

Complex behavior and perception in *Drosophila* emerges from iterative feedback-regulated reflexes

Thesis by

Floris van Breugel

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy



California Institute of Technology

Pasadena, California

2014

(Submitted November 26, 2013)

To my parents.

Acknowledgments

Many people provided invaluable advice, encouragement, support, criticisms, inspiration, and teaching that contributed to this thesis. I am especially grateful for Michael Dickinson's support, feedback, and above all, patience. When I joined Michael's lab, I was excited to study insect flight control, inspired by my undergraduate research in robotics with Hod Lipson. However I did not understand why I should do so with fruit flies when other insects seemed far more intriguing to me. I spent hours chasing hoverflies in the gardens around Pasadena, and rolling a portable 3D tracking system around the Caltech campus in search of honeybees. Ultimately, I conceded that having an incubator full of well-behaving flies was invaluable if I was to collect meaningful data. Without this experience, I would never have appreciated why flies were such ideal model organisms for studying behavior in a quantitative way. Having accepted defeat in my dream of working outside, I could not abandon my other interests: cameras, and robotics. Michael patiently let me spend several months building and programming an automated pan/tilt/focus system, which allowed me to point and focus a camera at a flying fly, with no immediate applications to research. Here, I owe an incredible debt to Andrew Straw, a postdoc in our lab at the time. Andrew opened me up to the world of programming in Python, not to mention generously providing the 3D tracking software, Flydra, which I used throughout my research program. This experience in cameras and programming eventually allowed me to build several novel systems that paid off in spades. As I spent more and more time in Michael's lab, surrounded by neurobiologists, Richard Murray and his lab group, asin particular Andrea Censi, provided valuable engineering-related feedback. I also would like to thank Kristi Morgansen, for helping me understand nonlinear control

theory such that I could answer a question that had been nagging me since soon after I joined Michael's lab: how do flies measure velocity (and distance)? In describing this general principle I introduce the term *idiokinometry*, which Eunice Kim helped me come up with. When I stumbled into the field of olfaction and odor plume tracking, Jeff Riffell provided invaluable feedback and encouragements for my experiments and data interpretations. I would especially like to thank Michael Dickinson, Richard Murray, Jeff Riffell, and Joel Burdick for being on my committee. When it came time to get my hands dirty with genetics, Anne Sustar and Ainnul Huda helped keep things simple for an engineer like me. Finally, in the heat of scientific excitement, it is easy to lose track of other aspects of life. I would like to thank my friends for understanding, and keeping me company in town, and in tents.

Abstract

For a hungry fruit fly, locating and landing on a fermenting fruit where it can feed, find mates, and lay eggs, is an essential and difficult task requiring the integration of both olfactory and visual cues. Understanding how flies accomplish this will help provide a comprehensive ethological context for the expanding knowledge of their neural circuits involved in processing olfaction and vision, as well as inspire novel engineering solutions for control and estimation in computationally limited robotic applications. In this thesis, I use novel high throughput methods to develop a detailed overview of how flies track odor plumes, land, and regulate flight speed. Finally, I provide an example of how these insights can be applied to robotic applications to simplify complicated estimation problems. To localize an odor source, flies exhibit three iterative, reflex-driven behaviors. Upon encountering an attractive plume, flies increase their flight speed and turn upwind using visual cues. After losing the plume, flies begin zigzagging crosswind, again using visual cues to control their heading. After sensing an attractive odor, flies become more attracted to small visual features, which increases their chances of finding the plume source. Their changes in heading are largely controlled by open-loop maneuvers called saccades, which they direct towards and away from visual features. If a fly decides to land on an object, it begins to decelerate so as to maintain a stereotypical ratio of expansion to retinal size. Once they reach a stereotypical distance from the target, flies extend their legs in preparation for touchdown. Although it is unclear what cues they use to trigger this behavior, previous studies have indicated that it is likely under visual control. In Chapter 3, I use a nonlinear control theoretic analysis and robotic testbed to propose a novel and putative mechanism for how a fly might visually estimate distance by

actively decelerating according to a visual control law. Throughout these behaviors, a common theme is the visual control of flight speed. Using genetic tools I show that the neuromodulator octopamine plays an important role in regulating flight speed, and propose a neural circuit for how this controller might be implemented in the flies brain. Two general biological and engineering principles are evident across my experiments: (1) complex behaviors, such as foraging, can emerge from the interactions of simple independent sensory-motor modules; (2) flies control their behavior in such a way that simplifies complex estimation problems.

Contents

Acknowledgments	iv
Abstract	vi
1 Introduction	1
2 Plume tracking behavior of flying <i>Drosophila</i> emerges from a set of distinct sensory-motor reflexes	6
2.1 Abstract	6
2.2 Introduction	7
2.3 Methods	9
2.3.1 Animals	9
2.3.2 Flight arena	9
2.3.3 Odor stimuli	11
2.3.4 Odor delivery	12
2.3.5 Odor plume calibration	12
2.3.6 Trajectory reconstruction and analysis	14
2.3.7 Model	15
2.3.8 Statistics – Fisher’s exact test	18
2.3.9 Statistics - bootstrapping	19
2.4 Results	20
2.4.1 Surging behavior	24
2.4.2 Casting behavior	26
2.4.3 Odor induced visual saliency	31

2.5	Discussion	33
2.5.1	Casting, surging, and odor induced object salience constitutes a stigmergic iteration	34
2.5.2	The role of odor induced visual saliency	38
2.5.3	Plume tracking, and visual saliency, do not diminish over time	39
2.5.4	How does a fly cast crosswind?	39
2.5.5	Conclusion	40
2.6	Supplementary figures	41
3	The visual control of landing and obstacle avoidance in the fruit fly, <i>Drosophila</i>	46
3.1	Abstract	46
3.2	Introduction	47
3.3	Methods	49
3.3.1	Animals	49
3.3.2	Flight arena	50
3.3.3	Experiment protocol	50
3.3.4	Trajectory reconstruction and analysis	52
3.3.5	High speed imaging	55
3.3.6	Analysis of saccades	56
3.3.7	Procedures for analyzing landing behavior	59
3.3.8	Statistical analysis	60
3.4	Results	60
3.4.1	Description of landings and flybys	60
3.4.2	Saccade results	61
3.4.3	Landing behavior	67
3.4.4	Crash landings	70
3.4.5	Post texture	71
3.5	Discussion	73
3.5.1	Attractive and aversive saccades	73

3.5.2	Landings	74
3.5.3	To land, or not to land?	82
3.5.4	Crashes	84
3.5.5	Post texture	85
3.5.6	Summary	86
3.6	Supplemental figures	88
4	Monocular distance estimation from optic flow during active landing maneuvers	92
4.1	Abstract	92
4.2	Introduction	93
4.2.1	Review of visual distance estimation in biological systems . . .	95
4.3	Modeling and observability analysis	97
4.3.1	Problem statement	98
4.3.2	Analysis	98
4.4	Implementation	106
4.5	Results and discussion	107
4.5.1	Applications to robotic systems	108
4.5.2	Implications for landing insects	110
4.5.3	Summary	111
5	Octopaminergic modulation of the visual flight speed regulator of <i>Drosophila</i>	113
5.1	Abstract	113
5.2	Introduction	114
5.3	Methods	117
5.3.1	Animals	117
5.3.2	Flight arena	117
5.3.3	Experiment protocol	118
5.3.4	Trajectory reconstruction and analysis	119
5.3.5	Immunohistochemistry and imaging	119

5.4	Results	120
5.5	Discussion	123
6	Conclusion	129
	Bibliography	133
	Bibliography	134

List of Figures

2.1	Wind tunnel and odor plume calibration.	10
2.2	Flight trajectories in the presence of an attractive odor plume.	21
2.3	Characterization of flies' responses to an ethanol plume.	22
2.4	Time flies spend inside an odor plume, and between plume encounters.	23
2.5	Flies turn upwind and increase flight speed 270 ms after entering an attractive odor plume (surge behavior).	25
2.6	Surge behavior is a function of both visual and wind stimuli.	26
2.7	Flies cast crosswind within 600-800 ms after leaving an attractive odor plume.	27
2.8	Casting behavior is a function of visual, but not wind, stimuli.	28
2.9	Flies cast with a range of slip angles.	29
2.10	Slip angles during casting are a function of flight speed.	30
2.11	Flies are attracted to high contrast visual features.	32
2.12	Flies are attracted to high contrast visual features (II).	32
2.13	A simple model of odor tracking behavior suggests that surge, cast, and odor-induced object saliency are sufficient, and each of them necessary, for efficient plume tracking.	36
2.14	Summary diagram, indicating the three independent sensory motor modules that describe foraging in the fruit fly.	40
2.15	Behavioral responses to a plume of ethanol and Vector960 (fruit fly attractant) are qualitatively similar.	41
2.16	Behavioral responses to a plume of ethanol and Vector960 (fruit fly attractant) are qualitatively similar.	42

2.17	Cast and surge behaviors are qualitatively similar in the presence of a high contrast checkerboard floor and a low contrast floor with high contrast spots.	43
2.18	Visually salient results presented in Figs. 2.5-2.8, 2.11-2.12 are statistically significant.	44
2.19	Surge, cast, and object saliency behaviors do not diminish over time.	45
3.1	Wind tunnel and high speed camera used for landing experiments.	51
3.2	Landing and fly-by trajectories.	54
3.3	Changes in flight heading are primarily accomplished using body saccades (defined as flight sequences with angular velocity (ω) $>$ 300° s^{-1}).	57
3.4	Changes in flight heading are primarily accomplished using body saccades (defined as flight sequences with angular velocity (ω) $>$ 300° s^{-1}), part II.	58
3.5	The post influences saccades performed in the vicinity of the post.	63
3.6	Saccades near the post are described by two independent behaviors, corresponding to landing and not landing.	64
3.7	Comparison of saccade features for landings and fly-bys.	66
3.8	Flies make targeted body saccades away from the post and saccade amplitude is independent of turn direction.	68
3.9	The motor programs for flight deceleration prior to landing and leg extension are distinct.	69
3.10	Image sequence from a high speed digital video recording of landing sequence.	70
3.11	Crash landings correspond to sequences in which flies do not decelerate fast enough.	72
3.12	Slip angle during flight has a wide range.	75
3.13	Retinal size-dependent expansion threshold compared to the time-to-contact model.	77
3.14	Potential control models for deceleration.	81

3.15	Time elapsed between leg extension and touchdown is less than 100 ms for the majority of landings (n=36).	83
3.16	Landing and fly-by behaviors are distinct, and controlled by retinal position, retinal size, and retinal expansion velocity cues.	87
3.17	Flight trajectories in the absence of the post, plotted as in Fig. 3.2 with a virtual post.	88
3.18	Comparison of saccade features for pseudo landings and fly-bys, with respect to a virtual post.	89
3.19	Object contrast subtly influences landing and fly-by behavior.	90
3.20	Influence of post texture on leg extension and evasive saccade maneuvers.	91
4.1	Theoretical errors in stereoscopic absolute distance estimates for humans (blue) and fruit flies (red), calculated using 4.3.	96
4.2	Geometric relationship of the moving camera and a reference object, as referenced in (4.4).	99
4.3	Condition number of the empirical local observability Gramian for a constant velocity trajectory (black) approaches infinity, whereas a constant acceleration trajectory (blue) or a constant optic flow trajectory (red) achieves significantly lower condition numbers.	102
4.4	Visual target and optic flow estimation.	106
4.5	Performance of the dynamic peering estimation algorithm.	109
5.1	Feedback control diagrams describing the three potential control model architectures under consideration for the flight speed regulator.	116
5.2	Wind tunnel with spatial frequency stimulus, and Confocal image of octopamine neurons.	118
5.3	Octopamine null flies respond to regressive visual motion during free flight with lower accelerations than wild type flies.	122
5.4	Preferred visual motion set point is modulated by identical changes in gain as the LPTC network (H2a), or it enters the visual sensory-motor cascade upstream of the large field motion network (H2b).	125

5.5 Model predictions for acceleration behavior and steady state flight speed. 126

5.6 Neural circuit summarizing the block diagram for H2b from Fig. 5.4A
in the context of a fly's hypothetical neuroanatomy. 127

List of Tables

2.1	Detailed statistics of trajectory data.	16
3.1	Trajectory statistics.	52
3.2	Flight speed under different conditions.	61

Chapter 1

Introduction

For thousands of years, since before Plato's time, humans have wondered how our brains work. Where do our thoughts come from? How are we able to perceive the world? Why do we behave the way we do? And for nearly as long, we have dreamed of building automated "robots," capable of performing tasks as complex as we can ourselves. At present, our best attempts do not even come close to the elegance and robustness of behaviors exhibited by any animal, let alone a human being. Our best hope of building such machines is to further our understanding of how our own brains function. The research presented in this thesis represents a (very) small step forward towards realizing both of these dreams.

Understanding the human brain is a daunting task. Today, we know that our brain consists of approximately 100 billion neurons, connected to one another in a complex network with over 1000 trillion connections. Understanding how sophisticated behaviors emerge from this complicated network of neurons is only now beginning to come within reach thanks to advances in genetic manipulations that make it possible to dissect neural circuits. Attempting to do so within the human brain, however, raises not only ethical issues, but would be an insurmountable task given the sheer size of the brain. Fortunately, pioneering comparative studies across a wide range of model organisms ranging from worms (*C. Elegans*), to sea slugs (*Aplysia*), fruit flies (*Drosophila melanogaster*), mice, and monkeys, have found that the majority of the underlying principles that govern neural interactions are conserved across species. Thus, research efforts to understand any organism will build valuable knowledge to-

wards understanding how our own brains function.

In order to understand how interactions between groups of neurons can give rise to the behaviors we observe, first requires a comprehensive and quantitative input/output understanding of these behaviors. Konrad Lorenz (1903-1989) and Niko Tinbergen (1907-1988) are credited with establishing this field, called neuroethology: the study of the neural basis for behavior. We are now in a golden age where it is finally possible to study behavior, and simultaneously “listen” to neurons *in vivo*. As a result, we have been able to identify the neural circuits that control individual behaviors. There is, however, a large gap between understanding individual reflexive behaviors, and how these reflexes interact with one another to give rise to more complex behaviors. Moving forward requires that we develop a more comprehensive, and quantitative, understanding of animals’ behavioral repertoires. This means studying animals while they perform complex behaviors with as much freedom as possible, while still controlling their sensory inputs.

To address these research goals, the fruit fly, *Drosophila melanogaster*, has emerged as an ideal model system for neuroethologists thanks to its reliable laboratory behavior and the wealth of genetic tools that have made it possible to begin dissecting their neural circuits. Despite having a brain no larger than a poppy seed, which consists of only 100,000 neurons, flies are capable of an impressive array of behaviors including rapid aerial acrobatics, tracking odor plumes in flight, and complex courting rituals. Possibly the most difficult, and important, of these is finding food. For flies, locating food is not only required for nourishment, but all other aspects of its life cycle as well including mating and oviposition, which also occur on food sources. Surprisingly, despite the ubiquity of *Drosophila* in laboratories, kitchens, and fruit stands, many questions regarding how they accomplish this critical task—from initial plume detection to landing—remain unanswered.

In this thesis, I describe in comprehensive detail the algorithm flies use in order to track an odor plume to its source (Chapter 2), and safely land on it while avoiding obstacles (Chapter 3). The final stage in landing requires a fly to extend its legs at an appropriate distance, yet flies are unable to directly measure distance using

their sensory system. In Chapter 4, I describe a bio-plausible, testable, and novel mechanism for actively estimating distance. All of the behaviors in Chapter 2-4 rely on a fly's ability to control and maintain an appropriate flight speed. In Chapter 5, I describe the role of the neuromodulator octopamine in the flight control system, and use a control theoretic framework to propose a putative neural circuit for the flight speed regulator. Together, the models presented in Chapters 2-4 describe a complete algorithm for how a fly locates, and lands on, a food source. Surprisingly, this seemingly complex behavior emerges from a set of just six simple, and independent, sensory-motor modules.

The idea that complex behavior can arise through simpler behavioral modules is not new. As early as 1911, researchers have recognized that complex limb maneuvers are accomplished through combinations of modular reflexes [25], for a more recent review see [12]. In 1959, French zoologist Pierre-Paul Grassé combined this principle with Tinbergen's concept of sign-stimulus triggered fixed-action-patterns to describe how behaviors which change the environment (and thus an animal's sensory perception) could be used to trigger subsequent behaviors, ultimately leading to a seemingly coordinated behavioral cascade. To describe this principle, Grassé introduced the term *stigmergy*, derived from the Greek words stigma (meaning "mark" or "sign") and ergon (meaning "work" or "action"). Grassé used this principle to explain how a group of termites could collectively build a complex nest structure without the need for an internal blueprint or direct communication between individuals [68]. Although the term was initially introduced in the context of social insects, the principle is far more general. Perhaps the most elegant example of stigmergy can be found in the nest building behavior of the solitary mud wasp [132]. In a series of elegant experiments, Smith showed that the wasps built their nests in a series of five distinct steps. By altering their structures at key points in the building cycle, he showed that the wasps do not use a blueprint of the final structure, but rather, that their actions are cued by specific aspects of the structure. In its most basic form, stigmergy describes how complex, coherent, and seemingly intelligent behavior can emerge from the iteration of simple independent sensory-motor modules. The key insight is that previous

actions will determine future actions by changing the environment, and therefore an animal's sensory perception, creating a closed-loop cycle. This interplay provides a mechanism for indirect communication between sensory-motor modules both within a single animal, and between multiple animals. In this way, complex sequences can be indirectly "programmed" without explicitly wiring the sequences together within the brain, resulting in a robust and "evolvable" architecture.

A natural extension of stigmergy is to relax the requirement that actions must *physically* change the environment, and instead allow for the possibility that actions can change an animal's *perception* of the environment simply by moving through it. These new perceptions can trigger new behavioral actions, ultimately leading to complex behavioral cascades similar to those that allow wasps and termites to build elaborate nests, but in this case, the cascades result in the construction of behavioral trajectories. The generalization from physical changes to perceived changes has intriguing consequences, because an animal can actively move so as to simplify, and even enrich, its sensory perception, a principle that I will refer to as idiokinometry. Although similar to "Active Sensing," the distinction is important. This mode of sensory simplification (or enrichment) requires that the animal move, physically changing its position (or more generally, its sensors position) within the environment according to some sensory feedback regulator, thus sampling from an entirely new point of view. This is in contrast to active sensing (for a review, see [129]), such as the classic examples of echolocation in bats and marine mammals [146], electrolocation in fish [156], and active eye movements observed in mammals [99, 104], in which the sensors do not change position (only orientation).

The principle of idiokinometry has been formalized in control theory under the name of nonlinear observability, which draws on Lie algebraic tools to describe how aspects of control can help expand sensory perception, an idea that I elaborate on in Chapter 4. However, these tools only provide a measure of *which* controls must be actuated, not *how*. Studying biological systems can give us inspiration as to how controls can be actuated to make estimation problems easier. To gain an intuitive understanding of this principle, imagine trying to run upwind by holding a windsock

while running in an open field. To construct a traditional control law to regulate this behavior, you would need a measurement of the wind direction relative to your direction of travel. If you stand still, determining the direction the wind is blowing is easy - just look at the windsock. While running, however, the apparent wind created by your motion will confound this measurement, and you can no longer determine the wind direction by looking at the angle of the windsock. One thing you can be certain of, however, is that if the windsock is not aligned with your direction of travel, you are not headed upwind. This observation implies a simple solution: while running, slowly turn until the sock is aligned with your direction of travel. Following this algorithm will result in you heading directly upwind. Note that running downwind faster than the wind speed would also result in the windsock being aligned with your direction of travel. This solution, however, is unstable given a simple proportional controller. In Chapters 2-4, I describe three instances of how flies may use this principle to simplify the task of tracking an odor plume, and landing on it. The first example, similar to the windsock problem, explains how flies are able to find the upwind and crosswind directions during plume tracking by using visual feedback (an idea that was first proposed for mosquitoes, under the name *visual anemotaxis*, by Kennedy [82]). The second example, is a detailed description of the body saccades flies make, which simplifies their visual experience by removing the rotational components of optic flow except for brief moments in time (where vision appears to be ignored [9]). Finally, I describe an example of how flies' visual control of flight speed can allow them to measure absolute distance to a target in order to trigger leg extension at the appropriate time. This principle of idiokinometry is likely general throughout biology, but its importance has not been strongly recognized thus far. The lessons we learn about what sensory-motor reflexes animals, such as flies, use may help design simple yet robust control algorithms for resource limited robots.

Chapter 2

Plume tracking behavior of flying *Drosophila* emerges from a set of distinct sensory-motor reflexes¹

2.1 Abstract

For a hungry fruit fly, locating fermenting fruit where it can feed, find mates, and lay eggs, is an essential and difficult task requiring the integration of both olfactory and visual cues. Here, I develop an approach to correlate flies' free flight behavior with their olfactory experience under different wind and visual conditions, yielding new insight into plume tracking based on over 70 hours of data. To localize an odor source, flies exhibit three iterative, reflex-driven behaviors, which remain constant through repeated encounters of the same stimuli: (1) 270 ms after encountering a plume, flies increase their flight speed and turn upwind, using visual cues to determine wind direction. Due to this substantial response delay, flies pass through the plume quickly after entering it. (2) 640 ms after losing the plume, flies initiate a series of vertical and horizontal casts, using visual cues to maintain a crosswind heading. (3) After sensing an attractive odor, flies exhibit an enhanced attraction to small visual features, which increases their probability of finding the plume's source. A fly's own flight dynamics, which are dominated by sensory-motor delays, play a larger role in determining the time history of its olfactory experience than does plume structure. As a consequence,

¹As of Nov. 1st 2013, the material presented in this chapter is currently under review at Current Biology.

delays in the onset of crosswind casting and the increased attraction to visual features are necessary behavioral components for successfully locating an odor source. These results provide a quantitative behavioral background for elucidating the neural basis of plume tracking using genetic and physiological approaches.

2.2 Introduction

Finding food is perhaps the most important task a motile animal faces on a routine basis, which has driven the evolution of efficient strategies that rely on spatial and temporal integration of multiple sensory modalities such as olfaction and vision. The fruit fly, *Drosophila melanogaster*, is a convenient organism in which to study foraging behavior thanks to its reliable laboratory behavior and the wealth of genetic tools that have made it possible to begin dissecting the neural circuits encoding both olfaction and vision [58, 128, 153]. However, our knowledge of their free flight responses to attractive odor plumes remains limited due to the challenges associated with correlating individual behaviors with the invisible odor plumes [59]. Understanding how flies use their sensory information to guide flight behaviors during foraging would provide a comprehensive ethological context for the expanding literature on cellular and molecular mechanisms.

Fruit flies feed on fermenting fruit, which emit a cocktail of odors including alcohols, acids, and CO₂ [168]. In their natural habitat, such odors are carried by wind through fields and orchards in turbulent plumes resulting in a complex odor landscape consisting of clean air interspersed with packets of high odor concentration. Due to the importance of advection relative to diffusion, this intermittent plume structure can persist for great distances downwind before dispersing below detectable levels [110, 122]. Plumes of this nature present both an opportunity and a challenge, as odors can be detected from far away, yet the chaotic spatial distribution means there are no smooth concentration gradients or continuous plume segment that the insects could follow to the source [109].

Over a century ago, Barrows [6] observed that the plume tracking behavior of

fruit flies, *Drosophila ampelophila*, consists of a simple algorithm consisting of two distinct behaviors: surging upwind and zigzagging (casting) crosswind. Subsequent studies of plume tracking behavior in insects have focused on pheromone tracking of male moths, which exhibit similar behavior [4, 3, 29, 41, 81, 89]. The emerging model is that upon encountering an attractive plume, moths surge upwind by visual anemotaxis, and at the same time an internal triggering mechanism causes the moth to make characteristic casting maneuvers [3, 81, 83]. *D. melanogaster* exhibit similar plume tracking behavior, however, their casting does not appear to be controlled by an internal clock, as evidenced by their maintaining an upwind heading in a wind tunnel with a homogenous odor mixture [26]. An alternative to the internal clock model is that casting is triggered directly by plume loss. This is similar to the basic strategy of maintaining a straight course in the presence of attractive cues and altering course in their absence that is shared by a diverse array of organisms operating at different scales including *Drosophila* larvae [64], bacteria [10], and humans [116].

In addition to determining which of these two models best describes the plume tracking behavior of the fruit fly, many other key issues need to be answered. For example, the success in tracking a plume will likely depend critically on response dynamics, but we have no precise measurements for the sensory-motor delays associated with the initiation of surges or casts. Although it is likely that flies use visual cues to maintain upwind flight during surges and cross wind flight during casts, this has not yet been demonstrated experimentally. Another critical unknown is the degree to which flies alter their behavior in the face of different wind speeds, an environmental condition that varies extensively in the field. All these details have proven difficult to study because the variable and invisible nature of odor plumes makes it challenging to simultaneously visualize the odor stimulus and an animal's flight path at the same time. To overcome this fundamental limitation, I developed an experimental system that delivers a temporally- and spatially-calibrated odor plume within a laminar flow wind tunnel equipped with a multi-camera, 3-D tracking system. These experiments offer a comprehensive and detailed overview of the algorithm flies use to localize an odor source, and can be described in the context of a simple stigmergic model.

2.3 Methods

2.3.1 Animals

Experiments were performed on 2-to 3-day-old fruit flies, *Drosophila melanogaster* Meigen, Canton-S background. Flies were deprived of food, but not water, for 6-8 hours prior to the start of the experiment in order to motivate flight. Flies were on a 16 hr/8 hr light/dark schedule, with lights out at 11pm local time. For each experimental trial, I introduced a group of 12 female flies to the corner of the arena within a small test tube between 6 and 8 pm. The flies were then free to move throughout the flight arena for a period of 12-18 hours, during which time data were collected automatically.

2.3.2 Flight arena

I performed all experiments in a 1.5 m x 0.3 m x 0.3 m working section of a wind tunnel (Fig. 2.1A) that has been described previously [26, 100, 141, 150]. In these experiments, except where otherwise noted, the wind tunnel was set to 0.4 m s⁻¹, which was chosen based on previously published measurements of wind levels in an apple orchard [26]. On the two long walls and floor of the arena, I projected different visual stimuli using a Lightspeed Designs DepthQ (Oregon City, OR) projector with the color wheel removed (120 Hz update rate, 360 Hz frame rate, mean luminance of 50 cd/m²). I generated the stimuli using the VisionEgg open-source image-rendering software [139]. For the purposes of tracking, the arena was backlit with an array of near-infrared (640 nm) LEDs. The cameras were equipped with long-pass filters (Hoya R-72) so that the camera images were not contaminated by the pattern that was displayed in visible wavelengths.

I tracked the 3-dimensional position of individual flies within the chamber using a camera based real-time tracking system that is described in detail elsewhere [140]. The 10-camera (Basler Ace 640-100 gm, Basler, Exton, PA) system generated an estimate of fly position at 100 frames per second with a median latency of 39 ms

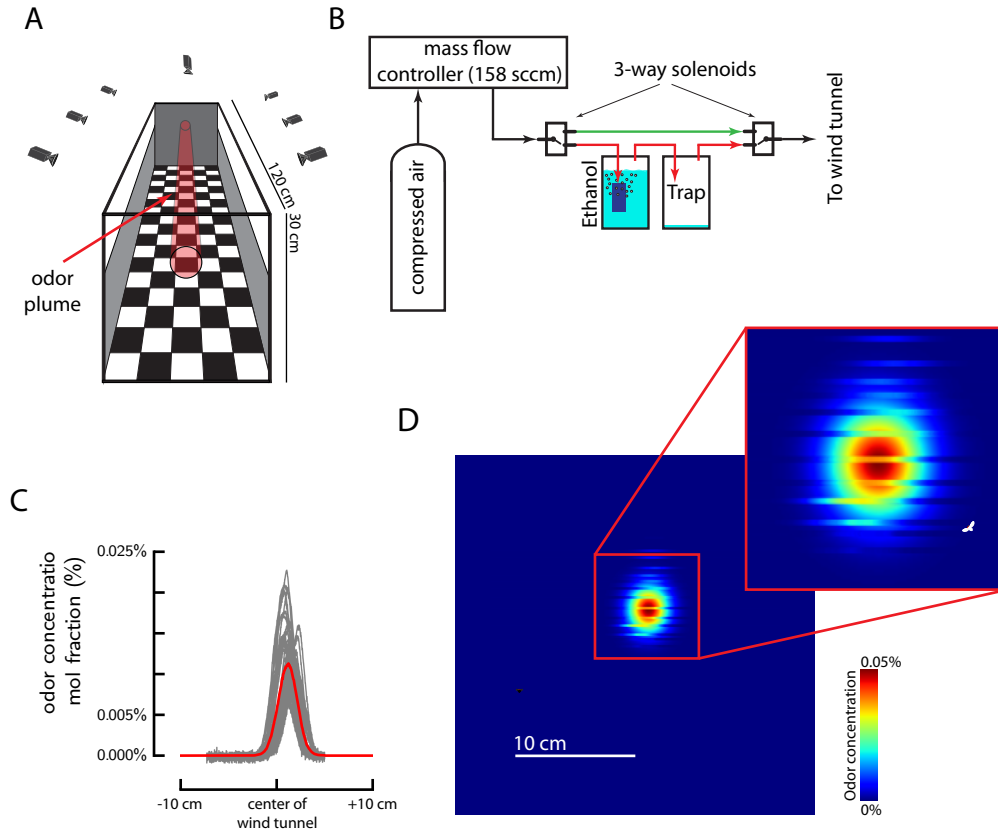


Figure 2.1: (A) Experiments were performed in a $1.5 \times 0.3 \times 0.3$ m³ wind tunnel, equipped with a camera based automated 3D tracking system to observe flies as they interacted with a narrow ribbon plume of ethanol odor. (B) Odor delivery system. (C) Raw (gray) and mean (red) measurements of odor concentration across a horizontal crosswind cross-section of the wind tunnel in the presence of a constant ribbon plume of ethanol (0.4 m s^{-1} wind). The measurements were collected with a photo-ionization detector that was automatically translated through the wind tunnel for 50 cycles over the course of 15 min. The original voltage traces were converted to odor concentration values by first principles, see Methods. (D) Gaussian model of the ethanol plume (0.4 m s^{-1} wind), overlaid with the twelve original mean odor concentration measurements (see C for an example) used to fit the model. In the enlarged view, a flying fly is shown to scale.

by triangulating the fly’s position from 2-dimensional tracking data contributed by two or more synchronized cameras. The 3-dimensional position was estimated with an extended Kaman filter, using a constant velocity motion model. Because of the high frame rate relative to flight behavior, I found that this simplification worked well in practice. The resulting trajectories were smoothed to remove digitization errors, and to estimate velocity, using a simple forward/reverse, non-causal Kaman filter. I excluded non-flying trajectories, as well as trajectories that were less than 1 second in length, from my analysis.

2.3.3 Odor stimuli

In previous studies, apple cider vinegar or a banana and yeast cocktail have been used to attract flies [26, 45, 54, 58]. Unfortunately these odors are difficult to work with, as they are prone to sticking to tubing used to deliver the odors, making it impossible to obtain a clean pulsing signal. To circumvent this issue I chose to use pure ethanol, a common component of rotting fruit, which is less likely to stick to tubes and other parts of the wind tunnel. Although the details of how ethanol is detected by flies are not well understood, it is attractive to flies [6, 119, 167, 168]. Photo-ionization detectors are also particularly sensitive to ethanol, making it possible to measure the same plume that flies experienced during my experiments without using a tracer gas that may behave differently.

To confirm that the behaviors I observed were due to the detection of an attractive odor, rather than physiological changes due to my choice of ethanol as an attractive odor, I repeated my experiments with the commercially available fruit fly attractant similar to balsamic vinegar, Vector960 (Pest Control Solutions, St. Louis, MO). Although the responses were less obvious (likely due to differences in concentration in the air), I did not find any qualitative differences (Fig. 2.15, 2.16). Between switching odors, I thoroughly cleaned all the potentially contaminated tubes and vials with Hexane, and ran the experiment with no odor to ensure that all the system was clean of any residual odors.

2.3.4 Odor delivery

To compare the responses of flies to an attractive odor plume and clean air I developed a controlled odor delivery system (Fig. 2.1B). Breathable air from a compressed air tank (Praxair, Seattle, WA) was sent through a mass flow controller (Part No. 902C-P5BM-I1, Sierra Instruments, Monterey, CA) set to 158 sccm. I used an Arduino Nano and solid-state relay to switch a three-way solenoid valve (Part No. LHDA0531215H, Lee Corp, Essex, CT) that directed the airflow either through an aquarium stone submerged in liquid ethanol contained in a glass vial, or through uncontaminated tubes. These two pathways passed through another 3-way solenoid valve controlled in parallel with the first, such that the desired flow of either odorous air or clean air entered the wind tunnel through a 1.67 mm (I.D.) Polyethylene tube from the same point with minimal switching latency.

For pulsing experiments, the Arduino was programmed to switch the odor on for 0.4 seconds, allowing clean air to pass through the system for 4 seconds between each pulse. The timing of the pulses was sent over USB to the computer and saved to disk for analysis. For the constant plume experiments, the odor delivery was turned on at midnight, and persisted for 4 hours, after which time the solenoids were switched off to deliver clean air for the remainder of the night.

2.3.5 Odor plume calibration

To determine the three-dimensional odor landscape, I scanned the wind tunnel with a miniature photoionization detector (PID) in the presence of an ethanol plume. To scan the pulsing plume, I positioned the PID at one point and saved the time course of the Arduino's control signals together with the output of the PID. I repeated these measurements for 28 pulses at each of 61 positions, and constructed a 3-dimensional time-varying Gaussian model based on the data using a least squares fit. Since I focused my analysis on the results collected with the continuous plume, rather than the pulsing plume, the results from these measurements are not shown for the sake of space.

To calibrate the continuous plume, I attached the PID system to a frame that was actuated with a stepper motor and timing belt. The motor was controlled so that the PID was slowly (0.01 m s^{-1}) driven back and forth along the horizontal crosswind direction through the plume for a period of 15 minutes (50 round trips) while data was streamed to the computer via an Arduino. This was repeated for 12 different altitudes, each at three different positions along the wind line. For each run I calculated the baseline subtracted mean (Fig. 2.1C), and used these means to calculate a least squares fit to a two-dimensional Gaussian model of the crosswind concentration profile at each of the three positions along the wind line (Fig. 2.1D). The three fits were very similar, so to simplify my analysis I modeled the odor plume as a constant Gaussian 2D model that stretched down the length of the wind tunnel as a cylindrical plume. To compare flies' odor plume tracking behavior under different wind speed conditions, I repeated the entire calibration process for three wind speeds (0.3 , 0.4 , and 0.6 m s^{-1}). For the slowest wind speed, I reduced the mass flow of air to 77 sccm to obtain a stable plume. The mass flow rates for 0.4 and 0.6 m s^{-1} wind speeds were both 158 sccm .

The photo-ionization detector provides data in arbitrary units, which must be calibrated in order to provide a measure of the actual odor concentration. Unfortunately, the calibration shifts significantly over the course of a few hours. Since each calibration routine took over 6 hours, I took a first principles approach to calibrate the magnitude of the odor signal. I assumed that by bubbling air through ethanol, the ethanol and air mixture in the vial is at steady state. This assumption potentially yields an over estimate of the concentration of ethanol. Under this assumption, the mol fraction of ethanol in the air is described by the ratio of the vapor pressure of ethanol to the atmospheric pressure:

$$\text{Mol fraction of ethanol} = \frac{6.66 \text{ kPa}}{101.325 \text{ kPa}} = 0.065. \quad (2.1)$$

Given the mass flow rate of air through the system, 158 mL/min , I calculated the flow of mols of ethanol per minute using the ideal gas law:

$$n = \frac{PV}{RT} = \frac{(101325Pa)(0.000158m^3)}{(8.314JK^{-1}mol^{-1})(293.15K)} = 0.00656 \text{ mols air}/\text{min}, \quad (2.2)$$

$$\text{mols ethanol}/\text{min} = 0.00656 * 0.065 = 4.26x10^{-4}. \quad (2.3)$$

The wind speed in the wind tunnel was set to 0.4 m s^{-1} , or 24 m min^{-1} , thus over the course of one minute, $4.26x10^{-4}$ mols of ethanol are distributed across 24 meters along the wind direction, or $1.775x10^{-5}$ mols per meter. The 2-dimensional Gaussian model from my PID measurements describes the concentration profile that these molecules must be distributed with in the plane perpendicular to the wind direction. Normalizing the Gaussian model such that its integral is $1.775x10^{-5}$ yields a function that describes the number of mols of ethanol in a $1*dy*dx \text{ m}^3$ volume. Using the ideal gas law again, I calculated the number of mols of total gas in that same volume: 41.57 mols. Thus, to calculate the molar fraction of ethanol at any point in the wind tunnel, I take the ratio of the output from the normalized Gaussian function and 41.57. This yields a peak concentration of 0.0476% mols of ethanol in clean air.

2.3.6 Trajectory reconstruction and analysis

By combining the three-dimensional trajectories from the tracking system with my model of the odor landscape I was able to reconstruct each flies' olfactory experience synchronized to its behavior. Visual inspection of the trajectories indicated that in the presence of both the pulsing and continuous odor plumes the flies were casting and surging as has been described previously (Fig. 2.2). To uncover the details of this behavior, I developed custom analyses that allow us to present the results of all of the trajectories together. All my analysis code was written in Python using the open-source software packages Scipy and Matplotlib.

The tracking software used to collect the 3-dimensional flight trajectories was unable to maintain the identity of individual flies over the entire course of my ex-

periments (12-18 hours). Therefore, I was unable to test whether individual flies behaved consistently different from one another, and each trajectory was treated as an independent sample. The values of “N” reported throughout the figures in this paper represent the number of trajectories that contributed data to each Figure panel. This, however, is only one interpretation of “N.” In many cases, individual trajectories contributed more than once to my analysis, if the fly entered and left the odor plume multiple times. Furthermore, the number of trajectories should not be confused with the number of animals (which range between 36 and 108 per experiment). These alternate values of “N” are reported in a supplementary Table 2.1.

2.3.7 Model

In order to quantitatively assess the importance of the three sensory-motor reflexes we describe, we created a simple 3-dimensional simulation. In our simulation, a single visual feature randomly emitted 3 cm diameter spherical “packets” of attractive odor at a mean rate of 5 per second. We reasoned that this simplified model (as opposed to modeling actual aerodynamics) would be sufficient to explore the implications of flies’ plume tracking behavior because their olfactory experiences are largely determined by their own flight dynamics, rather than the plume dynamics. The values were chosen arbitrarily such that the plume-tracking problem would be difficult, but not impossible. Different values resulted in consistent behavioral differences between the tracking algorithms we tested. The odor packets were advected by a 0.4 m s⁻¹ wind that randomly changed direction according to a uniform distribution between $\pm 100^\circ$ /second. This resulted in a slow random walk of wind speed direction. Once the first odor packet reached a distance of 1 meter from the source, a virtual fly was spawned inside of that odor packet. This allowed us to circumvent the problem of initial plume finding, which is beyond the scope of this study.

The virtual fly was programmed to surge and cast with the same sensory-motor delays we observed (270 ms, and 640 ms, respectively). To simplify programming, we allowed the virtual fly to have access to the absolute wind direction, rather than

Visual pattern	Wind speed (m/s)	Odor	# Exps (# Flies)	# Trajecs	Clean air control			Odor stimulus			Figures	
					Mean length (mean \pm std) (sec)	# Trajecs found plume	# Odor encounters per trajec (mean \pm std)	# Trajecs	Mean length (mean \pm std) (sec)	# Trajecs found plume		# Odor encounters per trajec (mean \pm std)
Checkers	0.4	Ethanol	9 (108)	3743	5.3 \pm 4.4	1368	1.5 \pm 1.0	5542	7.6 \pm 8.4	3182	3.0 \pm 3.0	2A-D-E; 3A-F; 4A-F; 5A; 6A-B; 7
Stripes wind	0.4	Ethanol	3 (36)	1136	6.5 \pm 6.0	313	1.4 \pm 0.9	1425	5.1 \pm 4.4	609	2.0 \pm 1.4	3E; 4E
Stripes \perp wind	0.4	Ethanol	3 (36)	1082	5.2 \pm 4.6	315	1.5 \pm 1.2	929	6.7 \pm 6.4	497	2.4 \pm 1.9	3E; 4E
Dots	0.4	Ethanol	5 (60)	2508	5.6 \pm 5.2	1090	1.8 \pm 1.7	2705	7.4 \pm 7.5	1795	3.0 \pm 2.9	5B-D; S2
Checkers	0.3	Ethanol	5 (60)	5453	5.4 \pm 5.5	1158	1.6 \pm 1.1	4509	6.3 \pm 6	2606	3.4 \pm 3.6	3F; 4F; 6B
Checkers	0.6	Ethanol	5 (60)	1400	5.4 \pm 4.7	433	1.4 \pm 0.9	2137	2.2	1056	3.2 \pm 2.7	3F; 4F; 6B
Checkers	0.4	Vector960	4 (48)	2819	6.0 \pm 5.3	888	1.7 \pm 1.3	3185	6.8 \pm 6.5	1635	2.2 \pm 1.9	S1
Dots	0.4	Vector960	7 (84)	5732	6.1 \pm 5.0	3879	2.1 \pm 1.7	4619	6.3 \pm 5.0	3580	2.7 \pm 2.3	S1
Checkers	0.4	Pulsing Ethanol	4 (48)	3417	6.6 \pm 6.8	1036	1.5 \pm 1.0	1421	6.4 \pm 7.0	173	1.2 \pm 0.8	2D,E

Table 2.1: Detailed statistics of trajectory data.

implementing the programmatically more complex solution of visual anemotaxis. Our experiments did not provide accurate insight as to how flies trigger casting reversals, either in the horizontal or vertical directions. For our model, we initially chose 0.5 second casting intervals for both the horizontal and vertical aspects of casting (these values are similar to what we observed for flies in our wind tunnel). However, we quickly noted that by using an identical reversal frequency for horizontal and vertical casts, the virtual flies would follow a highly periodic (X-shaped) pattern in the crosswind plane. Whereas this strategy may be ideal in the case of a pulsing plume in constant wind, following this algorithm in randomly shifting (or turbulent) winds could prevent a fly from ever re-locating the plume because the plume may have shifted outside of the flies periodic search trajectory. In this case, a more strategic algorithm may be to use a less periodic trajectory, accomplished by adding a noisy element to the reversal timing, or by using an irrational relationship between the horizontal and vertical cast timing. For our simulation, we arbitrarily chose 0.5 second for the horizontal reversals, and 0.3090169 seconds for the vertical reversals. Note that with these values, the ratio of horizontal to vertical casting periods is approximately equal to the golden ratio, an elegant irrational number found throughout nature and art [95]. This was an arbitrary choice, and future experiments will be necessary to determine what ratio of horizontal to vertical casting periods flies use, and whether this value changes under different environmental conditions (e.g., wind and visual). Although our simulation provides the framework for exploring what values may be theoretically optimal under different conditions, without a biological basis for environmental conditions and behavioral actions we are hesitant to draw any conclusions. Once the fly came within 20 cm of the visual feature, we programmed it to approach the feature, and if it sensed an odor, land on it. The 20 cm choice was again arbitrary, and will likely depend on the size, and contrast, of the actual feature. We allowed each simulation to run for 20 seconds, after which time flies had either landed on the feature, or moved past it without hope of relocating it. We ran three simulations with different behavioral algorithms (each with 1000 repetitions). In the first, the virtual fly followed the algorithm outlined above (similar to what we observed real flies do).

Next, we set the surge and cast delays to be equal (270 ms), and found that the time flies required to locate the food source increased. Finally, we removed the visual attraction component, and found that nearly none of the flies successfully located the source (due to the shifting winds, and the sensory-motor delays). Although simple, our simulation allows for easy comparisons of different behavioral and environmental parameters. We invite other researchers to explore this space using our code, which is made freely available here: <https://github.com/florisvb/FlyPlumeTracking>.

2.3.8 Statistics – Fisher’s exact test

Odor plume tracking behavior in *Drosophila*, as in other insects, is a complex time-dependent and multi-dimensional behaviors, with clear qualitative differences. To determine which aspects of flies’ behavior were due to the presence of an attractive odor, I compared their behavior in the presence of an attractive odor plume to their behavior in the presence of a pseudo-plume of clean air. I found several elucidating behavioral parameters that showed clear differences between the two treatments, including flight heading in the horizontal and vertical planes, airspeed, and altitude.

To reduce such behaviors to a single descriptive statistic for performing statistical tests could hide interesting and important details of the process, as well as running the risk of giving misleading results. Rather than reduce the data to a single descriptive statistic, I developed a method whereby I could assign a p-value to each individual pixel of the density maps shown in figures 2.5-2.8 and 2.11 (Fig. 2.18). These p-values give a quantitative statistical sense of the importance of the behavioral differences between the clean air and attractive odor cases.

To determine these p-values, I use a non-parametric resampling approach, Fisher’s exact test [51, 52]. For each resampling, I randomly reshuffled the labels “clean air” and “attractive odor” assigned to the trajectories, resulting in two new test groups. These new groups are run through my analysis, resulting in density maps (e.g., Fig. 2.5B). Each row of the density maps is normalized such that the integral for that row equals one. Then, on a pixel-by-pixel basis, I calculated the difference between the

two groups and record this value. After repeating this process for 1000 resampling's, I constructed a distribution these differences for each pixel. Comparing the difference between the actual control and attractive odor cases to this distribution allows us to calculate a two-tailed p-value that describes the probability that my result is due to the random sampling process, rather than an actual difference in the underlying mechanisms driving the behavior. Larger numbers of resampling's provide more resolution, and confidence, on the p-value. The computations involved in resampling data at the trajectory level, and calculating p-values on a per pixel basis, are not trivial. I found that 1000 resampling's gave consistent results while providing sufficient resolution to calculate p-values as small as 0.001. All of the clear behavioral differences I present in this paper are significant, in the statistical sense, with a p-value of ≤ 0.001 across the relevant time and parameter space (Fig. 2.18).

2.3.9 Statistics - bootstrapping

Many of the analyses presented in this paper rely on the distributions (D^*) of behaviors observed in large numbers of trajectories to draw conclusions on flies' stereotypical behavior. In order to provide a sense of variability in my data due to random sampling processes, I used a basic non-parametric bootstrapping method to calculate 95% confidence intervals for these distributions [47]. The general approach is to resample the original set of trajectories (Y), with replacement, to obtain a new set of trajectories (Y^*), with the same (or similar) sample size as the original dataset. Note that Y^* may include duplicates. Then the analysis is performed on Y^* to determine the distribution D^* . This is repeated many times, and the resulting values of the distributions D^* are sorted so that a mean and 95% confidence interval can be calculated. In my analysis, I found that 500 iterations provided repeatable measures of confidence intervals. For a more detailed explanation, including several examples, see [30].

2.4 Results

I measured the 3D trajectories of fruit flies in a laminar flow wind tunnel as they interacted with a controlled plume of ethanol, a compound which is strongly attractive [6, 167, 168, 119]. The system was equipped with a high speed projection system that allowed us to present different visual scenes to the floor and walls of the tunnel (Fig. 2.1A-B). To correlate the flies' behavior with their olfactory experience, I used a photo-ionization detector to construct a three dimensional map of the plume (Fig. 2.1C-D). Hungry flies were released into the wind tunnel between 6-8 pm, and the odor was injected into the wind stream between midnight (their subjective dusk) and 4 am. This allowed us to control for other potential sensory cues by comparing the flies' behavioral response to the odor plume, and their response to a pseudo-plume of clean air.

Our automated data collection system allowed us to collect more than 50,000 trajectories (mean length > 5 sec) of flies entering and leaving the plume under different visual conditions and wind speeds. In the presence of both a pulsing and continuous plume, flies exhibited stereotypic surge and cast behavior (Fig. 2.2). As expected, in the presence of the attractive odor, flies spent significantly more time within the region of the odor plume (Fig. 2.3A). To present the results from a large dataset in an informative manner I developed a graphical format that shows the time history of various behavioral parameters aligned to the moment when each fly either entered or exited the plume (Fig. 2.3B). To align the flies' behavioral responses with their olfactory experiences, I first empirically estimated the threshold odor concentration that triggered a plume cast (Fig. 2.3C). Based on these results, I chose an odor concentration threshold of 3σ from the plume maximum, which corresponds to 5×10^{-4} % ethanol in clean air. However, none of the results presented in subsequent figures were sensitive to changes in this threshold choice within a range of 1 to 4σ .

In the presence of the continuous laminar ethanol plume, flies spent a median of 250 ms within the plume during each encounter (Fig. 2.4A). By comparison, flies spent 210 ms (median value) within the confines of the identically-sized pseudo-plume

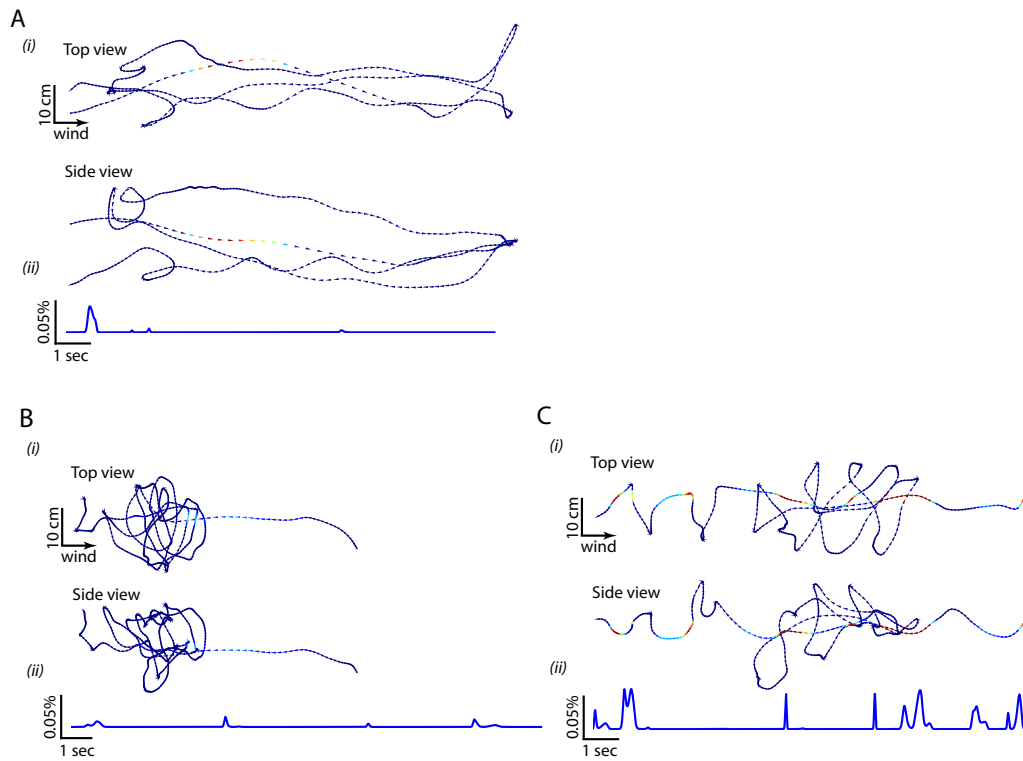


Figure 2.2: (A-C) (i) Example flight trajectories, color-coded for instantaneous odor concentration, (A) for a pseudo-plume of clean air, (B) a pulsing plume of ethanol, and (C) a constant ribbon plume of ethanol. (ii) Odor experience vs. time for each trajectory.

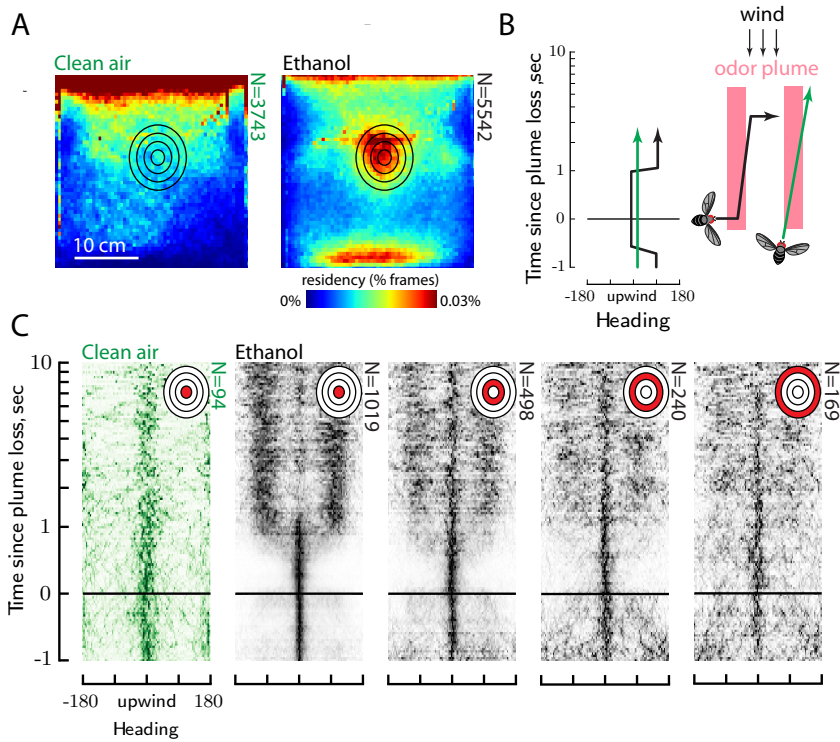


Figure 2.3: Characterization of flies' responses to an ethanol plume. (A) Residency heat map of flight trajectories in clean air and ethanol for the crosswind cross-section of the wind tunnel. The black ovals indicate so-concentration lines of 1-4 standard deviations from the peak of the 2D Gaussian model of the odor plume shown in Fig. 2.1D. (B) Cartoon of two flight trajectories, and their respective mapping onto a plot showing heading relative to the time when they exit the odor plume. The black trajectory depicts the stereotypical behavior I observed for a fly in the presence of an attractive odor plume; the green trajectory depicts the stereotypical behavior of a fly in the absence of any odors. I use these colors consistently throughout the paper. (C) To show the responses of all the trajectories together, I plot the heading response of flies relative to each time they exit the odor plume. These trajectory snippets are overlaid, and shown as a density map in which each row is normalized such that it spans the entire color range (higher color density indicates more trajectories). In the presence of the ethanol plume, approximately 0.4-1 seconds after leaving the plume, flies begin to fly crosswind (casting). To align the trajectories in this way requires that I define a behavioral concentration threshold of ethanol. To determine at what concentration flies begin to show a behavioral response, I selected for trajectories that pass through four different regions of the plume corresponding to the level sets of $1-4\sigma$, and set the behavioral threshold to the minimum of that region. Based on these results, I chose the generous threshold of 3σ . Responses to a pseudo-plume of clean air did not show any clear changes based on my choice of threshold within these ranges (not shown).

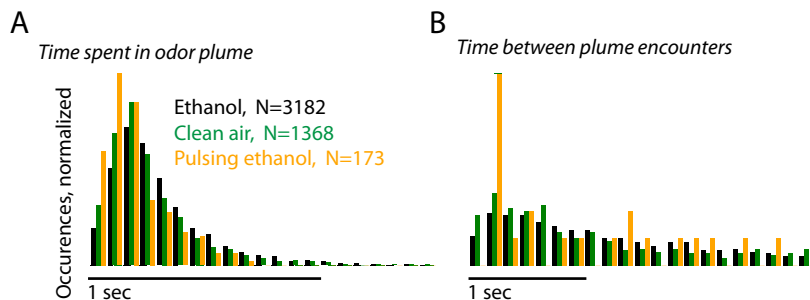


Figure 2.4: (A) Flies spent similar amounts of time in a constant ribbon plume of ethanol (0.25 sec) as they did in a pseudo-plume of clean air (0.21 sec), and a pulsing plume (0.16 sec). Though small, these differences were statistically significant (max p-val $< 10^{-6}$, MWUB). (B) The timing between odor encounters, after the initial encounter, was similar across the three plumes with no statistically significant differences (min p-val > 0.37 , MWUB). For E and D, the distributions show the mean value of times across plume encounters for each individual trajectory, thus each trajectory contributed once to E and once to D.

of clean air. The similarity of these values indicate that flies possess very limited capacity to actively remain within the 5 cm wide odor plume once they encounter it, despite that fact the plume is roughly 10 times the wingspan of a fly. In the presence of the pulsing plume (400 ms on, 4 sec off), flies spent a median of 160 ms in the plume. Although subtle, these differences were statistically significant (max p-val $< 10^{-6}$, Mann-Whitney U test with Bonferoni Corrections (MWUB)). I did not find significant differences between the three groups (continuous, pulsed, pseudo-plume) for the time between plume encounters within a given trajectory (Fig. 2.4B), which ranged from 1.1 to 1.4 seconds (min p-val > 0.37 , MWUB). In interpreting these values it should be noted that the tracking system often did not allow us to follow flies for durations as long as the period between successive odor packets in the pulsed plume experiments. Thus, the data in Fig. 2.4B primarily reflect the time between encounters of the odor within a single packet as it flowed through the tunnel.

In designing my system, my initial goal was to collect data using a pulsed plume with dynamics similar to those observed in an open field [110]. However, as described above, I found little difference in the statistics of plume interaction (time spent in plume and time between plume encounters) in the continuous and pulsed-plume cases

(Fig. 2.4). The lack of effect of plume structure is explained by the fact that, on each encounter, flies remained with the approximately 5 cm diameter plume for a relatively brief period of time (see Fig. 2.4A). Thus, even in a pulsed plume, flies experienced the odor intermittently not because they lost contact due to filament truncation, but rather because their lateral and vertical motion carried them through the plume. Because flies encountered the plume more frequently in the continuous plume (and thus I collected interactions at a higher rate), I focused my efforts on this paradigm. When encountering the plume, flies exhibited stereotypic surging and casting behavior as has been described qualitatively in previous studies. In the following sections I describe the details of these behaviors, organized in the ethologically relevant sequence of events that take place after a fly first encounters the attractive odor.

2.4.1 Surging behavior

In the presence of an attractive odor, flies exhibited a clear preference for flying upwind compared to downwind (Fig. 2.5B), consistent with prior studies. The spatially calibrated plume made it possible to accurately estimate the sensory-motor delay of this behavior. Within approximately 270 ms after entering the ethanol plume, flies turned upwind and increased their ground speed relative to control flies (Fig. 3B,C). During this upwind surge, flies exhibited tighter distribution about the upwind direction than did control flies, which also showed a preference for flying upwind (Fig. 2.6A). To test whether the ability to surge upwind after encountering an odor involved visual feedback, I repeated the experiment under three visual conditions: a checkerboard floor, a floor with stripes parallel to the wind-line, and a floor with stripes perpendicular to the wind-line. The visual anemotaxis model of Kennedy [82], posits that a flying insect regulates its flight heading such that the direction of the visual flow it experiences is aligned with its direction of motion, thereby minimizing visual side slip. Under this control scheme, we would expect the error between flight heading and the upwind direction to be smaller if the visual stripes run parallel to the wind direction and larger if the stripes run perpendicular to the wind direction. These

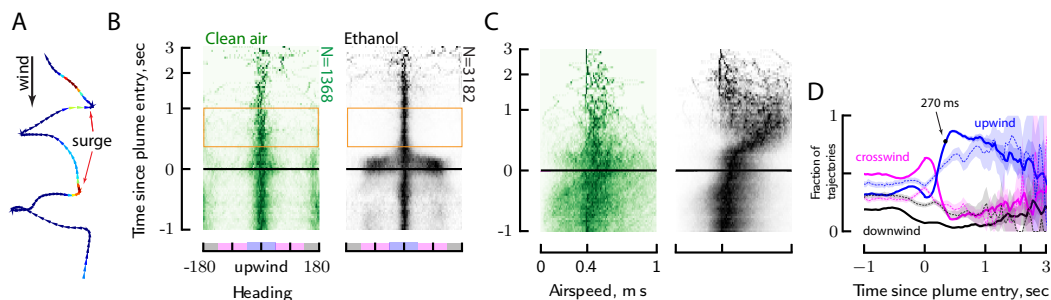


Figure 2.5: Within 270 ms after encountering an attractive odor plume, flies turn upwind using visual cues, and simultaneously increase their forward flight speed. (A) A trajectory (repeated from Fig. 2.2B) demonstrating the stereotypical upwind turn shortly after the fly enters the plume (color encodes odor concentration, see Fig. 2.1). (B) Heading (in the horizontal plane) of each trajectory relative to each time the fly enters the odor plume, for the duration of time that it remains inside the plume. These trajectory snippets are overlaid, and shown as a density map (higher color density indicates more trajectories). To make the distributions more visually apparent, the colors of each row were normalized to span the entire color range. The control plot (left) was generated using a “pseudo-plume” of clean air with the same shape as the actual ethanol plume. Time is plotted on a log scale between 0 and 3. (C) Airspeed of each trajectory relative to each time a fly enters the odor plume, plotted in the same manner as B. In B and C the visually apparent differences are statistically significant (see Fig. 2.18, p-val 0.001-0.01, Fisher’s exact test). (D) Fraction of trajectories flying upwind, crosswind, or downwind relative to plume entry for a pseudo-plume of clean air (dashed) and ethanol (solid). Values were calculated by binning the normalized heading shown in B into four 90° sections corresponding to upwind, crosswind (left/right combined), and downwind (see color bars on the abscissa of B). Light colored shading indicates bootstrapped 95% confidence intervals. The arrow indicates the 95% rise time to the peak fraction of trajectories flying upwind.

predictions were confirmed by my experimental results, which show a significantly tighter distribution of upwind heading for flies surging in the presence of stripes parallel to the wind-line ($p \leq 0.001$, Fischer’s exact test) (Fig. 2.6B). Flies also showed a significantly tighter distribution in the presence of faster wind speeds ($p \leq 0.001$, Fischer’s exact test), and a reduced accuracy in surging upwind at slow wind speeds (Fig. 2.6C), results that are also consistent with visual anemotaxis.

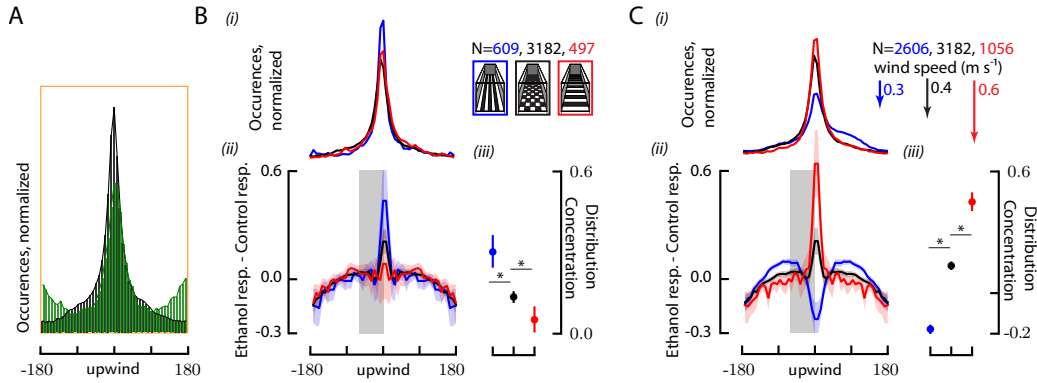


Figure 2.6: (A) Histogram of headings, binned across the time range outlined in orange in Fig. 2.5B. (B) (i) Comparison of heading histograms (calculated as in A) for three visual environments. (ii) Difference between heading histograms in Bi and their respective clean air control distributions, with enforced symmetry prior to subtraction to minimize artifacts. The light colored shading shows bootstrapped 95% confidence intervals. (iii) Concentration (peak sharpness) of each distribution in Bii, defined by the difference between the peak value and the mean across the gray-shaded region in Eii (mean and bootstrapped 95% confidence intervals). Statistically significant differences were calculated by the two-tailed Fisher's exact test (***) indicates $p \leq 0.001$). (C) Same as B, for different wind speeds, each with a checkerboard floor.

2.4.2 Casting behavior

Despite the unnatural constancy of the ethanol plume in my experiments, flies remained within the plume for only 250 ms. Within approximately 640 ms after leaving the plume, flies began casting crosswind (Fig. 2.7A-B). These maneuvers were not confined to a horizontal plane, but rather the flies tended to make casting movement oriented at oblique angles in the crosswind plane corresponding to $\pm 45^\circ$ and $\pm 135^\circ$ from vertical (Fig. 2.7C). This suggests that the casts are tuned to probe for the plume equally in the horizontal and vertical dimensions. The timing of cast initiation was independent of the time flies spent in the plume prior to leaving it (Fig. 2.7D). This suggests that casting is triggered reflexively by each plume loss event and is not strongly influenced by the flies' prior experience within the plume.

In addition to its role in upwind surges, flies might also employ visual anemotaxis during crosswind casts. To test this hypothesis, I again made use of stripes aligned either parallel or perpendicular to the wind line. If flies use visual slip to control

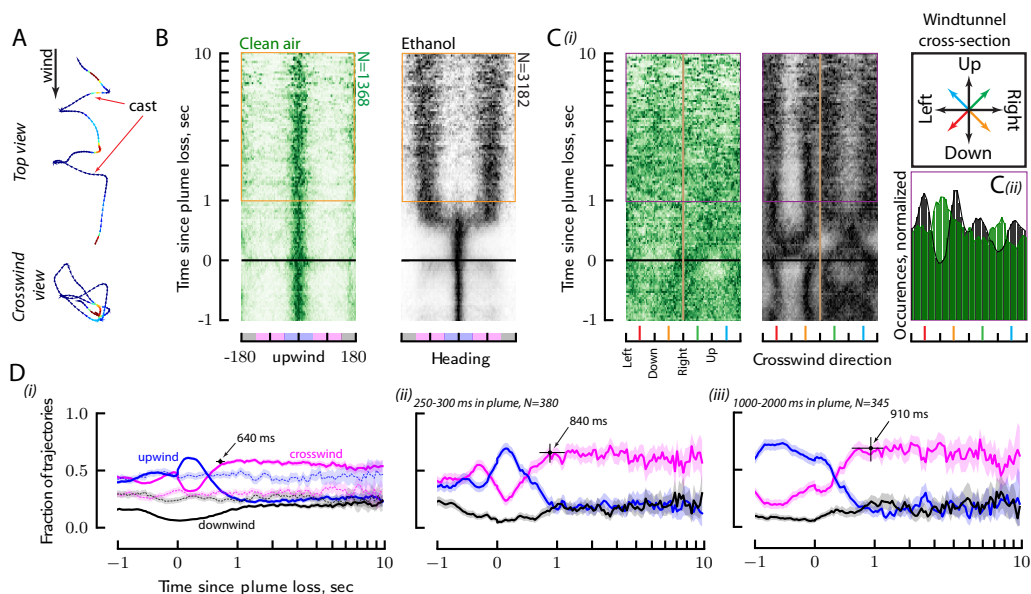


Figure 2.7: Within 600-800 ms after leaving an attractive odor plume, flies begin casting crosswind using visual cues, irrespective of wind speed. (A) A sample trajectory repeated from Fig. 2.2B that shows the stereotypical crosswind cast soon after the fly leaves the plume (color encodes odor concentration, see Fig. 2.1). (B) Heading (in the horizontal plane) of each trajectory relative to each time the fly leaves the odor plume, for the duration of time that it remains outside of the plume. The figures were generated from the same trajectory snippets and in the same manner as Fig. 2.5B. (C) (i) Heading of each trajectory in the cross-wind plane perpendicular to the wind direction relative to each time the fly leaves the odor plume, plotted in a similar manner as B. In B and C the visually apparent differences are statistically significant (see Fig. 2.18, p-val 0.001-0.01, Fisher's exact test). (ii) Normalized histogram of the headings for C_i, binned across the time range outlined in purple. (D) Fraction of trajectories flying upwind, crosswind, or downwind relative to plume entry for a pseudo-plume of clean air (dashed) and ethanol (solid). Values were calculated by binning the normalized heading shown in B in a similar manner to Fig. 2.5D. The arrows indicate the 95% rise time to the peak fraction of trajectories flying upwind (bar shows 95% confidence interval). Results are shown for (i) all the data in B, and two subsamples of the data corresponding to trajectories that spent (ii) 250-300 ms and (iii) 1-2 seconds inside the plume prior to departure.

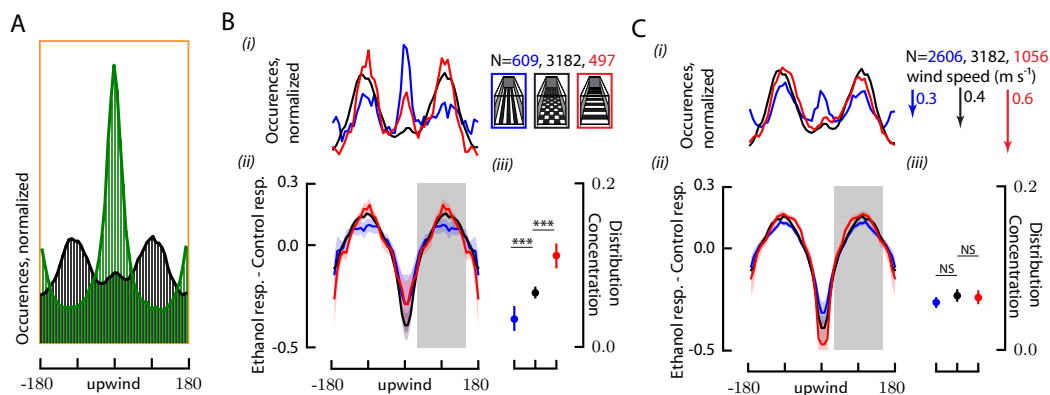


Figure 2.8: (A) Histogram of headings, binned across the time range outlined in orange in Fig. 2.7B. (B,C) Comparison of heading during casting behavior for different visual environments and wind speeds, plotted in the same manner as Fig. 2.6B,C.

their cast heading as they do during surging maneuvers, we would expect to find a tighter distribution about the crosswind direction if the lines are perpendicular to the wind, whereas we would expect a much broader distribution if the lines are parallel to the wind. This hypothesis was confirmed by my experimental results, which show a significantly ($p \leq 0.001$, Fisher's exact test) tighter distribution of heading in the crosswind direction when the lines are perpendicular to the wind (Fig. 2.8B). Unlike the upwind surges, however, I did not find any significant effect of wind speed on the flies' casting accuracy (Fig. 2.8C). The slightly increased proportion of flies that fly upwind in the slow wind case can be explained by the reduced odor injection flow, which was necessary to obtain a laminar plume structure under these conditions (see Methods).

The simplest mechanism by which a fly could maintain a crosswind flight heading independent of wind speed would be for it to orient its body into the wind by visual anemotaxis, and control its flight force vector so as to generate pure visual side slip, thereby maintaining a body orientation pointed directly upwind. To test whether flies use this approach, I installed an 11th camera directly above the wind tunnel looking straight down with a sufficiently narrow field of view to resolve flies' body orientation for a small sub-section of the wind tunnel (Fig. 2.9). I found that rather than maintaining a constant body orientation into the wind, flies instead cast with

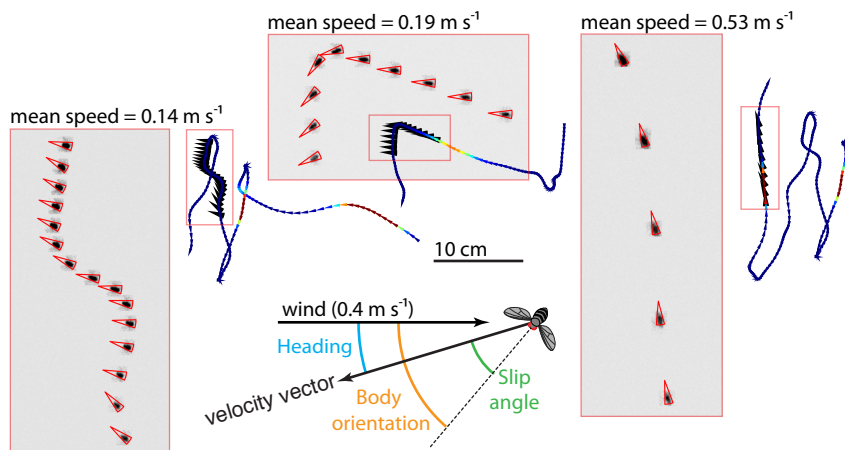


Figure 2.9: Flies cast with slip angles that range between 0-90°, and the slip angle is proportional to their flight speed. Three example trajectories of flies that pass through the ethanol plume prior to entering the volume over which I was able to track their body orientation with an auxiliary camera. The images show overlays of sequential camera images (recorded at 100 fps, displayed at 25 fps), together with the body angle estimates shown by the red triangle (the point indicates the head). The body orientation estimates were smoothed with a forward/reverse non-causal Kalman filter, and are shown as black triangles overlaid on the complete trajectories, which are color-coded as in Fig. 2.3.

a broad distribution of body orientations. To explain the variability in behavior, I examined the correlation between their flight speed and slip angle (where the slip angle is defined as the difference between their direction of travel and their body orientation). I found a strong correlation such that fast moving flies align their bodies with their flight direction during casts, whereas slow moving flies tend to fly sideways (Fig. 2.10). At intermediate speeds, flies employ an intermediate slip angle between 0° and 90°. Further, at slow wind speeds flies cast with a faster ground speed and a smaller slip angle, whereas in fast wind speeds flies cast with a slower ground speed and cast primarily by side slip. The fact that flies can fly crosswind by maintaining different slip angles suggests that the circuitry underlying optomotor anemotaxis is sophisticated enough that zero side slip can be regulated about an arbitrary optic flow axis with respect to the body and head.




Figure 2.10: From all the segments where I was able to collect body orientation data, I selected those that occurred 1-10 seconds after the fly encountered the odor plume (and is thus likely casting) and used these segments in plots i-iii. (i) Histogram of flight heading, similar to Fig. 2.8A. The difference in the histograms shown here and in Fig. 2.8A are expected since the auxiliary body orientation camera was centered and could only capture a small portion of the wind tunnel's width. (ii) Distribution of heading vs. body orientation. Each column of pixels is normalized such that it has a maximum and a minimum. (iii) Ground speed vs. slip-angle for the portions of ii that correspond to a heading within $\pm 20^\circ$ from crosswind (shown as a pink back drop in i).

2.4.3 Odor induced visual saliency

In a natural setting, tracking a chaotic and sparse odor plume is not only challenging, but it may never lead a fly to the actual source of the plume, although it will likely get them close. To pin point the source of the odor it seems sensible that a fly would use additional visual information to guide local search behavior. In my experiments with a checkerboard pattern on the floor of the arena I observed that flies often flew towards the bottom of the arena soon after leaving the attractive odor plume, and spent a disproportionate amount of time near the floor compared to their behavior in clean air (Fig. 2.11A). I hypothesized that this behavior was evidence of attraction to high contrast visual features following the detection of an attractive odor. To test the hypothesis further, I projected a small high contrast circle on an otherwise low contrast checkerboard floor, and projected two additional circles on the vertical walls of the tunnel. Under these conditions, flies did not exhibit as strong preference for exploring the floor of the arena, presumably because the contrast of the checkerboard pattern had been reduced (Fig. 2.11B). However, a close examination of the trajectories revealed that flies approached and hovered in the vicinity of the three small visual features after encountering the odor plume (Fig. 2.11C). These explorations of the visual features are obvious in heat-maps of residency time in the tunnel when compared to the clean air case (Fig. 2.12B).

The initial upward/downward symmetry in the distribution of trajectories leaving the plume ($t=0$, Fig. 2.11A) indicates that flies did not leave the plume by flying down in order to explore the checkerboards. Further, I did not find a correlation between time spent in the plume and the time course of the fraction of trajectories that explored the floor of the arena after leaving the plume. These results suggest that odor-induced visual saliency is triggered by the loss of an attractive plume, much like casting behavior. The accumulation of trajectories near the floor of the arena in the presence of a checkerboard floor raises the question of whether the aspects of casting behavior presented in Fig. 2.7-2.8 might be confounded by this behavioral effect. My results, however, show that this is not the case (Fig. 2.17). Instead, casting and the

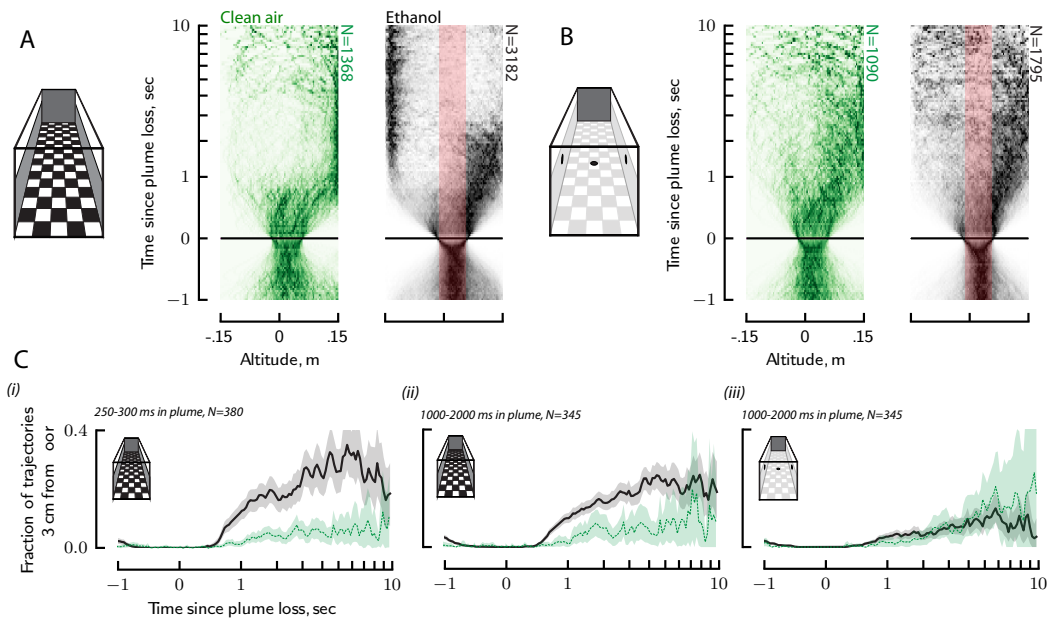


Figure 2.11: Odor increases visual saliency of high contrast objects. (A-B) Altitude response of flies relative to the time they exit the odor plume for different visual environments, for the duration of time prior to plume re-entry. The light red overlay in the ethanol case indicates the extent of the odor plume. In A and B, the visually apparent differences are statistically significant (see Fig. 2.18, p-val 0.001-0.01, Fisher's exact test). (C) Time course of object attraction is independent of time spent in plume. Light colored shading indicates bootstrapped 95% confidence intervals.

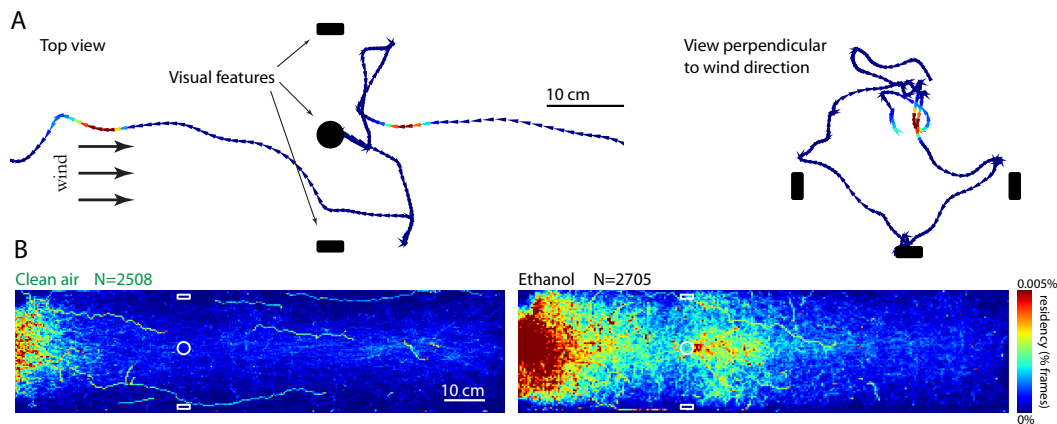


Figure 2.12: (A) Example trajectory of a fly that approaches all three of the high contrast features after encountering the odor plume. (B) Residency heat maps of flies trajectories in clean air, and in the presence of an ethanol plume. For both of these heat maps I only include frames where the fly is below the mid-line of the wind tunnel (which corresponds, approximately, to the lower extent of the odor plume).

exploration of high contrast features are distinct behaviors that are simultaneously triggered after plume loss.

2.5 Discussion

Although I set out to study plume tracking responses of flies to a pulsed plume, I noted that their behavior was quite similar to that elicited by a narrow continuous ribbon plume. After an initial contact, flies re-encountered the continuous plume with nearly the same interval (approximately 1 sec) as they did in the pulsed plume and pseudo-plume of clean air. Initial contact with the pulsing plume, however, was far less frequent, resulting in fewer trajectories. Flies spent a similar amount of time inside the bounds of the plume in each of the three cases (0.16-0.25 seconds), demonstrating that a fly's flight trajectory plays a larger role in its perceived olfactory experience than the dynamics of the odor plume itself. The density of trajectories above and below the plume at the time of departure in Fig. 2.11, together with the upwind heading at the time of departure shown in Fig. 2.7, indicate that flies left the plume more often through changes in altitude than horizontal heading.

These results are surprising considering that previous studies with tethered flies have shown that they have the sensory capacity to detect differences in odor concentration across their antenna, and thus should be able to track an odor gradient in flight [45, 58]. Although the circular laminar plume should present a best-case scenario for the use of osmotropotaxis during flight, I did not find any evidence that flies used this mechanism. This is not surprising, considering that in a natural setting concentration of odor is not a reliable guide to the direction of the source except when the source is decimeters away [109]. However, flies might use osmotropotaxis to decide in which direction they should make their cast after leaving the plume. For example, a fly might be able to detect whether it exited the plume going left or right, and thus initiate casting in the opposite direction (towards the plume). For flies leaving the plume with both small heading angles (10° - 20°) and large heading angles (45° - 135°) relative to upwind, I did find that flies were slightly more likely to turn

in the direction of the plume after leaving it (56% and 69%, respectively). However, this preference was not significantly different from the response of control flies in a pseudo-plume of clean air ($p > 0.4$, Fischer's exact test), suggesting that their decision is more likely based on visual cues from the walls of the wind tunnel, rather than olfactory ones.

2.5.1 Casting, surging, and odor induced object salience constitutes a stigmergic iteration

After encountering an attractive odor plume, flies initiated a surge within approximately 270 ms. This delay is generally consistent with previously reported values for flies [26] and other insects [29]. Presumably the delay is a limitation of the system (which includes olfactory processing, integration with visual stimuli, motor planning, transduction, biomechanics, and aerodynamics) rather than a programmed delay, and it has intriguing consequences. Often, the flies' flight trajectory prior to contact will take them through to the other side of the plume before they initiate a surge, resulting in only a brief encounter with the odor (insects have been found to react to short bursts of odor of less than 10 ms [167]). If the animals initiated their casting behavior with the same delay, they would make very slow headway towards the source. One possible solution to this problem is to incorporate a delay before initiating a cast in response to leaving the plume, guaranteeing some upwind progress for each plume encounter. This is indeed what I observed; the measured delay of approximately 640 ms seems longer than the minimum required for sensory-motor processing (and twice the delay associated with the surge). The additional delay also provides a low pass filter, which would serve to smooth their behavioral responses to the high frequency dynamics of natural odor plume filaments.

To quantitatively assess the importance of the casting delay, we constructed a simple computer simulation to test the performance of several plume-tracking algorithms (Figure 2.13). In our simulation, a single visual feature randomly emitted 3 cm diameter odor packets (mean five per second), which were advected by

a 0.4 m s⁻¹ wind that randomly shifted direction (according to a uniform distribution between -100° and 100° per second). For details, and source code, please visit: <https://github.com/florisvb/FlyPlumeTracking>. A virtual fly was released 1 m downwind of the source, and programmed to follow a simple algorithm consisting of the three independent behavioral modules we observed in our experiments: (1) surging upwind upon encountering an attractive odor with a delay of 270 ms; (2) casting crosswind within 640 ms of losing the plume; and (3) simultaneously approaching the visual feature after coming within 20 cm of it. We found that these three behaviors are sufficient, and each of them necessary, to guide the fly to an odor source. Furthermore, artificially reducing the delay before flies initiated a cast significantly increased the time before the fly located the food source (Fig. 2.13B).

Following Grassé, I use the term stigmergy to describe how a complex behavior, such as the trajectories I observed, can emerge from an iterative sequence of simple sensory-motor reflexes without the need for a plan or memory [68]. The delay between plume loss and casting initiation is independent of the amount of time the flies spent inside the plume, suggesting that casting behavior is most likely initiated in response to plume loss, rather than some internal mechanism. These observations are consistent with the previously published result that in the presence of a homogeneous plume, flies continue to surge upwind, never initiating casting behavior [26]. Together, these results provide convincing evidence for a stigmergic model. This model is distinct from what has been proposed to explain pheromone tracking of moths, which will initiate casting maneuvers in the presence of a homogenous plume rather than a continuous surge. To explain this behavior, Kennedy and Marsh proposed that casting is controlled by an internal clock, which is temporarily suppressed at the onset of odor [3, 81, 83]. Recent evidence, however, suggests that moths' casting dynamics are not simply a function of an internal clock, but rather that plume dynamics influence the amplitude of their casts [164]. An alternative to the internal clock mechanism is that moths exhibit the same stigmergic iteration that I propose for flies, and that their casting maneuvers in homogenous plumes are explained by a perceptual loss of the plume due to sensory adaptation, a phenomenon that is consis-

Figure 7

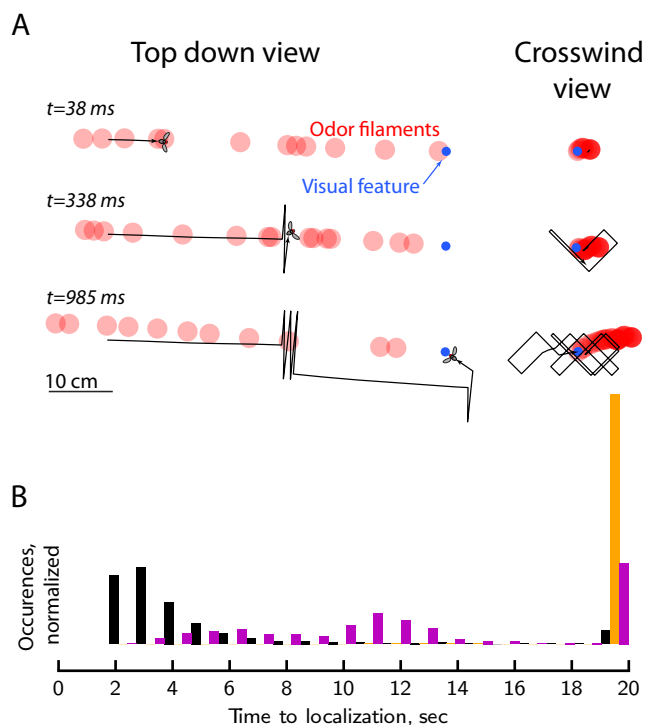


Figure 2.13: Surge, cast, and odor-induced object saliency are sufficient, and each of them necessary, for efficient plume tracking. (A) Three frames from a simple simulation that show a virtual fly interacting with randomly generated 3 cm diameter odor packets that are advected by a randomly shifting 0.4 m s⁻¹ wind. (B) Distribution of odor source localization times for three different algorithms each run 1000 times; the simulations were truncated at 20 sec. Black: the fly is programmed to follow the algorithm we observe real flies do, including surging with 270 ms latency, casting with 640 ms latency in both horizontal and vertical directions (switching of directions is arbitrarily programmed to occur every 500 ms in the horizontal direction, and 308.44 ms in the vertical directions), and approaching the visual feature (a behavior we arbitrarily trigger when the fly is within 20 cm of the object). Purple: the fly is programmed with the same algorithm as Black, except the latency for surging and casting are both set to 270 ms. Orange: same algorithm as Black, except without the attraction to visual features (the virtual flies fly past the feature due to the shifting winds and the sensory-motor delays).

tent with some experimental studies [5]. To determine whether moths might use the same mechanism that I propose for flies will require very careful behavioral assays where the olfactory experience is known for entire flight sequences. This may prove more difficult in moths, because their large flapping wings have a significant effect on the plume structure [85, 98, 149].

Casting in moths is typically characterized by increasingly wider lateral crosswind zigzags [41, 81, 89], however, I did not find any evidence for increasing width of casting maneuvers in my experiments with flies. This may be a result of the geometry of my wind tunnel (0.3x0.3 m² cross section), however casts in an open environment are typically no more than approximately 25 cm in amplitude (unpublished observations of freely flying flies in the laboratory). To minimize the influence of the tunnel geometry, I chose a visual stimulus for the walls of the wind tunnel—a single horizon line—to minimize collision avoidance behaviors [145], while still providing a visual reference for altitude control [141]. However, additional studies in a less constrained environment are necessary to determine precisely how flies control the timing between cast reversals, which is likely mediated by an internal mechanism. Because casting maneuvers are often generated by pure side slip maneuvers, this behavior will prove difficult to study in tethered preparations that are currently prevalent throughout the field, which detect changes in body-yaw [121].

Previous studies of plume tracking have restricted their analysis to the horizontal plane, with a few exceptions in plume tracking in moths [126, 155, 164]. The results, however, indicate that the oft-overlooked vertical component of casting behavior in flies is of the same magnitude as the horizontal component for flies. This is not surprising, considering the 3-dimensional nature of tracking an odor plume in flight. The relative timing of the vertical and horizontal aspects of casting has important implications on the overall search pattern flies execute in the crosswind plane. Equal timing would cause the fly to pass through the point where casting was first initiated in every cycle. A timing ratio close to an irrational number, however, would (over time) cause the fly to explore the entire crosswind plane. In the case of constant wind directions, where the highest likelihood of plume re-encounter is in the place where

the fly last encountered it, identical timing is optimal. In shifting winds, however, plume re-encounter may best be served by unequal timings. To determine which strategies flies use—and if they are environment dependent—is a question that can only be answered in a larger tunnel space.

2.5.2 The role of odor induced visual saliency

Flies show a greatly increased attraction to the downwind side of high contrast objects in the presence of an attractive plume, even though the objects were separated from the plume by more than 10 cm. This behavior could allow a fly to locate the source of a patchy plume more quickly than if it relied solely on plume tracking behavior, a hypothesis supported by my simulations (Fig. 7B). The attraction to visual features, like casting behavior, is triggered upon plume loss. By triggering both of these distinct behaviors simultaneously causes flies to cast in the crosswind direction just downwind of high contrast visual features, which is precisely where one would expect to find the odor plume emitted by a piece of fermenting fruit. If the fly senses an odor, a surging behavior would bring it towards the fruit and initiate the cascade of behaviors that ultimately lead to a successful landing.

In a previous study of free flight odor plume tracking in *Drosophila*, Budick (2006) found that flies exhibit a reduced attraction to a visual object (a post) in the presence of wind compared to still air. This result is interesting in light of the odor-induced visual saliency that I observed. Little is known about how insects localize odor sources in still air [29]. In the absence of wind, there can be no odor plumes to track with a surge and cast algorithm. Instead, it may be more beneficial for an insect to explore visual features in the hopes of encountering a food source. In the presence of wind, however, a more efficient strategy may be to focus on intercepting an attractive odor plume, a behavior best served by ignoring visual features until the presence of an attractive odor is identified. Although theoretical studies have been done to determine optimal search strategies for initial plume interception [127], behavioral evidence supporting them is lacking.

2.5.3 Plume tracking, and visual saliency, do not diminish over time

In my experiments with the continuous odor plume, odor was injected into the wind tunnel continuously for a period of 4 hours. It is conceivable that the flies' behavior could change during this time either due to their circadian rhythm, learning, or exhaustion. However, in comparing their surge, cast, and visual attraction behaviors for the first 30 minutes to the last 30 minutes I did not find any substantial behavioral differences (Fig. 2.19). These findings provide additional support for the simple reactionary stigmergic model.

2.5.4 How does a fly cast crosswind?

The observation that flies can accurately cast crosswind across a range of wind speeds has intriguing implications. For a fly to fly crosswind, it must have knowledge of the direction that the wind is coming from relative to its current direction of travel. A flying fly, however, has no known sensory mechanism for determining the component airspeed velocity it perceives that is due to ambient wind speed. To do so would require a measurement of absolute ground speed. A recent model suggest that insects could theoretically estimate absolute if they accelerated (or rotated) by a known amount while simultaneously measuring the change in its perceived wind speed and direction (see Chapter 4), however there is no evidence that they are capable of this computation.

A simpler solution would be for a fly to first orient itself upwind during the surge maneuver, and cast by making a pure side slip maneuver. This is not, however, what I observe. Instead, flies cast with side slip angles ranging from 0 to 90°, depending on their ground speed (Fig. 2.10). The next most parsimonious solution is again for the animal to begin by flying upwind, and then keep track of the angle by which it rotates its body while turning, either by storing the control signal used to generate the open loop maneuver, or through sensory mechanisms such as the halteres [8, 43] or antenna [102]. Knowledge of the turn angle would allow a fly to remember the angle




Figure 2.14: Summary diagram, indicating the three independent sensory motor modules that describe foraging in the fruit fly.

of the upwind direction relative to its body, and thus calculate at what angle it should experience visual flow to maintain a cross wind flight heading. Once it has turned the desired amount, it could use visual feedback to maintain the desired heading, making it robust to changes in wind speed and direction (even lulls in wind), which are common in natural flows. Although I do not have the data to test this hypothesis directly, evidence in experiments with moths where the wind flow was stopped shortly after they began casting showed that moths will continue to cast in the “crosswind” direction during lulls in wind [4].

2.5.5 Conclusion

I used an automated 3D tracking system to observe fruit flies, *Drosophila melanogaster*, as they interacted with a calibrated laminar plume of ethanol under a different visual environments and wind speeds. The flies’ tracking behavior can be described by a simple stigmergic model consisting of three distinct reflexes: (1) flies surge upwind within 270 ms of entering an odor plume: (2) cast crosswind 640 ms after losing the plume, and (3) explore high contrast visual features in the vicinity of the plume. All of these behavioral modules involve an integration of olfactory and visual information. Plume re-entry leads to another surging episode, and thus results in an iterative cycle that allows the fly to efficiently locate fermenting fruits in order to feed, find mates,

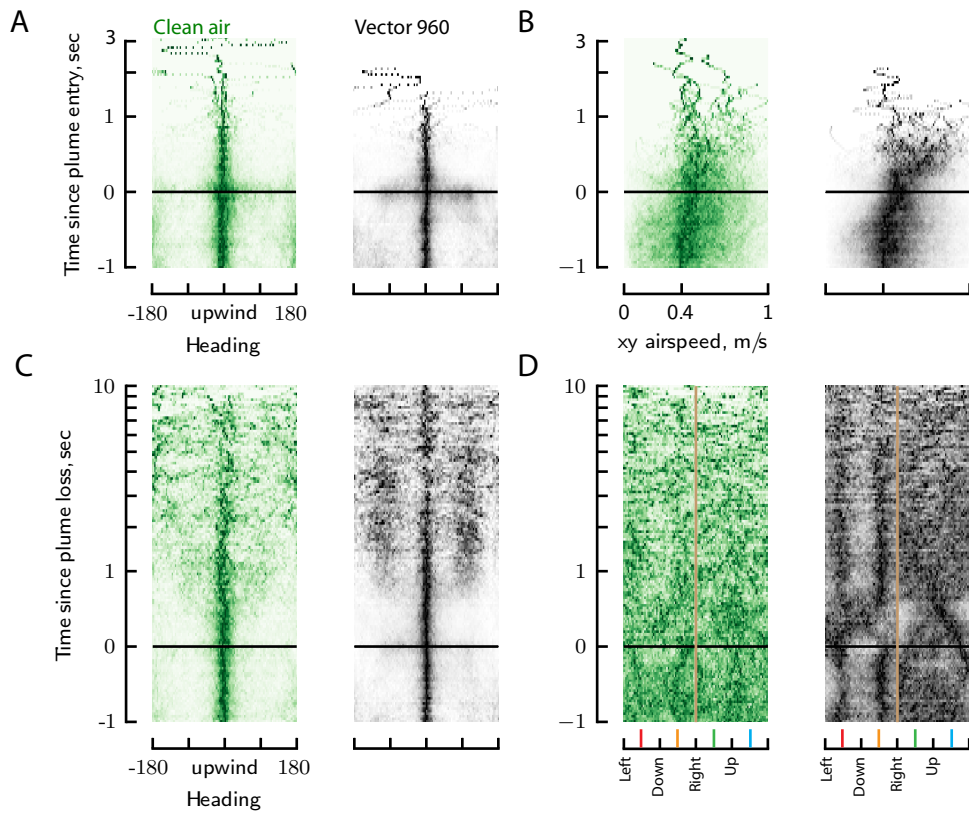


Figure 2.15: Behavioral responses to a plume of ethanol and Vector960 (fruit fly attractant) are qualitatively similar. (A) Surge heading, see Fig. 2.5B. (B) Surge airspeed, see Fig. 2.5C. (C) Cast heading, see Fig. 2.7B. (D) Cast heading in crosswind plane, see Fig. 2.7C.

and lay eggs.

2.6 Supplementary figures

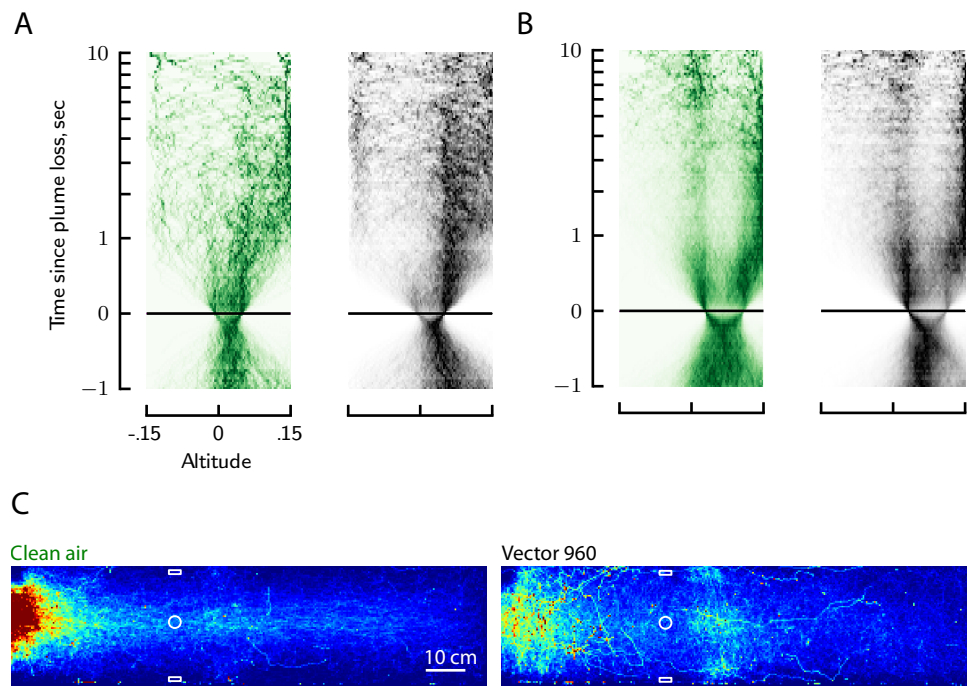


Figure 2.16: (A) Altitude response with checkerboard floor, see Fig. 2.11A. (B) Altitude response with low contrast floor and dots, see Fig. 2.11B. (C) Residency heatmap, see Fig. 2.12B.

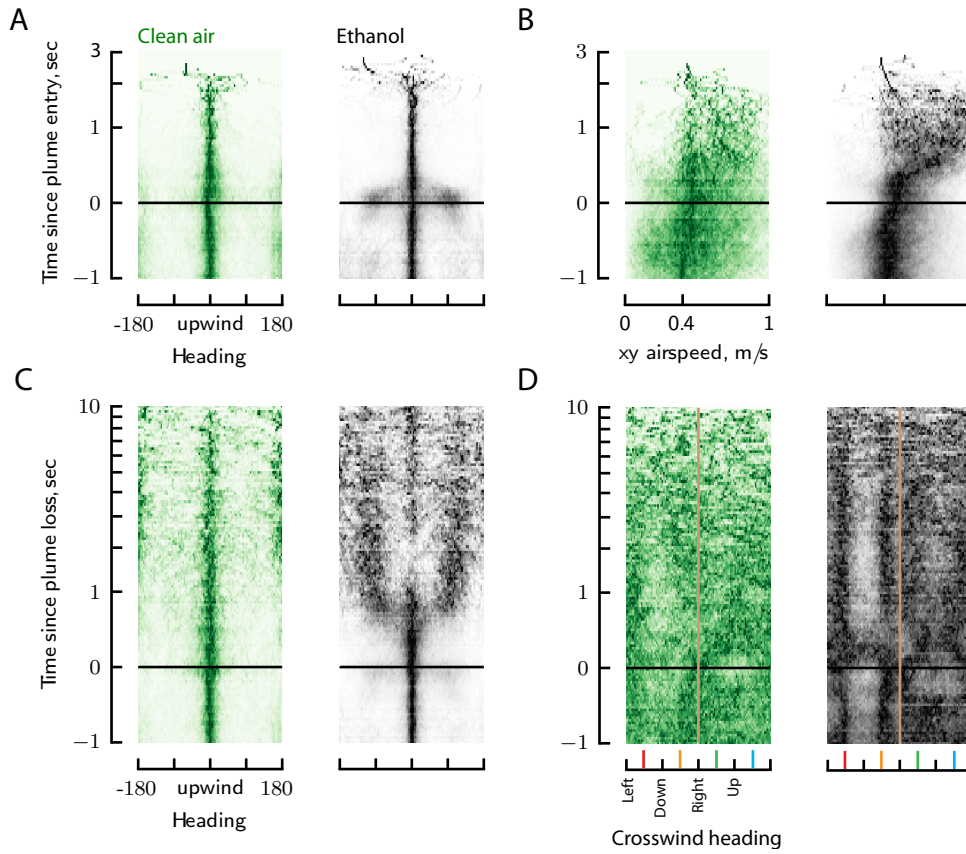


Figure 2.17: Cast and surge behaviors are qualitatively similar in the presence of a high contrast checkerboard floor and a low contrast floor with high contrast spots. After performing the experiments described in Fig. 2.11, I repeated the analysis presented in Figs. 2.5-2.8 with my data collected with the visual environment shown in Fig. 2.11B. I did not find any qualitative difference. (A) Surge heading, see Fig. 2.5B. (B) Surge airspeed, see Fig. 2.5C. (C) Cast heading, see Fig. 2.7B. (D) Cast heading in crosswind plane, see Fig. 2.8C.

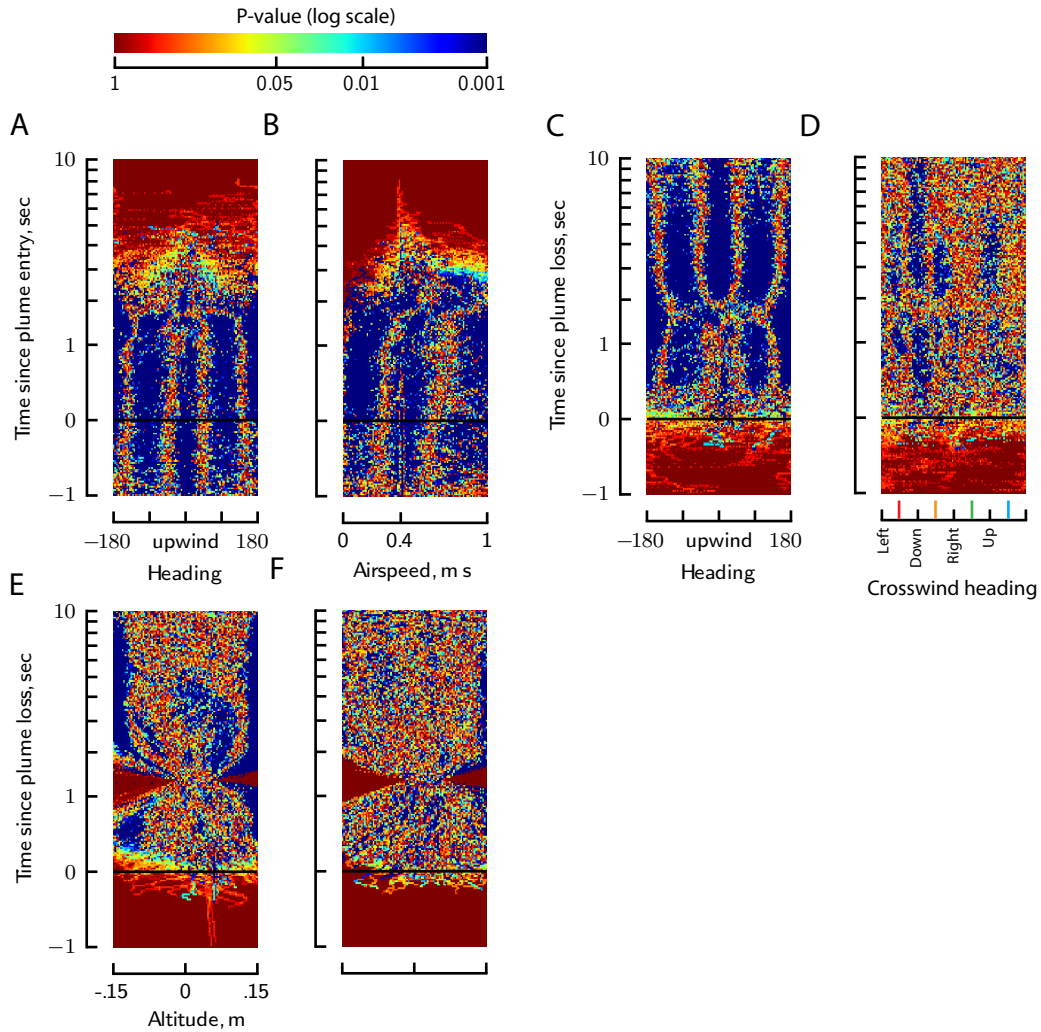


Figure 2.18: Visually salient results presented in Figs. 2.5-2.8, 2.11-2.12 are statistically significant. Each panel shows a heat-map (log color scale) representation of the per-pixel p-values that describe the probability that the differences in behavior that I observed in the presence of ethanol vs. clean air are due to random sampling. P-values were calculated by Fisher's exact test, a non-parametric permutation test. See Experimental Procedures for details. (A) Surge heading, see Fig. 2.5B. (B) Surge airspeed, see Fig. 2.5C. (C) Cast heading, see Fig. 2.7B. (D) Cast heading in crosswind plane, see Fig. 2.7C. (E) Altitude response with checkerboard floor, see Fig. 2.11A. (F) Altitude response with low contrast floor and dots, see Fig. 2.11B.

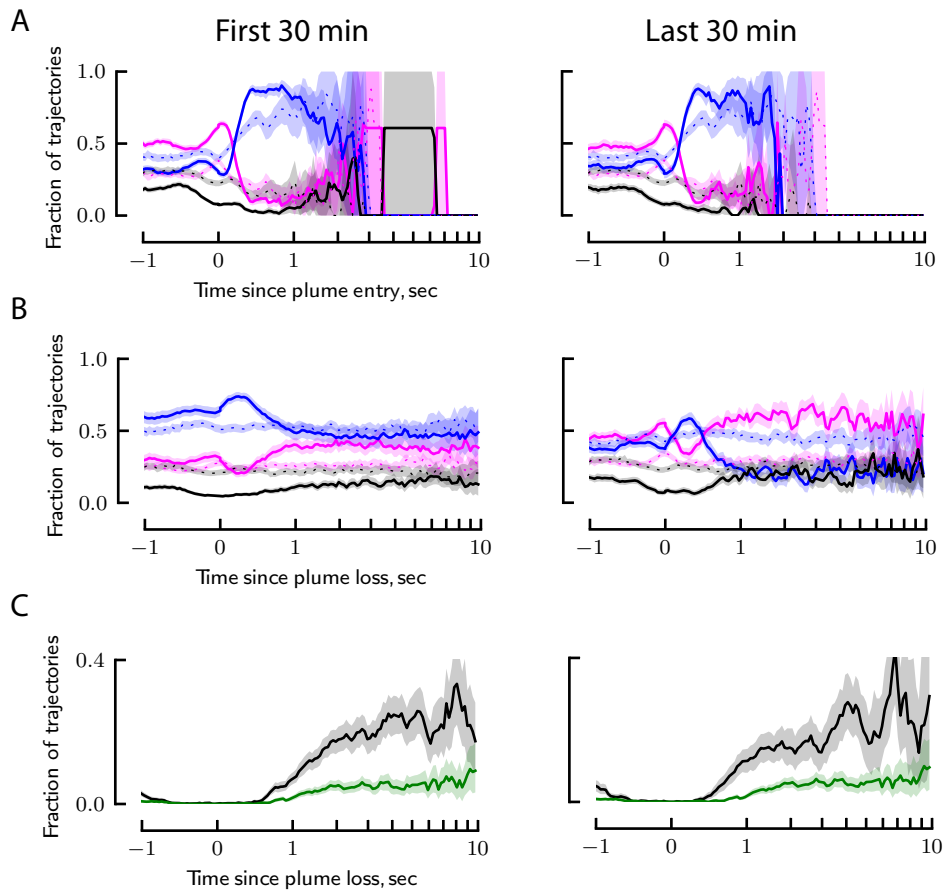


Figure 2.19: Surge, cast, and object saliency behaviors do not diminish over time. (A) Surge behavior, plotted in the same manner as Fig. 2.5D for trajectories collected in the first 30 minutes of the experiment (left), and the last 30 minutes of the experiment (right). (B) Casting behavior, plotted in the same manner as Fig. 2.7D. (C) Visual saliency behavior, plotted in the same manner as Fig. 2.11C.

Chapter 3

The visual control of landing and obstacle avoidance in the fruit fly, *Drosophila*¹

3.1 Abstract

Landing behavior is one of the most critical, yet least studied, aspects of insect flight. In order to safely land, an insect must recognize a visual feature, navigate towards it, decelerate, and extend its legs in preparation for touchdown. Although previous studies have focused on the visual stimuli that trigger these different components, the complete sequence has not been systematically studied in a free-flying animal. Using a real-time 3D tracking system in conjunction with high speed digital imaging, I was able to capture the landing sequences of fruit flies (*Drosophila melanogaster*) from the moment they first steered toward a visual target, to the point of touchdown. This analysis was made possible by a custom-built feedback system I built that actively maintained the fly in the focus of the high speed camera. The results suggest that landing is comprised of three distinct behavioral modules. First, a fly actively turns towards a stationary target via a directed body saccade. Next, it begins to decelerate at a point determined by both the size of the visual target and its rate of expansion on the retina. Finally, the fly extends its legs when the visual target reaches a threshold retinal size of approximately 60°. My data also let me compare landing sequences to

¹The material presented in this chapter is almost exactly transcribed from my publication [150].

flight trajectories that, although initially directed toward a visual target, did not result in landing. In these fly-by trajectories, flies steer toward the target but then exhibit a targeted aversive saccade when the target subtends a retinal size of approximately 33° . Collectively, the results provide insight into the organization of sensory motor modules that underlie the landing and search behaviors of insects.

3.2 Introduction

Deciding where and when to land, and subsequently performing a successful landing, is arguably the most critical aspect of flight behavior for most flying animals. Despite this importance, landing has been subject to much less research compared to other aspects of flight behavior. This is likely due to the difficulties of enticing insects to land in specific locations where they can be carefully observed. Unlike take-off, which begins at a fixed location and takes place within a relatively small spatial volume, landing consists of a complex sequence of different behavioral modules separated in space and time. An animal must orient towards a particular sensory feature, alter its wing motion to change its speed and posture, extend its legs, and then make contact with the substratum.

Previous studies of landing have focused on the visual stimuli that trigger and regulate landing behavior in a variety of insects including hoverflies (*Syrirta pipiens*) [38], houseflies (*Musca domestica*) [15, 17, 20, 143, 16, 157], blowflies (*Calliphora eurythrocephala*) [65, 46] (*Apis mellifera*) [134, 49], the milkweed bug (*Oncopeltus fasciatus*) [35], and fruit flies (*Drosophila melanogaster*) [144, 165]. Furthermore, electrophysiological recordings exist for neurons believed to be associated with landing and collision avoidance (locusts (*Locusta migratoria*): [123, 56, 57]; blowflies: [143]). The fruit fly, *Drosophila melanogaster*, is well suited for studies of landing behavior because its small size permits the analysis of long flight sequences within a controlled laboratory setting. This makes it possible to study landing behavior from the moment a visual target first influences an animal's flight motion to the point of touchdown. In addition, the visual system of *Drosophila* has been extensively studied [18], largely

due to the genetic tools that are uniquely available in this species [131].

Although the free flight landing sequences of *Drosophila* have not been studied explicitly before, several behavioral modules that are likely to represent components of the behavior have been studied extensively. First identified in mosquitoes by Kennedy [82], flying *Drosophila* exhibit a robust orientation reflex to vertical contrast edges known as fixation ([67, 66]; see also [120]). Although the functional relevance of this tethered flight phenomenon is not clear, free flight experiments suggest that fixation might serve to lead flies towards salient visual objects [100] and thus could represent the earliest component of a landing sequence. Further, *Drosophila* and other flies often exhibit a flight pattern characterized by relatively long segments of straight flight, interspersed with rapid turns called saccades [38, 161, 145]. If these saccades are directed toward certain features, they too should be considered an early component of the landing sequence. Although past studies have shown that visual cues play a large role in triggering saccades and influencing their course direction (i.e., left vs. right) [73, 145, 9, 135], there is not yet definitive evidence that *Drosophila* make directed saccades, such that they turn precisely to the angular position of a particular visual feature.

Once a fly starts approaching an object, it will receive an expansion cue, which has been shown to elicit a robust collision avoidance response [144, 9, 135]. In order to land, however, the animal must override this reaction and maintain a collision course, reducing its flight speed such that it can touch down safely. One free flight study of houseflies, suggests that visual cues are responsible for triggering this deceleration phase prior to landing [157].

Tethered flies exhibit a robust leg extension reflex in response to an expanding visual stimulus (green bottle flies: [65]; houseflies: [15, 17]; fruit flies: [144, 165]). Although this reflex has been previously termed “the landing response,” it only represents one component of landing behavior and its position in the complete free flight landing sequence is not known. Each of the elements of the landing sequence, including saccade generation, deceleration, and leg extension, have been shown to be visually mediated [157, 65].

In this study, I examine the landing behavior of freely flying *Drosophila* from the point when they can initially resolve a visual target to the moment of touchdown. The analysis was possible in part because of an automated 3D tracking system that allowed me to collect a large number of landing trajectories, as well as a real-time focus-following system that permitted the capture of high temporal and spatial resolution images during the final stages of this behavior. My dataset also enabled me to explicitly compare landing sequences with sequences in which flies initially flew towards the target, but then steered away from it. The results suggest that landing sequences begin with a body saccade directed toward the vertical edge of a visual target. Flies then begin to decelerate at a point determined by both the retinal size of the target, and its rate of expansion. Finally, flies extend their legs when the visual target subtends a critical angle. The nature of these triggering mechanisms normally ensures that a fly begins to slow its approach and then extends its legs in time for touchdown. In cases in which flies steer toward a target, but then do not land, the initial orienting behavior is followed by an aversive saccade that is triggered when the target subtends a critical angle. Collectively, the results indicate how a complex behavioral cascade may emerge from a temporal sequence of separated sensory-motor modules.

3.3 Methods

3.3.1 Animals

Experiments were performed on 3-5 day-old female fruit flies, *Drosophila melanogaster* Meigen, from a laboratory stock descended from a wild-caught population of 200 mated females. Flies were deprived of food, but not water, for 2-6 hours prior to the start of the experiment in order to motivate flight. For each experimental trial, I introduced a group of 12 flies to the corner of the arena within a small test tube. The flies were then free to move throughout the flight arena for a period of 12-24 hours, during which time data were collected automatically. A wet KimWipe in one

of the corners of the flight arena provided the flies with water for the duration of the experiment.

3.3.2 Flight arena

I performed all experiments in a 1.5 m x 0.3 m x 0.3 m working section of a wind tunnel (Fig. 3.1) that has been described previously [26, 100, 141]. In these current experiments, the wind tunnel was switched off, so that the internal air was still. To provide the flies with visual contrast, I projected the floor and two long side-walls with a static black and white checker pattern (checker size 3 cm, 11.5° retinal size at 15 cm distance) using a Lightspeed Designs DepthQ projector with the color wheel removed (120 Hz update rate, 360 Hz frame rate, mean luminance of 50 cd/m²). The two shorter walls of the chamber, consisting of the fine mesh screens of the upstream and downstream ends of the wind tunnel, were not illuminated. The ceiling of the chamber was transparent acrylic.

I tracked the 3D position of individual flies within the chamber using a real-time tracking system that is described in detail elsewhere [140]. The 6 camera system generated an estimate of fly position at 100 frames per second with a median latency of 39 ms. For purposes of tracking, the arena was backlit with an array of near-infrared (850 nm) LEDs. The cameras were equipped with long-pass filters (Hoya R-72) so that the camera images were not contaminated by the checkerboard pattern that was displayed in visible wavelengths.

3.3.3 Experiment protocol

In order to observe the flies' response to a conspicuous visual object, I placed a 15 cm tall, 1.9 cm diameter post in the center of the arena. The top of this post reached the half height of the arena. I used two different types of posts: a solid black post (matte black spray paint on a smooth aluminum cylinder) and a black and white checkered post (5 mm checkers, 5° retinal size at a distance of 5.7 cm, printed on white paper and tightly wrapped around an aluminum cylinder). I also collected data in the

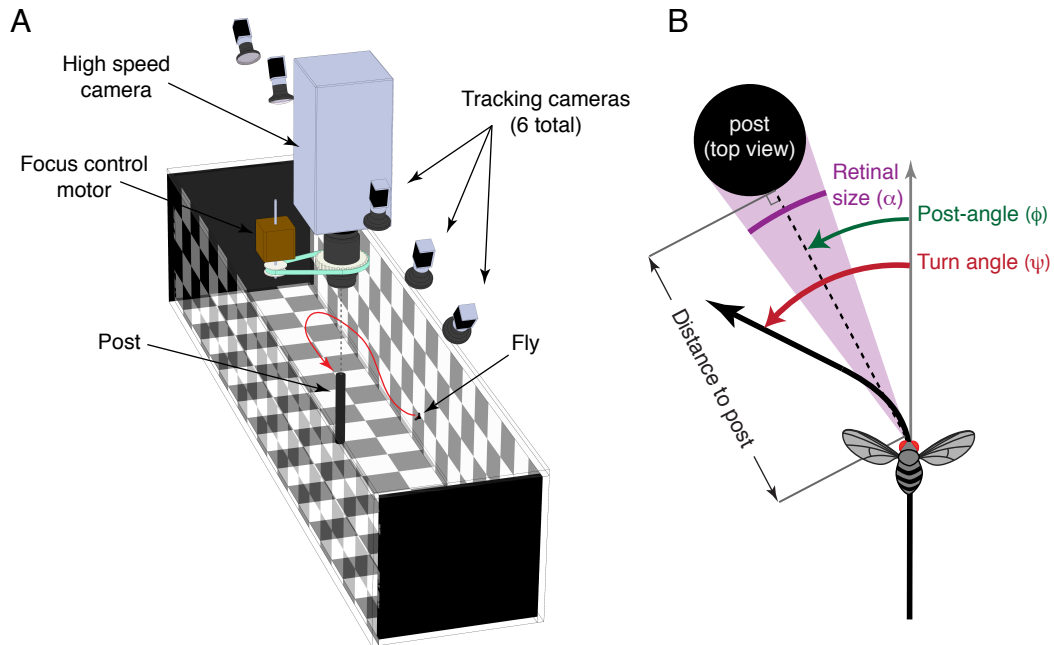


Figure 3.1: (A) Free flight arena equipped with 3D tracking system and high speed video camera with automated focus system. (B) Definition of visual parameters used throughout paper. Retinal size (α) is defined as angle subtended by the post on the fly's retina (see (3.1)). Post-angle (ϕ) is defined as the azimuthal angle of the center of the post from the fly's flight path. Turn angle (ψ) is defined as the angular change in the fly's heading over the course of a saccade (see Fig. 3.3-3.4). This term is synonymous with saccade amplitude.

Post type	N trials	Total trajecs	Mean num trajecs / trial	Min num trajecs / trial	Max num trajecs / trial	Total num trajecs after culling	Percent Landing
Black	14	25393	1813	4	5611	578	32%
Checkered	15	45908	3060	269	14499	608	6.00%
No Post	3	9951	3317	275	9951	266	NA

Table 3.1: Trajectory statistics.

absence of a post. Because both visual inspection and statistical comparisons of the data suggested that the results were only marginally influenced by post texture (solid black vs. checkered, see Fig. 3.19-3.20), I pooled both data sets for the bulk of my analysis unless otherwise noted.

3.3.4 Trajectory reconstruction and analysis

All analyses of flight trajectories were done using custom software I wrote in Python. Each trajectory was treated as an independent sample, as the tracking software was not able to maintain fly identities over the extended period of my experiments. As a result, it is impossible to tell how many individual flies contributed to my analysis (the possible range is 29-348). The tracking system logs all locomotor behavior within the arena including bouts of flying and walking. The data presented are derived from 29 experimental runs, and 3 controls, which generated a total of 81,252 trajectories. Further details of the experimental database are provided in Table 3.1.

Each trajectory was smoothed to remove digitization errors using a simple forward/reverse, non-causal Kalman filter. Unless otherwise noted, the control data from the no post trials were culled and analyzed in the same way as the data collected with the post, using an imaginary post with the same dimensions and in the same position as the real post. For my analysis, I only considered trajectories that started at a distance greater than 10 cm away from the post, and approached to within at least 3 cm of the post. Furthermore, except for the no post controls, I removed trajectories in which the flies flew above the top of the post at any point within the region of interest. This procedure was taken to maximize the likelihood

that I was examining flight behavior that was influenced by the presence of the post. These criteria removed the large majority of the original trajectories, which were often short sections of either flying or walking. As I did not collect as much data under the no-post control arrangement, I did not require the control trajectories to be below the (not present) post height. Because of the visual symmetry of the upper and lower halves of the wind tunnel in these experiments, this difference in processing of the no post data is unlikely to influence my results. Further, very few of my conclusions depend on an explicit comparison with the no post data.

I used a simple algorithm to automatically classify trajectories as either ‘landings’ or ‘fly-bys’. Trajectories that ended within 1 cm (approximately 4 body lengths) of the post with a velocity of less than 6 mm/sec were labeled as landings; all others were labeled as fly-bys. These soft criteria helped in preventing erroneous classifications. Furthermore, visual inspection of the raw data in Fig. 3.2 shows that my criteria were sufficient to properly segregate landings and fly-bys. Data from the no-post controls were also identified as pseudo-landings (trajectories that intersected the 3D volume of the an imaginary post with identical dimensions to the real one) and pseudo-fly-bys (all other trajectories). All landing sequences were analyzed from start to finish. Fly-bys were only analyzed from the beginning of the trajectory until the point just prior to the first saccade following the closest approach to the post (see Fig. 3.2). This procedure was necessary to eliminate the portion of the sequence in which the fly was flying away from the post, and thus unlikely to be responding to it.

Flight trajectories were described using a number of variables calculated in each frame, including distance to the post (measured to the post surface), ground speed, acceleration, heading (calculated as the tangent to the velocity vector), and angular velocity. Except where otherwise noted (e.g., as in Fig. 3.2, right hand column), all values were calculated as projected in the x-y plane, thus ignoring changes in altitude. In order to focus my analysis on experimental parameters that can provide insight into possible sensory processes, I use the following variables that are defined relative to the fly’s position (see Fig. 3.1B). Turn angle (ψ) is the angular change in the fly’s heading over the course of a saccade. Post-angle (ϕ) is the angle between the fly’s

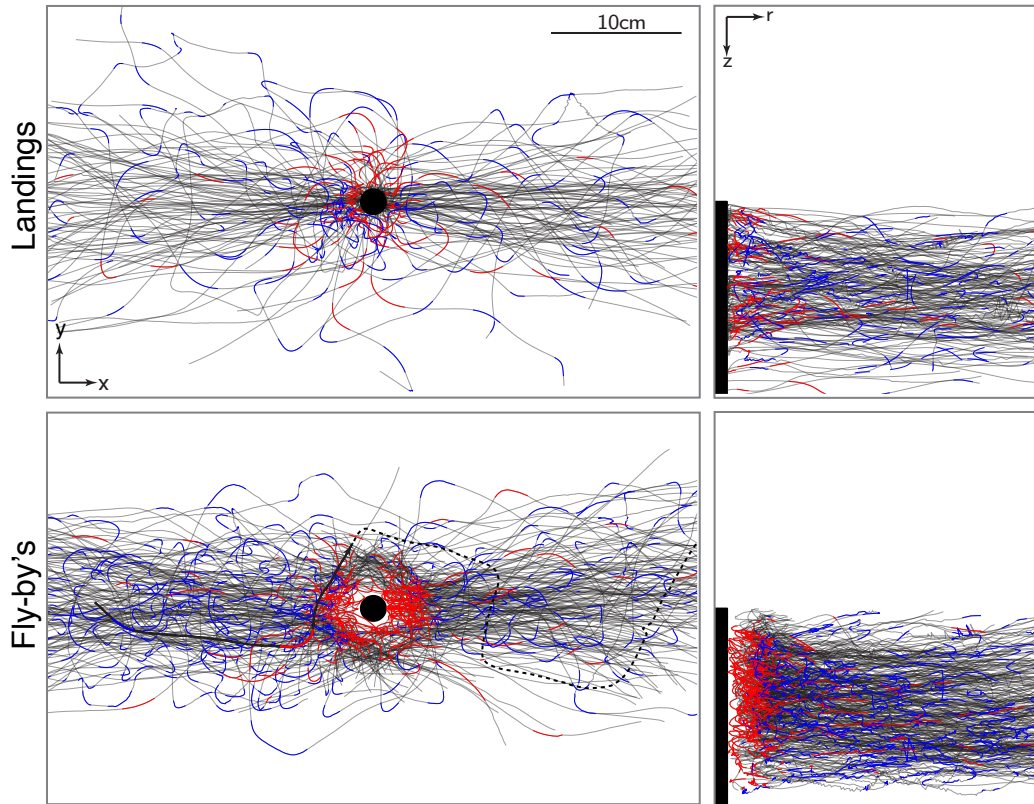


Figure 3.2: Flight trajectories in the presence of a post, classified as fly-bys ($n=300$) and landings ($n=177$). The trajectories for fly-bys represent a subset of the total data set of $n=1065$. The thick red overlays on each trace indicate the portions of the trajectories that were classified as the last saccade each fly made prior to its nearest approach to the post, which is the focus of my subsequent analysis. See Fig. 3.3-3.4 and the Methods section for my definition of a saccade. The thick blue overlays indicate all other saccades. The dotted line in the bold fly-by trace indicates the portions of trajectories after the flies' closest approach to the post. These portions were omitted from the rest of my analysis because it was unlikely a fly's behavior was influenced by the post after this point.

heading and the vector to the post. Retinal size (α) is the angle subtended by the post on the fly’s retina, which is calculated as:

$$\alpha = 2 \arcsin \left(\frac{r}{d} \right), \quad (3.1)$$

where r is the radius of the post, and d is the distance from the fly to the center of the post.

3.3.5 High speed imaging

To examine the landing and fly-by behavior near the post with greater spatial and temporal resolution, I mounted a high speed camera (Photron SA1, San Diego, CA, USA) looking down, approximately 30 cm above the post (see Fig. 3.1), equipped with a 105mm Nikon macro lens (Nikon USA, Melville, NY, USA). I operated the camera at 5,000 frames per second with a resolution of 1024x1024 pixels. To overcome the depth of field limitations imposed by the limited available lighting and the large magnification needed to resolve the flies’ legs, I designed an automated motorized focus-following system. The custom built system used the real-time elevation measurements from the 3D tracking system to automatically adjust a friction belt connected to the manual focus ring of the lens. This system allowed us to capture sharp video of flies at any elevation in the flight arena. The camera was post-triggered whenever a fly came to within 1 cm of the post (capturing both landings and close fly-bys).

One initial limitation of my focus-following system is that due to the distortion from the lens I could not focus through the top of the cylindrical post to visualize a clean circular section at any point along the height of the post. To solve this problem, I machined a gently tapering post (1.9 cm at the base and 1.3 cm at the top). This tapered shape made it possible to see the moment of touchdown along the entire length of the post without obstructions from the top. To account for the effect of the slight change in diameter of the post in these high speed experiments, I used the diameter of the post at the altitude of the fly, when calculating the retinal size of the

post. This slight modification in my analysis is reasonable, given that trajectories did not vary significantly with altitude. Furthermore, I replicated all analysis using a fixed value for the diameter equal to the mean diameter of the post, but found no significant difference in the results.

3.3.6 Analysis of saccades

The free flight trajectories of fruit flies are characterized by relatively long straight segments and short rapid turns, called saccades (for an example, see Fig. 3.3B-C). I focused much of my analysis on these saccadic flight segments, which I define as the portions of a trajectory during which the angular velocity exceeds a threshold of 300° s^{-1} . To justify both the particular threshold I used to define saccades as well as my decision to focus on them, I performed an extensive analysis of the angular velocity of all the trajectories I collected (in the presence of the post), and a subsequent analysis of the segments I labeled as saccades (Figs. 3.3, 3.4).

The distribution of angular velocities of freely flying flies is well approximated by the sum of a Gaussian distribution ($\mu = 0 \text{ s}^{-1}$, $\sigma = 85^\circ \text{ s}^{-1}$) and a log-normal distribution ($\mu = \pm 300^\circ \text{ s}^{-1}$, $\sigma = 2.3^\circ \text{ s}^{-1}$) (Fig. 3.3A). These results are similar to those published on tethered flies [105, 9] and as in those studies I take this as evidence for two distinct flight modes: noisy straight flight (described by the Gaussian distribution) and active saccadic turns (described by the log-normal distribution). Although the precise shape of the distribution is to some extent a function of the filtering used to calculate the angular velocity, I did not find any significant difference in the shape when angular velocity was calculated directly from heading (Fig. 3.3C, blue trace) or a Kalman estimate (Fig. 3.3C, black trace). The distribution shown in Fig. 3.3C comes from the Kalman-estimated angular velocities, which I also use in all subsequent analysis to eliminate measurement noise, such as the transient seen in Fig. 3.3C.

I define saccades as the portions of trajectories in which there is a $>95\%$ probability that the point lies within the log-normal distribution shown in Fig. 3.3A, a

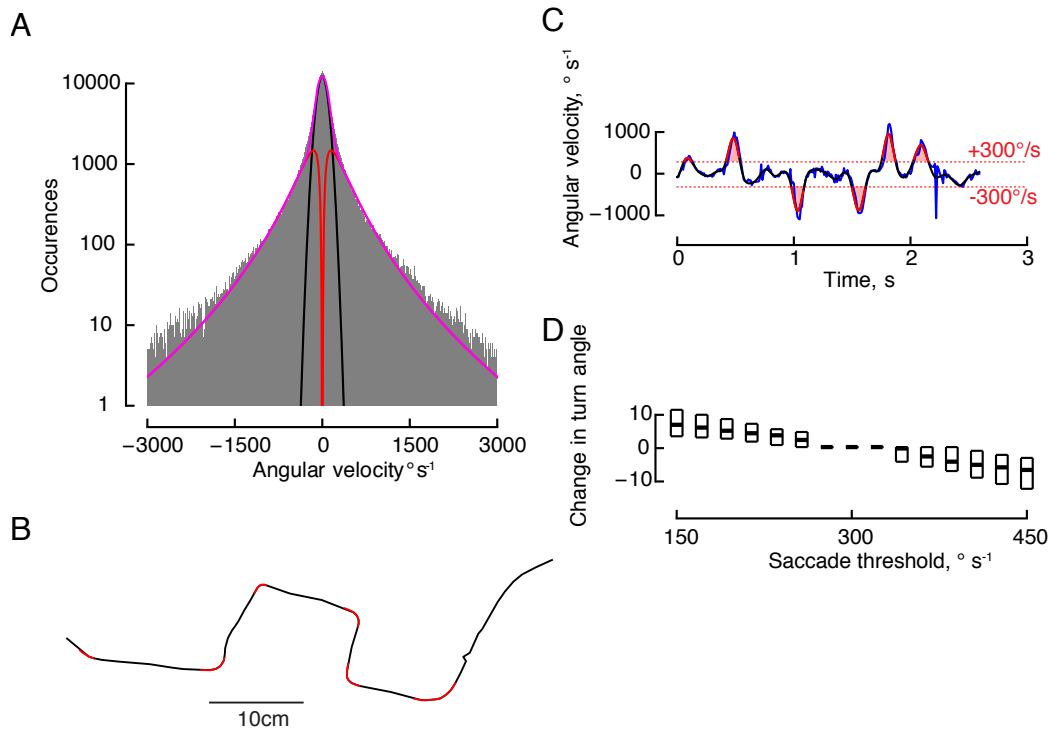


Figure 3.3: Changes in flight heading are primarily accomplished using body saccades (defined as flight sequences with angular velocity (ω) $> 300^{\circ} \text{s}^{-1}$). (A) Histogram of angular velocity (ω) at all time points for all trajectories ($n=1224$, $N=352,680$). The distribution is well approximated by the weighted sum (S, magenta) of a Gaussian distribution (G, black, $\mu=0^{\circ} \text{s}^{-1}$, $\sigma=85^{\circ} \text{s}^{-1}$) and a log normal distribution (L, red, $\mu=\pm 300^{\circ} \text{s}^{-1}$, $\sigma=2.3^{\circ} \text{s}^{-1}$) such that $S = 0.8G + 0.2L$. (B) A typical free flight trajectory, seen from above, with red portions indicating saccades. (C) Angular velocity of the trajectory in B, without smoothing (blue) and as a Kalman estimate (black). The sharp transient in angular velocity at $t=2.2$ s is due to a measurement error, such events are eliminated by the Kalman filter. (D) Changes in the ω –threshold used to classify saccades have little effect on measurements of turn angle (ψ). The relative change in ψ is plotted as a function of the saccade threshold; box plots indicate mean and first and last quartiles.

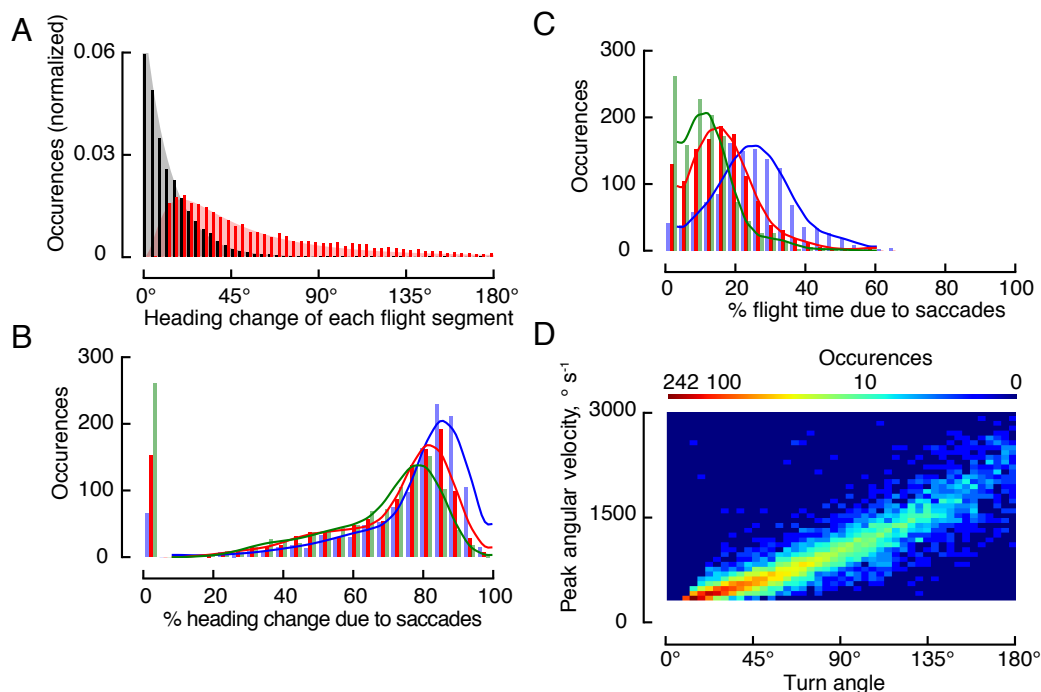


Figure 3.4: Changes in flight heading are primarily accomplished using body saccades (defined as flight sequences with angular velocity (ω) $>$ 300° s⁻¹). (A) Distribution of net change in heading of each saccade segment (ΔH_s) (red, N=5635) compared to each non-saccade segment (ΔH_n) (black, N=6697) for all trajectories (n=1224). The distributions are well approximated by a log-normal distribution (red, $\mu=40^\circ$, $\sigma=2.3^\circ$), and an exponential distribution (gray, $\lambda=0.075$). (B) Approximately 80% of all net changes in heading occur during saccades. The distributions show the ratio of the sum of ΔH_s to the sum of $\Delta H_s + \Delta H_n$ for each trajectory. I tested three definitions for the ω -threshold used to define a saccade: 200° s⁻¹ (blue), 300° s⁻¹ (red), and 400° s⁻¹ (green). The associated curves are smoothed representations of the distributions calculated using a 3rd order 0.4 Hz Butterworth filter. (C) Saccades account for approximately 20% of the flight time. The histograms show the ratio of time a fly was saccading to total trajectory time for each trajectory. The color scheme and curves are defined as in F. (D) Peak angular velocity vs. turn angle for all saccades, plotted as a heat-map with a logarithmic color scale.

classification which corresponds to an angular velocity threshold of 300° s^{-1} . This is similar to the thresholds used in prior work in free flight (300° s^{-1} , [145]) and magnetically tethered flight (350° s^{-1} , [9]). The amplitude of these saccades, referred hereto after as turn angle, does not vary by more than $\pm 10^\circ$ within a range of saccade threshold choices from 150° s^{-1} to 450° s^{-1} (Fig. 3.3D).

The saccades I observed ranged in amplitude from 20° - 270° (Fig. 3.4A), and are well described by a log-normal distribution with a mean of 40° and scaling factor of 2.3° . Although these turns only comprise approximately 20% of the flight time (Fig. 3.4C), they account for roughly 80% of all net navigational changes in heading (Fig. 3.4B). The turn angle is also tightly correlated with peak angular velocity (Fig. 3.4D), which closely matches prior results found in tethered flies [9]. This shows that flies do not perform turns of arbitrary amplitude at constant velocity, and suggests that each turn is an isolated maneuver that can be analyzed independently of the sequences that precede and follow the turn. Although it is possible that flies actively adjust their heading during the remaining flight segments, analyzing these aspects of flight is beyond the scope of this paper. I did not find a significant correlation between turn angle and horizontal flight speed ($R^2=0.02$), nor between angular velocity and flight speed ($R^2<0.001$).

3.3.7 Procedures for analyzing landing behavior

To study the onset of deceleration in landing flies I followed each trajectory backwards in time, starting from the point of landing, and defined the first negative to positive crossover of acceleration (defined as the derivative of flight speed in the x-y plane) as the point of deceleration initiation. For the subset of my landing data for which high speed imaging was available, I scored leg extension manually as an all or nothing event and defined the point of leg extension as the first point in time when the legs appeared to be fully extended.

3.3.8 Statistical analysis

Throughout the paper, I used multiple linear regression to determine whether relationships between two continuous variables, y and x_1 (e.g., speed and deceleration initiation, see Fig. 3.9) were influenced by a third categorical variable, x_2 (e.g., whether or not a fly turned after initiating deceleration, see Fig. 3.9). In these analyses, x_2 was represented as a boolean variable (0 or 1). I used the statistical program R [118] and the Python wrapper rpy2 [60] to fit the linear model to the data (R command `lm`):

$$y = a_1x_1 + b_1 + b_2 + a_{1:2}(x_1 : x_2), \quad (3.2)$$

where a_1 and b_1 are the slope and intercept of the regression between the primary variables (y and x_1), b_2 is the intercept due to the third variable (x_2), and $a_{1:2}$ is the slope due to the interaction between x_1 and x_2 . Essentially, this is analogous to running an analysis of covariance, but allowing for interactions between the covariate x_1 and the categorical variable x_2 . To show whether the impact of x_2 was significant, I report a pair of p-values ($p(x_2) = b_2, p(x_1 : x_2) = a_{1:2}$), where b_2 corresponds to the intercept due to x_2 and $a_{1:2}$ corresponds to the slope due to the interaction between x_1 and x_2 , respectively. T -tests tests were performed using the statistics sub-module of SciPy [78], a Python module. I define a threshold of $\alpha = 0.01$ to be significant, and $0.01 < \alpha < 0.05$ to be marginally significant.

3.4 Results

3.4.1 Description of landings and flybys

Excluding trajectories that were above the level of the post, I analyzed a total of 1224 flight trajectories (Fig. 3.2) (in addition to 194 no post controls (Fig. 3.17)). Flight speeds (measured at a distance of 10 cm from the post) for landing and fly-by trajectories were similar in all my experimental conditions with a mean of 0.33 ± 0.12

Behavior	Post Type	N trajecs	Speed, m s ⁻¹	Total length, s
			m±s	m (min, max)
Fly-by's	all	1047	0.32±0.12	3.03 (0.37, 24.99)
	black	439	0.31±0.13	3.28 (0.37, 24.99)
	checkered	608	0.34±0.12	2.84 (0.47, 16.13)
Landings	all	177	0.37±0.13	1.76 (0.52, 8.98)
	black	139	0.39±0.13	1.75 (0.52, 7.25)
	checkered	38	0.32±0.10	1.79 (0.66, 8.98)
Pseudo landings	none	97	0.27±0.16	7.61 (0.82, 83.88)
Pseudo fly-by's	none	97	0.31±0.17	6.26 (0.92, 30.12)

Table 3.2: Flight speed under different conditions.

m s⁻¹ (see Table 2 for details).

In the trajectories that were classified as fly-bys (Fig. 3.2, n=1065), flies exhibited a turn away from the post when they were roughly 3 cm away from the object. This avoidance behavior is clearly manifest as a zone of exclusion around the post in the top down view of Fig. 3.2. Within the trajectories classified as landings (Fig. 3.2, n=177), I observed two general patterns: flies that flew straight towards the post, and those that performed a saccade at some point within 10 cm of the post. The flies that landed did not show any preference for flying near the top of the post (on average flies landed 6.3 ± 3.4 cm below top of the post; see Fig. 3.2 right column). This shows that under these experimental conditions flies are not more likely to land at the top of an object as previously suggested based on their preference for flying at the level of horizontal edges [141].

3.4.2 Saccade results

Because the most significant changes in direction are made during body saccades (see Methods, Fig. 3.3-3.4), I focused my initial analysis on these behaviors. Of particular interest is the last saccade flies perform prior to their nearest approach to the post, which for landing flies corresponds to the moment of landing. For all of the following analysis I focus exclusively on these saccades (the red overlays in Fig. 3.2). For my initial analysis I grouped landings and fly-bys together. In the absence of the post,

relatively few flies performed a saccade when the retinal size of the imaginary post exceeded 25° , whereas in the presence of the post I saw a distinct peak near 35° (Fig. 3.5C). This suggests that the saccades flies make when in the vicinity of the post are likely a response to the post, in particular when the retinal size exceeds 25° . Prior to making a saccade, the distributions of post-angle in the experiments done in the presence of the post and in the absence of the post are similar (Fig. 3.5A). The fact that the mean post-angle is near 0° in both cases is presumably a result of the geometry of the rectangular flight tunnel, which favors longitudinal flight. The post-angle after turning, however, is clearly different in the post vs. no-post conditions (Fig. 3.5B), suggesting that the presence of the post influences the turn angle of the flies. I further confirmed this by convolving the distribution of saccade turn-angles (Fig. 3.4A) with the post-angle prior to turning, yielding a unimodal distribution (data not shown) that is clearly different from the multimodal distribution shown in Fig. 3.5B. In order to examine the influence of the post on saccade behavior in more detail, I will focus only on those saccades made when the retinal size of the post exceeded 25° .

The distributions of post-angle after flies make a turn suggest that there are two independent behaviors governing these actions (Fig. 3.6). One group turns such that they are headed away from the post (described by Gaussian distributions with $\mu = \pm 140^\circ$, $\sigma = 50^\circ$), the other group turns such that they are headed towards the post (described by a Gaussian distribution with $\mu = 0^\circ$, $\sigma = 40^\circ$). These distributions are each separated by more than twice the common standard deviation, confirming that this is indeed a multi-modal distribution [7]. These two groups correspond surprisingly closely with the independently-classified landing and fly-by trajectories (Fig. 3.6C). This result justifies my subsequent analysis in which I treat the behavior of saccades made by landing and non-landing flies independently.

To determine in what way the post influences the saccades made by landing and non-landing flies, I examined the relationship between the post-angle prior to performing a saccade and the subsequent turn angle. If flies always made saccades directly towards the post I would expect them to make a turn roughly equal in magnitude to

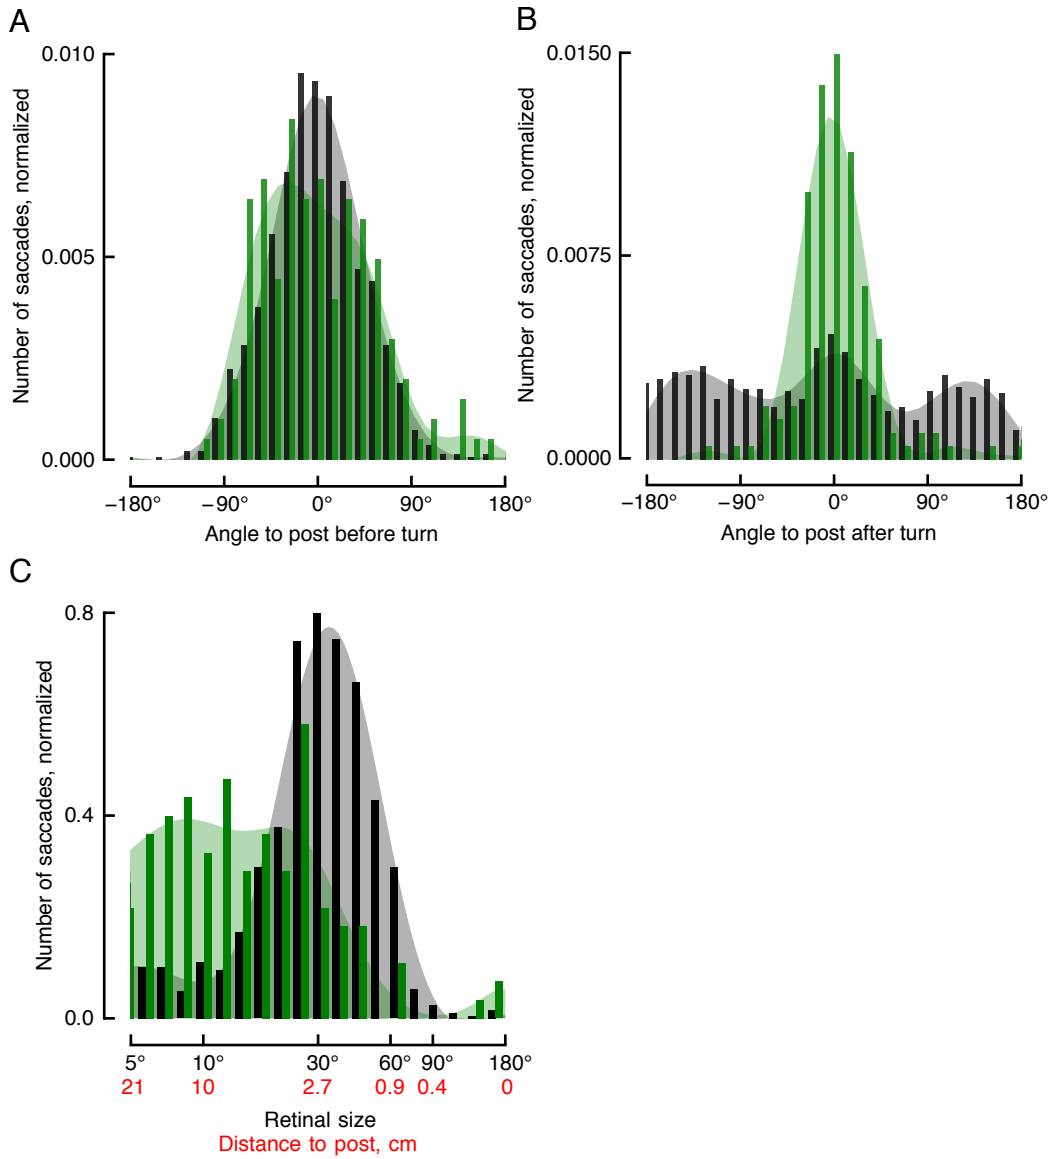


Figure 3.5: The post influences saccades performed in the vicinity of the post. (A) Prior to turning, flies show a qualitatively similar distribution of post-angle in the presence (black, $n=1012$) and absence (green, $n=148$) of the post. (B) After turning, flies show a qualitatively different distribution of post-angle in the presence (black) and absence (green) of the post. (C) Flies show a qualitatively different distribution for the retinal size at which they make their last saccade in the presence (black) and absence (green) of the post. In all panels the associated shading shows smoothed representations of the distributions calculated using 3rd order 0.3 Hz Butterworth filters.

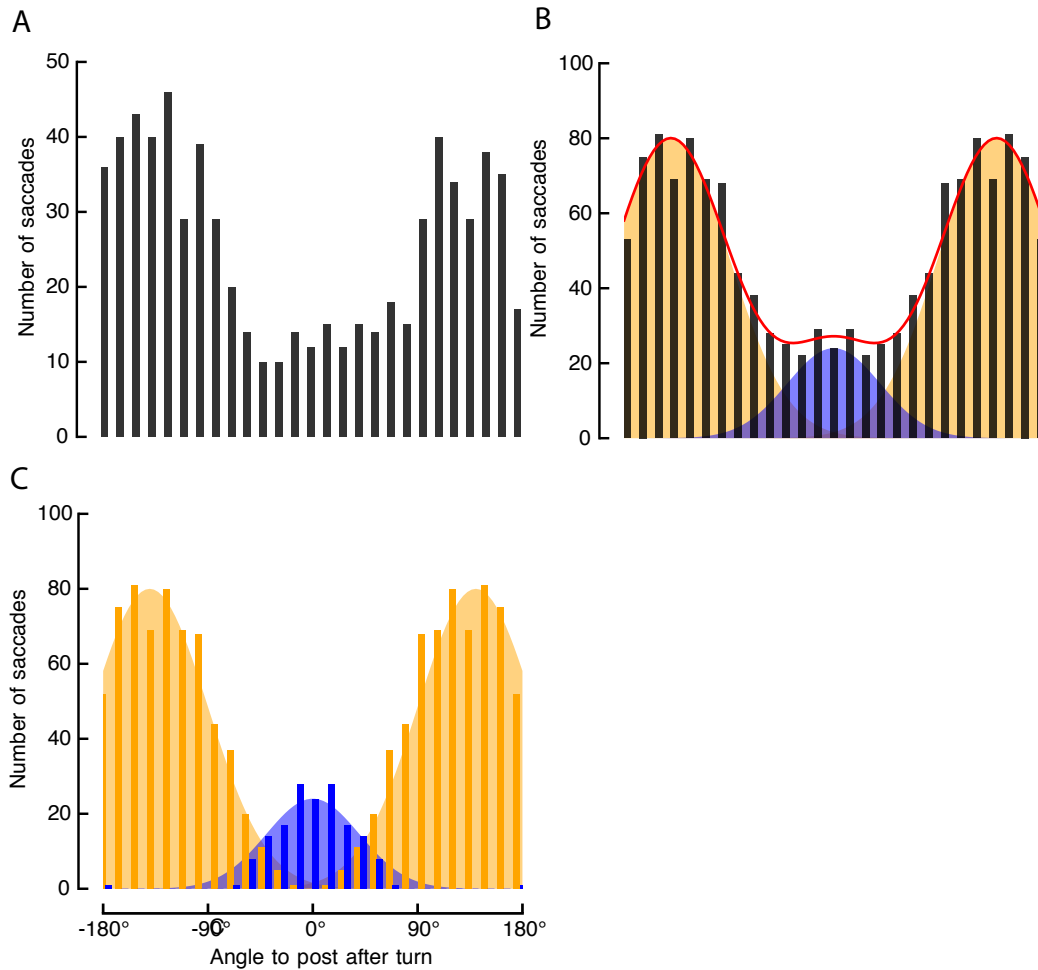


Figure 3.6: Saccades near the post are described by two independent behaviors, corresponding to landing and not landing. (A) Post-angle after turning for all saccades made in the presence of the post when the retinal size exceeded 25° . (B) The histogram shows the same data as in A, but after the distribution has been reflected about the central axis and added to the original. The symmetric distribution is bi-modal and is well approximated by a linear sum of a two Gaussian distributions shown in blue ($\mu=0^\circ$, $\sigma=40^\circ$) and orange ($\mu=\pm 140^\circ$, $\sigma=50^\circ$). (C) The two Gaussian distributions from B are repeated, and the histograms from B are segregated independently into landings (blue) and fly-bys (orange).

the post-angle, whereas turns directed away from the post would fall above or below that line. For the saccades prior to landings, these points lie along a line with slope close to one (Fig. 3.7A), suggesting that the flies were turning towards the post. This is not surprising, because in order to land on the post the flies' last saccade prior to landing must necessarily have been directed towards the post. Indeed, I see a similar regression in the pseudo-landing flies from the no-post control experiments (Fig. 3.18A). In the presence of the post, however, flies tended to make their last saccades much closer to the post than the control flies performing a 'pseudo-landing' in the absence of a post (Figs. 3.7C, 3.18C). Because I have already demonstrated that the saccades made when the retinal size of the post exceeded 25° were likely influenced by the post (Fig. 3.5), it is conceivable that the landing flies were in fact directing their saccades towards the post; I will revisit this hypothesis more rigorously later in my analysis.

In contrast, the non-landing flies tended to make aversive saccades when they were close to the post (Fig. 3.7B), at a mean retinal size of $33^\circ \pm 17^\circ$ (Fig. 3.7D). As expected, there are no obvious trends in the no-post control data for pseudo-non-landing flies (Fig. 3.18C). Although some of the non-landing flies exhibited a final saccade towards the post, these were all done at a much greater distance from the post. To examine the aversive saccades more closely, I again set a threshold on the retinal size of the post at 25° (which corresponds to a distance of approximately 3cm). These saccades – the last saccades non-landing flies made prior to their nearest approach to the post – are strongly correlated with post-angle, as evidenced by the two nearly parallel linear regressions for left and right turns (Fig. 3.8). The data are colored blue and red according to left and right turns, respectively. The small clusters of lightly shaded points (with a post-angle close to $\pm 90^\circ$) correspond to shallow turns towards the post when the flies were flying past the post. The separation between these clusters and the rest of the turns made in the same direction suggest that they are not governed by the same behavioral algorithm, and were thus excluded when I calculated the regressions. In fact, these clusters appear to be associated with the data corresponding to turns in the opposite direction. For example, the light red

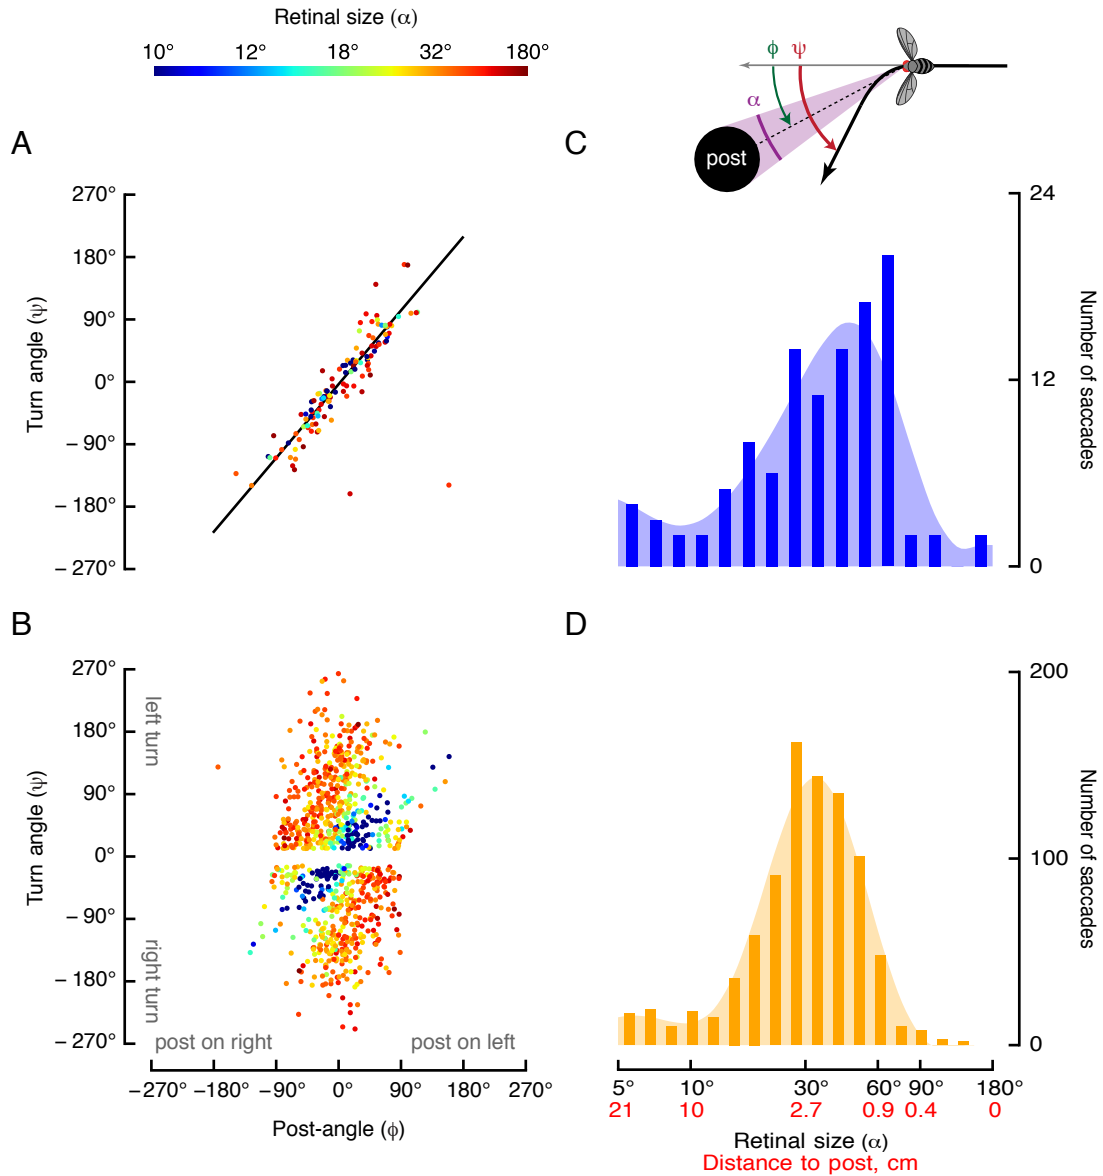


Figure 3.7: Comparison of saccade features for landings and fly-bys. Saccade turn angle (ψ) is plotted as a function post-angle (ϕ) for landings (A) and fly-bys (B) with the color of each point indicating the retinal size of the post at the time of the saccade. The black line in A shows a linear regression through the data, $\psi=1.18\phi-3.45^\circ$ ($R^2=0.88$, $n=126$). (C) The distribution of saccades in A (landings) as a function of retinal size and distance to post. (D) The distribution of saccades in B (fly-bys) as a function of retinal size and distance to post. In panels C and D, the associated shading shows smoothed representations of the distributions calculated using 3rd order 0.3 Hz Butterworth filters.

points (shallow left turns) may be an extension of the dark blue points (right turns).

3.4.3 Landing behavior

Of all the trajectories in which a fly came to within 3 cm of the post, 14% landed on the post ($n=1224$). In order for a fly to safely land it must decelerate to a safe flight speed and extend its legs in time for touchdown (although not necessarily in that order). To examine the deceleration behavior I estimated the point at which a fly started to decelerate prior to landing on the post (see Methods section). I initially restricted my analysis of deceleration to those trajectories where flies did not perform a saccade after initiating deceleration. When these points are plotted as a function of each fly's instantaneous speed and the log of the retinal size of the post, the data fall along a straight line (purple points, Fig. 3.9A-B). Thus, flies that are flying fast begin to slow down when they are farther away from the post compared to flies that are flying slowly. The data show that some flies began to decelerate as far away as 10-20 cm from the post, a retinal size of 5-10°, which is equivalent to one or two ommatidial acceptance angles. In contrast, I did not find a significant correlation between deceleration initiation and retinal size for fly-bys ($R^2=0.031$, $n=142$) nor pseudo-landings in the absence of the post ($R^2=0.153$, $n=82$). This suggests that the deceleration behavior seen in Fig. 3.9A is a unique hallmark of landing behavior.

If I now relax the requirement that flies not turn after initiating deceleration I see that these points also fall along a straight line (blue points, Fig. 3.9B). Furthermore, the regressions for turning and non-turning flies are very similar (Fig. 3.9B). Statistically the slopes of the regressions are not significantly different from one another ($p(x_1 : x_2) = 0.10$) and the differences in intercepts are only marginally significant ($p(x_2) = 0.06$), (see Methods). Because this deceleration pattern is uniquely indicative of landing behavior, the saccades performed after initiating deceleration were likely made with the intent to land, and thus intentionally directed towards the post.

To study leg extension in free flight I used a high speed video camera mounted directly above the post equipped with an automated motor driven focus-following




Figure 3.8: Flies make targeted body saccades away from the post and saccade amplitude is independent of turn direction. Saccade turn angle is plotted as a function of post-angle for saccades performed by non-landing flies. The data shown are the same as in Fig. 3.7C, but with the omission of data points where the retinal size is less than 25° (see Fig. 3.5 for justification). The two lines show linear regressions for left turns: $\psi = 2.68\phi + 155^\circ$ ($R^2 = 0.25$, $n = 342$) and right turns: $\psi = 2.47\phi - 149^\circ$ ($R^2 = 0.27$, $n = 270$). Assuming bilateral symmetry in the behavior, the regression becomes: $\psi = 2.58\phi - \text{sign}(\phi)152^\circ$ ($R^2 = 0.259$, $n = 565$). The light colored points represent shallow saccades towards the post when the post was close to 90° on either side of the fly. For reasons discussed in the text, these points were omitted when calculating the regressions. The histograms at the top show the distribution of left and right turns. The associated shading shows a Gaussian distribution ($\mu = \pm 25^\circ$, $\sigma = 40^\circ$). The histograms on the right show the distribution of turn angle for left and right turns. The associated shading shows a Gaussian distribution ($\mu = \pm 90^\circ$, $\sigma = 60^\circ$). The units are occurrences, and the black bars show the distributions associated with the light colored points, which were omitted when calculating the colored histograms.

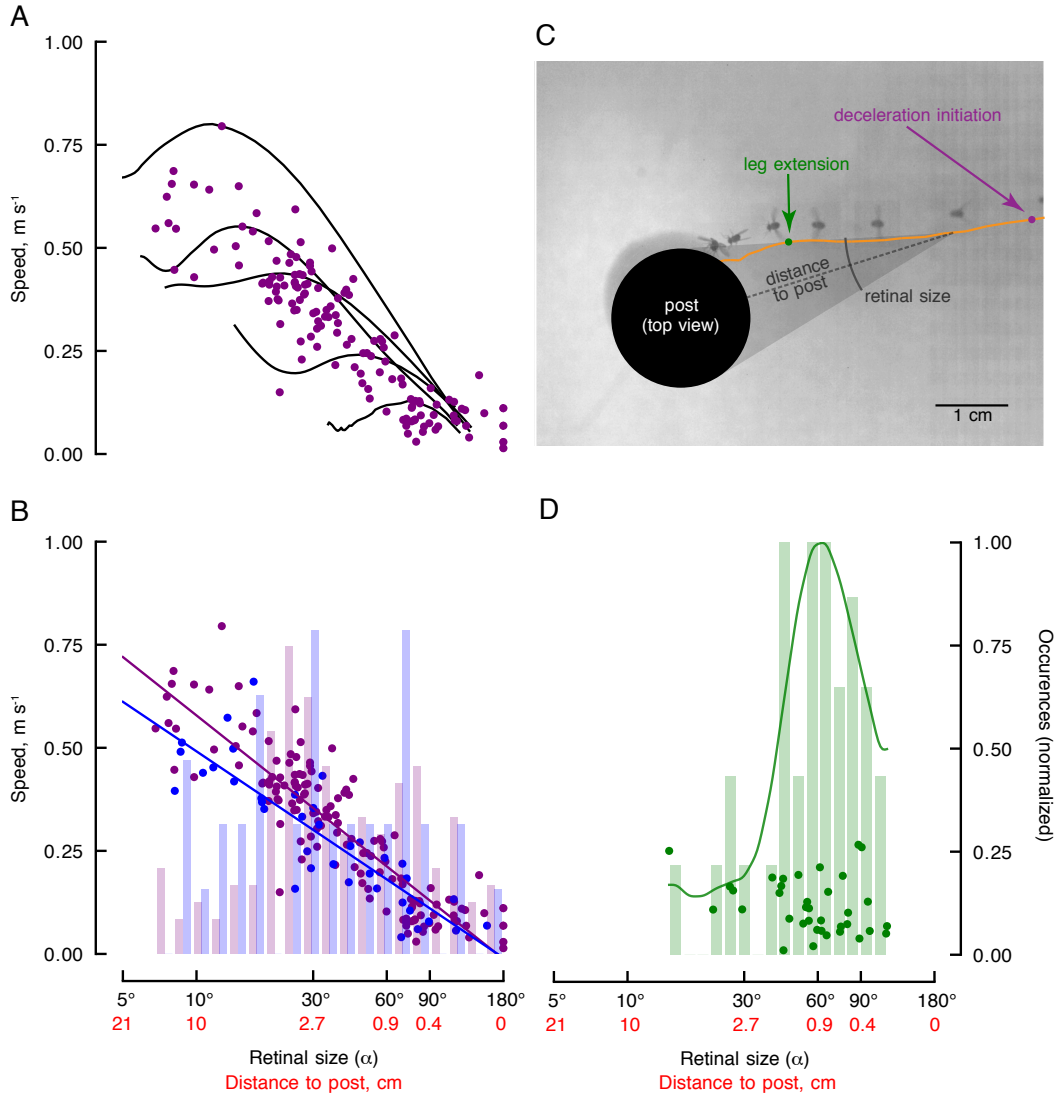


Figure 3.9: The motor programs for flight deceleration prior to landing and leg extension are distinct. (A) Points of deceleration initiation, plotted as horizontal speed vs. the retinal size of the post, for all landing trajectories for which the flies did not perform any saccades after initiating deceleration. The black traces show four sample trajectories across a range of initial flight speeds. (B) The purple points are repeated from A, and the purple line shows a linear regression through these data (speed = $-0.20\alpha + 0.22$, $R^2=0.8$, $n=140$). The blue points show the point of deceleration initiation for all landing trajectories in which flies performed a saccade after initiating deceleration, and the blue line shows a linear regression through these data (speed = $-0.17\alpha + 0.19$, $R^2=0.75$, $n=37$). The histograms show the distribution of the data across retinal size. (C) A sample landing trajectory captured on high speed video showing the point of deceleration initiation (purple) and leg extension (green). (D) The point of leg extension for all landing trajectories for which high speed video data were available. The light green histogram shows the distribution of leg extension events as a function of retinal size. The associated curve is a smoothed representation of the distribution calculated using a 3rd order 0.3 Hz Butterworth filter.




Figure 3.10: Image sequence from a high speed digital video recording of landing sequence.

system (see Methods section). Due to the limitations of high speed image capture (e.g., long downtimes between trigger events to allow for data transfer), I was only able to collect high speed imagery for a small subset of landings ($n=36$). As suggested in previous tethered flight studies, I observed leg extension to be a highly stereotyped, all-or-nothing behavior (Fig. 3.10). I determined the point of leg extension and initial touchdown for each video by visual inspection (see Methods section). In contrast to the start of deceleration, leg extension appears to be independent of flight speed, and unimodally distributed about retinal size ($61 \pm 22^\circ$, $n=36$). The final stage of deceleration took place after touchdown (which I defined as the first point of contact between a fly's legs and the surface of the post). After touchdown, the flies decelerated from an average of 7.1 cm/s ($\pm 3.2 \text{ cm/s}$, $n=30$) over the course of roughly half a body length (1 mm). This would correspond to a constant deceleration of 2.9 m/s^2 or roughly $0.3g$.

3.4.4 Crash landings

In observing the high speed sequences, it is clear that a sizable percentage (35.7%) of the landings were crashes in which the head or wings collided with the post before the fly extended its legs. In the cases of wing crashes (Fig. 3.11C), the fly's wing struck the surface of the post and in nearly all cases the fly rotated towards the post, extended its legs, and ended up landing successfully. Head-on crashes (Fig. 3.11D) also resulted in leg extension in nearly all cases, but the end result appeared to depend on flight speed. Slower flies tended to recover quickly and land, whereas the faster ones

bounced off of the surface of the post, tumbled through the air, and either managed to turn and land on the post or fell to the floor of the flight arena. I performed the same analysis for the crash landings as the normal landings to qualitatively compare the two (Fig. 3.8A-B). In 76% of the crash landings, the flies began to decelerate at the appropriate retinal size for their flight speed, based on the regression of deceleration initiation for successful landings from Fig. 3.9B. They did not, however, decelerate fast enough to be at an appropriate flight speed at touchdown, and they did not extend their legs in time. The flight speed of the trajectories that result in crashes is not any greater or lower than the typical flight speeds of successful landings.

3.4.5 Post texture

As a part of my experimental design, I tested two different types of post surfaces: solid black and checkered (5 mm squares, 5° retinal size at a distance of 5.7 cm). For the primary results of this paper (which are summarized in Fig. 3.16), I did not find any substantive differences in the behavior of flies landing on the two different post types both by visual inspection and statistical analysis, with two exceptions (Fig. 3.19-3.20). The largest difference in behavior is in the distance at which non-landing flies saccade away from the post. On average flies turn away from the checkered post at a retinal size of 10° less (thus earlier) compared to the black post ($p < 0.001$, Fig. 3.20B). The second substantial difference is in the intercept of the linear regression associated with the turn angle of aversive saccades (Fig. 3.19A-B). The regressions suggest that in the presence of the checkered post flies make a smaller turn (by approximately 40°) compared to those made in the presence of the black post ($p(x_2) = 0.03, 0.001$, for left and right turns, respectively).

Although neither deceleration nor leg extension are strongly correlated with post contrast (Fig. 3.19C-D), the percentage of trajectories that ended in the fly landing in the presence of the checkered post (7.4%, $n=662$) was much lower than for the solid black post (29%, $n=637$). Of these landings (for which I had high speed data) a much higher percentage were crash landings in the presence of the black post (38.5%

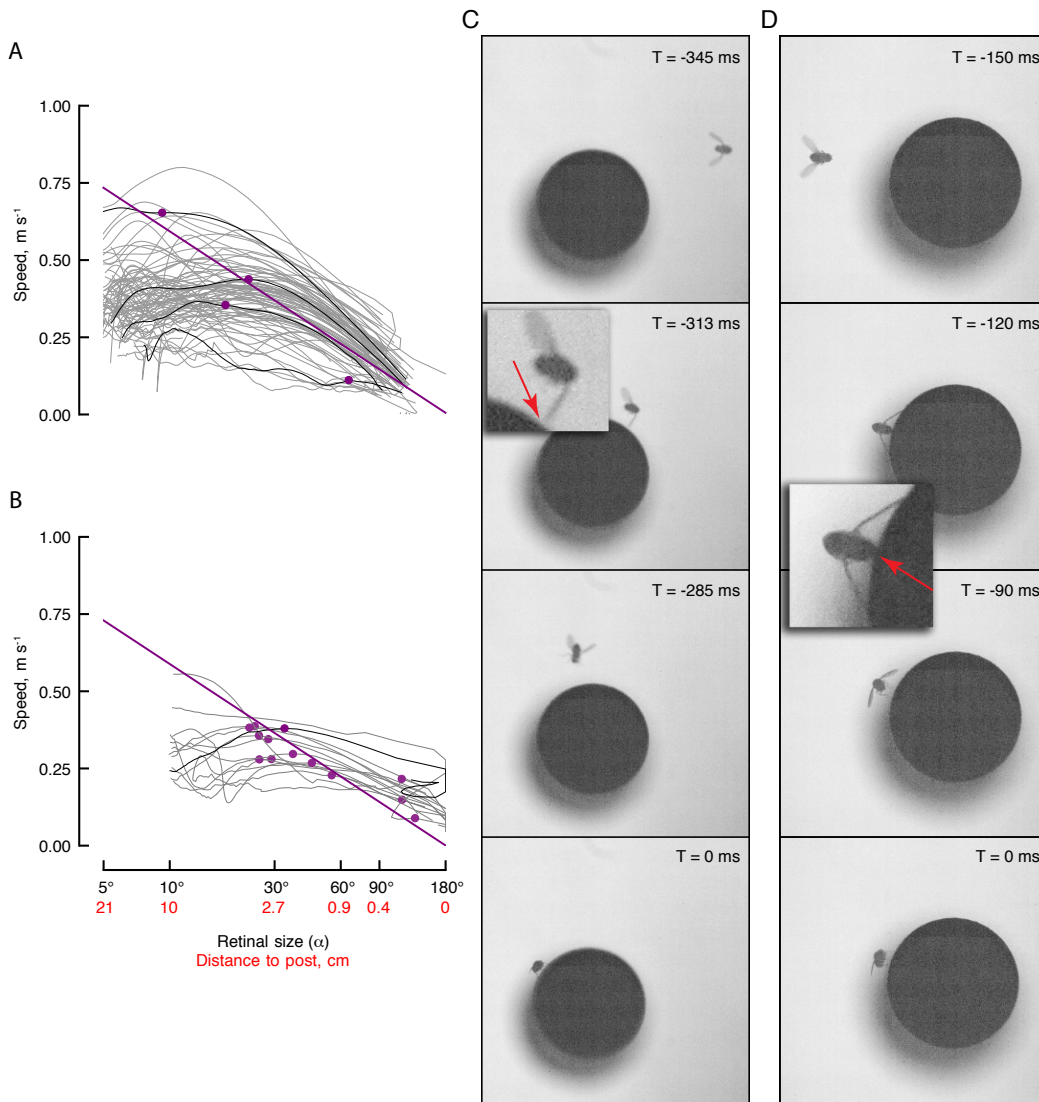


Figure 3.11: Crash landings correspond to sequences in which flies do not decelerate fast enough. (A) Examples of successful landing trajectories plotted as in Fig. 3.9A. (B) Speed vs. retinal size plots for crash landings. The purple line shows the linear regression from the landing data in Fig. 3.9A. (C-D) Image sequence from high speed digital video recordings for (C) a head on crash and (D) a wing impact crash. The point of impact is highlighted by the red arrow in the enlarged inset. Since successfully landing flies touchdown with their legs first (Fig. 3.9E) the distance at which their speed falls to zero is greater than zero. Crashing flies, however, collide with the post (often head first) with velocities much greater than zero. This explains why the trajectories of successfully landing flies terminate earlier than those of the crashing flies.

of landings on the black post, $n=29$, were crashes) compared to the checkered post (18.5% of landings were crashes, $n=27$).

3.5 Discussion

In this study, I examined landing behavior of *Drosophila* by analyzing their flight behavior in the presence of a conspicuous vertical post. Landings appear to consist of a temporal sequence of three distinct behaviors. First, flies actively turn towards the visual target via saccades that are directed toward the center of the post (Fig. 3.7A). The flies next begin to decelerate at a point that is a function of both their speed and the retinal size of the post (Fig. 3.9A-B), although in some cases deceleration is initiated prior to turning towards to the post (Fig. 3.9B). Finally, flies extend their legs just prior to touchdown when the post subtends an absolute size of approximately $61 \pm 22^\circ$ on their retina (Fig. 3.9D). In contrast, flies that do not land make a targeted saccade away from the post when it subtends a retinal size of approximately 33° (Figs. 3.7-3.8).

3.5.1 Attractive and aversive saccades

The tendency for flies to fly towards prominent visual features, such as high contrast edges and posts, was first documented by Kennedy for tethered mosquitoes, and has been studied thoroughly in tethered [66, 67] and free-flying *Drosophila* [100]. Over time, this so-called fixation behavior could be achieved via either smooth tracking, according to the model proposed by Reichardt and Poggio [120], or a series of directed saccades (or both, as proposed by Land [91] and observed in humans). Achieving target tracking through saccadic maneuvers has been proposed for hoverflies when visual targets lie outside of their optical fovea [38] as well as for houseflies chasing conspecifics [14, 13]. To my knowledge, I believe my data provide the first evidence that *Drosophila* make targeted body saccades towards and away from visual features, although many tethered flight studies have suggested that they saccade in the direction of a visual target [72]. Although subtle, this is an important distinction in that

in order to perform a targeted saccade the magnitude of the turn – and not just the direction – must depend on the retinal position of the object prior to the turn.

Whereas the saccades directed towards the post do not seem to be triggered by specific target size, the distribution of aversive saccades suggests a trigger threshold of roughly 33° (Fig. 3.7D). This corresponds quite closely with data from a previous study in the same apparatus, using a different-sized post (1.27 cm diameter, 30 cm tall) [100]. Another study using magnetically-tethered *Drosophila* (free to rotate about their yaw axis) reported that flies exhibit aversive saccades in response to expanding squares when the retinal size reaches approximately 60° , with an estimated neural processing delay of about 50 ms [9]. Electrophysiological recordings in locusts [56, 57] as well as hawkmoths [163] also supports a triggering mechanism for aversive maneuvers that depends upon an absolute angular threshold. This retinal size threshold model, which is independent of contrast, also would explain why I did not observe a significant difference in the aversive saccades in experiments using the black and checkered posts.

The simplest implementation of a neural mechanism for achieving targeted body saccades would make use of the assumption that the fly’s body and head are oriented in the direction that it is flying. By further analyzing the high speed video that I collected during my experiment I found that flies do not, in fact, always orient their bodies in the direction that they are flying, even during straight flight segments, resulting in a slip angle that ranges between 0° and 50° (Fig. 3.12). The resolution of my video was not sufficient to resolve head motion, so it is possible that the flies partially corrected for this slip by adjusting the position of their head. Alternatively, the slip angle could be partly responsible for the variance in the observed turn angles I measured (Figs. 3.7A, 3.8).

3.5.2 Landings

The combination of free flight 3D tracking in a large volume and with a focus-following optics on a high speed camera, enabled us to observe the complete landing sequence




Figure 3.12: Slip angle during flight has a wide range. Slip angle is defined as the angle between a fly's body orientation and the direction it is flying. The data shown are drawn from the straight portions (excluding saccades and portions of flight very close to the post) of 71 trajectories for which high speed video sequences were available. Occurrences are measured as individual frames from the videos, which were filmed at 5,000 frames per second ($N=963$).

of freely flying *Drosophila* from approach and deceleration to leg extension and the final moments of touchdown. Despite the wide variety of approaches (Fig. 3.2), many of which show intermittent targeted body saccades towards the post, flies exhibit a clear pattern in when they initiate deceleration prior to landing (Fig. 3.9A-B). This pattern depends on a combination of their speed and distance to the post. However, it is very unlikely that the fly can accurately measure either its ground speed or its distance to the post directly. Given the constraints of the visual and mechanosensory systems available to the fly, it is more likely that they use a measure of the retinal size of the post (which is correlated with distance) and the rate of expansion (which is correlated with ground speed). For this reason, I re-plotted the linear fit from Fig. 3.9B in terms of the retinal size of the post, α (as defined by (3.1)), and its retinal expansion velocity, Ω (Fig. 3.11):

$$\Omega = \frac{d\alpha}{dt} = \frac{-2s (r/d^2)}{\sqrt{1 - (r/d^2)}}, \quad (3.3)$$

where s is the flight speed in the x-y plane, r is the radius of the post, d is the distance from the fly to the center of the post. This retinal size dependent expansion threshold model (RSDET) can be visualized by considering the approach trajectories of flies flying at constant velocity toward the post, plotted in the α - Ω plane (Fig. 3.13).

A fast flying fly crosses this threshold, and thus starts to decelerate, at a greater distance from the post than a slow flying fly. From the perspective of the fly, a relatively low rate of expansion is sufficient to trigger deceleration when an object size is small, whereas a higher rate of expansion is required when the retinal size of an object is large. Note that this model is independent of the actual physical size of the object, but by combining measures of α and Ω , deceleration would not, for example, be triggered by a physically large, but distant object (because Ω would be too small). This principle is nearly identical to the rate of relative expansion velocity model (RREV), previously proposed by Wagner for houseflies (Wagner, 1982). The RREV model states that insects should begin decelerating when the ratio of retinal expansion velocity and the retinal size of an object reaches a critical threshold:

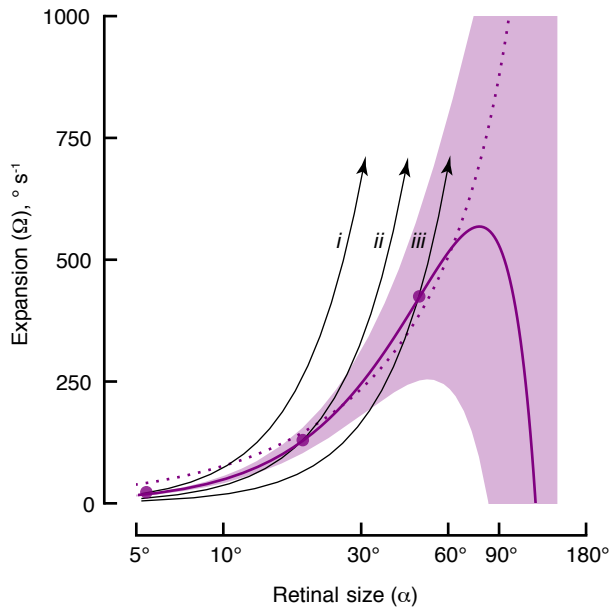


Figure 3.13: Retinal size-dependent expansion threshold compared to the time-to-contact model. The solid purple line and shading shows the regression and standard deviation from Fig. 3.7C, calculated as expansion vs. retinal size using (3.3). The model predicts that flies initiate deceleration once their trajectory passes above the purple line. Three constant velocity trajectories are shown along with the associated predicted points of deceleration (0.6, 0.4, 0.2 m/s, corresponding to i, ii, iii, respectively). The dotted purple line shows an example threshold curve for the time-to-contact model that closely matches the RSDET model up to a retinal size of about 80°. This curve was calculated using 3.7 and a time-to-contact value of 120 ms.

$$RREV = \frac{1}{\alpha}\Omega, \quad (3.4)$$

Wagner further simplified this definition as:

$$RREV' \approx -2s\frac{1}{d}, \quad (3.5)$$

where d is the distance to the center of the post, and s is the flight speed. Note, however, that this simplified definition is derived by making some approximations that are only valid when the distance to the object is much greater than the object's radius. However, in many of the trajectories I recorded, flies initiated deceleration only after they were quite close to the post, at distances of the same order as the radius. Thus, although the RSDet and RREV models are fundamentally identical (both are a measure of the expansion relative to retinal size) and are probably experimentally indistinguishable, I will keep the terminology distinct for the sake of mathematical clarity.

Using the above approximation, the RREV model is formally the inverse of the time-to-contact (τ), that is, the time before the fly will collide with an object assuming that it is indeed on a collision course and flying at constant velocity:

$$RREV' = \tau^{-1}, \quad (3.6)$$

For reference, without any approximations the time-to-contact can be calculated as:

$$\tau = \frac{-2 \sin(\alpha/2)}{\Omega \sqrt{1 - \sin^2(\alpha/2)}}. \quad (3.7)$$

Because of this relationship, the RREV model has been referred to as the time-to-contact model [56]. The time-to-contact model, using a time-to-contact threshold value of about 120 ms, can be made to fit my RSDet model quite closely up to retinal sizes of 80° (Fig. 3.11). Note that once flies are close enough to the post for it to reach a retinal size of 80° they have typically already initiated deceleration (Fig. 3.9B), so comparing performance of either model beyond this point is not particularly

meaningful.

Although the time-to-contact model may be an intuitive way of thinking about the time course of a trajectory, I prefer not to use this terminology because neither of two critical assumptions, constant velocity and a direct collision course, is valid for the flies' behavior. Furthermore, there is no need for the fly to actually calculate a time to contact when a simple retinal size dependent expansion threshold calculation is sufficient. The three models: RSDET, RREV, and time-to-contact, are all fundamentally the same with regards to the sensory inputs they use and the behaviors they predict.

Once the fly has started to decelerate it needs to control its speed such that it will reach a safe touchdown velocity prior to landing. This is not a trivial calculation, as there are no experimentally supported hypotheses for how a fly might accurately measure the distance between itself and a physical object. Previous work on honeybees landing on flat horizontal surfaces suggests that they use a simple controller that holds the rate of retinal expansion constant, elegantly guaranteeing that they reach a flight speed of nearly zero just before touchdown [134]. However, my data for fruit flies landing on a vertical post are not consistent with this model (e.g., when speed is plotted with respect to distance in a linear scale, I do not see a linear relationship between these two parameters). This should not be surprising, because the two experimental paradigms are fundamentally different (both in geometry and animals).

As the retinal size of an object and its derivatives are the only optical cues directly available to the fly, the most biologically plausible controller would use some combination of retinal size, rate of expansion, and the rate of change of expansion. In Fig. 3.14, I show that both rate of retinal expansion (Ω) as well as the rate of change of the rate of retinal expansion ($\dot{\Omega}$) could, in principle, be used to safely decelerate prior to impact without requiring a measurement of ground speed, object distance, or a priori knowledge of how large the object is. In both cases, Ω is first put through a threshold function (3.8), where the threshold (ξ) is defined by the solid purple curve in Fig. 3.11:

$$\Omega_{\xi,t} = \max(\Omega_t - \xi, 0). \quad (3.8)$$

The discrete time control models for Ω (Fig. 3.14B) and $\dot{\Omega}$ (Fig. 3.14C) are defined by (3.9, 3.10), respectively:

$$s_{t+1} = s_t - K_{\Omega}\Omega_{\xi,t}, \quad (3.9)$$

$$s_{t+1} = s_t - K_{\dot{\Omega}}\dot{\Omega}_{\xi,t}, \quad (3.10)$$

where s is the flight speed in the x-y plane, K_{Ω} is the gain for Ω , and $K_{\dot{\Omega}}$ is the gain for $\dot{\Omega}$. In order for (3.10) to be stable in the final moments of landing, I need to add the requirement that $\dot{\Omega}_{\xi,t}$ be greater than or equal to zero. These models, which match the general trend of the observed data, are provided only as proofs of concept that such controllers could be employed by the fly given the limited sensory information available to it. Future experiments will be needed to test whether or not flies actually use such algorithms.

The next stage of landing – leg extension – has been subject to more prior work than the rest of the landing sequence due to the ease of studying it in tethered flight preparations. Several hypotheses for the underlying neural mechanism that trigger leg extension have been proposed. One idea is that flies use the same sort of retinal size threshold trigger that is used for the expansion avoidance response. This model is supported by several studies on tethered flying *Drosophila*, which extend their legs at a fixed retinal size of about 50° with a neural processing delay of approximately 50 ms [165, 144]. Leg extension behavior in houseflies and blowflies, however, is inconsistent with the retinal size threshold, because they will extend their legs in response to sinusoidal gratings expanding within a grating mask of fixed size [17, 161]. The authors of these studies proposed a spatio-temporal integration model, in which flies integrate motion energy until a certain threshold is reached [17]. Note that because multiple pathways could trigger leg extension, these two hypotheses are not mutually

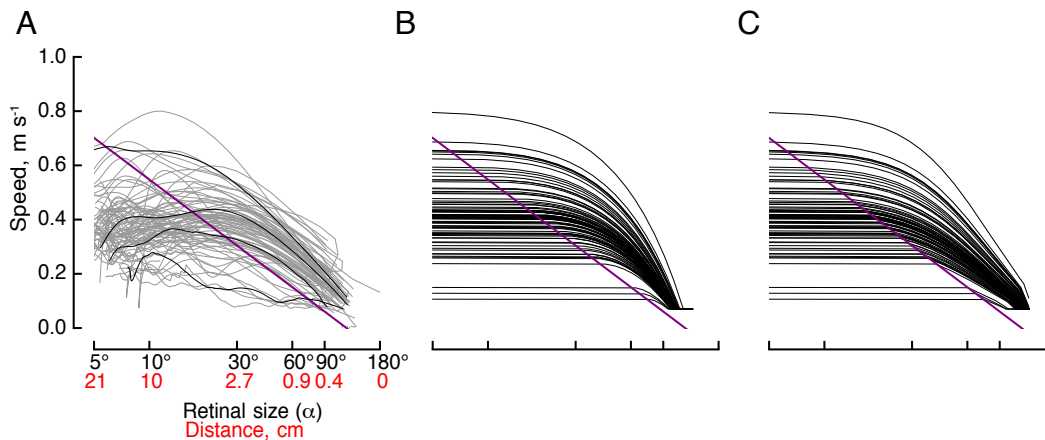


Figure 3.14: Potential control models for deceleration. The purple trace shows the observed point of deceleration, repeated from Fig. 3.9A. (A) Observed speed vs. retinal size traces, duplicated from Fig. 3.11A. Four arbitrary trajectories at different initial velocities are highlighted for clarity. Using the same initial conditions from the observed data in A, I simulated the time course of trajectories using two kinds of deceleration controllers: (B) $s_{k+1} = s_k - 3500(\max(\Omega_k - \tau, 0))$, (C) $s_{k+1} = s_k - 170000(\max(d(\Omega_k - \tau)/dt, 0))$, where τ represents the expansion threshold drawn from Fig. 3.13. The $\max()$ function simply ensures that the threshold subtraction returns zero when Ω_i is less than τ .

exclusive. Alternatively, a fly could react to a sudden change in luminosity in its visual field, referred to as the temporal contrast model. This model has been proposed for triggering escape responses in stationary *Drosophila* [76, 148] as well as leg extension responses in blowflies [65], but does not elicit turning or leg extension in *Drosophila* [144]. Finally, flies could use the “time-to-contact” model, or the more general RSDET model that I proposed above as the trigger for deceleration.

Because I do not see any correlation between leg extension and flight speed, and thus rate of expansion (Fig. 3.9D), it is unlikely that either the spatio-temporal integration model or the RSDET model are responsible for triggering leg extension in free flight. The most parsimonious hypothesis for explaining the leg extension behavior I observed is a retinal size threshold model. The value I measured for this threshold ($61 \pm 22^\circ$, $n=36$), is remarkably close to a previously published result from tethered flies of about 50° [144].

However, it is unlikely that the leg extension trigger is as simple as a retinal size

threshold, as many of the non-landing flies experience retinal sizes of approximately 60° , yet very few extended their legs (we did, however, observe some examples of non-landing flies extending their legs as if in preparation for landing). The additional requirement that the extent of the object needs to be centered on the field of view might prevent unnecessary leg extension during fly-bys, as they will typically have turned away from the post by the time it subtends 60° (Fig. 3.7D). This idea is consistent with psychophysical tuning curves from tethered flies which show that leg extension is triggered by expansion centered on the fly's retina, but not by lateral expansion [144].

The final stage of landing – touchdown – appears to be a stereotyped open-loop process, likely done without visual feedback. I did not find any indication that flies orient their bodies (and thus legs) in relation to the surface of the post, a behavior that has been observed in honeybees [49]. Because honeybees weigh on the order of 200 times more than a fruit fly, and they often need to make more challenging landings, such as on flowers swaying in the breeze, it is not surprising that their landing sequence is more complex. The time between leg extension and touchdown is less than 50 ms for about 1/3 of the landings I observed (Fig. 3.15). This is such a short time frame (equal to the visual processing delay found in other experiments described earlier) that it is unlikely that this aspect of landing is under tight visual control. In the majority of touchdowns I observed that one of the two front legs would touch the post before the other. Because the legs were spread out, this asymmetric touchdown creates a moment arm that automatically orients the fly to the normal of the post surface.

3.5.3 To land, or not to land?

What factors determine a fly's choice to land or not? Is the decision dictated solely by sensory experience, or does the animal's internal state play a role? In other words, if two different flies started with the same exact initial sensory conditions, would they be likely to follow similar trajectories? A simple explanation for whether a fly




Figure 3.15: Time elapsed between leg extension and touchdown is less than 100 ms for the majority of landings (n=36).

lands or not would be that certain visual cues ultimately lead to landing while others lead to aversive maneuvers. I term this deterministic scheme the sensory attractor hypothesis. Based on my analysis, the most relevant visual cues for landing and obstacle avoidance are retinal size, flight speed (or, alternatively, rate of expansion), and post-angle. Thus, if the decision to land were governed by a sensory attractor, I would expect it to be possible to construct a predictor based on some combination of these sensory stimuli. Despite my efforts, the best predictor I found (based on post-angle and retinal size) exhibited a 57% accuracy at predicting whether a fly would land or not at a distance of 6 cm. Although it is possible that a better predictor exists, my analysis suggests that the decision to land or not is either largely random, determined by sensory parameters other than retinal size, rate of expansion, and post-angle, or influenced by the internal state of the fly. It is certainly feasible that unaccounted for sensory cues could play a role in this behavioral decision. For example, it has already been shown that attractive odors have a significant impact on post approach and landing behavior [54, 34, 26, 135]. However, because my experiments were performed

in the absence of such cues, a more plausible explanation is that the decisions I observed were influenced by some combination of internal states such as hunger, exhaustion, the need to lay eggs, or defecate.

Assuming that the decision is not random, and that the internal state model is correct, at what point in their trajectory do flies make a decision to land on or turn away from an object? Because flies begin to decelerate when the post subtends as little as 5 to 10° (equivalent to subtending 1 to 2 ommatidia), they probably can detect the presence of an object at those distances. Why then, do non-landing flies wait until the post subtends about 33° before they turn away from it? Perhaps they are using the visual cue of the post as a means to navigate, and thus until there is a more enticing visual feature to fly towards there is no need to turn away until a collision becomes imminent. Alternatively, this observation could be an artifact of the rectangular flight arena I used. If a fly were to turn away from the post when it was only 5-10° it would not be able to travel very far before having to turn again.

3.5.4 Crashes

Our focus-following system, which enabled us to capture the final moments of landing, revealed that 35.7% of landings were crashes (these crashes were excluded from the data shown in Fig. 3.9). Due to my tracking system being unable to maintain fly identity over the course of my experimental runs, I cannot rule out there are simply some individuals that are bad flyers, and I plan to address this question in the future. For the analysis in this paper I assumed that all flies behave similarly.

The result that the majority of crashing flies initiate deceleration at the appropriate time according to the RSDDET model is of particular interest. It is not that the flies failed to decelerate; the problem appears to be that they did not decelerate fast enough. According to my proposed models for velocity control prior to landing, deceleration is controlled by some combination of the rate of expansion and the rate of change of expansion (Fig. 3.12). If the flies are able to measure expansion sufficiently well to trigger deceleration, why were they unable to use that measure to control

their deceleration? One hint comes from the observation that flies crashed into the solid post more frequently than they crashed into the checkered post. It would make sense for the measure of expansion from the checkered post to be more accurate, as there are more contrast edges to use for calculating motion. However, perhaps a more important error in the crash landings than insufficient deceleration is the failure of these flies to extend their legs. If they were able to detect the presence and size of the post in order to initiate deceleration, why did they not extend their legs in time? One possible explanation for this is that leg extension may only be triggered if the flight speed is below some threshold. This hypothesis is supported by Fig. 3.9D and Fig. 3.11B, which suggest a threshold of approximately 0.2 m/s. I plan to address these hypotheses in future experiments.

3.5.5 Post texture

Although subtle, the differences in behavior in the presence of the checkered and solid black posts deserve comment. The most apparent difference in behavior correlated with post texture was the fact that far fewer flies landed on the checkered post compared to the solid black post (Tables 3.1, 3.2). More experiments will be necessary to get a full understanding of why this is the case, but my current hypothesis is that under the experimental conditions, at distances beyond which the flies could resolve the individual checkers on the checkered post (approximately 5-10 cm), the checkered post may not have been as conspicuous as the black post. This would have caused fewer flies to be attracted to, and thus land on, the checkered post. Furthermore, since flies show such a strong preference for navigating towards conspicuous vertical patterns, it would not be too surprising if they adjust their behavior in the presence of such visual features by increasing flight speed. Together with the reduced saliency of the checkered post hypothesis, this would explain the slight difference in mean flight velocities I observed in landing flies in the presence of the black post ($0.39 \pm 0.13 \text{ m s}^{-1}$) compared to the checkered post ($0.32 \pm 0.13 \text{ m s}^{-1}$).

I also found that non-landing flies turn, on average, at a retinal size of 10° earlier

in the presence of the checkered post ($p(x_2) = 0.03, 0.001$, for left and right turns, respectively; see Fig. 3.19-3.20 for complete statistical details). This suggests that the aversive saccade maneuvers are, to some degree, a function of object contrast. I found the same trend for leg extension, although the statistics are less convincing ($p=0.08$). These observations are in agreement with the similarly subtle results found in tethered flies [9], and favors a model in which the internal contrast of the object, and not just the position and expansion rate of its edges does influence the underlying visual processing. This may, at first, appear to be at odds with the contrast independent retinal size threshold that I and others propose as the trigger for both aversive saccades and leg extension, but there could easily be two (or more) pathways for triggering these behaviors. For example, aversive saccade maneuvers can be elicited by wide field expansion [145]. Alternatively, the circuitry that somehow tracks the expansion of edges might nevertheless be sensitive to internal motion. Indeed, a system using exclusively a retinal size threshold would have functional limitations, as it would cause a fly to either extend its legs or initiate an aversive saccade when a physically large, but very distant object reached the appropriate retinal size. By incorporating some sort of expansion threshold in addition to the retinal size threshold, these unnecessary triggers could be prevented. This addition to the model could explain the slight dependence on contrast I observed.

3.5.6 Summary

In conclusion, my findings suggest that flies have an internal state that guides their decision of whether to land or not. In both cases flies actively turn towards the post when the post is far away. Flies that do not land make targeted body saccades away from the post once it reaches a critical retinal size threshold of $33^\circ \pm 17^\circ$. Landing flies, however, continue to make targeted body saccades towards the post and eventually start to decelerate when a retinal size-dependent expansion threshold is reached, followed by leg extension, which is triggered by a fixed retinal size threshold of $61 \pm 22^\circ$, and finally, touchdown (Fig. 3.16).

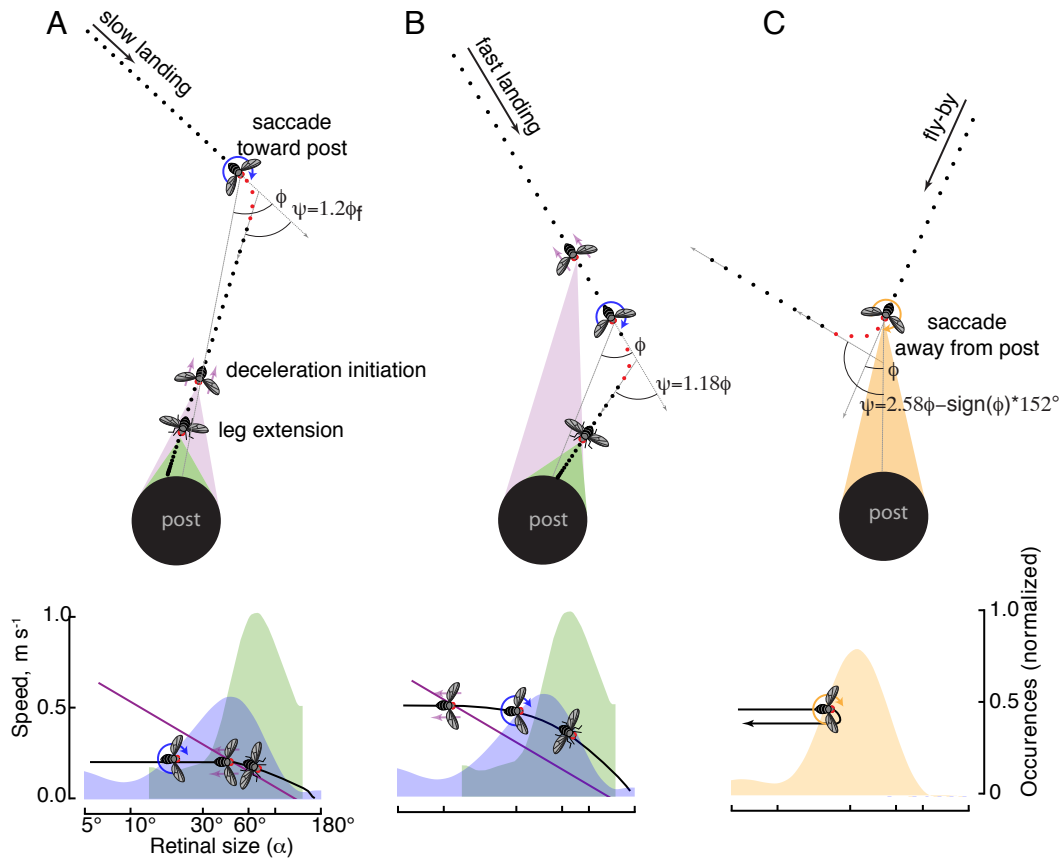


Figure 3.16: Landing and fly-by behaviors are distinct, and controlled by retinal position, retinal size, and retinal expansion velocity cues. The colors are consistent with the related figures earlier in the paper. The purple line indicates the point at which flies initiate deceleration (Fig. 3.9B), the green distribution shows when flies extend their legs (Fig. 3.9D), the blue distribution shows when flies make turns towards toward the post (Fig. 3.7C), and the orange distribution shows when flies make evasive saccades (Fig. 3.7D).

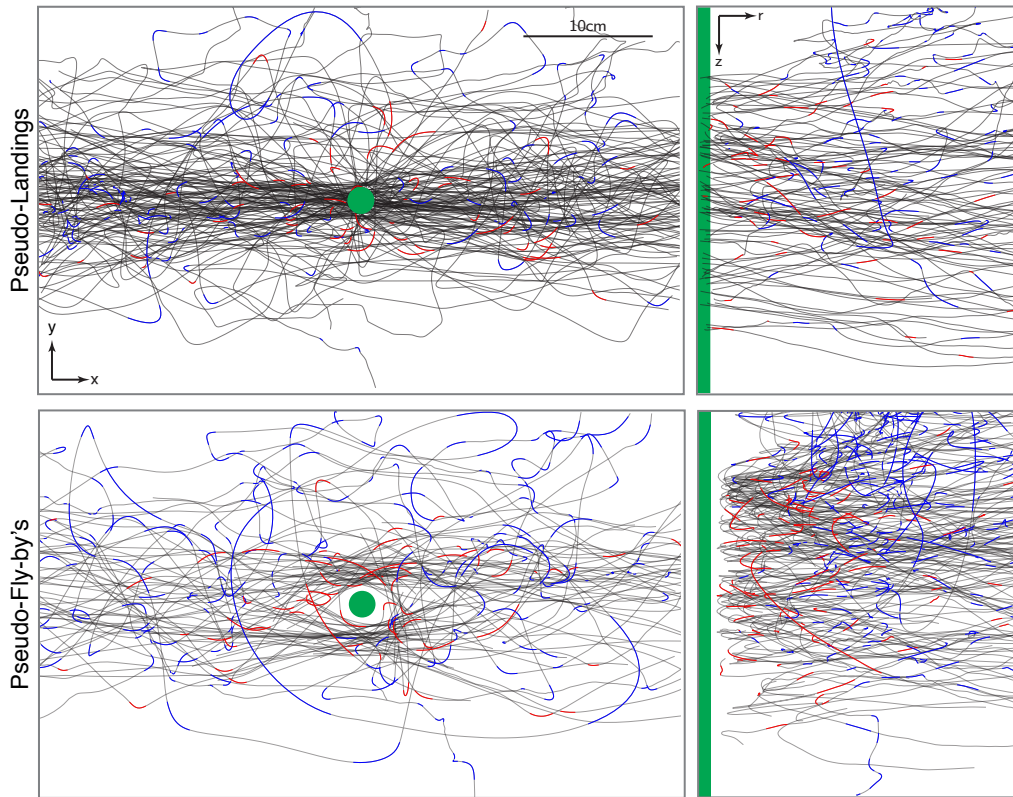


Figure 3.17: Flight trajectories in the absence of the post, plotted as in Fig. 3.2 with a virtual post. Trajectories were classified as if there were a post into pseudo-landings ($n=97$) and fly-bys ($n=97$). In order to increase the number of trajectories in these control datasets I extended the imaginary post to the top of the arena.

3.6 Supplemental figures

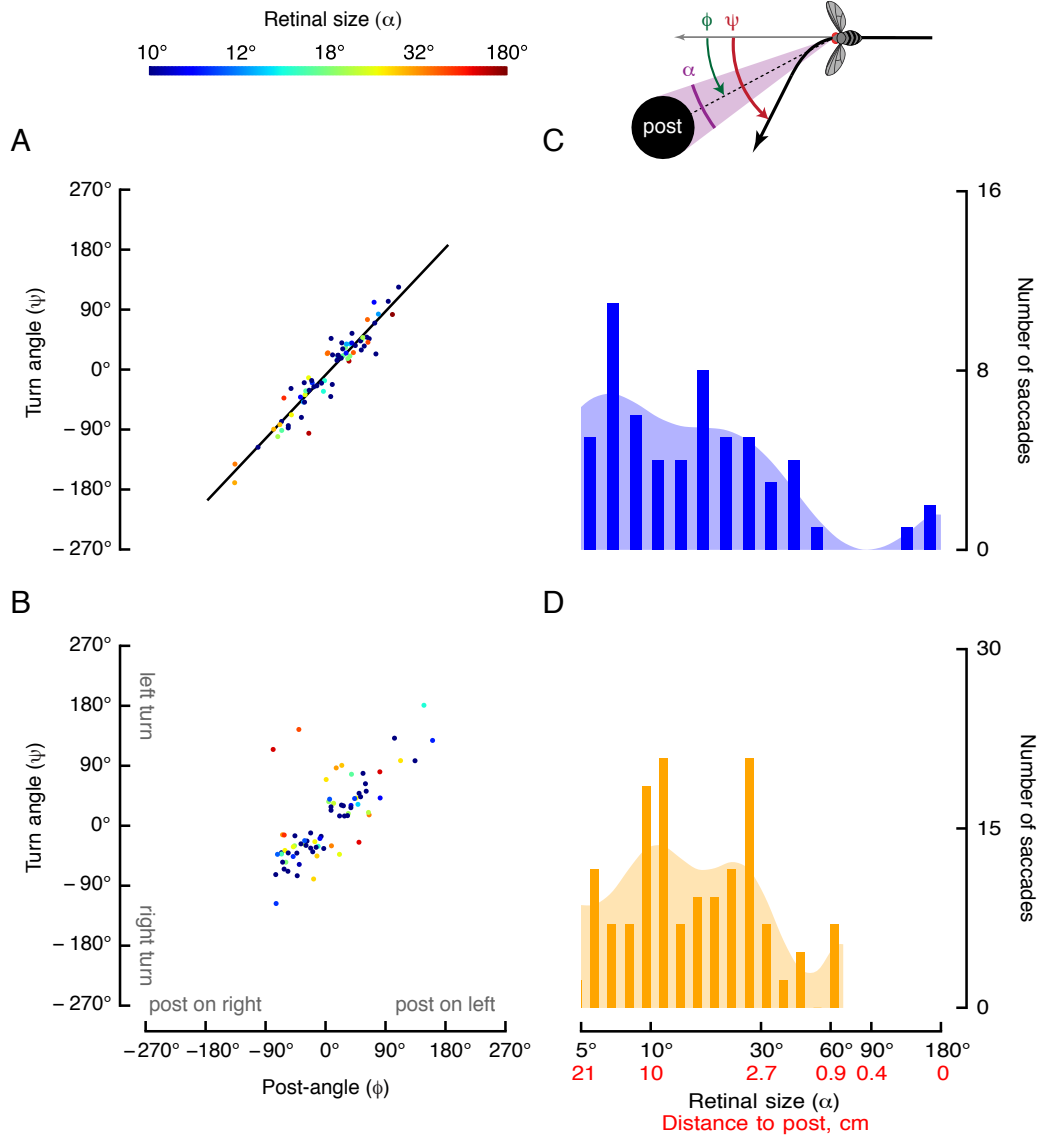


Figure 3.18: Comparison of saccade features for pseudo landings and fly-bys, with respect to a virtual post. Saccade turn angle (ψ) is plotted as a function post-angle (ϕ) for landings (A) and fly-bys (B) with the color of each point indicating the retinal size of the post at the time of the saccade. The black line in A shows a linear regression through the data, $\psi=1.06\phi-4.06^\circ$ ($R^2=0.90$, $n=68$). (C) The distribution of saccades in A (landings) as a function of retinal size and distance to post. (D) The distribution of saccades in B (fly-bys) as a function of retinal size and distance to post. In panels C and D the associated shading shows smoothed representations of the distributions calculated using 3rd order 0.3 Hz Butterworth filters.

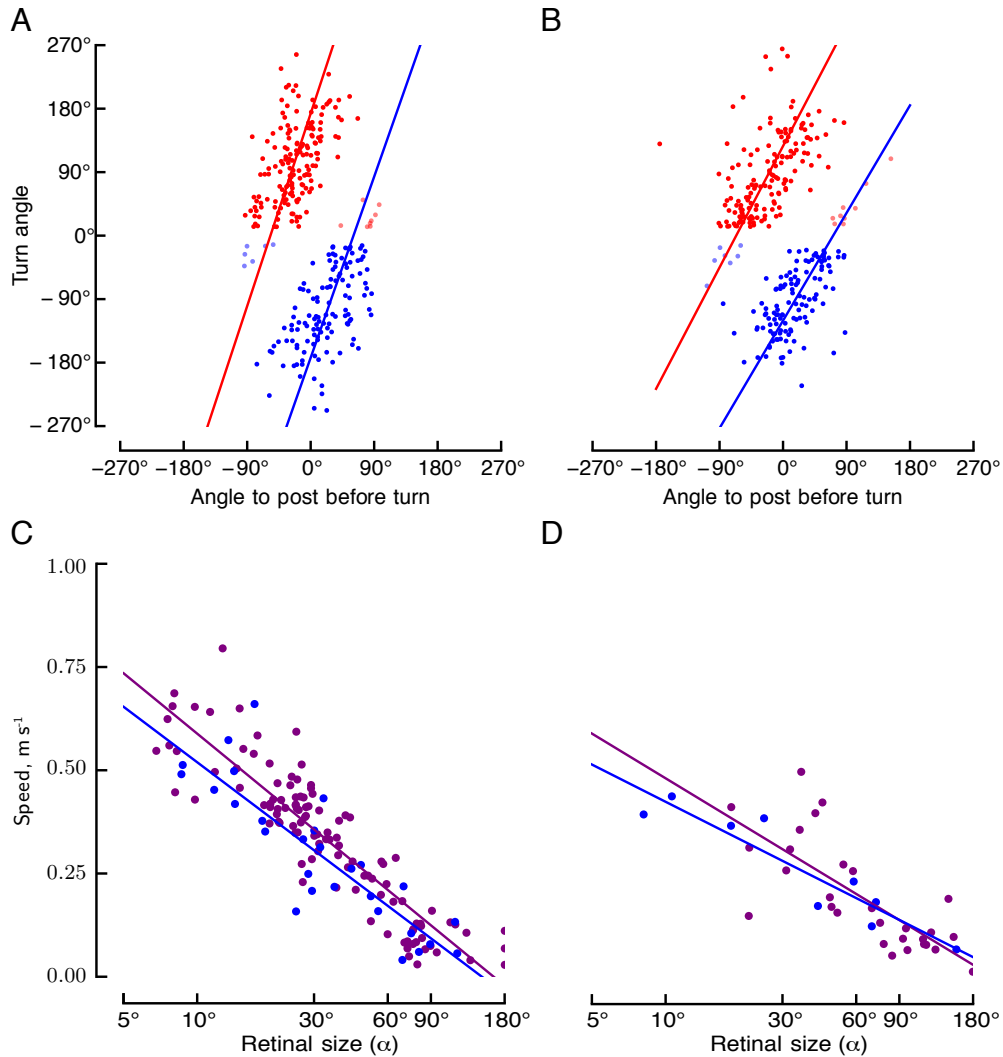


Figure 3.19: Object contrast subtly influences landing and fly-by behavior. Figure 3.8 is repeated, but with the data segregated for the black post (A) and checkered post (B). The linear regressions are given in Table S. I did not find any significant differences for left turns (blue points, $p(x_2)=0.03$, $p(x_1:x_2)=0.58$) or right turns (red points, $p(x_2)=0.001$, $p(x_1:x_2)=0.11$), except for the intercept ($p(x_2)=0.001$) for right turns, and marginal significance for left turns ($p(x_2)=0.03$). In panels (C) and (D) Fig. 3.9B is repeated with the data segregated for the black and checkered post, respectively. The linear regressions are given in Table S. Although the statistics suggest that there may be significant differences in deceleration behavior due to the post texture for the flies that did not turn after initiating deceleration (purple points, $p(x_2)=0.66$, $p(x_1:x_2)=0.011$), visual inspection of the data does not show any notable difference. The flies that did turn after initiating deceleration do not show any statistically significant difference (blue points, $p(x_2)=0.67$, $p(x_1:x_2)=0.30$). The details of this statistical analysis are given in the Methods. In both cases, visual inspection does not suggest that there are any substantive differences, and there is insufficient data for statistical comparisons to be taken at face value.

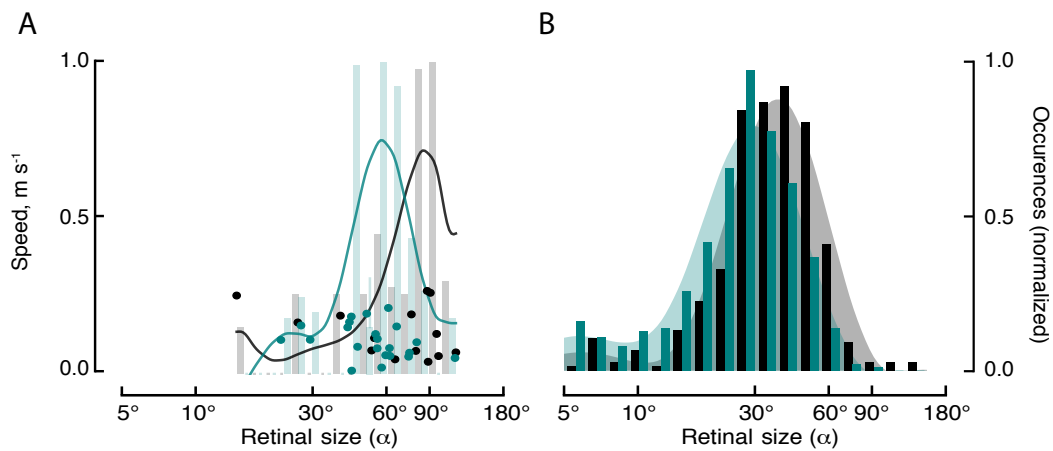


Figure 3.20: (A) Leg extension behavior is not likely influenced by post texture. The data from Fig. 3.9D is repeated for the black post (black) and checked post (teal). The difference in these distributions is insignificant (t-test, $t=1.79$, $df=34$, $p=0.08$). (B) Retinal size at which non-landing flies perform an evasive saccade is subtly influenced by the contrast of the post (t-test, $t=8.86$, $df=1222$, $p<0.001$). Fig. 3.7D is repeated, but with the data segregated for the black post (black) and the checked post (teal).

Chapter 4

Monocular distance estimation from optic flow during active landing maneuvers¹

4.1 Abstract

Vision is arguably the most widely used sensor for position and velocity estimation in animals, and it is increasingly used in robotic systems as well. Many animals use stereopsis and object recognition in order to make a true estimate of distance. For a tiny insect such as a fruit fly or honeybee, however, these methods fall short. Instead, an insect must rely on calculations of optic flow, which can provide a measure of the ratio of velocity to distance, but they cannot determine either parameter independently. Nevertheless, flies and other insects are adept at landing on a variety of substrates, a behavior that inherently requires some form of distance estimation in order to trigger distance-appropriate motor actions such as deceleration or leg extension. Previous studies have shown that these behaviors are indeed under visual control, raising the question: how does an insect estimate distance solely using optic flow? In this paper I use a nonlinear control theoretic approach to propose a solution for this problem. My algorithm takes advantage of visually controlled landing trajectories that have been observed in flies and honeybees. Finally, I implement my algorithm, which I term *dynamic peering*, using a camera mounted to a linear stage

¹As of Nov. 1st 2013, the material in this chapter is in press at *Bioinspiration and Biomimetics* [151].

to demonstrate its real-world feasibility.

4.2 Introduction

Many animals rely heavily on vision to gather information about their position and velocity relative to objects in the world around them. As an animal moves, it perceives apparent motion of these objects, with nearby objects moving faster than distant ones. This phenomenon is referred to as optic flow and essentially provides a measure of the ratio of forward movement to the distance of surrounding objects [86]. This coupled relationship intuitively suggests that estimating either absolute distance or velocity from optic flow is challenging. Indeed, it is not immediately clear whether an accurate estimate of both position and velocity from optic flow is possible at all.

Optic flow is only one of many sensory modalities available to most animals, and given additional sensory information it is possible to derive accurate estimates of velocity and position. For example, terrestrial animals might count strides, as supported by experiments in desert ants [166]. Another theoretical possibility for walking animals is to use visual odometry based on ventral optic flow, with the system calibrated by proprioceptive information that accurately measures the distance between the eye and the ground [63]. Animals with high acuity vision and long-term memory might estimate distance to recognizable objects by remembering their typical size [62]. Perhaps the most straightforward strategy is stereopsis, which can provide a distance estimate based on parallax [11, 62]. As will become clear in my subsequent discussion, however, none of these approaches are plausible for small flying insects. How is it then possible for a fruit fly or honeybee to avoid some objects and land on others, without knowing either how fast they are going or how far away the obstacles are located?

Growing evidence from studies of insects and birds reveals a clever algorithm that solves this problem, at least in the context of landing. Although the details vary across the studies, the core strategy is that by maintaining a constant rate of optic flow, an animal can decelerate to a safe touchdown velocity without needing to directly

measure either its velocity or distance to target [134, 160, 150]. Because optic flow essentially provides a measure of the ratio of velocity to distance (units of s^{-1}), it is often instead referred to as the inverse of time-to-contact (the time it would take for the animal to hit an object assuming it is on a collision course and maintains its initial speed and heading). A controller implementing this simple law would cause an animal to decelerate in such a way that its velocity would be inversely proportional to the distance to nearby objects. These observations have inspired the implementation of similar algorithms on robotic aircraft to achieve smooth automated landings [32].

Although the constant optic flow algorithm offers a robust strategy for deceleration, other components of landing behavior may require different sensory mechanisms. For example, at some point before contact, the animals need to extend their legs to help touch down safely. For animals that tuck their legs tightly during flight, this motion would ideally happen at a short and consistent distance from the landing target. Indeed, both flies and honeybees appear to have this capacity [49, 150]. Furthermore, tethered flight experiments with fruit flies have demonstrated that leg extension behavior is triggered by visual cues [17, 144]. Honeybees have also been shown to rely exclusively on visual cues to accurately determine the distance to targets such as artificial flowers [31, 94, 93]. These observations suggest that somehow insects are able to measure some feature that is tightly correlated with distance using only vision, yet the precise details of how this might be accomplished are not known.

In this paper I present a novel algorithm, termed dynamic peering, that I propose as a potential model for how insects accomplish this task as well as being a useful algorithm for small scale and computationally limited robotics applications where traditional forms of distance estimation such as stereo, sonar, and laser rangefinders are too large, heavy, costly, or computationally intensive. Image based sensing also has the advantage that it is a passive system, rather than needing to send out or receive active signals.

4.2.1 Review of visual distance estimation in biological systems

Before beginning the derivation of the dynamic peering algorithm, I present a brief review of experimentally confirmed mechanisms for vision based distance estimation in biological systems (see also [37]). Perhaps the most familiar mechanism for estimating distance from visual information, on which humans heavily rely, is object recognition. Given a recognizable object and knowledge of its typical size, it is possible to estimate the distance to the object. This process is, however, a cognitively complex task that relies on high acuity vision, many layers of visual processing, and a large memory of objects. It is thus an unlikely general mechanism for insects, which have neither the high acuity vision nor the neural capacity necessary for this approach. Another strategy is to use an image-forming lens with a shallow depth of field, which makes it possible to calculate depth from the level of defocus [115]. Such a mechanism, however, is incompatible with the optics of compound eyes [92]. The simplest, and most widely used mechanism for estimating distance—both in biology as well as in computer vision and robotics—is stereopsis. Stereopsis works by calculating the parallax between two (or more) simultaneous images of the same object taken from different viewpoints to triangulate the absolute distance to the object. Many factors play a role in the accuracy of this method, however a simplified distance estimate error (e_d) from a stereoscopic camera pair is given by:

$$e_d \geq \frac{d^2 \xi}{lf}, \quad (4.1)$$

where d is the distance to the target, l is the interocular distance, f is the focal length of the lens, and ξ is the pixel resolution [154]. To use this equation with biological systems, I rewrite the equation using the relationship:

$$\frac{\xi}{f} = 2 \tan \left(\frac{\theta}{2} \right), \quad (4.2)$$

where θ is the angular resolution of the visual system. Substituting 4.2 in 4.1 yields:




Figure 4.1: Theoretical errors in stereoscopic absolute distance estimates for humans (blue) and fruit flies (red), calculated using 4.3.

$$e_d \geq 2 \tan\left(\frac{\theta}{2}\right) \frac{d^2}{l}. \quad (4.3)$$

To gain intuitive insight into this relationship, consider the human visual system, which has an interocular distance of approximately 65 mm and a stereoscopic angular resolution of 0.003° [159]. Given these parameters, the minimum theoretical error in distance estimates is described by the blue curve in Figure 4.1. Next, consider the visual system of the fruit fly, which has an interocular distance of approximately 0.3 mm. Although the precise angular resolution of the flies' visual system is not known, I approximate it by the ommatidial acceptance angle of 5° , yielding the red curve in Figure 4.1. Given these results, stereopsis might provide useful information for behaviors that involve operating at very close distances, for example when a male chases a female during courtship. However, it is unlikely that flies could use stereopsis to trigger leg extension during flight, which occurs when they are approximately 1 cm from the target [150], at which distance the minimum error using stereopsis would be approximately 3 cm. Although details vary, such limitations are general for most insects given their small values of interocular distance.

A related, time domain approach to stereopsis is to use sequential images from a single moving eye that has traversed some known distance between acquisitions. Some insects, such as locusts, use this approach to estimate distances before jumping or attacking prey by moving their heads back and forth in a regular fashion, a behavior known as “peering” [36, 133, 158]. For this approach to work, the distance between the positions at which the two images are acquired must be known. This distance could, for example, be calculated if the velocity and time interval are known, or determined directly via proprioceptive sensory feedback [117]. However, because a flying animal has no accurate measure of its true groundspeed (only its airspeed), this method is not feasible for estimating distance for landing behaviors.

Suppose that rather than moving an eye or camera at some known velocity, it is accelerated at a known rate. The acceleration could either be produced along the direction of travel in the case of a straight trajectory, or by changes in direction. As I will show in the following section, knowledge of the acceleration and the time interval between image acquisitions is sufficient to estimate both velocity and absolute distance. This approach, which I term dynamic peering, has the intuitive functionality of the nonlinear observer I will develop more formally in the following sections.

4.3 Modeling and observability analysis

Of the approaches for vision based distance estimation described above, only dynamic peering has the potential for use by a small flying insect, and thus also by a similarly-scaled robot. In this section, I will take a control theoretic approach to formalize this concept and construct a nonlinear observer that can estimate distance if, and only if, certain controls (actuation patterns) are applied. In order to simplify the system so that the principles are as transparent as possible, I will focus on the problem of landing on a large flat target given a system with simple linear dynamics limited to a single translational degree of freedom. In the final section, I describe a physical implementation of the proposed observer, which demonstrates its feasibility in the real world. Finally, I will propose some methods that would allow

my algorithm to be implemented by higher degree of freedom systems in complex environments. Throughout the following sections, I will use the term “camera” and “robot” to represent a simple imaging device and some moving agent, but the terms “eye” and “animal” could be used as well.

4.3.1 Problem statement

Given a single camera and control over forward acceleration, an agent flying toward a flat surface must decelerate to a safe speed and estimate the distance in order to prepare for touchdown.

4.3.2 Analysis

As the camera approaches the surface, from its perspective, textures on the wall will appear to move across its visual field at some angular velocity that is a function of its forward flight speed, the distance to the wall, and the heading angle between the texture and the agent’s trajectory. If I assume perfectly spherical optics, this relationship can be described as follows,

$$\frac{d\alpha}{dt} = -\frac{v_h}{d_h} \tan(\alpha), \quad (4.4)$$

where d_h is the distance to the object, v_h is the velocity in the direction of d_h , α is the position of the object on the camera’s retina, and $\dot{\alpha}$ is the angular velocity (e.g., optic flow) of the object relative to the camera, see Figure 4.2.

Assuming spherical lens geometry, α corresponds directly to the heading of the object relative to the camera. Thus, for each direction α , the ratio v_h/d_h is directly proportional to $\dot{\alpha}$ by a constant of $-1/\tan(\alpha)$. In the case that the camera is moving directly towards a flat wall, I can relate all such measurements for different α ’s and rewrite 4.4 in terms of the forward velocity (v) and the distance to the wall (d):

$$-\frac{\dot{\alpha}}{\tan(\alpha)} = \frac{v_h}{d_h} = \frac{v}{d}. \quad (4.5)$$




Figure 4.2: Geometric relationship of the moving camera and a reference object, as referenced in (4.4).

By only using measurements corresponding to small α (e.g., from the center of its field of view), the equation further simplifies to:

$$\frac{v}{d} = -\frac{\dot{\alpha}}{\tan(\alpha)} \approx -\frac{\dot{\alpha}}{\alpha}. \quad (4.6)$$

These simplifications make it possible to estimate v/d as the mean of $\dot{\alpha}/\alpha$ across all (small) α . In the following sections I use this relationship to describe optic flow as v/d . Note that although this relationship appears to be poorly defined at $\alpha = 0$, $\dot{\alpha}$ also tends toward zero in this direction, and the limit remains well-defined. In a real-world implementation where noise is unavoidable, this relationship is prone to producing large errors. In my implementation section, I discuss a simple solution to this problem.

Next consider a robot equipped with a single camera flying straight towards a static object. The equations of motion can be written as:

$$\begin{bmatrix} \dot{d}(t) \\ \dot{v}(t) \end{bmatrix} = \begin{bmatrix} v(t) \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(t) = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} d(t) \\ v(t) \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(t), \quad (4.7)$$

where d is the distance to the object, v is the forward flight speed, and u is the control input (which is equivalent to acceleration with these dynamics). With optic flow as the system's only sensory input, I can write the observations as:

$$y(t) = [v(t)/d(t)]. \quad (4.8)$$

Although the dynamics are linear, the observation equation, y , is nonlinear. In order to use linear systems analyses, I begin by linearizing the system about a nominal trajectory, $(d_t(t), v_t(t))$. This choice allows us to write the system in the canonical state space form:

$$x(t) = Ax(t) + Bu(t), \quad y(t) = Cx(t) + Du(t), \quad (4.9)$$

where, for my system, the terms are defined as follows:

$$x(t) = \begin{bmatrix} d(t) \\ v(t) \end{bmatrix}, \quad (4.10)$$

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad C = \begin{bmatrix} -v_t/d_t^2 & 1/d_t \end{bmatrix}, \quad D = \begin{bmatrix} 0 \end{bmatrix}.$$

To address the question of whether or not it is possible to estimate distance and velocity using only optic flow, I examine the system's observability (a measure for how well the states of a system, such as position and velocity, can be inferred given the available sensory measurements [112]). First, I check the linear observability condition for an arbitrary nominal trajectory by calculating the rank of the observability matrix:

$$\begin{bmatrix} C \\ CA \end{bmatrix} = \begin{bmatrix} -v_t/d_t^2 & 1/d_t \\ 0 & -v_t/d_t^2 \end{bmatrix}. \quad (4.11)$$

This matrix is full rank along any trajectory provided that neither the velocity nor distance is zero, suggesting that the system is observable. Because in practice I am concerned with non-zero velocities at non-zero distances, this limitation is not critical. This conclusion, however, goes against the intuition presented in the introduction, which suggested that it should be difficult, if not impossible, to extract either distance or velocity from optic flow. To explore this discrepancy, I use an alternative check for calculating the observability of a linear time-varying system, the observability Gramian:

$$P(t_0, t_f) = \int_{t_0}^{t_f} e^{A^T t} C^T C e^{A t} dt \quad (4.12)$$

where C (defined in 4.10) is time varying. The advantage of this approach is that the condition number (the ratio of the minimum and maximum singular values of a matrix), termed the *local estimation condition number*, is a direct measure of the well-posedness of the estimation problem [88]. The smaller the condition number, the better posed the estimation problem. Calculating 4.12 analytically is often challenging for complex systems. Instead, it is possible to numerically estimate the observability Gramian, termed the *empirical local observability Gramian*, by simulating the system and comparing the outputs y for perturbations $\pm\epsilon$ of the initial condition [75, 88]. It can be shown that as $\epsilon \rightarrow 0$, the estimate converges to the result from 4.12. Using this approach on my system shows that the condition number for a constant velocity trajectory approaches infinity for any time interval (Figure 4.3, black trace), suggesting the system is not in fact observable in a linear sense. However, significantly smaller condition numbers exist if I consider non-constant velocity trajectories (Figure 4.3).

For a more direct confirmation that non-zero control inputs are required for the system to be observable, I can employ a nonlinear observability analysis, which draws

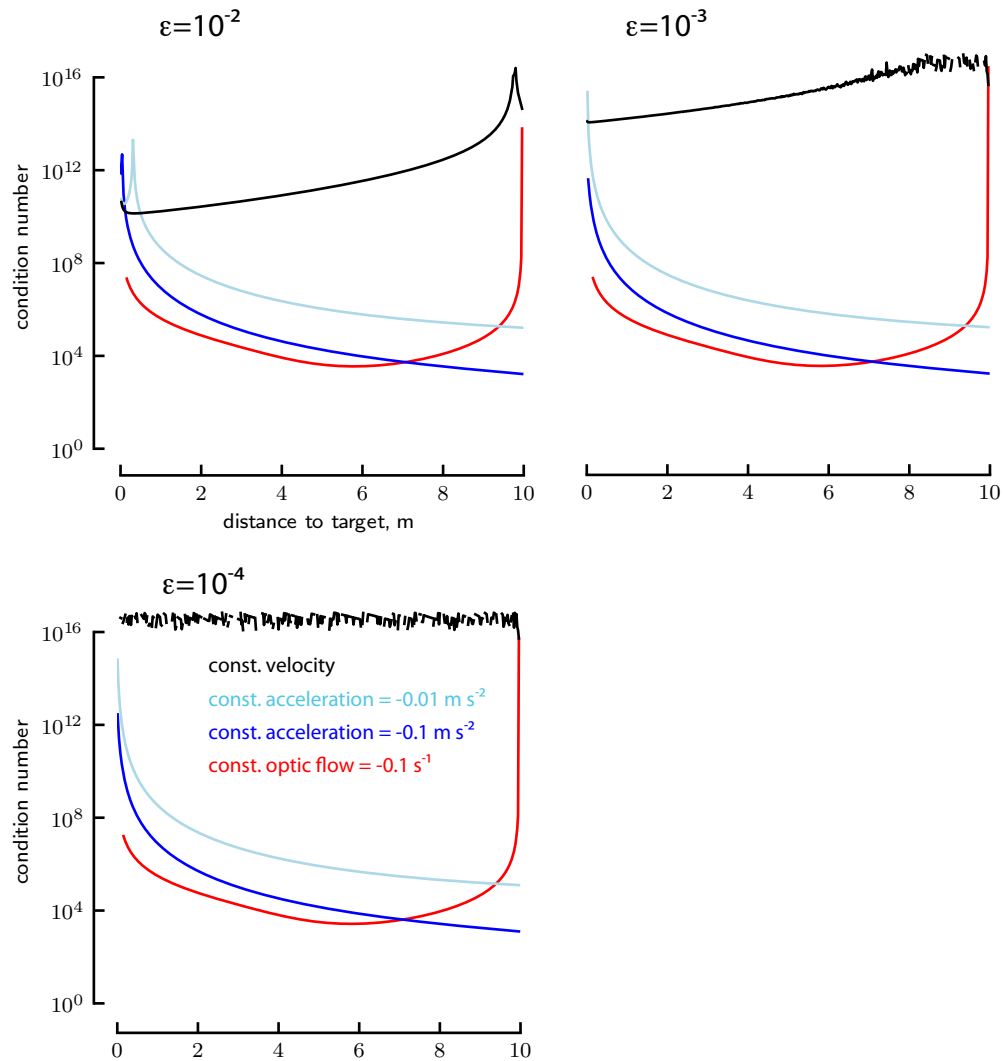


Figure 4.3: Condition number of the empirical local observability Gramian for a constant velocity trajectory (black) approaches infinity, whereas a constant acceleration trajectory (blue) or a constant optic flow trajectory (red) achieves significantly lower condition numbers. The constant optic flow trajectory is obtained by using the controller given by 4.16, and will be discussed at length later in the paper.

on Lie algebraic tools to account for the contribution of active controls.

A brief review of the key elements will be presented here; for a more detailed discussion see [112]. The Lie derivative of the observation equation, $y = h(x)$, with respect to a vector field, $f_i(x) \in \mathbb{R}^n$, where $x \in \mathbb{R}^n$, is defined as:

$$L_{f_i}h = \frac{\partial h}{\partial x} f_i. \quad (4.13)$$

Intuitively, the Lie derivative represents the change in a function or vector field along a vector field. Applied to the observability problem presented in this paper, $L_{f_i}h$ is the change in the observations ($y = h(x) = \text{optic flow}$) along either the drift dynamics, $f_0 = Ax$, or the control direction, $f_1 = B$. Because the drift dynamics cannot be turned off, I take a repeated Lie derivative with respect to both f_0 and f_1 in order to calculate the change in the observations with respect to the control. This repeated Lie derivative is defined as:

$$L_{f_0 f_1} h = L_{f_0} L_{f_1} h = \frac{\partial}{\partial x} (L_{f_1} h) f_0. \quad (4.14)$$

Next, define the *observability Lie algebra*, \mathcal{O} , which is the collection of Lie derivatives of the observations, h , with respect to the drift dynamics and each of the controls. In the case of my problem:

$$\mathcal{O} = \{h(x), L_{f_0}h(x), L_{f_0 f_1}h(x)\} = \left\{ \frac{v}{d}, \frac{-v^2}{d^2}, \frac{v}{d^2} \right\}. \quad (4.15)$$

If the Jacobian of \mathcal{O} is full rank (i.e. if the number of linearly independent terms is equal to the number of states in the system) at all states, the system is said to be observable. If \mathcal{O} must include Lie derivatives between h and a control vector, f , in order to become full rank, this control dimension must be actuated in some way in order for the system to be observable. In my system, the terms $\{v/d, -v^2/d^2\}$ result in linearly dependent terms within the Jacobian, so \mathcal{O} reduces to $\{v/d, v/d^2\}$, which has a full rank Jacobian if and only if the term $L_{f_0 f_1}h(x)$ is included (and for $v \neq 0$ and $d \neq 0$). The presence of the term $L_{f_0 f_1}h(x)$ suggests that in order for the system

to be fully observable, some acceleration must be applied. That is, if the camera-equipped robot were approaching the wall at constant velocity, I would not be able to observe d and v separately. If the camera accelerates, however, such an observation is possible. These results are consistent with my conclusions from the observability Gramian analysis in Figure 4.3, and is the critical feature of my proposed method.

In principle, I could conclude my analysis at this point and simply implement an unscented Kalman filter or particle filter to estimate the unobserved states [39] and ensure that the system accelerates sufficiently often. Unfortunately, however, the observability Lie algebra calculation does not give any indication of *what* that control needs to be; only that it cannot be identically zero. To address this limitation, previous studies have developed methods that use the condition number of the observability Gramian to rank potential trajectories [75]. In this paper, I take a bio-inspired approach as described below.

Consider the landing algorithm presented in the introduction, where an animal maintains a constant rate of optic flow in order to guarantee a safe landing velocity at touchdown. This motion can be accomplished by using a simple proportional feedback controller on the velocity:

$$u = k(r - r_d), \quad (4.16)$$

where $r = v/d$, and r_d is the desired ratio of v/d . To understand the implications of such a controller I write:

$$\text{optic flow} \propto r = v/d, \quad (4.17)$$

$$\frac{d}{dt}r = \frac{d}{dt}vd^{-1} = \frac{\dot{v}}{d} - \frac{v^2}{d^2} = \frac{\dot{v}}{d} - r^2 = \frac{u}{d} - r^2. \quad (4.18)$$

Rewriting this equation yields the following equation for distance:

$$d = \frac{\dot{v}}{\dot{r} + r^2} = \frac{u}{\dot{r} + r^2}. \quad (4.19)$$

Thus, given non-zero control, it is possible to estimate distance directly from the control input and a measurement of r along with its derivative, both of which can be extracted from measurements of optic flow. It is important to note, however, that this relationship has several pitfalls, particularly for a noisy system. As I will describe in the implementation section, current computer vision methods for calculating optic flow are indeed noisy, so these pitfalls present a real problem. In the case that $\dot{r} + r^2 = 0$, the distance estimate approaches infinity. Although this will theoretically not happen if the acceleration is sufficiently large, there is no such guarantee in a noisy system. Given the bio-inspired controller that maintains a constant rate of optic flow ($\dot{r} = 0$), however, the distance estimate in 4.19 can be simplified to:

$$d = \frac{u}{r_d^2}, \quad (4.20)$$

where r_d is the desired rate of optic flow, u is the control input (acceleration), and d is the distance to the target. This equation evolves with time according to:

$$d = d_0 e^{r_d t}, \quad (4.21)$$

where t is time, and d_0 is the (unknown) initial distance. In this way, I have removed the potential for dividing by zero, and all the effects of noise reside within the relative safety of the numerator. Furthermore, by using 4.21 I have reduced the estimation problem to that of estimating a single parameter—the initial distance—allowing us to calculate a clean distance estimate with a sequential least squares filter [39, pg. 19]². This approach is, of course, limiting for general applications. However, in the context of landing, where a constant rate of optic flow controller is ideal, these constraints do not present a critical problem. On the contrary, they provide an exceptionally elegant method for triggering behaviors such as leg extension based on the internal state of control (or the measured value of acceleration). In the discussion, I explore the possibilities for expanding this approach to more general cases.

²The sequential least squares filter makes it possible to update a least squares estimate as new measurements arrive, rather than batch processing the entire dataset at once. The filter essentially works as a Kalman filter without dynamics.




Figure 4.4: Visual target and optic flow estimation. (A) Example camera image showing the visual target and region of interest (red box). (B) Optic flow as a function of camera pixel from two successive frames, calculated using OpenCV’s Lucas Kanade algorithm. For the purposes of control, I calculated a linear fit of the data (red line) over the region of interest indicated in (A).

4.4 Implementation

As a proof-of-concept physical demonstration of the distance estimation algorithm presented in the previous section, I implemented it using a camera (Basler Ace 640-100 gm) equipped with a 1.4mm fisheye lens (Fujinon C Mount 1.4mm CCTV Fish-Eye) mounted to a linear stage. The camera was driven by a computer controlled servo along a 1.5 m track towards a panoramic image of a forested scene. All the image processing and estimation were done on a desktop computer running Ubuntu Linux. Images were acquired with the open source camera aravis driver, and the rectified images were published on a Robot Operating System (ROS fuerte) network. All subsequent processing was done in Python. I used OpenCV’s implementation of the Lucas Kanade algorithm to calculate optic flow over a region of interest that corresponded to an approximately 45° field of view in the horizontal direction. For simplicity, I restricted my analysis to the optic flow along the horizontal dimension.

Figure 4.4 shows a representative measurement of optic flow along this axis. Recall from 4.5 that optic flow ($\dot{\alpha}$) is equal to $-\tan(\alpha)v/d$. In order to extract v/d , I fit a line to the central portion of this curve using a RANSAC algorithm [50], and used the slope of this line as the estimate for v/d . Note that this approach circumvents the potential for dividing by zero at small α and also makes the algorithm robust to poor

camera alignment and noisy optic flow estimates.

In order to use these optic flow estimates for my estimation problem, I first had to calibrate the system, since the Lucas Kanade algorithm provides normalized values between ± 1 . To calibrate the system, I drove the camera at various known rates of v/d and recorded the associated slope. After collecting several of these points, I did a least squares fit to determine the relationship between my slope estimate and the true v/d .

Next, I implemented a simple proportional controller 4.16 with gain $k = 6$ to adjust the acceleration of the camera such that it maintained the desired value of optic flow. The acceleration commands were turned into updated velocity commands based on the operating frequency of the control loop and were sent to an Arduino Uno board over USB. The Arduino continuously generated step and direction commands that were sent to the stepper motor controller, which in turn moved the camera. These step commands served as my ground truth of the actual distance and velocity of the camera to which I could compare my algorithms' estimates. The rate limiting step of my implementation was the communication to the Arduino board, resulting in a 50 Hz operating frequency, which was sufficient for my demonstration purposes.

To estimate the distance, I used a two-step process. For each cycle of the 50 Hz control loop, I estimated distance using 4.20. These distance estimates were then run through a sequential least squares (SLS) filter [39] to estimate the initial distance d_0 in 4.21. Because 4.20 is only valid when the system is moving with the desired rate of optic flow, I kept the covariance in the SLS filter artificially high until the desired rate of optic flow was reached. My software, as well as the data collected using this system, are freely available online at www.github.com/florisvb/dyneye.

4.5 Results and discussion

The camera started out at zero velocity 1.5 m away from the target, quickly accelerated up to the desired optic flow set-point, and subsequently began gently decelerating as it approached the target so as to maintain the desired constant rate of optic flow.

Figure 4.5 shows a comparison between the actual position, the nonlinear observer estimates at each time point calculated using (18), and the sequential least squares (SLS) filtered estimates. I also implemented a square-root unscented Kalman filter (UKF) [152] and found very similar performance to the SLS results, but with much greater computational overhead (not shown for the sake of graphical clarity).

Because my nonlinear observer 4.20 only provides valid information when the system is close to the desired trajectory, the position and velocity estimates show large initial errors. These initial errors are an artifact of my system starting out at zero velocity and initially needing to accelerate to reach the desired optic flow rate. The initial errors are more pronounced in experiments with higher rates of optic flow, as the system took longer to reach the target level. This left little space on the limited track to operate at the final desired rate. In a freely moving system, deceleration could be triggered when the optic flow reaches the desired threshold, and subsequent distance and velocity estimates would be accurate so long as the system maintained the target level of optic flow. For very slow rates of optic flow ($r > -0.01$), the poor estimates are most likely due to insufficient changes in pixel values between sequential image acquisitions to calculate accurate measures using the Lucas Kanade algorithm. This limitation could potentially be solved by using longer delays between image acquisitions.

4.5.1 Applications to robotic systems

The approach presented here is best suited for landing applications due to the trajectory choice; however, by integrating the algorithm into a more complex trajectory it could provide more general utility in navigation tasks. For example, a flying robot could periodically approach the ground below it with constant optic flow to estimate its altitude and use this measurement as a calibration for other optic flow estimates. Recall that optic flow alone can provide relative measurements to different objects, so if the distance to any one of these objects is known, the others can be calculated as well. Between these bouts, the robot could use other bio-inspired visual cues

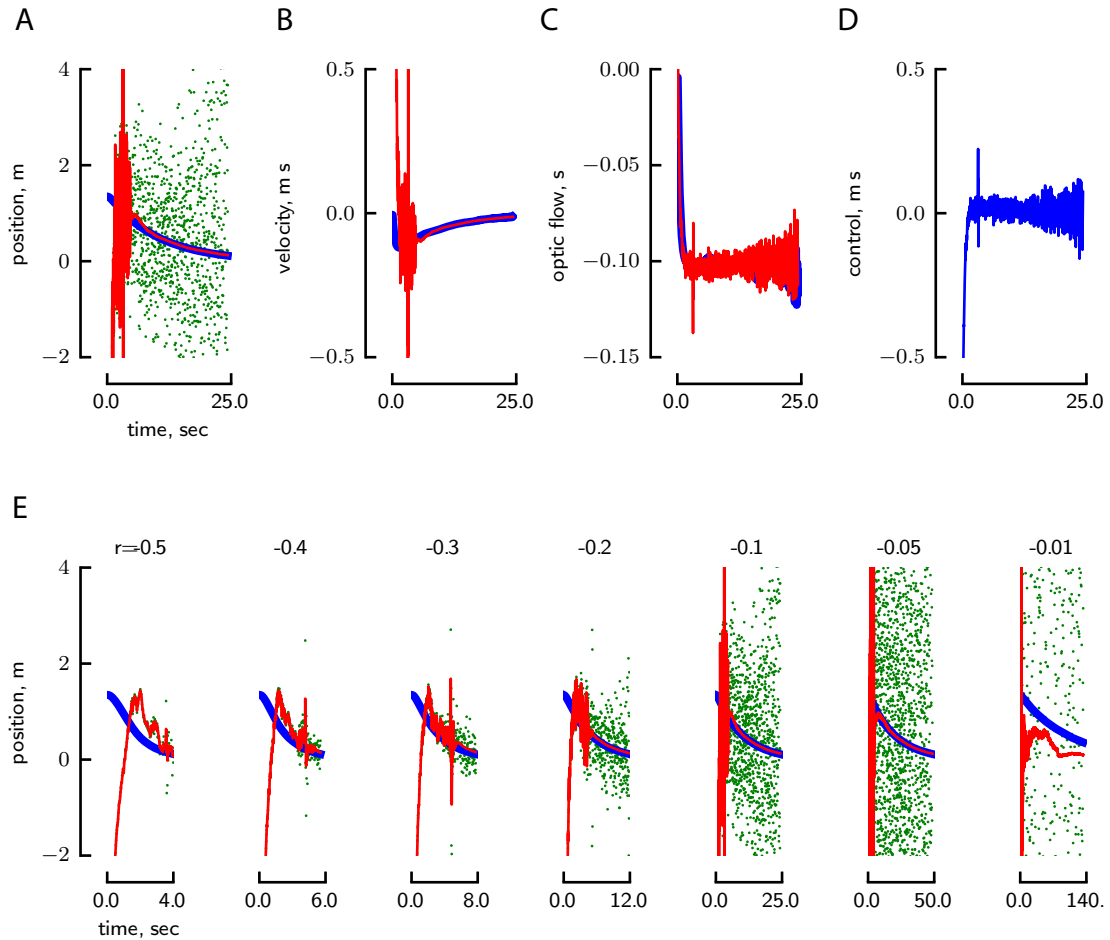


Figure 4.5: Performance of the dynamic peering estimation algorithm. Panels A-D show results for a desired optic flow rate of -0.1 s^{-1} . (A) Actual position (blue), raw dynamic peering estimate (green), and the sequential least squares filtered dynamic peering estimate (red). (B) Actual velocity (blue), dynamic peering with sequential least squares filtering (red). (C) Actual velocity/distance (blue) and optic flow measurement based on Lucas Kanade calculations (red). (D) Control commands, equivalent to acceleration (blue). (E) Same Figure as (A) repeated for different desired optic flow rates (indicated above each plot).

to maintain a constant altitude, such as local horizon following[141]. Although my implementation made use of a full desktop computer, this use was purely for convenience. Optic flow can be calculated with simple parallel analog circuits, without the need for extensive memory [136].

I restricted my analysis to a single degree of freedom system for simplicity, however the principles presented here can be generalized to three-dimensional motion. In three dimensions, the non-zero acceleration requirement can also be satisfied by non-zero angular accelerations, i.e., changes in direction [1].

4.5.2 Implications for landing insects

The dynamic peering algorithm presented here provides a plausible mechanism for how insects might estimate distance in order to trigger leg extension or any other behavior during landing maneuvers. Once the animal starts decelerating in preparation for landing by maintaining a constant rate of optic flow, it could trigger leg extension when either its internal control signal, or a measured value of acceleration, reaches a particular threshold. This threshold would always correspond to the same distance if the animal uses the same rate of optic flow for regulating its deceleration, a hypothesis supported by behavioral observations [134, 150]. Such a threshold calculation could be easily implemented by the nervous system with as little as a single neuron. To my knowledge, this is a novel and plausible principle. Insects, such as flies [53, 106, 125], are known to use optic flow to regulate their velocity by adjusting acceleration, suggesting that such an internal control signal is indeed present.

To experimentally test whether insects use dynamic peering, one could construct an experiment in which animals are tricked into decelerating faster or slower than usual by using a virtual reality system. If the insects use an approach similar to what I describe, their leg extension behavior would be correlated with their internal control commands, and thus the observed acceleration behavior. Unfortunately this type of experiment is challenging to perform, as it requires free flight observations of leg extension behavior and simultaneous control of the visual stimulus in closed loop. As

reported previously, flies landing on a target typically extend their legs after having first initiated a deceleration [150]. However, flies that approached the target but did not decelerate, did not extend their legs. These results suggest that leg extension is not simply triggered by object size or image velocity, but rather it depends on some aspect of the flies' internal state, or their deceleration behavior, an observation that is consistent with my model for distance estimation. Tethered flight experiments in fruit flies, as well as other insects, have shown that visual stimuli are sufficient to elicit strong leg extension responses [17, 144]. These results from tethered flies seem contradictory to my hypothesis because the flies extend their legs without physically decelerating. However, in these tethered flight experiments it is impossible to know if the flies were attempting to decelerate. Assuming they were, this behavior could be explained if they used an internal signal (e.g., efferent copy) of the control output. It is also possible that there are multiple sensory-motor pathways that can elicit leg extension.

Our analysis focused on a single degree of freedom trajectory, which is consistent with the landing behavior of fruit flies which do not make significant changes in heading after initiating deceleration prior to landing [150]. Additional degrees of freedom in the system would make it possible for an insect to extract a true depth estimate from optic flow simply by changing direction by a known amount. Previously published results from experiments with honeybees are consistent with this option, having shown that they do actively change direction in a stereotypical fashion in order to determine the height of objects above the ground using optic flow information [93].

4.5.3 Summary

To summarize, in this paper I used tools from control theory to show that non-zero acceleration is necessary and sufficient to estimate distance from optic flow with a single camera. There are, however, an infinite number of possible trajectories that satisfy these requirements, many of which will fail in real-world noisy implementations. In order to choose a trajectory that provides an accurate estimate of distance,

I turned to biological inspiration from landing insects, which decelerate so as to keep their optic flow at a constant value. This choice of trajectory simplifies the estimation problem to a single parameter (initial distance), resulting in an accurate estimate of distance from sequential least squares filtering. My dynamic peering algorithm provides a plausible, and testable, mechanism for how insects might trigger leg extension prior to landing, as well as being a novel method for estimating distance with a single camera in robotic systems.

Chapter 5

Octopaminergic modulation of the visual flight speed regulator of *Drosophila*¹

5.1 Abstract

Recent evidence suggests that flies' sensitivity to large field optic flow is increased by the release of octopamine during flight. This increase in gain presumably enhances visually-mediated behaviors such as the active regulation of forward speed, a process that involves the comparison of a vision-based estimate of velocity with an internal set point. To determine where in the neural circuit this comparison is made, I selectively silenced the octopamine neurons in the fruit fly, *Drosophila*, and examined the effect on vision-based velocity regulation in free flying flies. I found that flies with inactivated octopamine neurons accelerated more slowly in response to visual motion than control flies, but maintained nearly the same baseline flight speed. My results are parsimonious with a circuit architecture in which the internal control signal is injected into the visual motion pathway upstream of the interneuron network that estimates ground speed.

¹As of Nov. 1st 2013, the material in this chapter is under review at the Journal of Experimental Biology.

5.2 Introduction

Many animals modulate the properties of their neural networks according to behavioral state in order to increase their functionality [103], a principle that presumably applies to control circuits that regulate behavioral actions. Studies in mice [2, 111], monkeys [108, 147], and flies [79, 101], show that modulation of the visual processing system in particular is a common feature across taxa, however, the behavioral implications of these modulations are unknown. Presumably, increasing the sensitivity of the visual system during certain behaviors allows these animals to react more quickly to visual disturbances through a sensory motor feedback control loop.

Feedback control works by comparing a perceived sensory signal with a desired output, and adjusting motor actions to minimize their difference. Within such a circuit, there are two fundamentally different locations where modulation (e.g., changes in gain) might take place: in the sensory system (before the comparison is made), or in the controller (after the comparison is made). Our current understanding of the underlying neuroanatomy in monkeys, mice, and flies, would suggest that this gain modulation is situated on the sensory side of the feedback comparison (Fig. 5.1, H1). However, gain changes in sensory signals can lead to unpredictable behavioral responses because of a large bias introduced to the error signal (when the desired value is not equal to zero). One solution would be to link changes in sensory gain with simultaneous and identical changes in the branch including the command input so that the error signal is balanced (Fig. 5.1, H2a).

A simpler alternative is to modulate the gain of the feedback controller that calculates the error signal (Fig. 5.1, H2b). To incorporate this approach for a visual-motor task would necessitate that an animals' desired reference signal enters the visual stream early, upstream of the neuromodulatory inputs that alter sensory gain. Although topologically distinct, H2a and H2b are mathematically identical. Designing neurobiological experiments that can distinguish between H1 and H2 requires the ability to artificially modulate a specific neural circuit involved in a well-characterized feedback control task without affecting other aspects of behavior. The fruit fly,

Drosophila melanogaster, is particularly well-suited for addressing this question, because genetic tools make it possible to selectively label and manipulate the activity of small groups of neurons in intact behaving animals [113, 131, 153].

In flies and other species of insects, estimates of self-motion are extracted from patterns of optic flow. Local optic flow is estimated by a two-dimensional array of so-called elementary motion detectors [18, 21, 48, 70] and then integrated across visual space by large interneurons in the lobula and lobula plate. The lobula plate tangential cells (LPTC) are particularly well-characterized and many exhibit receptive fields that make them sensitive to different patterns of self-motion, such as those created by rotation and translation during flight [80, 84, 87]. This LPTC network is thought to serve a critical role in flight control, a hypothesis supported by their connections to descending pathways controlling wing and neck motor neurons [69, 137, 138]. Additional experiments in which lobula plate neurons are ablated by either physical or genetic means add further support to this hypothesis [61, 71, 74].

Recent studies in *Drosophila* suggest that gain modulation of the LPTC network occurs during both walking [33] and flight [101]. In the case of flight, this modulation appears to be mediated by octopamine neurons which become active during flight and cause an increase in the gain of the visual responses in at least one class of LPTCs, the vertical system (VS) cells [142]. Evidence from other flies [79, 97] suggest that gain modulation during flight is not restricted to VS cells, but may represent a ubiquitous feature of the visual motion system [27]. Thus, modulation by octopamine probably increases the sensitivity to many large field visual cues, perhaps allowing the fly to react more quickly to visual disturbances during flight. In this paper I test this hypothesis directly by observing flies' velocity and acceleration responses to translational visual motion with and without octopamine by genetically silencing their octopamine producing neurons. Then, I use my results to distinguish between the two models proposed in Fig. 5.1.

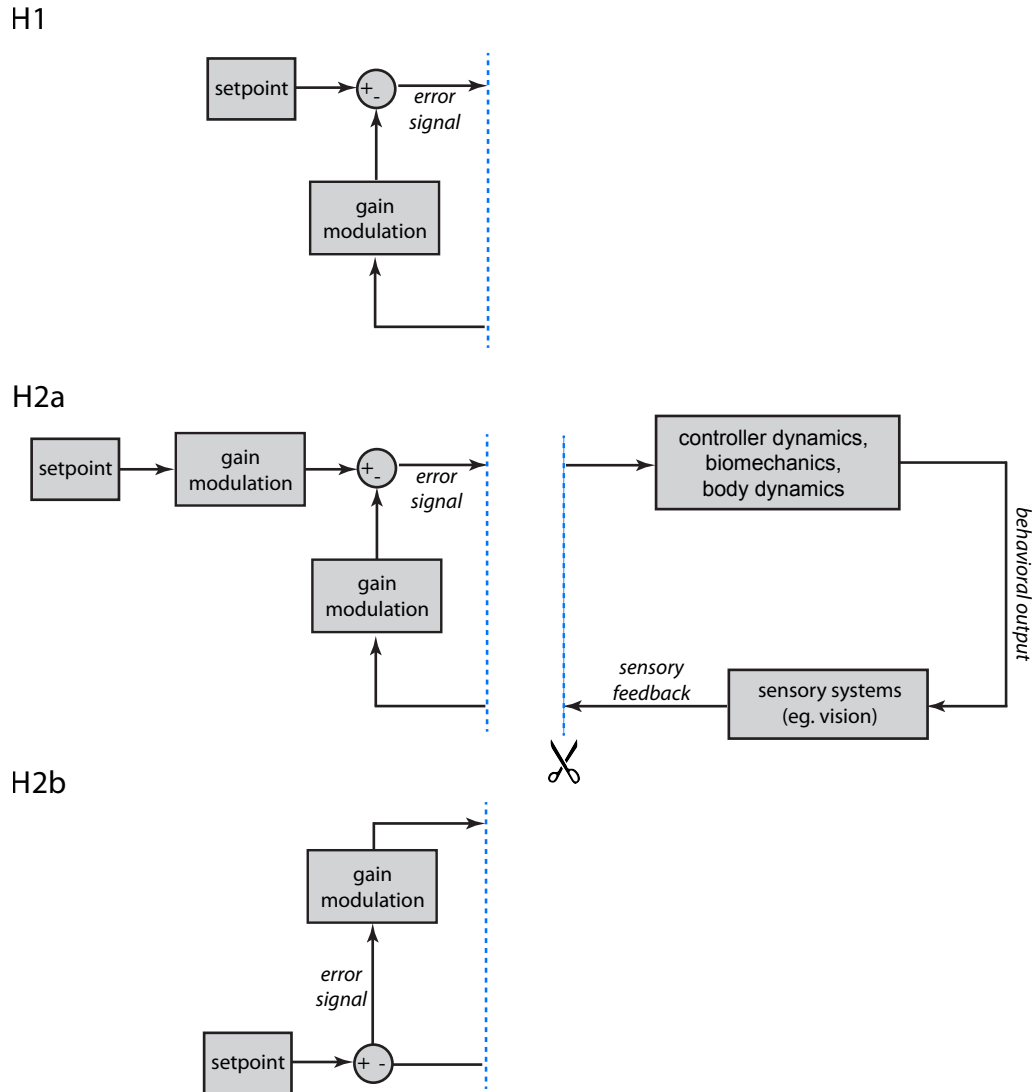


Figure 5.1: Feedback control diagrams describing the three potential control model architectures under consideration (left), each of which is paired with the system dynamics and sensory feedback shown on the right.

5.3 Methods

5.3.1 Animals

Experiments were performed on 2-3 day-old fruit flies, *Drosophila melanogaster* Meigen, using the following transgenic constructs in a w[+] Canton-S genetic background: Tdc2-Gal4 (FBst0009313) and UAS-Kir2.1-EGFP (FBti0017552). In the text, the parental controls are always referenced in the following order: UAS-Kir2.1, Tdc2-Gal4. Flies were deprived of food, but not water, for 6-8 hours prior to the start of the experiment in order to motivate flight. For each experimental trial, I introduced a group of 12 flies to the corner of the arena within a small test tube. The flies were then free to move throughout the flight arena for a period of 12-18 hours, during which time data were collected automatically. With the exception of Fig. 5.3B, all experiments were done with female flies.

5.3.2 Flight arena

I performed all experiments in a 1.5 m x 0.3 m x 0.3 m working section of a wind tunnel (Fig. 5.2A) that has been described previously [26, 100, 141, 150]. In these experiments the wind tunnel was switched off, so that the internal air was still. On the two long walls and floor of the arena I projected a sine grating perpendicular to the length of the wind tunnel with a constant linear spatial frequency of 4.2 m^{-1} using a Lightspeed Designs DepthQ (Oregon City, OR) projector with the color wheel removed (120 Hz update rate, 360 Hz frame rate, mean luminance of 50 cd/m^2). The spatial frequency was chosen because it lies within the range of maximum response for similar visuomotor behaviors in a similar arena [53, 141]. I generated the sine grating, and controlled its temporal frequency, using the VisionEgg open-source image-rendering software [139].

I tracked the 3D position of individual flies within the chamber using a real-time tracking system that is described in detail elsewhere [140]. The 11-camera (Basler Ace 640-100 gm, Basler, Exton, PA) system generated an estimate of fly position at




Figure 5.2: (A) Wind tunnel in which the experiments were performed, with a properly scaled representation of the 4.2 m^{-1} spatial frequency used during the experiments. (B) Confocal image of a *Drosophila* brain and thoracic ganglion showing the GFP labeled octopamine neurons (green), which I genetically silenced with Kir2.1. The imaging for B was done by my colleague, Dr. Marie Suver.

100 frames per second with a median latency of 39 ms. For the purposes of tracking, the arena was backlit with an array of near-infrared (850 nm) LEDs. The cameras were equipped with long-pass filters (Hoya R-72) so that the camera images were not contaminated by the pattern that was displayed in visible wavelengths.

5.3.3 Experiment protocol

To automate the data collection I used a position and velocity dependent trigger near both ends of the wind tunnel. When a fly passed through either trigger volume while flying towards the opposite end of the wind tunnel, the visual display of sine gratings began to move randomly at one of eight specified temporal frequencies (0-16 Hz), in the same direction as the flies initial motion. The stimulus continued to move for 12 seconds, which was more than sufficient time for the fly to reach the opposite end of the tunnel. I restricted my analysis to the first 0.5 seconds of these trajectories. After each visual presentation there was a refractory period of 25 seconds during which time the trigger remained off regardless of the flies' behavior.

5.3.4 Trajectory reconstruction and analysis

All analyses of flight trajectories were done using custom software written in Python using the open-source software packages Scipy and Matplotlib, as well as Matlab. Each trajectory was treated as an independent sample, as the tracking software was not able to maintain fly identities over the extended period of the experiments. Therefore I could not test whether individual flies behaved consistently differently from other flies. I tested 24 individual flies for each of the parental controls, and 108 Tdc2-Gal4, UAS-Kir2.1-EGFP flies, resulting in 42-106 trajectories for each temporal frequency and each of the three genetic lines. Trajectories that came within 5 cm of the wind tunnel ceiling, or turned around part way through the trial, were left out of my analysis. Each trajectory was smoothed to remove digitization errors, and to estimate velocity and acceleration, using a simple forward/reverse, non-causal Kalman filter. I measured a 19 ± 1 ms delay between the computer generating the visual stimulus and the projector [55] (Fuller, 2011). Adjustments were made to the trajectories to align them with the true estimated timing of the stimulus post facto. Means and standard deviations of the accelerations in response to the visual stimulus were calculated by averaging the Kalman estimates of acceleration for the time range of 200-400 ms after the onset of visual motion.

All statistical comparisons were performed with the two-tailed Mann-Whitney U test. Statistically significant differences at the $p = 0.001$ level were calculated using the Bonferroni method for multiple hypothesis testing [130]. I used open-source python software for computing the statistically homogeneous groups, which is described in detail elsewhere [124]. This software is made freely available by Andrew Straw: <http://astraw.github.com/pairs2groups/>.

5.3.5 Immunohistochemistry and imaging

Note: this section was performed by a co-author and colleague, Dr. Marie Suver. She dissected brains in 4% paraformaldehyde in PBS and fixed for a total of 30 min. She then incubated them overnight at 4°C in a primary antibody solution containing 5%

normal goat serum in PBS-Tx, mouse anti-nc82 (1:10, DSHB) and rabbit anti-GFP (1:1000, Invitrogen). Brains were then incubated for two hours at room temperature in a secondary antibody solution containing 5% normal goat serum in PBS-Tx, goat anti-mouse Alexa Fluor 633 (1:250, Invitrogen) and goat anti-rabbit Alexa Fluor 488 (1:250, Invitrogen). She then mounted the brains in Vectashield and imaged them on a Leica SP5 II confocal microscope under 20x magnification and scanned at 1 μ m section intervals. She adjusted intensity and contrast for single channels for the entire image using ImageJ 1.46r.

5.4 Results

To study the functional role of octopamine-mediated modulation in the visual system, I examined the velocity control system of flies in free flight. Previous studies showed that flies use large field visual motion to regulate their flight speed about a fixed visual motion set point [53, 106]. Flies maintain a constant groundspeed over a large range of headwind velocities [40], demonstrating that this vision-based feedback system is quite robust. Assuming that the LPTC network plays a role in this behavior by estimating the groundspeed from optic flow, then the flies' ability to regulate forward flight speed should be compromised by silencing the octopamine neurons, thereby reducing or eliminating the gain boost in the network during flight. I tested this hypothesis by presenting regressive (back-to-front) visual motion to flying flies and recording their flight trajectories with a 3D tracking system (Fig. 5.2A, see Methods for additional details). The visual stimulus consisted of a moving contrast grating (spatial frequency = 4.2 m⁻¹) presented at a temporal frequency that varied between trials from 0 to 16 Hz (corresponding to 0-3.8 m s⁻¹ linear pattern velocity). To determine the role of octopamine in this behavior, I silenced putative octopaminergic neurons by expressing the inwardly-rectifying potassium channel Kir2.1 [77] using the Tdc2-Gal4 driver line (Fig. 5.2B).

Flies responded to the regressive visual motion by accelerating in the direction of the stimulus after an initial delay of approximately 100-150 ms. The magnitude of

the acceleration increased monotonically with the temporal frequency of the motion stimulus (Fig. 5.3). However, for temporal frequencies of 1 Hz (0.23 m s⁻¹ linear pattern velocity) and higher, flies with inactivated octopamine neurons showed significantly lower accelerations when compared to parental controls, whereas there were no significant differences in acceleration responses exhibited by the two parental controls at any temporal frequency with the exception of 8 Hz (Fig. 5.3D). Across all genetic lines, and all temporal frequencies, over 92% of the flies' velocity was in the direction of visual motion for the duration of the trials. Baseline flight speeds were just slightly higher for the flies with inactivated octopamine neurons (median 0.26 m s⁻¹) compared to parental controls (medians 0.23, 0.21 m s⁻¹) (Fig. 5.3E).

In addition to its role in gain modulation of the visual pathway, octopamine is known to influence a wide range of physiological effects. In my experiments, female flies with Kir2.1 expression in their octopamine neurons exhibited a swollen abdomen that is consistent with tonically low octopamine levels and the inability to oviposit [107]. To test whether the results could be explained by the additional mass of the gravid female flies, I repeated the experiment with males of each genetic line (Fig. 5.3B). The results for male flies were very similar, indicating that the results for females were not due to differences in abdominal mass.

Octopamine is also known to play a role in body tissues such as muscle [114]. Thus, another possibility is that the inhibition of the octopamine system could negatively affect the flies' ability to accelerate via its action on the motor system. The maximum acceleration I observed in the Tdc2-Gal4/UAS-Kir2.1 flies was 1.5 m s⁻², elicited by a temporal frequency of 16 Hz. This magnitude of acceleration would be sufficient to match the accelerations generated by control flies at temporal frequencies of 1-2 Hz. However, even at these low temporal frequencies, the flies with inactivated octopamine neurons accelerated significantly less than the parental controls. It is thus more parsimonious that the influence of octopamine acts on groundspeed regulation via the visually-mediated feedback control pathway, rather than through effects on the motor system. Although I cannot be sure that the behavioral effect of octopamine I observed is mediated by the LPTC network, the result is consistent with electrophysiological

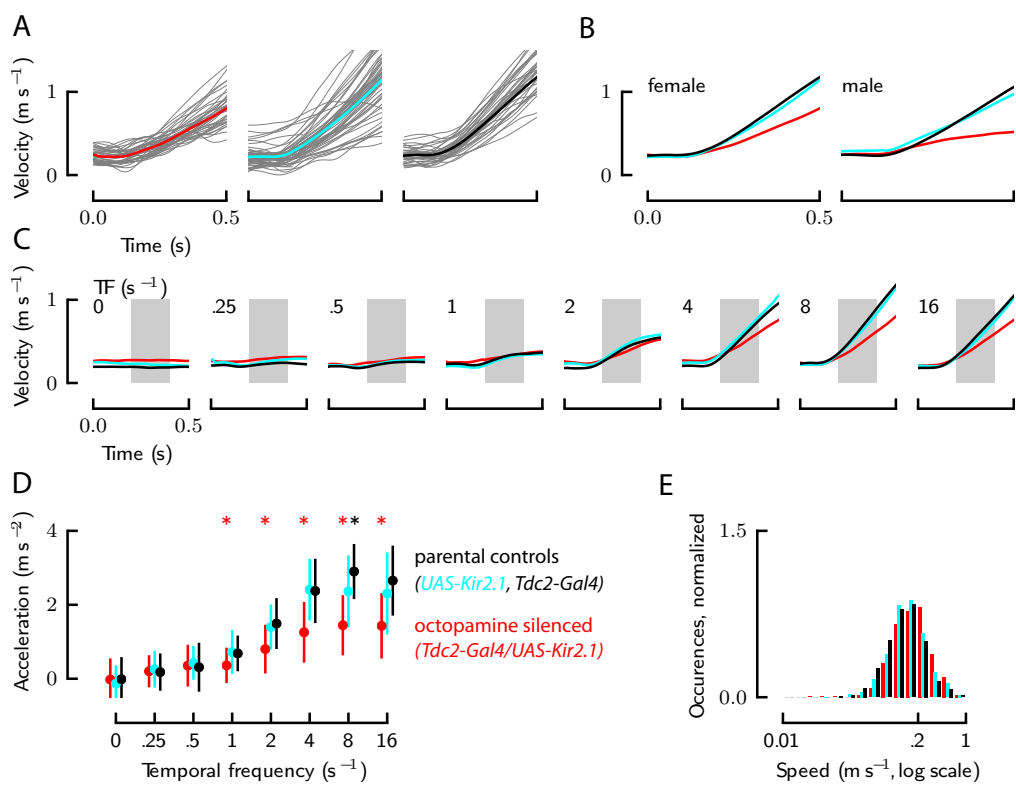


Figure 5.3: Octopamine null flies respond to regressive visual motion during free flight with lower accelerations than wild type flies. (A) Individual (gray) and mean (colored) velocity vs. time responses to regressive visual motion with a temporal frequency of 8 Hz for flies with inactivated octopamine neurons (red), and the parental controls (blue, black). I use consistent colors in subsequent panels. The visual motion started at $t=0$. (B) Mean velocity responses to an 8 Hz temporal frequency stimulus for males and females. (C) Mean velocity responses to regressive motion at different temporal frequencies (for each genotype of each temporal frequency, $42 \leq n \leq 106$). For the sake of graphical clarity, I chose not show the variance in these traces. The raw traces from panel A are a representative of the data from other temporal frequencies. (D) Acceleration responses to regressive motion vs. temporal frequency (mean \pm std dev). Accelerations for each trajectory were calculated as their mean acceleration during the 200-400 ms window after the stimulus was triggered, shown as a gray background in C. Significant differences in acceleration behavior between the flies with inactivated octopamine neurons (red) and the parental controls (blue, black) are indicated for each temporal frequency with red asterisks ($p < 0.001$). Significant differences between the two parental controls are indicated with a black asterisk ($p < 0.001$). (E) Histogram of steady state velocities for trajectories 1 ms prior to the onset of the regressive motion stimulus.

recordings that show that octopamine modulates the gain of neurons within this system [79, 96, 97, 142].

5.5 Discussion

To test the models presented in Fig. 5.1, I constructed a control theoretic model of the flight speed regulator based on, but not identical to, earlier models [53, 55, 125] (Fig. 5.4A). Although the focus of my analysis is on the visual processing and flight control, in order to build a functional model it is necessary to incorporate an accurate prediction of a fly's passive flight dynamics as well. Following the example of Fuller (2011), I modeled the passive flight dynamics as a simple unity gain low pass filter with a time constant of 170 ms, which takes into account the aerodynamic drag on a fly's flapping wings as well as inertial effects of its body mass. At high temporal frequencies, the flies exhibited maximum accelerations of up to 2.9 m s^{-2} , which is similar to the maximum of 2.5 m s^{-2} found by [125]. This saturation presumably represents a biomechanical limit, and I modeled it by incorporating an acceleration saturation, although I present the results both with and without the saturation. Prior experiments have demonstrated that flies exhibit an antenna-mediated reflex that increases active damping in response to fast changes in air speed [55]. Although this component only marginally affects the dynamics in my model, I include it for completeness. The effect of the antenna is nearly identical to the passive flight dynamics, resulting in a unity gain low pass filter with a time constant of 170 ms and a 20 ms delay.

Previous models of the groundspeed regulator have described the visual processing dynamics of the fly as a pure delay [55, 125]. However, this treatment ignores the temporal frequency dependence introduced by early motion processing in the optic lobes, a feature that is thought to emerge from the fundamental properties of motion detectors [48, 70]. To describe these filter dynamics more completely, I estimated this function based on published electrophysiological recordings of VS cells within the LPTC network [142] (Fig. 5.4B). The recordings also provide an accurate estimate

of the processing delay accrued before the LPFC network (36 ± 12 ms), which is independent of temporal frequency (Fig. 5.4C). Prior models of groundspeed regulation used a pure delay of 100 ms; to account for this total delay I included a 64 ms delay in the process dynamics. As in previously proposed models of flight speed control [55], I chose an integral type controller, which is necessary to explain a fly's ability to maintain constant flight speed in variable wind conditions [40] (David, 1982).

Although it consists of many components, my dynamic model contains only one free parameter, a proportional gain term, which I chose using Matlab's implementation of Lagarias' simplex search method [90]. To account for the differences in visual motion gain between control flies and those with inactivated octopamine neurons, I chose a proportional gain with values of 5.5 and 2.2, respectively. To match the steady state behavior of the parental controls, I used a preferred visual set point of 0.22 m s^{-1} for the H2 model (which is equivalent to the observed median baseline flight speed), and set point of 1.2 m s^{-1} for the H1 model (which is equivalent to the product of the gain, 5.5, and the observed median baseline flight speed of 0.22 m s^{-1}).

Both models are equally accurate at predicting the observed temporal frequency-dependent acceleration responses, with an octopamine-mediated increase in gain of 150% during flight (Fig. 5.4D). The maximum acceleration of the flies with inactivated octopamine neurons lies well below the biomechanical saturation limit of 2.5 m s^{-2} , and their behavior is well captured by the low pass filter derived directly from physiology experiments. The behavior of the parental control lines is best explained by a model that includes a biomechanical saturation. The two models differ substantially, however, in their steady state predictions (Fig. 5.5B). According to the H1 model, a reduction in gain due to the absence of octopamine would result in a 150% increase in steady state flight speed, whereas the H2 model predicts that there would be no change in steady state flight speed. In my experiments the flies with inactivated octopamine neurons showed only a marginal increase in flight speed of 18%, despite the fact that these flies accelerated more slowly in response to visual motion than control flies did. Of these two models, my results are best explained by H2.

H2a is mathematically indistinguishable from H2b because the gain is simply

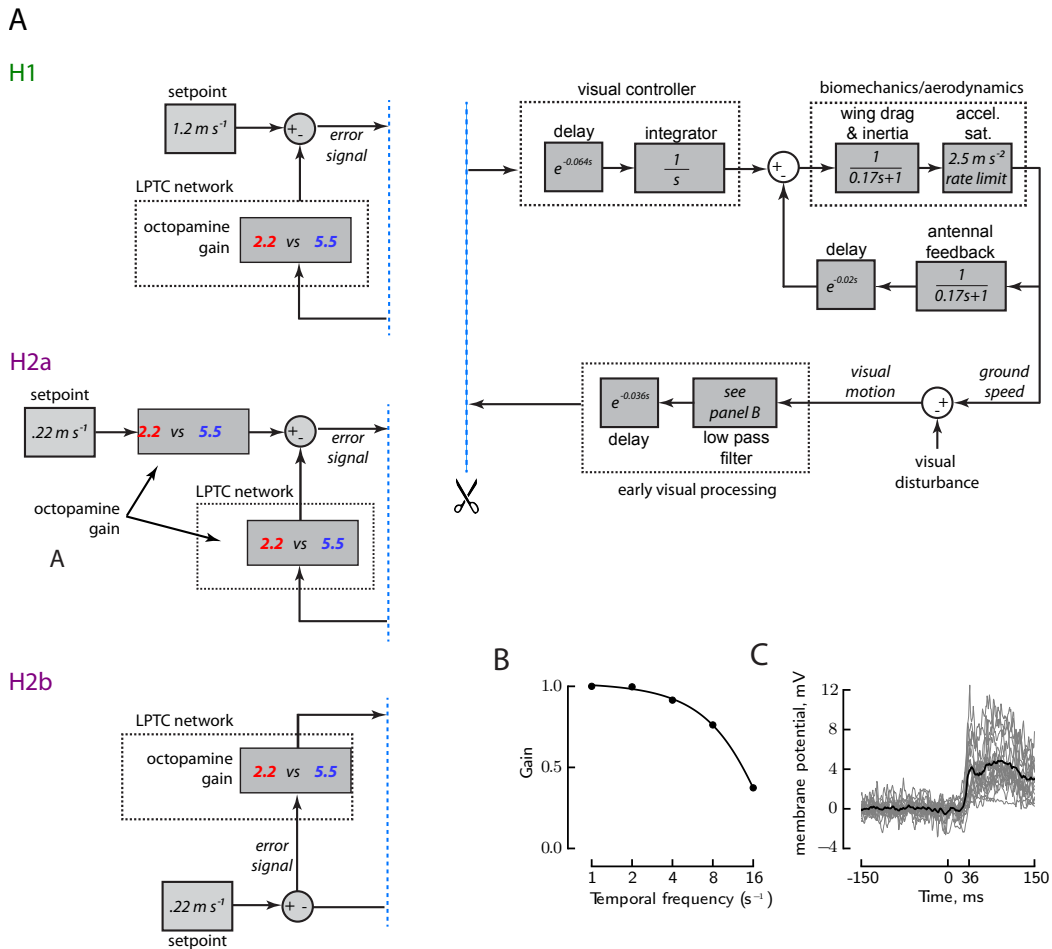


Figure 5.4: Preferred visual motion set point is modulated by identical changes in gain as the LPTC network (H2a), or it enters the visual sensory-motor cascade upstream of the large field motion network (H2b). (A) Block diagram showing the models under consideration. (B) Gain vs. temporal frequency curve used for the low pass filter in the visual sensory system of my model. The data points are drawn from the temporal frequency tuning curve given in Suver et al. (2012) Fig. 1D, which summarizes the responses of electrophysiological recordings of VS cells in response to vertical motion. The original data were scaled such that the gain at a temporal frequency of 1 Hz is 1. The line shows a 3rd order polynomial fit. Note that this results in a transfer function defined in the linear temporal frequency domain, rather than the oscillatory temporal frequency domain. In order to implement this type of filter in my control model, I calculate the gain based on the linear temporal frequency of the stimulus. (C) Baseline subtracted membrane potential of vertical system cells in response to a downward 8 Hz visual motion stimulus during flight; data repeated, and magnified, from Suver et al (2012). The gray traces show the mean responses each of 19 individual flies, and the bold trace shows the group mean.

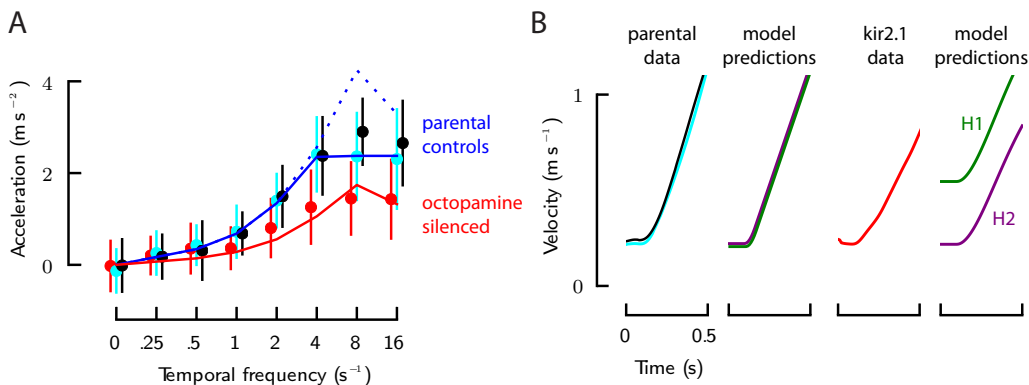


Figure 5.5: (A) Model predictions compared to the results from Fig. 5.3D. The solid blue line shows the model prediction for the parental controls (gain = 5.5) with the biomechanical saturation, whereas the dotted blue line shows the prediction without saturation. The solid red line shows the model prediction for the flies with inactivated octopamine neurons (gain = 2.2). (B) Model predictions compared to mean velocity vs time responses for parental controls (left) and flies with inactivated octopamine neurons (right). The data traces are repeated from Fig. 5.3A. Note that H2 is a better fit.

distributed to multiply the preceding blocks (the sensory signal, and the desired set point) independently. This does, however, have significant implications for the neural implementation. Any mismatch in the two gain blocks in H2a would result in a shift in baseline flight speed. This could explain the slight increase in baseline flight speed I observed. However, time dependent changes in the activity of octopamine neurons would result in unpredictable oscillations in flight speed. In H2b, on the other hand, time dependent changes in octopamine would serve to change the dynamics of the flies' response, rather than their baseline activity. Given these potentially detrimental effects, I suspect that the H2a is less likely than the H2b model, although there is currently insufficient data to distinguish between them.

The H2b model suggests that the preferred flight speed set point may enter the visual sensory-motor cascade upstream of the LPTC network, which implies that their functional role is not restricted to that of a set of matched filters for detecting self-motion. Although the LPTC network is considered as part of the visual system, they synapse directly onto descending interneurons that project to neck, leg, and wing motor centers in the thoracic ganglia [69, 137, 138]. Thus, the major output

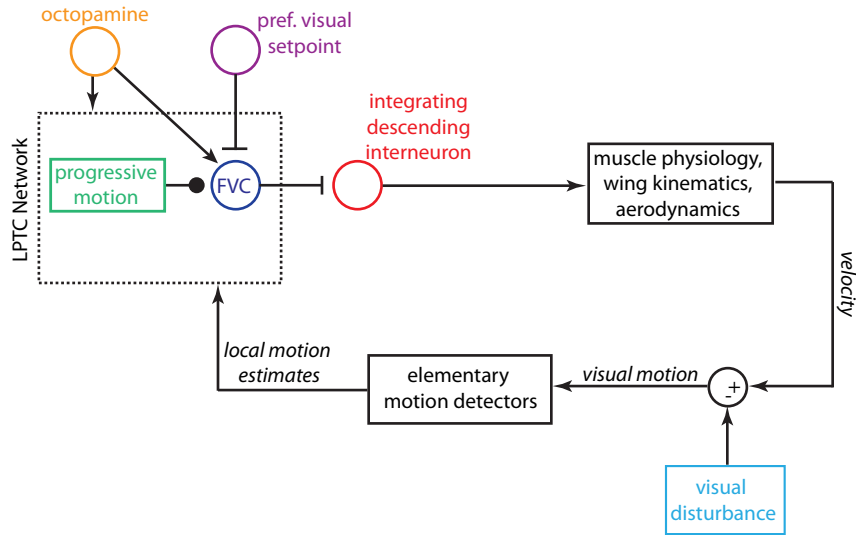


Figure 5.6: Neural circuit summarizing the block diagram for H2b from Fig. 5.4A in the context of a fly’s hypothetical neuroanatomy. The simplest possible implementation is that a neuron within the LPTC network acts as a forward velocity controller (FVC) neuron, which either receives a signed translatory motion input, or calculates it itself, and subtracts the preferred visual set point. The output of this cell is then modulated by octopamine. The FVC cell projects to descending interneurons, which integrate the signal before activating motor-neurons in the thoracic ganglion.

elements of the LPTC network are only 1-2 synapses away from motor neurons, and it is reasonable to propose that control signals, such as those determining the forward velocity set point, might enter the sensory-motor path within the LPTC network rather than downstream from it.

Octopamine neurons, however, project throughout the brain (Fig. 5.2B), resulting in a vastly more complex network of modulation than accounted for by either the H1 or H2 models. In addition to its role in the LPTC network, octopamine has been implicated in modulation of contrast sensitivity and motion adaptation in visual processing presynaptic to the LPTC network, presumably within the elementary motion detection circuits [42]. Since contrast has been shown to influence baseline flight speed [140], octopamine could have an indirect influence on a flies’ preferred visual set point, which could also explain the slight increase in baseline flight speed that I observed.

To place the H2b model within in the context of the flies' known neural architecture, I constructed a simple putative neural circuit (Fig. 5.6). The key conclusion from the experimental and modeling results is that the octopamine gain is likely applied after the flies' estimate of translational motion is compared to its preferred visual set point. I hypothesize that this comparison, and the subsequent octopaminergic gain modulation, is mediated by a neuron within the lobula plate, which I call the Forward Velocity Controller (FVC). Previous studies suggest that the flies' preferred visual set point may be a function of visual cues such as total luminance, contrast, and distance to objects, as well as olfactory cues. Thus, future electrophysiological studies should investigate the influence of such sensory cues on the response properties of visual interneurons within the LPTC network.

Chapter 6

Conclusion

In this thesis, I explore the question of how the complex behavior of foraging emerges from the numerically simple brain of the fruit fly, *Drosophila melanogaster*. In Chapters 2-4, I described in detail six sensory-motor modules in the fruit fly that are sufficient to produce robust foraging behavior: (1) surging upwind upon encountering an attractive odor; (2) casting crosswind upon plume loss; (3) exploring small visual features after encountering an odor; (4) maintaining constant optic flow during landing maneuvers; (5) turning away from obstacles if the retinal expanse (or possibly the optic flow) reaches a threshold; and (6) leg extension triggered by a threshold retinal size (Chapter 3), or an internal control signal (Chapter 4). Finally, in Chapter 5, I demonstrated the behavioral effect that the neuromodulator Octopamine has on a fly's flight control system. By extension, it is likely that neuromodulators play an important role in tuning the interaction between many sensory-motor modules. For example, neuromodulators may be responsible for governing the decision of whether to land or not, based on olfactory stimuli. Given the apparent complexity of plume tracking behavior that arises from such a small set of behaviors, the complexity that could arise from hundreds to thousands, or perhaps millions, of such modules, especially with added layers of neuromodulation, is beyond our imagination - perhaps, it is even what makes our imagination possible.

Many more sensory-motor modules that guide the behavior of flies have been described previously (for example [162, 28]), and understanding how they all interact with one another is a critical next step. In a recent article, Dickinson reviews these

modules, and describes how a stigmergic¹ iteration of these behaviors could give rise to long distance dispersal [44]. Many of the sensory-motor modules appear to be independent, as evidenced by inconsistent sequential ordering. For example, in Chapter 3, I showed that flies sometimes begin decelerating prior to turning towards an object, and sometimes vice versa. Furthermore, sometimes flies land, without having followed an odor plume [150], and they show attraction to visual features without experiencing odors [67]. The sensory-motor modules of the fly may likely be wired in parallel within the brain (while sharing certain components, such as basic visual processing), which would reduce sensory-motor delays and energy costs in information transmission, enhance evolvability, and reduce developmental complexity. Thanks to this apparent modularity, it has been possible to isolate these behaviors and study each one independently, which has indeed been the focus of neuroscientists for the past century. We have now come to a point, however, where we are flooded with data on individual behaviors, often in highly reduced preparations, with little understanding of the greater context. In order to understand how complex behaviors emerge, it is imperative that we discover how each of these sensory-motor modules interact with one another, and the effect that neuromodulators have on these interactions. Many of these interactions, such as those described in this thesis, do not become apparent until the impact they have on a freely moving animals' sensory system is considered. For example, in Chapter 4, the leg extension trigger hypothesis I proposed only emerged as a possibility because of the trajectory that results from a simple optic flow regulator. As another example, the importance of casting, surging, and visual attraction during odor plume tracking do not become clear until the moving fly is modeled in a realistic environment, and the impact of 200-600 ms delays becomes apparent.

Therefore, in order to understand how complex behaviors emerge from independent sensory-motor modules, I believe that neuroscientists, roboticists, and computer scientists, must come together to build an extensible model that incorporates each published sensory-motor module into a comprehensive embodied system. Such a

¹Stigmergy: the emergence of complex behavioral cascades through the iteration of independent sensory-motor reflexes. See the Introduction for more details.

model would have several clear and immediate uses: it would (1) help organize our growing knowledge into a comprehensive, and continuously updated, summary; (2) describe the emergent properties of independent sensory-motor modules; (3) elucidate gaps in our knowledge by forcing scientists to create a working model, and thereby producing testable hypotheses. With a carefully constructed framework, such a modeling environment would serve the neuroscientist community as a crowd sourced, self-validating, and continuously updating neuroethological review, while automatically generating well defined and testable hypotheses to help direct research towards gaps in our knowledge. Simultaneously, the approach would give roboticists a chance to test this bio-plausible brain structure on physical systems, revealing its potential advantages and disadvantages compared to modern control approaches, and producing testable hypotheses for biologists to consider.

Fortunately, thanks to the apparent independence and presumably parallel structure of these modules, this task is not as daunting as it may seem. By building a computational framework consisting of carefully considered protocols it will be possible for independent researchers to write simulations of the sensory-motor modules they observe in such a way that they are all compatible with one another. The basis for such frameworks already exists, for example, the Robot Operating System (ROS) makes it possible to write simple, parallel, and asynchronous modules that can communicate with one another indirectly through a custom messaging protocol - all on a single desktop computer. Through software version control, it will be possible to keep track of an indefinite number of hypotheses for how the modules are constructed. For example, the third sensory-motor module I describe in a fly's foraging behavior—odor-induced object saliency—could be implemented in one of several ways. Perhaps the visual signal itself is modulated, changing the perceptual sensory experience of the fly at the level of the lobula plate, or even earlier (similar to the effect that Octopamine has on wide field visual motion, as described in Chapter 5). Or perhaps there exists a circuit that determines visual features, and the operation of this circuit may be modulated to pick out features more readily. Or downstream circuits that trigger attending to visual features is modulated. These are examples of

testable hypotheses that could give rise to the behavior we observe, and each of them could be programmed as alternatives within the model's framework such that future theorists can test the implications of these models, and experimentalists can devise experiments to test between them given an understanding of the implications the different approaches would have on the structure and function of the overall network. This type of simulation would not, initially, be directed at the level of individual neurons. Instead, it represents an abstraction that is far simpler to explore, is immediately within our grasp, and has immediate practical applications to engineering and neuroethology.

The principle of using independent and layered sensory motor modules to design control architectures for robotic systems has not gone unnoticed by engineers. The possibilities of such an architecture were first elegantly described by Valentino Braitenberg in a series of thought experiments in his book *Vehicles* [19]. Turning these thought experiments into reality, however, has proven to be a challenge. Rodney Brooks took on this challenge twenty years ago, when he proposed his "subsumption architecture" for designing artificially intelligent robots [22, 23, 24]. Though initially promising, his approach has not been as successful as "traditional" control approaches based on central processing and full state estimation. This is largely due to the serial architecture of computer chips and significant advances in processing speeds. Modern computer chips are sufficiently fast and efficient to use central processing for robotic applications such as the self driving cars of the DARPA Grand Challenge, and Boston Dynamics' famous robots such as Big Dog. Furthermore, designing control architectures from independent behavioral modules is a non-intuitive task, generally requiring countless iterations of trial and error (sometimes aided through evolutionary design approaches). This approach is both time consuming, and results in poorly defined behavioral responses that require exhaustive testing to fully characterize. As a result, only relatively simple modern robots, such as the Roomba Vacuum Cleaner, are governed by a collection of independent reflexes, rather than a central state estimation based controller. By building a comprehensive model of the behavioral modules that have been discovered in flies (and other animals), and analyzing the interactions be-

tween modules, we will build a library of useful reflexes that can be used in robotic systems, potentially making Brooks' dream a reality.

The most intriguing application of the stigmergic approach is that by following simple control laws, it is possible to indirectly simplify challenging estimation problems, a principle which I introduced in the introduction as idiokinometry. For example, in Chapter 4, I describe a putative mechanism for distance estimation which reduces the problem to a single variable by following a trajectory of constant optic flow. A second example was given in Chapter 2, where I described how flies solve the problem of determining the crosswind direction without needing to measure the magnitude or direction of the wind. First, flies find the upwind direction by minimizing visual slip. Then, to fly crosswind, a fly can turn by a known amount (ϕ), and maintain a new optic flow axis equal to $(90^\circ - \phi)$. So long as the wind direction does not shift, the fly will maintain a crosswind trajectory. A third example is implied in Chapter 3, where I show that the majority of flies' changes in heading are accomplished through ballistic saccadic maneuvers. This vastly simplifies visual processing by removing the rotational component of motion. The key point is that rather than determining the control necessary to construct an explicit observer to estimate unobservable states, the controls are chosen to *simplify* the problem such that these observations are implicitly determined. This is a powerful principle, which has been subject to relatively little research in both engineering and biology. It seems likely, however, that this principle is pervasive throughout biological systems, and may prove to be a key component in building robust and efficient engineered systems with limited computational resources.

Bibliography

- [1] Atiye Alaeddini and Kristi A Morgansen. Autonomous State Estimation Using Optic Flow Sensing. In *American Control Conference*, 2013.
- [2] Mark L Andermann, Aaron M Kerlin, Demetris K Roumis, Lindsey L Glickfeld, and R Clay Reid. Functional specialization of mouse higher visual cortical areas. *Neuron*, 72(6):1025–39, December 2011.
- [3] T C Baker. Upwind flight and casting flight: complimentary phasic and tonic systems used for location of sex pheromone sources by male moths. In K B Doving, editor, *International Symposium on Olfaction and Taste X*, pages 18–25, Oslo, 1990. Graphic Communications Systems.
- [4] T C Baker and L P S Kuenen. Pheromone source location by flying moths: a supplementary non-anemotactic mechanism. *Science*, 216(4544):424–427, April 1982.
- [5] Thomas C Baker, Bill S Hansson, Christer Lofstedt, and Jan Lofqvist. Adaptation of antennal neurons in moths is associated with cessation of pheromone-mediated upwind flight. *Proceedings of the National Academy of Sciences*, 85:9826–9830, 1988.
- [6] W M Barrows. The reactions of the Pomace fly, *Drosophila ampelophila* loew, to odorous substances. *Journal of Experimental Zoology*, 4(4):515–537, 1907.
- [7] Javad Behboodian. On the Modes of a Mixture of Two Normal Distributions. *Technometrics*, 12(1):131–139, April 1970.

- [8] John A Bender and Michael H Dickinson. A comparison of visual and haltere-mediated feedback in the control of body saccades in *Drosophila melanogaster*. *Journal of Experimental Biology*, 209(Pt 23):4597–606, December 2006.
- [9] John A Bender and Michael H Dickinson. Visual stimulation of saccades in magnetically tethered *Drosophila*. *Journal of Experimental Biology*, 209(Pt 16):3170–82, August 2006.
- [10] H C Berg. *E. coli in Motion*, 2004.
- [11] Paul J. Besl. Active, optical range imaging sensors. *Machine Vision and Applications*, 1(2):127–152, June 1988.
- [12] E Bizzi, V C K Cheung, a D’Avella, P Saltiel, and M Tresch. Combining modules for movement. *Brain research reviews*, 57(1):125–33, January 2008.
- [13] Norbert Boeddeker and Martin Egelhaaf. A single control system for smooth and saccade-like pursuit in blowflies. *Journal of Experimental Biology*, 208(Pt 8):1563–72, April 2005.
- [14] Norbert Boeddeker, Roland Kern, and Martin Egelhaaf. Chasing a dummy target: smooth pursuit and velocity control in male blowflies. *Proceedings of the Royal Society B: Biological Sciences*, 270(1513):393–9, March 2003.
- [15] A. Borst. Time course of the Houseflies’ landing response. *Biological Cybernetics*, 54(6):379–383, September 1986.
- [16] Alexander Borst. How Do Flies Land? *BioScience*, 40(4):292–299, 1990.
- [17] Alexander Borst and Susanne Bahde. What Kind of Movement Detector is Triggering the Landing Response of the Housefly. *Biological Cybernetics*, 55:59–69, 1986.
- [18] Alexander Borst, Juergen Haag, and Dierk F Reiff. Fly motion vision. *Annual Review of Neuroscience*, 33:49–70, January 2010.

- [19] Valentino Braitenberg. *Vehicles: Experiments in synthetic psychology*. MIT press, 1986.
- [20] Valentino Braitenberg and Cloe Taddei Ferretti. Landing reaction of *musca domestica* induced by visual stimuli. *Die Naturwissenschaften*, 53(6):155–155, 1966.
- [21] Russell S A Brinkworth and David C O’Carroll. Robust models for optic flow coding in natural scenes inspired by insect biology. *PLoS Computational Biology*, 5(11):e1000555, November 2009.
- [22] R a Brooks. New approaches to robotics. *Science*, 253(5025):1227–32, September 1991.
- [23] Rodney A Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23, 1986.
- [24] Rodney A Brooks. A robot that walks; emergent behaviors from a carefully evolved network. *Neural Computation*, 1:253–262, 1989.
- [25] T Graham Brown. The Intrinsic Factors in the Act of Progression in the Mammal. *Proceedings of the Royal Society B: Biological Sciences*, 84(572):308–319, 1911.
- [26] Seth A Budick and Michael H Dickinson. Free-flight responses of *Drosophila melanogaster* to attractive odors. *Journal of Experimental Biology*, 209(Pt 15):3001–17, August 2006.
- [27] Sebastian Busch, Mareike Selcho, Kei Ito, and Hiromu Tanimoto. A map of octopaminergic neurons in the *Drosophila* brain. *Journal of Comparative Neurology*, 513(6):643–67, April 2009.
- [28] Gwyneth Card and Michael H Dickinson. Visually mediated motor planning in the escape response of *Drosophila*. *Current Biology*, 18(17):1300–7, September 2008.

- [29] Ring T Cardé and Mark A Willis. Navigational strategies used by insects to find distant, wind-borne sources of odor. *Journal of Chemical Ecology*, 34(7):854–66, July 2008.
- [30] James Carpenter and John Bithell. Bootstrap confidence intervals: when, which, what? A practical guide for medical statisticians. *Statistics in Medicine*, 19(August 1999):1141–1164, 2000.
- [31] B A Cartwright and T S Collett. How honey-bees know their distance from a near-by visual landmark. *Journal of Experimental Biology*, 82:367–372, 1979.
- [32] J S Chahl, M V Srinivasan, and S W Zhang. Landing Strategies in Honeybees and Applications to Uninhabited Airborne Vehicles. *The International Journal of Robotics Research*, 23(2):101–110, February 2004.
- [33] M Eugenia Chiappe, Johannes D Seelig, Michael B Reiser, and Vivek Jayaraman. Walking modulates speed sensitivity in *Drosophila* motion vision. *Current Biology*, 20(16):1470–5, August 2010.
- [34] Dawnis M Chow and Mark A Frye. Context-dependent olfactory enhancement of optomotor flight control in *Drosophila*. *Journal of Experimental Biology*, 211(Pt 15):2478–85, August 2008.
- [35] John C Coggshall. The landing response and visual processing in the milkweed bug, *oncopeltus fasciatus*. *Journal of Experimental Biology*, 57:401–413, 1972.
- [36] T S Collett. Peering - a locust behavior pattern for obtaining motion parallax information. *Journal of Experimental Biology*, 76:237–241, 1978.
- [37] T S Collett and L I K Harkness. Depth vision in animals. In D J Ingle, M A Goodale, and R J W Mansfield, editors, *Analysis of Visual Behavior*, pages 111–176. MIT Press, Cambridge, Massachusetts, 1982.
- [38] T S Collett and M F Land. Visual control of flight behavior in the hoverfly, *syritta pipiens* L. *Journal of Comparative Physiology*, 66:1–66, 1975.

- [39] John L Crassidis and John L Junkins. *Optimal Estimation of Dynamic Systems*. Chapman & Hall, London, second ed. edition, 2012.
- [40] C T David. Compensation for Height in the Control of Groundspeed by *Drosophila* in a New, 'Barber's Pole' Wind Tunnel. *Journal of Comparative Physiology*, 147:485–493, 1982.
- [41] C T David, J S Kennedy, and A R Ludlow. Finding of a sex pheromone source by gypsy moths released in the field. *Nature*, 303:804–806, 1983.
- [42] Roel de Haan, Yu-Jen Lee, and Karin Nordström. Octopaminergic modulation of contrast sensitivity. *Frontiers in Integrative Neuroscience*, 6(August):55, January 2012.
- [43] Michael H Dickinson. Haltere-mediated equilibrium reflexes of the fruit fly, *Drosophila melanogaster*. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 354(1385):903–916, May 1999.
- [44] Michael H Dickinson. Death Valley, *Drosophila*, and the Devonian Toolkit. *Annual Review of Entomology*, 59(1):131025101253002, February 2014.
- [45] Brian J Duistermars, Dawnis M Chow, and Mark A Frye. Flies require bilateral sensory input to track odor gradients in flight. *Current Biology*, 19(15):1301–1307, 2009.
- [46] Hendrik Eckert. On the landing response of the blowfly, *Calliphora erythrocephala*. *Biological Cybernetics*, 47(2):119–130, June 1983.
- [47] B Efron. Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, 7(1):1–26, 1979.
- [48] M Egelhaaf, A Borst, and W Reichardt. Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system. *Journal of the Optical Society of America*, 6(7):1070–87, July 1989.

- [49] C Evangelista, P Kraft, M Dacke, J Reinhard, and M V Srinivasan. The moment before touchdown: landing manoeuvres of the honeybee *Apis mellifera*. *Journal of Experimental Biology*, 213(Pt 2):262–70, January 2010.
- [50] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [51] R A Fisher. On the Interpretation of χ^2 from Contingency Tables , and the Calculation of P. *Journal of the Royal Statistical Society*, 85(1):87–94, 1922.
- [52] R A Fisher. *The design of experiments*. Oliver & Boyd, Oxford, England, 1935.
- [53] Steven N Fry, Nicola Rohrseitz, Andrew D Straw, and Michael H Dickinson. Visual control of flight speed in *Drosophila melanogaster*. *Journal of Experimental Biology*, 212(Pt 8):1120–30, April 2009.
- [54] Mark A Frye, Michael Tarsitano, and Michael H Dickinson. Odor localization requires visual feedback during free flight in *Drosophila melanogaster*. *Journal of Experimental Biology*, 206(5):843–855, March 2003.
- [55] Sawyer B Fuller. *Visual autocorrelators and antenna-mediated airspeed feedback in the control of flight dynamics in fruit flies and robotics*. PhD thesis, California Institute of Technology, 2011.
- [56] F Gabbiani, H G Krapp, and G Laurent. Computation of object approach by a wide-field, motion-sensitive neuron. *Journal of Neuroscience*, 19(3):1122–41, February 1999.
- [57] F Gabbiani, C Mo, and G Laurent. Invariance of angular threshold computation in a wide-field looming-sensitive neuron. *Journal of Neuroscience*, 21(1):314–29, January 2001.

- [58] Quentin Gaudry, Elizabeth J Hong, Jamey Kain, Benjamin L de Bivort, and Rachel I Wilson. Asymmetric neurotransmitter release enables rapid odour lateralization in *Drosophila*. *Nature*, 493(7432):424–8, January 2013.
- [59] Quentin Gaudry, Katherine I Nagel, and Rachel I Wilson. Smelling on the fly: sensory cues and strategies for olfactory navigation in *Drosophila*. *Current Opinion in Neurobiology*, 22(2):216–22, April 2012.
- [60] L Gautier, W Moreira, and G R Warnes. A simple and efficient access to R from Python., 2001.
- [61] Gad Geiger and Dick R Nässel. Visual orientation behaviour of flies after selective laser beam ablation of interneurons. *Nature*, 293(1):398–399, 1981.
- [62] James J Gibson. *The Perception of the Visual World*. Riverside Press, Cambridge, Massachusetts, 1950.
- [63] James J Gibson. Visually controlled locomotion and visual orientation in animals. *British journal of psychology*, 49(3):182–94, August 1958.
- [64] Alex Gomez-Marin and Matthieu Louis. Active sensation during orientation behavior in the *Drosophila* larva: more sense than luck. *Current Opinion in Neurobiology*, 22(2):208–15, April 2012.
- [65] Lesley J Goodman. The landing responses of insects I. The landing response of the fly, *Lucilia sericata*, and other Calliphoridae. *Journal of Experimental Biology*, 37:854–878, 1960.
- [66] Karl G Götz. Flight Control in *Drosophila* by Visual Perception of Motion. *Kybernetik*, IV(6):199–208, 1968.
- [67] Karl G Götz. Course-control, metabolism and wing interference during ultralong tethered flight in *Drosophila melanogaster*. *Journal of Experimental Biology*, 46:35–46, 1987.

- [68] Pierre-Paul Grassé. La theorie de la stigmergie: essai d'interpretation du comportement des termites constructeurs. *Insectes Sociaux*, 6:41–81, 1959.
- [69] W Gronenberg and Nicholas J Strausfeld. Descending neurons supplying the neck and flight motor of Diptera: Physiological and anatomical characteristics. *Journal of Comparative Neurology*, 302:973–991, 1990.
- [70] B Hassenstein and Werner Reichardt. Systemtheoretische Analyse der Zeit-, Reihenfolgen- und Vorzeichenauswertung bei der Bewegungsperzeption des Russelkafers *Chlorophanus*. *Zeitschrift Fur Naturforschung*, 11:513–524, 1956.
- [71] K Hausen and C Wehrhahn. Microsurgical Lesion of Horizontal Cells Changes Optomotor Yaw Responses in the Blowfly *Calliphora erythrocephala*. *Proceedings of the Royal Society B: Biological Sciences*, 219(1215):211–216, September 1983.
- [72] M Heisenberg and R Wolf. *Vision in Drosophila. Genetics of microbehaviour*. Springer Verlag, 1984.
- [73] Martin Heisenberg and Reinhard Wolf. On the fine structure of yaw torque in visual flight orientation of *drosophila melanogaster*. *Journal of Comparative Physiology*, 130:113–130, 1978.
- [74] Martin Heisenberg, R. Wonneberger, and Reinhard Wolf. Optomotor-blind H31 - A *Drosophila* mutan of the lobula plate giant neurons. *Journal of Comparative Physiology*, 124:287–296, 1978.
- [75] Brian T Hinson and Kristi A Morgansen. Flowfield Estimation in the Wake of a Pitching and Heaving Airfoil. In *Proceedings of the American Control Conference*, 2012.
- [76] M H Holmqvist and M V Srinivasan. A visually evoked escape response of the housefly. *Journal of Comparative Physiology A*, 169(4), October 1991.

- [77] D C Johns, R Marx, R E Mains, B O'Rourke, and E Marbán. Inducible genetic suppression of neuronal excitability. *Journal of Neuroscience*, 19(5):1691–7, March 1999.
- [78] E Jones, T Oliphant, P Peterson, and Others. SciPy: open source scientific tools for Python., 2001.
- [79] Sarah Nicola Jung, Alexander Borst, and Juergen Haag. Flight activity alters velocity tuning of fly motion-sensitive neurons. *Journal of Neuroscience*, 31(25):9231–7, June 2011.
- [80] K Karmeier, J H van Hateren, R Kern, and M Egelhaaf. Encoding of naturalistic optic flow by a population of blowfly motion-sensitive neurons. *Journal of Neurophysiology*, 96(3):1602–14, September 2006.
- [81] J S Kennedy. Zigzagging and casting as a programmed response to wind-borne odour: a review. *Physiological Entomology*, 8(2):109–120, 1983.
- [82] John S Kennedy. The visual responses of flying mosquitoes. *Proceedings of the Zoological Society of London*, 109(4):221–242, 1939.
- [83] John S Kennedy and D Marsh. Pheromone-regulated anemotaxis in flying moths. *Science*, 184(4140):999–1001, 1974.
- [84] Roland Kern, J H van Hateren, Christian Michaelis, Jens Peter Lindemann, and Martin Egelhaaf. Function of a fly motion-sensitive neuron matches eye movements during free flight. *PLoS Biology*, 3(6):e171, June 2005.
- [85] M A R Koehl, J R Koseff, J P Crimaldi, M G McCay, T Cooper, M B Wiley, and P A Moore. Lobster sniffing: antennule design and hydrodynamic filtering of information in an odor plume. *Science*, 294(5548):1948–51, November 2001.
- [86] Jan J Koenderink. Optic flow. *Vision Research*, 26(I):161–180, 1986.

- [87] Holger G Krapp and R Hengstenberg. Estimation of self-motion by optic flow processing in single visual interneurons. *Nature*, 384(6608):463–6, December 1996.
- [88] Arthur J. Krener and Kayo Ide. Measures of unobservability. In *Proceedings of the 48th IEEE Conference on Decision and Control*, pages 6401–6406. Ieee, December 2009.
- [89] L P S Kuenen and R T Cardé. Strategies for recontacting a lost pheromone plume: casting and upwind flight in the male gypsy moth. *Physiological Entomology*, 19(1):15–29, 1994.
- [90] Jeffrey C Lagarias, James A Reeds, Margaret H Wright, and Paul E Wright. Convergence properties of the nelder mead simplex method in low dimensions. *Society for Industrial and Applied Mathematics*, 9(1):112–147, 1998.
- [91] F Land. Visual tracking and pursuit: humans and arthropods compared. *Journal of Insect Physiology*, 38(12):939–951, 1992.
- [92] Michael F. Land and Dan-Eric Nilsson. *Animal Eyes*. Oxford University Press, Oxford, 2nd edition, 2012.
- [93] M Lehrer. Small-scale navigation in the honeybee: active acquisition of visual information about the goal. *Journal of Experimental Biology*, 199(Pt 1):253–61, January 1996.
- [94] M Lehrer, M V Srinivasan, and S W Zhang. Motion cues provide the bee’s visual world with a third dimension. *Nature*, 332:356–357, 1988.
- [95] Mario Livio. *The golden ratio: The story of phi, the world’s most astonishing number*. Random House Digital, Inc., 2008.
- [96] Kit D Longden and Holger G Krapp. State-dependent performance of optic-flow processing interneurons. *Journal of Neurophysiology*, 102(6):3606–18, December 2009.

- [97] Kit D Longden and Holger G Krapp. Octopaminergic modulation of temporal frequency coding in an identified optic flow-processing interneuron. *Frontiers in Systems Neuroscience*, 4(November):153, January 2010.
- [98] C Loudon and M A R Koehl. Sniffing by a silkworm moth: wing fanning enhances air penetration through and pheromone interception by antennae. *Journal of Experimental Biology*, 203:2977–2990, 2000.
- [99] Sean P Macevoy, Timothy D Hanks, and Michael A Paradiso. Macaque V1 Activity During Natural Vision : Effects of Natural Scenes and Saccades. *Journal of Neurophysiology*, 99:460–472, 2008.
- [100] Gaby Maimon, Andrew D Straw, and Michael H Dickinson. A simple vision-based algorithm for decision making in flying *Drosophila*. *Current Biology*, 18(6):464–70, March 2008.
- [101] Gaby Maimon, Andrew D Straw, and Michael H Dickinson. Active flight increases the gain of visual motion processing in *Drosophila*. *Nature Neuroscience*, 13(3):393–9, March 2010.
- [102] Akira Mamiya, Andrew D Straw, Egill To, and Michael H Dickinson. Active and passive antennal movements during visually guided steering in flying *Drosophila*. *Journal of Neuroscience*, 31(18):6900 – 6914, 2011.
- [103] Eve Marder and Dirk Bucher. Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. *Annual Review of Physiology*, 69:291–316, January 2007.
- [104] Susana Martinez-Conde, Rich Krauzlis, Joel M Miller, Concetta Morrone, David Williams, and Eileen Kowler. Eye movements and the perception of a clear and stable visual world. *Journal of Vision*, 8(14):85013, 2008.
- [105] M Mayer, K Vogtmann, B Bausenwein, R Wolf, and Martin Heisenberg. Flight control during free yaw turns in *Drosophila melanogaster*. *Journal of Comparative Physiology A*, 163(3):389–399, 1988.

- [106] V Medici and S N Fry. Embodied linearity of speed control in *Drosophila melanogaster*. *Journal of the Royal Society, Interface*, 9(77):3260–7, December 2012.
- [107] Maria Monastirioti, Charles E. Linn, and Kalpana White. Characterization of *Drosophila* tyramine beta-hydroxylase isolation of mutant flies lacking octopamine. *Journal of Neuroscience*, 16(12):3900–3911, 1996.
- [108] J Moran and R Desimone. Selective attention gates visual processing in the extrastriate cortex. *Science*, 229:782–784, 1985.
- [109] J Murlis, J S Elkinton, and R T Carde. Odor plumes and how insects use them. *Annual Review of Entomology*, 37(86):505–532, 1992.
- [110] J Murlis, M A Willis, and R T Carde. Spatial and temporal structures of pheromone plumes in fields and forests. *Physiological Entomology*, 25:211–222, 2000.
- [111] Cristopher M Niell and Michael P Stryker. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron*, 65(4):472–9, February 2010.
- [112] H Nijmeijer and A J van der Schaft. *Nonlinear Dynamical Control Systems*. Springer, New York, 1990.
- [113] Tyler a Ofstad, Charles S Zuker, and Michael B Reiser. Visual place learning in *Drosophila melanogaster*. *Nature*, 474(7350):204–7, January 2011.
- [114] Ian Orchard, Jan-Marino Ramirez, and Angela B Lange. A multifunctional role for octopamine in locust flight. *Annual Review of Entomology*, 38:227–249, 1993.
- [115] A P Pentland. A new sense for depth of field. *IEEE transactions on pattern analysis and machine intelligence*, 9(4):523–31, April 1987.
- [116] Jess Porter, Brent Craven, Rehan M Khan, Shao-Ju Chang, Irene Kang, Benjamin Judkewitz, Benjamin Judkewicz, Jason Volpe, Gary Settles, and Noam

- Sobel. Mechanisms of scent-tracking in humans. *Nature Neuroscience*, 10(1):27–9, January 2007.
- [117] Michael Poteser, Maria Anna Pabst, and Karl Kral. Proprioceptive contribution to distance estimation by motion parallax in a praying mantid. *Journal of Experimental Biology*, 201:1483–91, May 1998.
- [118] R Core Development Team. R: A Language and Environment for Statistical Computing, 2011.
- [119] M R Reed. The olfactory reactions of *Drosophila melanogaster* Meigen to the products of fermenting banana. *Physiological Zoology*, 11(3):317–325, 1938.
- [120] Werner Reichardt and Tomaso Poggio. Visual control of orientation behavior in the fly: Part 1. A quantitative analysis. *Quarterly Reviews of Biophysics*, 9(3):311–375, 1976.
- [121] Michael B Reiser and Michael H Dickinson. A modular display system for insect behavioral neuroscience. *Journal of Neuroscience Methods*, 167(2):127–39, January 2008.
- [122] Jeffrey A Riffell, Leif Abrell, and John G Hildebrand. Physical processes and real-time chemical measurement of the insect olfactory environment. *Journal of Chemical Ecology*, 34(7):837–53, July 2008.
- [123] F Claire Rind. A Directionally Selective Motion-Detecting Neurone in the Brain of the Locust: Physiological and Morphological Characterization. *Journal of Experimental Biology*, 19:1–19, 1990.
- [124] Alice A Robie, Andrew D Straw, and Michael H Dickinson. Object preference by walking fruit flies, *Drosophila melanogaster*, is mediated by vision and graviperception. *Journal of Experimental Biology*, 213(Pt 14):2494–506, July 2010.

- [125] Nicola Rohrseitz and Steven N Fry. Behavioural system identification of visual flight speed control in *Drosophila melanogaster*. *Journal of the Royal Society, Interface*, 8(55):171–85, February 2011.
- [126] Adam J Rutkowski, Roger D Quinn, and Mark A Willis. Three-dimensional characterization of the wind-borne pheromone tracking behavior of male hawk-moths, *Manduca sexta*. *Journal of Comparative Physiology A*, 195(1):39–54, January 2009.
- [127] M W Sabelis and P Schippers. Variable wind directions and anemotactic strategies of searching for an odour plume. *Oecologia*, 63(2):225–228, 1984.
- [128] B Schnell, M Joesch, F Forstner, S V Raghu, H Otsuna, K Ito, A Borst, and D F Reiff. Processing of horizontal optic flow in three visual interneurons of the *Drosophila* brain. *Journal of Neurophysiology*, 103(3):1646–1657, 2010.
- [129] Charles E Schroeder, Donald a Wilson, Thomas Radman, Helen Scharfman, and Peter Lakatos. Dynamics of Active Sensing and perceptual selection. *Current Opinion in Neurobiology*, 20(2):172–6, April 2010.
- [130] Juliet Popper Shaffer. Multiple hypothesis testing. *Annual Review of Psychology*, 46:561–84, 1995.
- [131] Julie H Simpson. *Mapping and manipulating neural circuits in the fly brain.*, volume 65. Elsevier Inc., 1 edition, January 2009.
- [132] Andrew P Smith. An investigation of the mechanisms underlying nest construction in the mud wasp paralastor sp. (hymenopter: eumenidae). *Animal behaviour*, 26:232–240, 1978.
- [133] Erik C Sobel. The locust’s use of motion parallax to measure distance. *Journal of Comparative Physiology A*, 167(5):579–588, 1990.

- [134] M V Srinivasan, S W Zhang, J S Chahl, E Barth, and S Venkatesh. How honeybees make grazing landings on flat surfaces. *Biological Cybernetics*, 83(3):171–183, August 2000.
- [135] Finlay J Stewart, Dean A Baker, and Barbara Webb. A model of visual-olfactory integration for odour localisation in free-flying fruit flies. *Journal of Experimental Biology*, 213(11):1886–900, June 2010.
- [136] Alan A Stocker. Analog Integrated 2-D Optical Flow Sensor. *Analog Integrated Circuits and Signal Processing*, 46(2):121–138, December 2005.
- [137] Nicholas J Strausfeld and U K Bassemir. Lobula plate and ocellar interneurons converge onto a cluster of descending neurons leading to neck and leg motor neuropil in *Calliphora erythrocephala*. *Cell and Tissue Research*, 240:617–640, 1985.
- [138] Nicholas J Strausfeld and H S Seyan. Convergence of visual, haltere, and prosternal inputs at neck motor neurons of *Calliphora erythrocephala*. *Cell and Tissue Research*, (240):601–615, 1985.
- [139] Andrew D Straw. Vision egg: an open-source library for realtime visual stimulus generation. *Frontiers in Neuroinformatics*, 2(November):4, January 2008.
- [140] Andrew D Straw, Kristin Branson, Titus R Neumann, and Michael H Dickinson. Multi-camera real-time three-dimensional tracking of multiple flying animals. *Journal of the Royal Society, Interface*, 8(56):395–409, March 2011.
- [141] Andrew D Straw, Serin Lee, and Michael H Dickinson. Visual control of altitude in flying *Drosophila*. *Current Biology*, 20(17):1550–6, September 2010.
- [142] Marie P Suver, Akira Mamiya, and Michael H Dickinson. Octopamine neurons mediate flight-induced modulation of visual processing in *Drosophila*. *Current Biology*, 22(24):2294–302, December 2012.

- [143] C Taddei-Ferretti and S Chillemi. Landing reaction of *Musca domestica*. VI. Neurones responding to stimuli that elicit the landing response in the fly. *Journal of Experimental Biology*, 94:105–118, 1981.
- [144] Lance F Tammero and Michael H Dickinson. Collision-avoidance and landing responses are mediated by separate pathways in the fruit fly, *Drosophila melanogaster*. *Journal of Experimental Biology*, 205(Pt 18):2785–98, September 2002.
- [145] Lance F Tammero and Michael H Dickinson. The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*. *Journal of Experimental Biology*, 205(Pt 3):327–43, February 2002.
- [146] Jeanette A Thomas, Cynthia F Moss, and Marianne Vater. *Echolocation in bats and dolphins*. University of Chicago Press, 2004.
- [147] Stefan Treue and John H. R. Maunsell. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature*, 382:539–541, 1996.
- [148] J R Trimarchi and A M Schneiderman. Flight initiations in *Drosophila melanogaster* are mediated by several distinct motor patterns. *Journal of Comparative Physiology A*, 176(3), 1995.
- [149] Shreejoy J Tripathy, Oakland J Peters, Erich M Staudacher, Faizan R Kalwar, Mandy N Hatfield, and Kevin C Daly. Odors pulsed at wing beat frequencies are tracked by primary olfactory networks and enhance odor detection. *Frontiers in Cellular Neuroscience*, 4(March):1, January 2010.
- [150] Floris van Breugel and Michael H Dickinson. The visual control of landing and obstacle avoidance in the fruit fly *Drosophila melanogaster*. *Journal of Experimental Biology*, 215(Pt 11):1783–98, June 2012.
- [151] Floris van Breugel, Kristi A Morgansen, and Michael H Dickinson. Monocular distance estimation from optic flow during active landing maneuvers. *Bioinspiration and Biomimetics*, In Press, 2013.

- [152] Rudolph van der Merwe and Eric A Wan. The square-root unscented Kalman filter for state and parameter-estimation. In *IEEE International Conference on Acoustics, Speech, and Signal Processing Volume 6*, pages 3461 – 3464, Salt Lake City, UT, 2001.
- [153] Koen J T Venken, Julie H Simpson, and Hugo J Bellen. Genetic manipulation of genes and cells in the nervous system of the fruit fly. *Neuron*, 72(2):202–30, October 2011.
- [154] A Verri and V Torre. Absolute depth estimate in stereopsis. *Journal of the Optical Society of America*, 3(3):297, March 1986.
- [155] N J Vickers and T C Baker. Latencies of behavioral response to interception of filaments of sex pheromone and clean air influence flight track shape in *Heliothis virescens* (F.) males. *Journal of Comparative Physiology A*, 178(6):831–847, 1996.
- [156] Gerhard von der Emde and Stephan Schwarz. Imaging of objects through active electrolocation in *Gnathonemus petersii*. *Journal of Physiology*, 96(5-6):431–44, 2002.
- [157] Hermann Wagner. Flow-field variables trigger landing in flies. *Nature*, 297:147–148, 1982.
- [158] G K Wallace. Visual scanning in the desert locust *Schistocerca gregaria* forskal. *Journal of Experimental Biology*, 36:512–525, 1959.
- [159] Gordon L Walls. Factors in Human Visual Resolution. *Journal of the Optical Society of America*, 33(9):487–505, 1943.
- [160] Yongchang Wang and Barrie J Frost. Time to collision is signalled by neurons in the nucleus rotundus of pigeons. *Nature*, 356:236–238, 1992.
- [161] Christian Wehrhahn, Tomaso Poggio, and Heinrich Bülthoff. Tracking and Chasing in Houseflies (*Musca*). *Biological Cybernetics*, (45):123–130, 1982.

- [162] Peter T Weir and Michael H Dickinson. Flying *Drosophila* orient to sky polarization. *Current Biology*, 22(1):21–7, January 2012.
- [163] M Wicklein and N J Strausfeld. Organization and significance of neurons that detect change of visual depth in the hawk moth *Manduca sexta*. *Journal of Comparative Neurology*, 424(2):356–76, August 2000.
- [164] Mark A Willis, E A Ford, and J L Avondet. Odor tracking flight of male *Manduca sexta* moths along plumes of different cross-sectional area. *Journal of Comparative Physiology A*, 199(11):1015–1036, October 2013.
- [165] W C Wittekind. The landing response of tethered flying *Drosophila* is induced at a critical object angle. *Journal of experimental biology*, 135:491–493, 1988.
- [166] Matthias Wittlinger, Rüdiger Wehner, and Harald Wolf. The ant odometer: stepping on stilts and stumps. *Science*, 312(5782):1965–7, June 2006.
- [167] R H Wright. The olfactory guidance of flying insects. *The Canadian Entomologist*, 90(02):81–89, 1958.
- [168] Junwei Zhu, Kye-Chung Park, and Thomas C Baker. Identification of odors from overripe mango that attract vinegar flies, *Drosophila melanogaster*. *Journal of Chemical Ecology*, 29(4):899–909, April 2003.