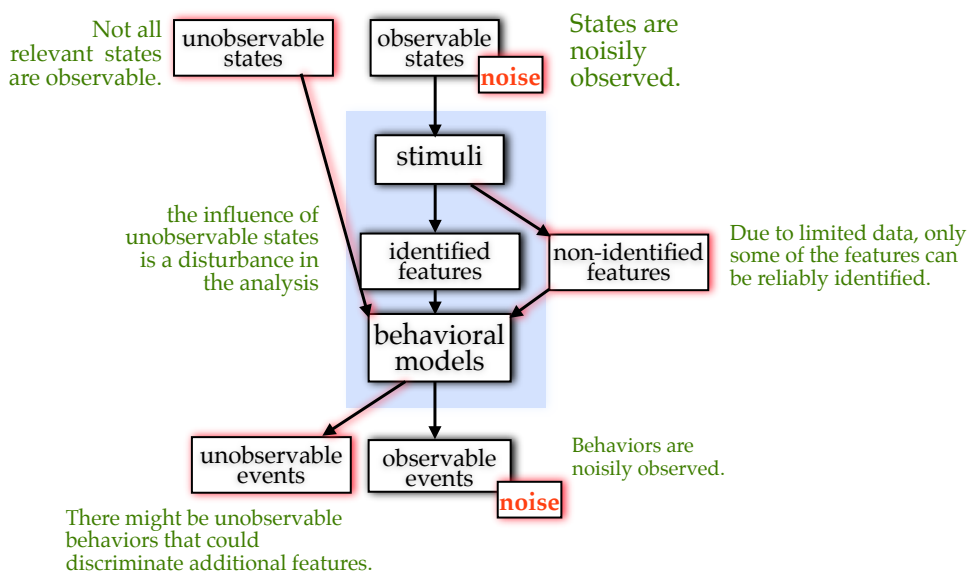


FIGURE 8.2. Nuisances in the analysis

9. COMPLETE PLOTS (GEOMETRIC SACCADDE DETECTOR)

The next pages show the complete statistics using the events detected by the GSD algorithm.

FIGURE 9.1. Time spent in each cell




FIGURE 9.2. Mean speed (whole trajectory)




FIGURE 9.3. Number of detected saccades (both left and right)




FIGURE 9.4. Number of detected left saccades




FIGURE 9.5. Number of detected right saccades




FIGURE 9.6. Observed saccade rate (both left and right)




FIGURE 9.7. Observed left saccade rate




FIGURE 9.8. Observed right saccade rate




FIGURE 9.9. Observed left saccade rate




FIGURE 9.10. Observed right saccade rate




FIGURE 9.11. Order of left saccade rate




FIGURE 9.12. Order of right saccade rate




FIGURE 9.13. Estimated feature




FIGURE 9.14. Uncertainty of estimated feature




FIGURE 9.15. Observed saccade rates as a function of the estimated feature




FIGURE 9.16. Observed saccade rates as a function of the estimated feature





FIGURE 9.17. Estimated event generation rates as a function of the estimated feature



10. COMPLETE PLOTS (ANGULAR-VELOCITY BASED DETECTOR)

The next pages show the complete statistics using the events detected by the AVSD algorithm.

FIGURE 10.1. Number of detected saccades (both left and right)




FIGURE 10.2. Number of detected left saccades




FIGURE 10.3. Number of detected right saccades




FIGURE 10.4. Observed saccade rate (both left and right)




FIGURE 10.5. Observed left saccade rate




FIGURE 10.6. Observed right saccade rate




FIGURE 10.7. Observed left saccade rate




FIGURE 10.8. Observed right saccade rate




FIGURE 10.9. Order of left saccade rate




FIGURE 10.10. Order of right saccade rate




FIGURE 10.11. Estimated feature




FIGURE 10.12. Uncertainty of estimated feature




FIGURE 10.13. Observed saccade rates as a function of the estimated feature




FIGURE 10.14. Observed saccade rates as a function of the estimated feature






FIGURE 10.15. Estimated event generation rates as a function of the estimated feature



REFERENCES

- [1] D. R. Cox and V Isham. *Point Processes*. Chapman & Hall, 1980. ISBN: 0-412-21910-7.
- [2] V. Guerriero, A. Iannace, S. Mazzoli, M. Parente, S. Vitale, and M. Giorgioni. “Quantifying uncertainties in multi-scale studies of fractured reservoir analogues: Implemented statistical analysis of scan line data from carbonate rocks”. In: *Journal of Structural Geology* 32.9 (2010), pp. 1271 –1278. ISSN: 0191-8141 [DOI](#) .