Steady as She Goes: Visual Autocorrelators and Antenna-mediated Airspeed Feedback in the Control of Flight Dynamics in Fruit Flies and Robotics

Thesis by

Sawyer Buckminster Fuller

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

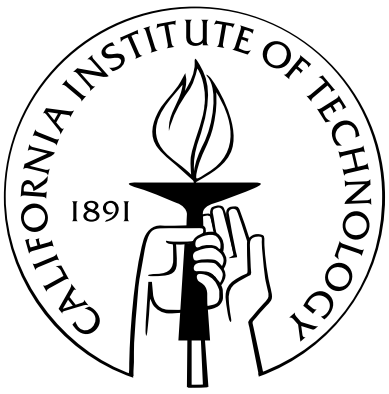

California Institute of Technology

Pasadena, California

2011

(Submitted May 26, 2011)

To Mom.

The key to a good thesis, I am convinced, is a well-selected opening quote.

- S. B. Fuller

Acknowledgments

Extreme thanks to Richard Murray for bringing me in, and for probably being my favorite advisor ever, and I've had a few. Thanks for letting me start on a great project, giving me flexibility to move around (and start working with flies), and pretty much everything else you could ask for, including saying how great my work is. To Michael Dickinson for letting me work on a crazy idea. To Joel Burdick for bringing me into Caltech. To Andrew Straw for getting me started thinking about correlators in the first place and for showing me how to do software the open-source way. This includes lots of tweaks on his software for my purposes, and for writing software good enough to do do things like measuring fast gust responses within milliseconds continuously for 24 hours. His software pervades this thesis. To Linda Scott for always having a smile. To Anissa Scott, Gloria Bain, Lilian Porter for making things all work. Thanks to colleagues and friends who were helpful sounding boards or gave invaluable feedback or just had a kind word, Akira Mamiya, Andrea Censi, Jasper Simon, Shuo Han, Dionysios Barmpoutis, John Dabiri, Ufuk Topcu, Elisa Franco, Ros Sayaman, Alice Robie, Will Dickson, Manuj Swaroop, David Gleason, Erik Schomburg, Jerry Cruz, Pratyush Tiwary, David Brown, Pelayo Dominguez, Aron Varga, Jessica Bastiaansen, Costas Anastassious, Celia Shiau, Olivier Delaire. And thanks Mom who always knew what to say when I was stuck somewhere in my work and knowing I could build great things, and to Dad for showing me how to build them.

Abstract

Achieving agile autonomous flight by an insect-sized micro aerial vehicle (MAV) will require improved technology that is radically smaller, lighter, and more power-efficient. One animal that has solved the problem is the fly, a virtuoso among insect flyers whose nervous system can perform sophisticated aerial maneuvers under severe computational constraints. This thesis is concerned with understanding and emulating the dynamics of the fly's feedback control system. Because vision is noisy and information rich, processing time may a problem for a fast-moving MAV or fly. By tracking the fruit fly *Drosophila melanogaster* in free flight in gusts of wind, I found that they incorporate feedback from wind-sensing antennae in a fast feedback loop that dampens the forward-flight dynamics. The slower dynamics are easier to control for long-delay visual feedback, making the fly more robust to the limitations of its visual system. Using the fly as inspiration, I designed a minimal, visual autocorrelation based controller that used a small array of visual sensors to stabilize a fan-actuated hovercraft robot in a narrow corridor. Using a model for correlators developed for the robot, I showed that a uniform array of visual correlators was sufficient to explain the free-flight velocity regulation behavior of flies, rather than a different model. In addition to illustrating the benefits of concurrent scientific analysis and engineering synthesis, the results give new insight into how to control small biological and man-made flying vehicles using limited, noisy sensors.

Contents

Acknowledgments **v**

Abstract **vi**

1 Introduction **1**

2 Responses of Free-flight Flies to Gusts of Wind **6**

2.1 Abstract 6

2.2 Introduction 6

2.3 Methods and apparatus 10

2.3.1 Realtime fly tracker 10

2.3.2 Wind gusting 10

2.3.3 Visual stimulus 13

2.3.4 Flies 13

2.3.5 Trial protocol 14

2.4 Results 15

2.4.1 The groundspeed of flies without aristae is highly variable 15

2.4.2 A long time delay in the visual response could explain the variability . 17

2.4.3 The antennae response is much faster, suggesting a stabilization mechanism 17

2.5 Conclusion 19

3 System Identification of the Fly's Visual and Mechanosensory Feedback Controller **21**

3.1 Abstract 21

3.2 Introduction 21

3.3 Methods . . . 22
3.4 Wing damping . . . 25
3.5 Antennae model . . . 26
3.6 Visual model . . . 30
3.7 Sensory Fusion . . . 33
3.8 Predictions of the model . . . 35
3.9 Discussion . . . 37
3.9.1 The gust response paradox . . . 40
3.9.2 Remarks on feedback architecture . . . 40
3.9.3 Relation to previous work . . . 42
3.9.4 Mechanisms . . . 43
3.9.5 Robustness to variability . . . 44
3.9.6 Robustness to environment geometry . . . 44
3.10 Appendix: Parameters used in simulations . . . 44

4 An Insect-inspired Autocorrelation Model for Visual Flight Control in a Corridor 48

4.1 Abstract . . . 48
4.2 Introduction . . . 48
4.3 Insect Flight Control . . . 51
4.4 Frequency-Domain Analysis of Correlators . . . 52
4.4.1 Correlator response to panoramic image motion . . . 52
4.4.2 Decomposing correlator response . . . 52
4.4.3 Incorporating the effect of spatial blurring . . . 53
4.4.4 Incorporating motion parallax and perspective . . . 54
4.5 A Controller That Uses Correlators to Approximate Retinal Velocity . . . 56
4.5.1 Tuning σ and $\Delta\phi$ angles for the environment to approximate retinal velocity ν_r . . . 57
4.5.2 Implementation in simulation . . . 58
4.6 Robotic Implementation . . . 59
4.7 A Controller Designed for Correlators . . . 63
4.7.1 Approximations to correlator model . . . 63

4.7.1.1 Aliasing factor (Equation 4.13) . . . 64
4.7.1.2 Temporal factor (Equation 4.12) . . . 64
4.7.1.3 Perspective approximation . . . 64
4.7.2 Finding state variables by square harmonics . . . 65
4.8 Conclusions and Future Work . . . 67

5 How Flies Use Correlators to Control Forward Velocity 70
5.1 Introduction . . . 70
5.2 Methods . . . 73
5.2.1 Model . . . 73
5.2.2 Drum case . . . 74
5.2.3 Tunnel case . . . 74
5.2.4 Apparatus . . . 76
5.2.5 Simulation . . . 77
5.3 Results . . . 79
5.3.1 Data . . . 79
5.3.2 Model fitting . . . 83
5.3.3 Velocity sensing . . . 85
5.4 Discussion . . . 86

6 Conclusion 89

Bibliography 93

List of Figures

1.1 Feedback and control block diagram of the fly 4

2.1 Multimodal sensory input to the fly 8
2.2 Diagram of experimental apparatus. 11
2.3 Photographs of the free-flight tracking and wind gusting apparatuses 12
2.4 Flies flying without their aristae are unstable, possibly because of a long delay in visual feedback 16
2.5 Flies' antenna-mediated response to wind gusts is much faster than their visual response and could eliminate the oscillations observed in arista-less flies. . . . 18
2.6 A hypothesis for how antennae could act to reduce oscillations. 19

3.1 A diagram of how tunnel-frame inputs and outputs are transformed into fly-frame inputs and outputs. 23
3.2 Aerodynamic wing drag is roughly proportional to airspeed. 27
3.3 System identification supports a proportional controller with a 20 ms delay as the antennae feedback controller. 29
3.4 A good fit for the visual controller is an integral controller with a 60 ms time delay 32
3.5 Vision and antennae response forces combine by summing 34
3.6 Addition of the effect of the antennae wind sense increases performance and robustness, with a small cost of increased susceptibility to wind gusts at intermediate frequencies 36
3.7 Comparison of response of model compared to data in the velocity domain . . . 38
3.8 Step response of the velocity-domain model reported in other work to our model 39
3.9 How distance to obstacles affects gain in the visual feedback loop 45
3.10 Block diagram of feedback system, rearranged for ease of analysis 47

4.1 Diagram of the correlator. 49
4.2 Diagram of fan-actuated robot with array of correlators in a corridor. 50
4.3 Simulated and analytic model correlator response R for different visual sensor blurring widths σ. 55
4.4 Simulation of robot using Humbert controller (sinusoid harmonics) and tuned blur width σ and correlator distance $\Delta\phi$. 59
4.5 Diagram of the fan-actuated hovercraft robot in its environment 60
4.6 Images of the robot and corridor 61
4.7 The luminance response profile of an individual infrared sensor 62
4.8 Trajectories of the robot in the corridor captured by the overhead vision system 63
4.9 Correlator response to state perturbations as a function of body-frame angle ϕ'. 65
4.10 The controller introduced in this work, designed explicitly for the correlators, has a larger basin of attraction in simulation and functions on non-sinusoid patterns as well. 68

5.1 Photograph of corridor with short-wavelength gratings back-projected onto the walls 76
5.2 Visualization of correlator simulation 78
5.3 Simulated and analytic correlator responses for different spatial frequencies . . 79
5.4 Drum ground speed trajectories 80
5.5 Tunnel ground speed trajectories 81
5.6 Fly accelerations in response to visual stimuli exhibit a temporal frequency tuning peak 82
5.7 Model correlator responses to drum and tunnel stimuli 83
5.8 By tuning the blur width in the correlator, it is possible to achieve an agreement between the relative responses at different spatial frequencies 84
5.9 Fly acceleration response plotted as a function of velocity 85
5.10 The model predicts that correlators could perform good-enough velocity estimation in the range of $F_s = 4$ to 16 cycles/m in a tunnel geometry 86

6.1 The flight trajectories of the fruit fly as they explore in the flight arena, as seen from above. Each dot is scaled according to flight speed, as if the animal was dribbling paint as it was flying. The stars indicate when the flies were subjected to a rapid gust of wind. Can we explain this complicated behavior? 92

List of Tables

3.1 Fit values for candidate antenna controllers 46
3.2 Fit values for candidate visual controllers 46
3.3 Parameters used in simulation 46

5.1 Simulation parameters for model of fly's forward-velocity correlator-based velocity controller 83

Chapter 1

Introduction

Man has always aspired to fly. But only very recently in history has he begun to do so, starting with the Wright Brother's first faltering flight in an airplane. Man-made craft have since left the atmosphere and even brought humans to the moon and rovers to Mars. But while the past century of flight innovation can be marked by milestones of bigger and faster, from the advent of jet power to supersonic flight to rocketry, today the focus of innovation has changed. New milestones are being measured by smallness, autonomy, and complexity [1]. Small-scale autonomous flight of a vehicle on the scale of an insect is an unsolved current research question. The difficulties are manifold: as sensors are reduced in size, their capability and precision must diminish as components are removed. Simultaneously, the available onboard computation power must use less power, limiting the choice of algorithm [2, 3].

One approach to building a small flying vehicle is to improve technology iteratively, building successively more sophisticated and smaller robots with each new generation by adapting existing technologies for reduced size and weight. In open spaces, small flying airplanes such as the Aerovironment Wasp use the global positioning system (GPS) to navigate between waypoints for military surveillance [1]. But in more confined and cluttered spaces, the problem becomes more difficult. GPS cannot indicate the position of obstacles, and even if the vehicle were equipped with a perfect onboard map of them, both the accuracy and bandwitdth of GPS become insufficient to avoid collisions below a certain scale [3]. Indoors, GPS is denied entirely, and other methods must be employed. Shen [4] reports an autonomous quadrotor helicopter that operates indoors without GPS by using a laser rangefinder and an inertial measurement unit to build a map. As impressive this technical feat, the robot is large, on the order of a meter across, so that it can carry the sensor suite

and powerful processor required. It also moves slowly. As vehicles get smaller, and greater maneuverability is desired, powerful sensors like laser rangefinders and computers must be done without because of their weight [2]. While vision can give precise estimates of self-motion, existing algorithms such as the one in use on the Mars rovers are slow computation heavy, requiring minutes for each step on Mars [5]. A new approach is needed [6].

Biology, which excels in the domain of smallness and sophistication, may suggest possible solutions. Hundreds of millions of years of evolution has lead to flies, butterflies, hummingbirds, all of which are able to perform admirably without fast serial computers or lasers [7]. Their flight capabilities exceed current human-engineered solutions [8], often exploiting unsteady fluid flow dynamics in their flapping flight [9, 7]. Insects are of particular interest because of their relatively simple nervous systems. Because their neurons are typically specified individually with characteristics that are conserved across individuals and even species [10, 11, 12, 13], it seems that an understanding of how their neurons operate to generate behavior is closer at hand [14, 15] than in vertebrates.

Engineers often take inspiration from biology. The original inspiration for feedback control lay in the term "cybernetics," coined by Norbert Wiener to refer to feedback control in biology [16]. An effort to emulate the brains of insects lead to an influential body of work by Brooks that challenged the foundations of robotic intelligence [17]. Instead a traditional approach to robot control, which modularized sensing, modelling, planning, and execution, he proposed tying sensors directly to actuators [18]. Higher level competence at behaviors like foraging were implemented by adding layers that modulated the lower-level locomotive behavior, mirroring the evolution of the brain [17]. Though these controllers were implemented on ambulatory and wheeled robots, they were never applied to aerial robots.

Findings about flying insects have been the inspiration for many robotic micro-flyers in other groups. Biological inspiration may be particularly valuable in flight, where realtime behavior, minimal computation, and sophisticated motion control are specialties of biological organisms [7]. The finding that bees center their flight in a corridor led to robots that balance lateral optic flow [19, 20, 21, 22]. Later, it was found that the lobula plate tangential cells integrate visual flow over wide-field patterns corresponding to states of self motion [11, 23, 24] and seem to be required for important insect motor control tasks [11]. This inspired the idea of decomposing visual flow into "matched filters" for self motion [25, 20, 26], including sinusoid harmonics, from which the state vector of the vehicle can be observed and

controlled [27, 2]. A small airplane inspired by the expansion-avoidance behavior observed in flies [28] was able to fly autonomously in a small room by making body-saccades away from walls [29]. A robot that regulates its attitude using ventral optic flow used a simple, insect-inspired controller that reproduces observations that that insects may descend in a headwind and ascend in a tailwind [30], though recent work has called into question whether that mode of flight control is active in flies [31]. Using Hebbian updates to learn and perform visual servoing demonstrates how correlators could arrive spontaneously during learning or evolution [32, 33]. For a review of bio-inspired engineering, particuarly in the domain of visual flight control, see [22].

Insight the other direction, from engineering into biology, often follows an effort in engineering synthesis of a biological system. Only then can we perceive the subtleties that invariably arise when a concept is reduced to practice. Braitenburg [34] showed how simple wheeled robots could show life-like behavior and even the illusion of free will by simple interconnections between sensors and motors. But understanding the principles required mentally constructing and running these little robots. Ijspeert [35] showed how regulation of central pattern generators could lead to transitions between swimming and walking gates on a robotic salamander. Biology's strongest trait is perhaps its ability to find practical and robust solutions. For example, [36] found by implementing a small flapping flyer that an automotive-like differential could compensate for wing damage by using a mechanical pivot to automatically flap the other wing farther. Only by constructing the device could he discover that a mechanical linkage could produce robustness that might ordinarily be presumed to arise from neurally-modulated mechanosensory feedback. This may inform future studies on the insect wing motor hinge [37].

It is the opinion of the author that understanding how to autonomously fly small vehicles and understanding how flight is performed by insects is best approached by performing engineering and biology simultaneously. In doing so, the benefits of both approaches will inspire bigger ideas.

This thesis considers the dynamic flight control of the fly from a control-theoretic perspective, an approach that has been lacking in the biological literature [38]. The fly's feedback controller is represented as a block diagram, with boxes representing components segregated by function or mechanism (Figure 1.1). I report contributions to the understanding of insect flight control by both studying fruit flies and building a robot that takes inspiration from

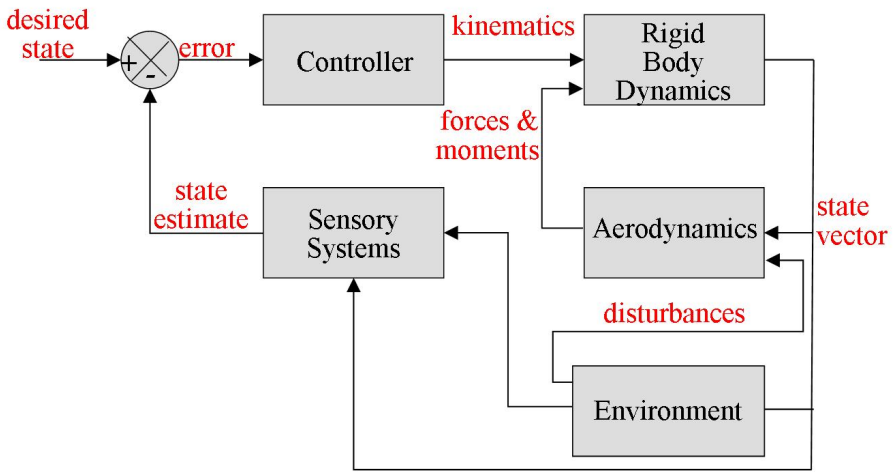

Figure 1.1: Feedback and control block diagram of the fly. This thesis is concerned with understanding characteristics of the "controller" and "sensory systems" blocks.

the fly's control system. Fruit flies make a compelling choice for study because of the ease of raising them, the short generation time (9 days), and most importantly, as a genetic model organism, genetic findings can be leveraged in behavioral studies and vice-versa [14]. In the study of the fly, I placed particular emphasis on studying sensory feedback systems that have potential application to robot flyers.

The thesis is organized as a series of self-contained chapters, each beginning with an abstract that summarizes the main results and an introduction that reviews the relevant background. chapters 2 and 3 detail how by studying dynamics of forward flight in flies, I uncovered a control challenge encountered by the fly that had not previously been addressed in detail: feedback from vision is relatively slow. By subjecting flies to sudden gusts of wind, I made the surprising finding that the antennae act as fast, proportional sensors that act to damp out and thus slow the fly's dynamics. The slowed dynamics mitigate the difficulty of long visual delays and give the fly robustness to variations in parameters ranging from fly-to-fly variability to environment geometry to non-idealities in visual correlators. And they may be applicable to autonomous flight control, in which computation-heavy vision is likely to be a slow sense as well. Next, I considered flight control from a synthesis perspective in chapter 4. Using a bio-inspired control algorithm with minimal computational requirements, I designed a controller that stabilized the motion a dynamic hovercraft robot using visual

flow estimates arriving from a small array of luminance sensors and correlators. In chapter 5, I found that the analysis of chapter 4 suggested a new study on free-flight behavior in flies to show that basic correlator behavior may be sufficient to explain the fly's forward speed regulation, rather than necessitating some more complicated model. In addition to illustrating the benefits of cross-disciplinary study, the results shed new light onto the flight control of small flying vehicles, biological or man-made.

Chapter 2

Responses of Free-flight Flies to Gusts of Wind

2.1 Abstract

In order to navigate through a complex and changing world, animals must rapidly combine information from sensory channels with different response bandwidths and modulate motor output accordingly. For example, to regulate their flight speed, insects are thought to employ both rapid mechanosensory signals and slower visual cues, although the means by which they combine information from these two modalities is unknown. To study this process of sensory-motor integration, we subjected free-flying fruit flies to impulsive gusts, visual gusts, and attenuated mechanosensory feedback by removing the aristae on their antennae. By observing the fly's velocity response to wind perturbations, we found that, surprisingly, the fast, antenna-mediated response acted in the same direction as the wind input, in effect augmenting the effect of drag on the wings. Aristae-ablated flies showed much higher variability and oscillations in flight speed, suggesting that the wind sense might not be for wind disturbance rejection but for "damping out" these flight velocity oscilations.

2.2 Introduction

Flies, among the most adept of insect fliers, display a sophisticated suite of aerial behaviors that require rapid sensorimotor processing. They can recover from tumbling surprise take-offs in tens of wingbeats [39] and can perform inverted landings on the ceilings with ease. But even the practical matter of navigating from one place to another in flight is a challenging task requiring sensorymotor specializations. Not only must they negotiate cluttered

environments abounding with obstacles such as trees and plants, but they must do so in wind, which may change direction and magnitude. How this is achieved has not yet been explained, although many clues have been found from studies on flies and other insects.

Flies are visual animals, devoting perhaps two-thirds of the neurons in their brain to visual processing [40]. They have two faceted eyes that sample nearly the entire visual sphere, with varying visual accuity and number of ommatidia depending on species [41, 42] (Figure 2.1). Insects' flight control relies heavily on visual flow, that is, the pattern of visual motion across the retina. Bees center their flight in a corridor, balancing lateral visual flow, and slow when the corridor narrows [43]. Both of these behaviors are consistent with a regulator that measures the angular velocity of visual flow across the retina (for a review, see [22]). Neurons that average visual motion across the retina have been found in the fly brain and homologs of such cells that may underly this behavior [11, 23]. A modeling effort suggests visual motion patterns could be used to orient upwind [44]. Flies also use preferences for simple horizontal or vertical features to structure their flight. They control altitude by keeping level with horizontal features [31] and are attracted to vertical features, using that attraction to fly straight and overcome the normally-aversive frontal visual expansion that accompanies forward flight [45]. Flies spend much of their time flying forward and straight. They organize their trajectories into linear segments of roughly constant velocity, punctuated by sudden turns known as body-saccades [46] that are often induced by visual expansion [28]. Their visual forward flight-speed regulator can be modelled as a delay followed by a second-order velocity controller [47]. A rich literature exists concerning the mechanism of visual motion detection, but evidence suggests that in forward flight, flies' response behavior acts proportionally to the difference between the fly's velocity and the velocity of background visual motion across a range of spatial frequencies [48].

Less well understood is the role played by the antennae (for a review, see [38]). For many flying insects, knowledge about the wind is vital for survival: to find food sources by smell, a fruit fly moves laterally to the wind to search for odor plumes and then turns upwind when one is found [51]. Perhaps the only example of known interaction between vision and wind sensing is that tethered but freely-rotating flies turn into a headwind and then use that percept to overcome the normally-aversive forward visual expansion [52]. Once facing upwind, flies use vision to maintain a constant groundspeed, independent of the speed of wind in a wind tunnel [53, 54].

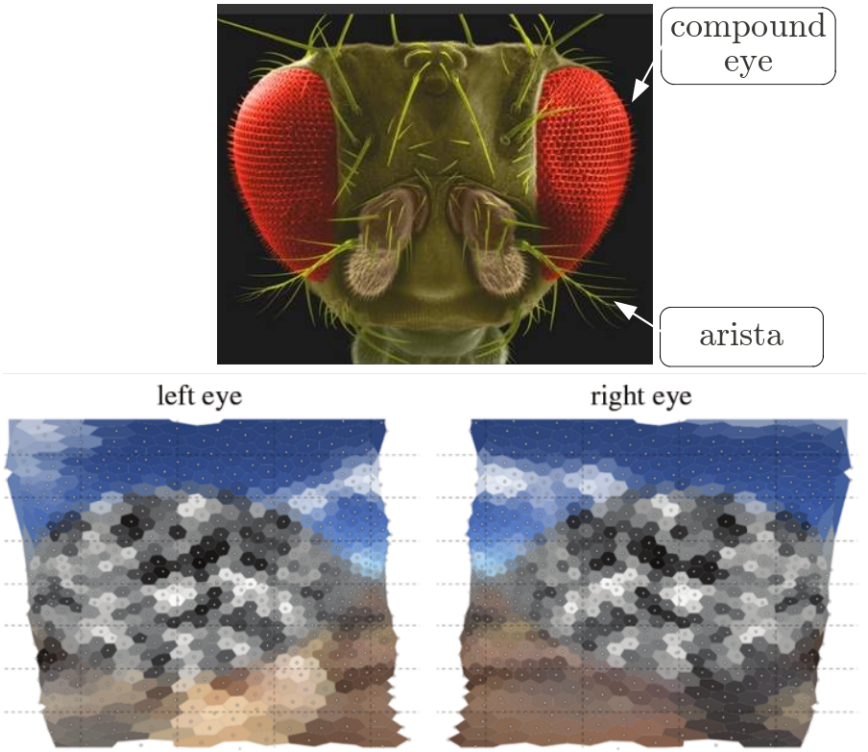

Figure 2.1: Multimodal sensory input to the fly. An electron micrograph of a fruit fly's head (top) shows the faceted eyes (red) and the antennae. The eyes take a low-resolution omnidirectional sample of the visual world, which would look something like this computer-generated rendition (below, from [49]). The aristae (top) are branched appendages of the antennae that protrude from the head and deflect in wind. Sensors at the base of the antenna, the Johnston's Organ, sense the wind-induced deflection [50]. In this and the next chapter, we present a quantitave analysis of the input-output behavior of these two senses simultaneously.

Flies detect wind using sensitive motion sensors in their antennae [55]. The antennae have also been shown to sense sound [56], the direction of gravity [57, 58], and may be directly involved in sensing and controlling wing kinematics [59]. The basic mechanism is that the arista, a branching fourth antennal sement, protrudes into the wind and is subject to a torque roughly proportional to airspeed [50]. The arista, together with the third antennal segment to which it is rigidly attached, can rotate passively around a joint with a roughly vertical axis coincident with the long axis of the third segment. The motion of this joint is sensed by the Johnston's organ, an array of chordotonal mechanosensory organs residing in the second antennal segment. An active antenna-positioning reaction carried out by the first antennal segment may increase the dynamic range of the wind sense [50].

Evidence suggests that flies can use their antennae to measure absolute airspeed, rather than, for example, its rate of change. In tethered flies, the steady-state wing beat amplitude varies with the steady-state airspeed [50]. Similar observations have been made for locusts [60] and dragonflies [61]. More recently, using a calcium-sensitive reporter gene, Johnston's Organ neurons were found to project to different portions of the brain depending on whether they responded to either absolute antenna deflection or its rate of change [62].

Whereas wing kinematics changes in response to antenna wind stimulation in the the steady-state on tethered animals has been explored, little is known about about how this reaction translates into changes in behavior in free-flight. Because of the range of different kinematic motions available to flapping insects, it is in general impossible to map changes in wing beat amplitude directly to forces and torques [63, 64]. In addition, tethered kinematics differ significantly from free-flight kinematics, perhaps because normal feedback signals are disrupted [64]. Hence, it remains unanswered how the documented wing kinematic changes affect flies' free-flight motions—do they give rise to thrust, lift, or torques, and in what directions? In addition, steady-state analyses do not examine dynamic behavior, which may be key to explaining how these organs affect the dynamic stability of the aloft fly.

To investigate the fly's sensor dynamics in straight segments of forward flight, we subjected them to wind and visual stimuli as they flew along a corridor while their position was tracked by cameras.

We found that the antenna-mediated wind response is much faster than the visually mediated response. But rather than acting to reject the gust of wind as might be expected, the antenna response does just the opposite: the velocity of flies with intact antennae is

more perturbed that for arista-ablated flies. Flies without their aristae showed significant variance in their forward velocities, nearly twice that of control flies. The variability appears to be sinusoidal, as if the flies are using a along-delay visual feedback loop that is at the edge of stability without the antennae. The antennae reflex function may be to act as an active damper, applying force as necessary to dampen these oscillations.

2.3 Methods and apparatus

2.3.1 Realtime fly tracker

We used a custom built multi-camera real-time fly tracker [65] (Figures 2.2 and 2.3) to record the three-dimensional positions of flies in free-flight. It consisted of five digital firewire videocameras (Basler A602f) taking images at 100 frames per second, five dedicated desktop computers performing image analysis to locate flies in each camera's view, and a central computer triangulating that information into (x, y, z) positions. Dimensions of the arena were 150 cm $\times$ 30 cm $\times$ 30 cm. Latency was approximately 50 ms. Tracking was performed with infrared backlight and infrared filters on each camera so that moving visual stimuli were not detected by the cameras.

2.3.2 Wind gusting

We constructed two different devices to generate sudden changes of wind velocity in the flight arena. The first was a set of motor-actuated wind-vanes that could open suddenly to allow air pulled by the fan to flow through the tunnel, giving a "step input" to the fly. Because the arena was in a wind tunnel designed for only uni-directional flow, the fan, and hence the vanes, could produce wind stimuli in only one direction. The second was an air piston that could move the mass of air in the tunnel in either direction very quickly. Though the apparatus was situated behind the fly during the trial (Figure 2.2), atmospheric pressure in front of the fly provided plenty of pressure to generate headwind gusts. To verify that the timing of the gust was the same along the length of the arena, two wind-measurement probes (described below) were placed 1 m apart at either end and simultaneously measured the gust. The two gust velocity traces were nearly identical with timing difference of 3 ± 1 ms, approximately the speed of sound ($\approx$350 m/s). Thus the gust was approximately the same along the length of the tunnel during a single 10 ms frame of the tracking cameras. Both

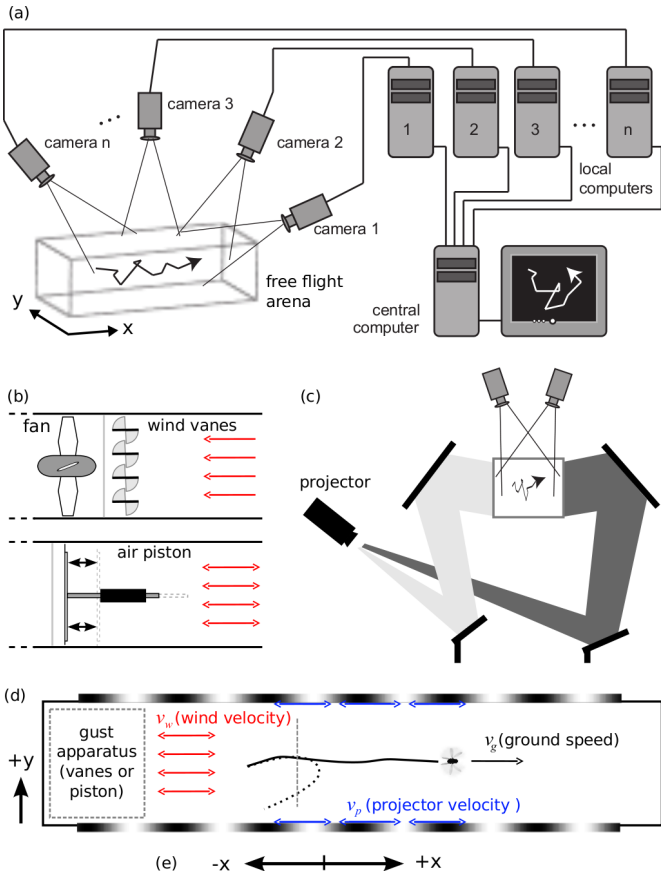

Figure 2.2: Diagram of experimental apparatus. An array of cameras and a cluster of computers processing video images triangulated the position of flies in real-time (a). Wind stimuli (b) were generated either by a fan and wind vanes opening or closing (top) or by an air piston spanning the cross section of the wind tunnel actuated by a linear motor (bottom). Because the wind tunnel is designed to provide unidirectional laminar flow, the fan could only pull air in one direction. Back-projected visual stimuli were generated by a high-speed monochrome projector reflected off of mirrors (c). A diagram of the wind tunnel seen from above (d) shows the locations of the gust apparatus and visual stimuli relative to the fly. In the inertial lab frame, stimuli to the fly are wind velocity v_w and projector velocity v_p. The fly's ground speed v_g is measured by the tracker as the output. A trial starts when a fly passes through the imaginary plane shown as a dotted line transverse to the axis of the wind tunnel. A drawing of an accepted trajectory in which the fly flew primarily along the axis of the tunnel is shown as a solid line; a rejected trajectory is shown as a dotted line. All velocities in the tunnel frame have the convention that positive is toward the right, according to (e).

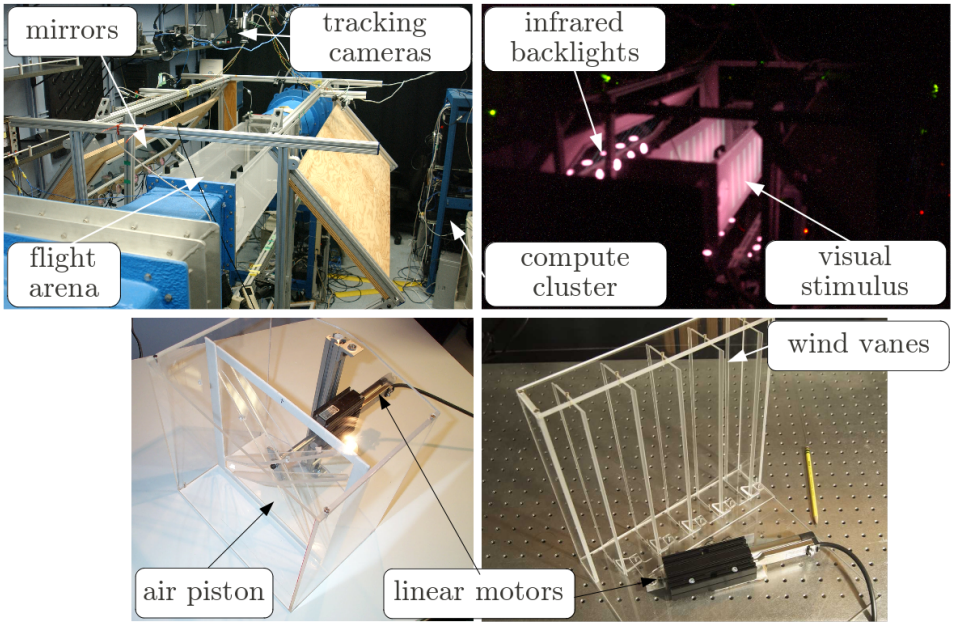

Figure 2.3: Photographs of the free-flight tracking (top) and wind gusting apparatuses (bottom). The free-flight arena is shown in light (top left) and under the experimental conditions of darnkess (top right). The pink glow on the right is the infrared backlighting used by the infrared-only tracking cameras, not visible to either human fly eyes, but detected by the CCD sensor in the digital camera used to take the picture.

gusting devices were actuated by a high-speed brushless linear motor (LinMot, Elkhorn, Wisconsin), and could perform short motions in as little as 35 ms or with a distance up to 125 mm. Because the motor servo controller could be programmed with only one trajectory at a time, a given population of flies in the tunnel was subjected to only one type of wind gust for the 24 hour period. The mechanism was commanded to gust by a single voltage pulse from a computer based on the real-time position estimate. The two devices are shown in Figures 2.2 and 2.3.

To minimize visual impact of piston or vanes motion, they were constructed from clear plexiglass and placed so that they were behind the fly during each trial. To test whether their motion could induce a visual response, a sham piston with holes spanning most of its cross section that did not generate any measurable wind disturbance when it moved was substituted for the real piston. When tested with flies, there was no discernible behavioral difference between trials in which the piston was moved and when it was not. This was expected because the physical motion of the gusting apparatus lay entirely in the rearward section of the visual sphere not sampled by the fly's eyes [66].

Using a hotwire anemometer (MiniCTA with p55 probe, Dantec Dynamics, Holtsville, NY), the time-course of the gust was measured at 1 kHz. Gusts were repeatable to within 0.02 m/s. The thickness of the boundary layer in the airflow in continuous-wind experiments was found to be less than 2 cm, so any trajectories passing within that distance of the walls were eliminated from consideration.

2.3.3 Visual stimulus

Visual stimuli were generated using the Vision Egg software on a PC running Ubuntu Linux with an nVidia Geforce 8500 GT graphics card [67]. A Lightspeed Designs DepthQ projector with color filter wheel removed was used to back-project the patterns (120 Hz update rate), and the mean luminance of the arena walls and floor when the projector displayed midgray was 50 cd/m2. The visual stimuli were sinusoid gratings on both walls with a spatial wavelength of 12 cm, moving in the x-direction. The 12 cm wavelength exhibited a strong response that was found to be roughly proportinal to forward velocity (Figure 5.6). This computer received 3-D coordinate estimates for all flies from the tracking computer over ethernet and orchestrated the automated experiment protocol. Because the visual frame rate of 120 Hz differed from the tracking frame rate of 100 Hz, the position of the visual stimulus was linearly interpolated to the time stamps of the camera frames.

2.3.4 Flies

Flies were 2–3 day old female *Drosophila melanogaster* Meigen descended from 200 wild-caught female specimens. Flies were manipulated on a cold stage under a dissection microscope. For aristae-ablated experiments, the aristae were removed using sharpened tweezers. Flies were kept on a 12h:12h light:dark cycle and experiments were started 5–9 hours before the end of their subjective day, continuing for 24 hours with 10 to 12 flies in the arena at a time. They were starved with water for 4-8 hours prior to each experiment to increase exploratory behavior. Yield was approximately 30 acceptable trajectories/day for control flies, and 15/day for arista-ablated flies. The mass of flies averaged 1.1 mg at the start of trials. After a 24 hour trial period, flies were recovered from the tunnel using a vacuum and were found to have lost approximately 20% of their mass, so the average fly mass m was estimated at 1.0 mg.

2.3.5 Trial protocol

A visual confinement protocol similar to that described in [48] was used to bring flies to the trigger plane at relatively high frequency, increasing experimental throughput. However, it was not possible to include a plume of odor because the air piston covered the entire cross section of the tunnel, blocking continuous flow. Because of flies' tendency to only fly in the presence of a headwind, it was necessary to have them flying forward in the quiescent air at the start of each trial (rather than be stationary but flying into continuous wind, as in [48]). The protocol used to bring flies to the trigger plane was executed as follows: (1) The walls were animated with velocity in the negative-x direction to induce any flying flies to move in that direction as well. (2) Once the chosen fly (the fly with the longest current trajectory) passed a threshold 5 cm away from the trigger plane, the animation direction was reversed again. (3) With animation now in the forward direction, the fly was brought forward until it passed the trigger plane. If either of the previous steps did not complete within 2.5 seconds, the protocol returned to step 1. Otherwise, once the fly passed the trigger plane, 4) a 1.2-second trial was initiated, consisting of a gust, a visual gust, or a combination of the two, in random order. At the end of this period, the state of execution was returned to step 1.

Because the linear motor interface was written in closed-source software, it could not be programmed and only one piston stroke type was available per population of flies and 24-hour period.

We eliminated trajectories in which flies performed sudden body-saccades, selecting only trajectories in which the fly flew essentially along the axis of the tunnel. Trajectories with non-axial velocity magnitudes that exceeded 0.25 m/s (after being filtered by a 10-sample box filter to eliminate transients) were eliminated from consideration. To insure uniform visual stimulus, trajectories not starting in the middle 1/3 (width-wise) of the tunnel were eliminated. Traverse measurements taken with the hotwire anemometer indicated that wind tunnel air flow was uniform farther than 2cm from the walls. Accordingly, trajectories in the upper 2/3 of the arena were accepted, with the exception of trajectories that came within 2 cm of the arena ceiling.

The fly tracker records trajectories and performs data association using a nonlinear extended Kalman filter that estimates velocities as well as positions. In parallel, it also estimates positions directly each time step using a least-squares estimate based on projection

onto the rays cast into each camera currently imaging the fly [65]. We used this noisier, unfiltered format because it provided the highest time-resolution, a necessity because of the high-speed dynamics in this study.

Registering the timing of all devices was of key importance. To collect the exact time that each frame from the cameras was taken, the tracking computer used feedback from the camera trigger device to create a model of offset and skew of pulses relative to time information collected over internet from NPT to collect timestamp of each frame. For the timing of the gust, the voltage pulse that induced the gust was recorded along with the wind velocity of the gust so that the timing of the gust relative to the pulse was recorded. The pulse was generated by the same usb trigger device that triggered each frame of the cameras. The device returned an acknowledgement of the command, and the middle of the roundtrip time (almost aways around 3 ms, and if longer, the trajectory was elimininated) was taken as when gust was started. To time the projector latency, we used a Texas Instruments opt101 light sensor sensor aimed at the projector and the camera trigger hardware to send a voltage pulse so that both could be observed on an oscilloscope simultaneously to find that the projector had a consistent 20 ms latency. Adjustments for these delays were made in the data *post facto*.

2.4 Results

2.4.1 The groundspeed of flies without aristae is highly variable

In comparison to control flies with intact antennae, flies with their aristae removed exhibited more variable trajectories and velocities. The variance of groundspeed within-trajectory was greater (4.6 cm/s vs 7.5 cm/s), as was the variance in mean velocities (7.1 cm/s vs 9.9 cm/s) (Figure 2.4). The velocity variability appears to be a result of greater degree of oscillations: flies without their aristae appear to be continually accelerating and decelerating in a sinusoidal pattern.

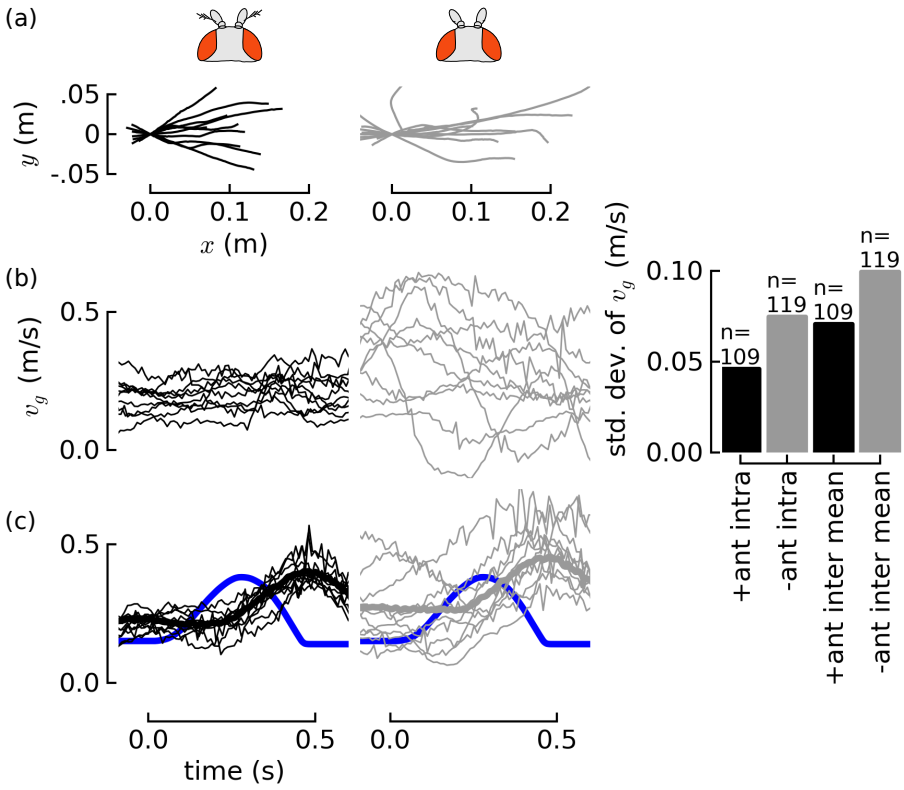

Figure 2.4: Flies flying without their aristae are unstable, possibly because of a long delay in visual feedback. Compared to the trajectories of contol flies with their antennae intact (a, left in black), a random sample of ten trajectories from flies with their aristae removed show (right, in grey) more variability. Trajectories were rendered as seen from above with $(0, 0)$ being the fly's position at the start of the trial. Arista-ablated flies exhibit more variable groundspeeds (b) that appear to be sinusoid-like. The variability (measured by standard deviation) of velocities of arista-ablated is higher within each trajectory, as the variability in mean velocities, compared to control flies (b, far right). When subjected to a visual stimulus provided by a change in projector speed at the start of the trial, shown in blue (c), flies respond by accelerating, but with a ≈ 100 ms delay. The visual delay could account for the velocity oscillations. Assuming feedback from the antennae is faster than that of the eyes, a parsimonious explanation is that the flies approach instability when forced to rely solely on longer-delay visual feedback alone.

2.4.2 A long time delay in the visual response could explain the variability

When subjected to a visual stimulus provided by a change in projector speed, flies respond by accelerating in a way that eliminated the visual slip, but with a delay of ≈ 100 ms. The visual delay could account for the velocity oscillations. If the feedback system had to only rely on slower-response visual feedback, it could be operating closer to the edge of stability, which would predict larger oscillations arising from the feedback, as can be observed in traces in Figure 2.4 (b).

2.4.3 The antennae response is much faster, suggesting a stabilization mechanism

We measured the groundspeeds of flies in a series of trials with different headwind gusts (Figure 2.5). In slow, step-like changes in wind velocity generated by suddenly opening the wind vanes, flies first decelerated and then recovered their initial velocity after 400-500 ms, consistent with the steady-state behavior first documented by [54] that flies maintain groundspeed independent of wind velocity. Further, we verified that this behavior was mediated by vision by animating the walls at the same time as the onset of the gust in such a way that the mean visual stimulus during the trial was near zero. Under this condition, flies decelerated to a new velocity and therafter maintained that velocity. A significant difference in responses between control and visually-abolished step gust responses only appeared after 200 ms, reinforcing the finding that vision operates with a relatively slow response.

In a fast impulsive gust, unlike in a step gust, an animal moving purely passively with no feedback (with a fixed amount of forward thrust to counteract the effect of aerodynamic drag on the wings) would recover its initial velocity after a short time because the wind velocity itself also returns to zero. However, in fast impulsive gusts generated by the air piston (15 mm gusts in 35 ms), we found a significant difference in velocity between flies with and without their aristae at 75 ms (Mann-Whitney U-test, $p < .001$). The divergence appears very fast, within 20 ms. These result suggest that flies sense rapid changes in airspeed with their antennae and respond by actively accelerating in the same direction as the gust. In the gentle, step gusts, the effect is so subtle it is hard to detect.

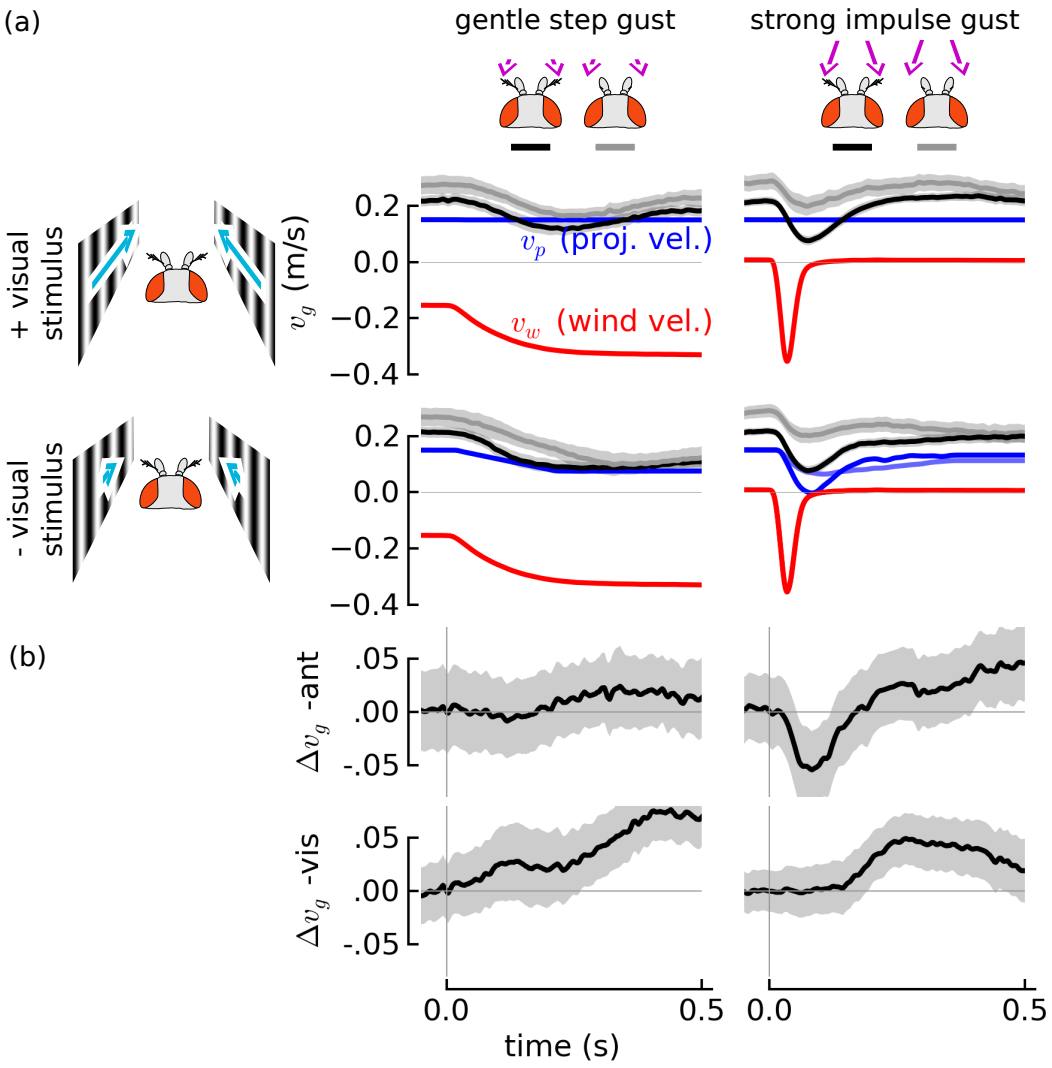

Figure 2.5: Flies' antenna-mediated response to wind gusts is much faster than their visual response. Fly groundspeeds v_g (mean +/- 95% confidence interval of mean) in gentle, step gusts (a, left) and strong, impulsive gusts (a, right) are compared against the groundspeeds of flies with their aristae removed (gray). In gentle step gusts (left), flies eventually recovered their initial velocity in spite of the increased headwind (upper left), This behavior is most likely mediated by vision. During a headwind gust, the fly's velocity drops, inducing visual slip in the negative direction, shown by the blue arrows on either side of a cartoon of the fly's head. To eliminate the visual slip normally associated with the headwind gust, the walls were animated at roughly equal velocity to the fly's mean velocity (collected from an earlier set of trials) (lower row). With visual animation, groundspeed recovery in the step gusts was abolished (lower left). In strong gusts there is a significant behavioral difference between flies with and without their aristae that arises soon after the gust (b, top right). But visual responses take much longer to occur (b, bottom right), supporting the hypothesis that there is a long visual delay. The antenna play and active role in forward velocity regulation and rather than acting to reject the gust, the response seems to abet the effect of the gust.

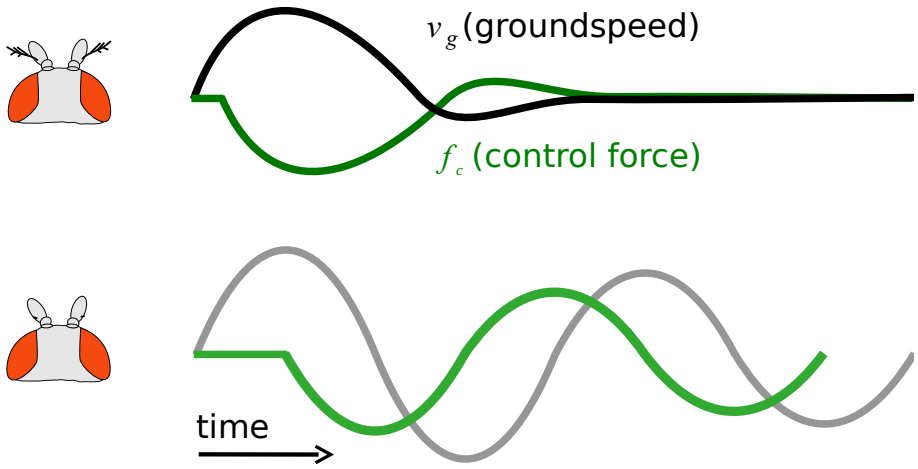

Figure 2.6: A hypothesis for how antennae could act to reduce oscillations. The time it takes for a fly to sense that its velocity is different from its desired forward velocity and exert a thrust force response f_c to eliminate the error (green) is much faster if the fly's antennae are intact (top). If instead there is only visual feedback (bottom), the response is so delayed that it leads to a continuation of the oscillations.

What use could there be to actively augmenting the disturbance caused by a gust of wind? The direction of the response from the antennae is in the same direction as would be required to abolish the ground speed oscillations observed in aristae-less flies in Figure 2.4. A representative diagram is provided in Figure 2.6.

2.5 Conclusion

In this chapter, we presented an apparatus to provide wind and visual stimuli to the fruit fly as it flew and its position in space was tracked. We found that the antenna-mediated wind sense is much faster than the visual sense. The direction of the response, however, was unexpected: rather than sensing the wind so that its influence could be "rejected," the antennae appear to be augmenting the passive drag force on the wings with active changes in wing motions. We hypothesized that the function could be to counteract, and thus damp out the oscillations observed in arista-less flies. These flies may be oscillating because they are forced to rely only on long-delay visual information and thus are operating at the edge of stability.

In the following chapter, we describe how a quantitative system identification procedure was applied to find a model for the fly's input-output behavior. We used the model to explain how the antennae function as airspeed dampers, and how this could be used to

provide robustness to parameter variability.

Chapter 3

System Identification of the Fly's Visual and Mechanosensory Feedback Controller

3.1 Abstract

This chapter concerns a system identification of the fly's free-flight behavior in response to wind and visual stimuli. A linear model fit to the flies' behavior suggests that the antennae act as a fast, proportional feedback regulator whereas vision provides longer-delay integral feedback, and that the two senses sum. The feedback from the antennae acts to augment the passive air drag of the wings, effectively doubling it, providing a velocity-proportional damper that slows the fly's dynamics and makes them easier to control for long-delay visual feedback. Flies without aristae exhibit oscillations in flight velocity that typify a feedback regulator at the edge of stability, as predicted by the model. The additional information provided by airspeed measurement may aid the fly in confined visual environments where the effective visual gain depends on the angular rate of motion across the retina which varies greatly depending on the distance to obstacles. Our results provide new insight into the functional architecture of flight control systems in insects.

3.2 Introduction

This chapter succeeds chapter 2 by finding a quantitative model of the fly's behavior so that statements can be made about its design and performance. We proposed a number of simple candidate linear models and fit them to the fly's response to a number of different wind and

visual gusts and selected the best model.

We found that the antennae functioned as fast active sensors with a feedback response proportional to airspeed. Since the visual response has a longer delay and hence is slower, the damping of the antennae slows the dynamics and makes forward motion of the fly easier to control with visual feedback. This has two interrelated consequences. First, there is less overshoot by flies with intact antennae, and thus velocity oscillation in flies with intact antennae, as shown in Figure 2.4 of the previous chapter. Second, the antenna effect increases robustness to changes in gain or time delay. It takes a much smaller increase in time delay (such as from cold weather or the effects of ethanol) or increase in gain to make a fly without aristae unstable with respect to forward velocity control. A number of effects can increase visual gain: odor [68], visual brightness or contrast [69, 70], state of excitation, fly-to-fly variability, or most importantly, in confined spaces the visual response will be higher for a given speed because of higher angular rate of visual motion across retina [71]. The antennae give the fly robustness to limitations in its visual system.

3.3 Methods

Since all inputs and outputs are received in the moving frame of the fly, we considered the feedback mechanism from that perspective. Thus neither were the fly's inputs directly observed nor its outputs directly measured, but inferred using both information about the fly's velocity and tunnel-frame inputs.

As the fly moves, it experiences visual motion on its eyes, which we will call v_v for visual velocity. It is necessary as well to define a direction convention, and so we defined positive to be the direction experienced by the fly when it is moving forward, known in the literature as "progressive" visual motion. We likewise defined positive airspeed v_a to be in the same direction (Figure 3.1).

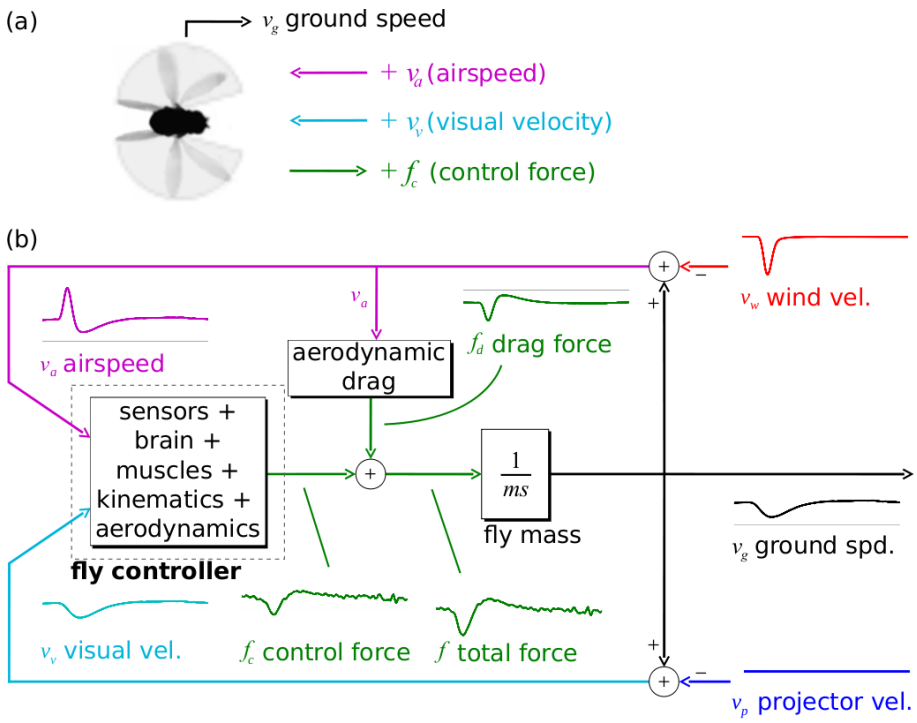

Figure 3.1: A diagram of how tunnel-frame inputs and outputs are transformed into fly-frame inputs and outputs. The two inputs, visual velocity v_v and airspeed v_a, and one control force output f_c are defined as positive along the direction of the arrows (a). A block diagram (b) indicates how tunnel-frame disturbances (right) are transformed into the moving frame of the fly, with example traces showing how a wind disturbance would be propagated. The goal is to find an abstract model of the box labeled "fly controller." Color conventions introduced for different quantities in this figure are used throughout the paper.

In the fixed geometry of the arena, we defined the visual velocity v_v as a linear term, but the fly probably measures angular rate across the retina [43], averaged over portions of the visual sphere. The mechanism for visual motion detection is likely a correlator [72] combined with spatial averaging by tangential cells [73, 23]. More detail is given in the introduction to Chapter 5. That the fly is likely measuring angular visual rates has implications for feedback control in varying geometries, to which we will return in the Discussion.

Because the dynamics of the fly can not be approximated as primarily viscous in nature at the timescale considered here, we performed modeling in the domain of forces. Previous studies on free-flight forward flight [48, 47] have modeled fly visual flight speed response in velocity domain only. But in this work, because the response from the antennae is much

faster, it is necessary to model the dynamics in the force domain. Equating forces and accelerations along the x-axis, the force-balance equation for the fly is

$$f_c + f_d = m\dot{v}_g, \tag{3.1}$$

where f_c is the active control force generated by the fly in response to sensory stimuli and f_d is the passive drag force arising from baseline wing kinematics . The fly changes its wing kinematics from baseline motions to alter the active force f_c, but we do not address the specific changes that occur in this work. Results from [74] indicate that passive wing damping drag force f_d is roughly proportional with airspeed v_a, as do our results and a simulation on an aerodynamic flapping-wing fruit fly model based on data taken from a dynamically-scaled model in a tow-tank [49] (and see Figure 3.2). A dynamic element with force proportional to velocity is known in control engineering as a "dashpot" or "viscous damper." Call this coefficient of proportionality b, so that $f_d = -bv_a$. If inertial forces dominate (short timescales), the force balance equation (Equation 3.1) reduces to $f_c = m\dot{v}_g$, and if viscous forces dominate (longer timescales), then it reduces to $f_c = bv_a$. The time constant of exponential relaxation to asymptote velocity in response to a steady force in the full force balance equation is equal to $\frac{m}{b} \approx .13$ s. Considering delay from vision reported in [48] is 0.081 - 0.1 s, the slower velocity-domain viscous dynamics may be sufficient for the purposes of purely visual modeling. But the faster antenna response we consider in this work necessitates inertial dynamics so the slower viscous domain dynamics approximation cannot hold.

Moving to the domain of active sensory feedback, we seek a model for both drag and active control forces. But rather than perform an exhaustive search among all possible nonlinear models, we considered a set of simple, plausible linear models. If a satisfactory one could be found, no further analysis would be necessary. And in general if they do not have strong nonlinearities, nonlinear systems can be approximated by linear ones in a large-enough neighborhood around equilibrium. Thus complicated nonlinear spiking ensembles of neurons performing computations in the nervous system of the fly may be abstracted as linear. Because the drag was roughly linear, this was considered as further evidence that linear models could be sufficient. The fitting approach was to perform least-squares fits on each candidate model. For clarity, the results are segregated by functionality into delay and

gain boxes, but most likely both operations are being performed by a single sensory-motor transduction cascade.

To estimate velocities and accelerations from position information, we filtered using causal linear filters and a first-order hold input model (linear interpolation between input data points). All filtering and fitting computations were performed using the python-control package version 0.3 (http://python-control.sourceforge.net/) and SciPy (http://www.scipy.org). Accelerations were calculated from position using a filter of the form $\frac{s^2}{(\tau s+1)^2}$, whereas velocities were estimated using the filter $\frac{s}{(\tau s+1)^2}$. When the information was already in final form, such as wind velocity measurements from the hotwire anemometer, it too was filtered, but with the filter $\frac{1}{(\tau s+1)^2}$ so that any phase lags induced by the filtering were replicated in all data. The time constant τ used in the filters was 3 ms.

3.4 Wing damping

It was first necessary to find a relation for the passive drag forces acting on the fly. To eliminate antenna feedback, the aristae of flies were removed, and to eliminate visual feedback, we considered only data coming from the early, pre-visual period of the gust. Based on experiments on tethered Drosophila performed by [74], as well on a quasi-steady model of flapping *Drosophila* wings based on data from a dynamically-scaled tow tank [49], we expected the drag force to be roughly proportional to the airspeed, or $f_d = -bv_a$. While airspeed drag at the scale of the fly would be expected to be primarily inertial, and thus proportional to airspeed squared, this does not take into account the effect of flapping wings. Consider this approximation: suppose wing angle trajectory can be aproximated as a sawtooth (rather than a more realistic sinusoid [9]), with v_a the free-stream airspeed and w the mean velocity of the wings relative to the fly, with $w \gg v_a$. Then drag on the downstroke (into the wind) is $f_d = -\alpha(v_a + w)^2$ according to the relation for intertial flow drag, where α incorporates coefficients of surface area, coefficient of drag, and air viscosity. On the upstroke, the drag reverses direction because the wings are moving much faster than the free-stream, and the drag force is $f_d = \alpha(w - v_a)^2$. Since upstroke and downstroke take equal time, the stroke-averaged force is just the average of these two, or $\overline{f_d} = \frac{1}{2}\alpha(-v_a^2 - 2v_aw - w^2 + v_a^2 - 2v_aw + w^2) = -2\alpha v_a w$. Thus the drag on the flapping wings is proportional to v_a assuming w is constant (assuming negligible effects from the

body [49]).

To perform a least-squares fit, we assumed the fly was in steady-state at the start of the trial with constant forward groundspeed. That is, drag and thrust forces were exactly balanced out. Then, during the gust, any change in drag force was detectable by a corresponding acceleration, or $\Delta f_d = m\dot{v}_g$, with the mass m of the fly known beforehand. Since both the groundspeed and windspeed are known, then $v_a = v_g - v_w$ and a fit can be performed for the data $-bv_a = m\dot{v}_g$. We found that this model fit the data reasonably well, as well as predicting a zero active force during the various gusts (Figure 3.2).

3.5 Antennae model

A model for antenna feedback was found by subjecting flies with intact antennae to different wind stimuli and only considering data during the early, pre-visual period of each trial so that visual feedback was eliminated. We proposed three different linear feedback models: derivative with delay, lag (low-pass), and proportional with delay. Because most mechanosensory input is fast-adapting (derivative-like), an integral feedback response, which would require the fly's nervous system to take two integrals of a derivative-like input, seemed unlikely and was not considered.

The three strengths of wind gusts were chosen so that their peak was approximately 0.5 m/s and ranged from the shortest time impulse the air piston could generate (a 15 mm gust in ~40 ms) up to a gust with the maximum throw of the piston, 125 mm.

For each candidate controller the goal was to find the best-fitting gain K_a and delay/lag T_a mapping the input airspeedv_a to the outupt control force f_c under all of the conditions. The frequency-domain transfer functions for the derivative, lag, and proportional controllers were $e^{-sT_a}sK_a$, $\frac{1}{T_a s+1}K_a$, and $e^{-sT_a}K_a$, respectively, where the exponential e^{-sT} represents a time delay of T sec. For each trajectory, the input was a vector of v_a values sampled at the frame rate of the camera and the output was a vector of f_c values at the same sampling rate. To find the best-fit K_a for a given T_a and candidate controller C_a, the following steps were performed:

1. For each trajectory, recenter initial inputs and outputs to zero by subtracting initial conditions v_{a0} and f_{c0}, assuming that at the start of trial the fly was at steady-state. This gives a recentered input vector $\tilde{v}_a$ and output vector $\tilde{f}_c$.

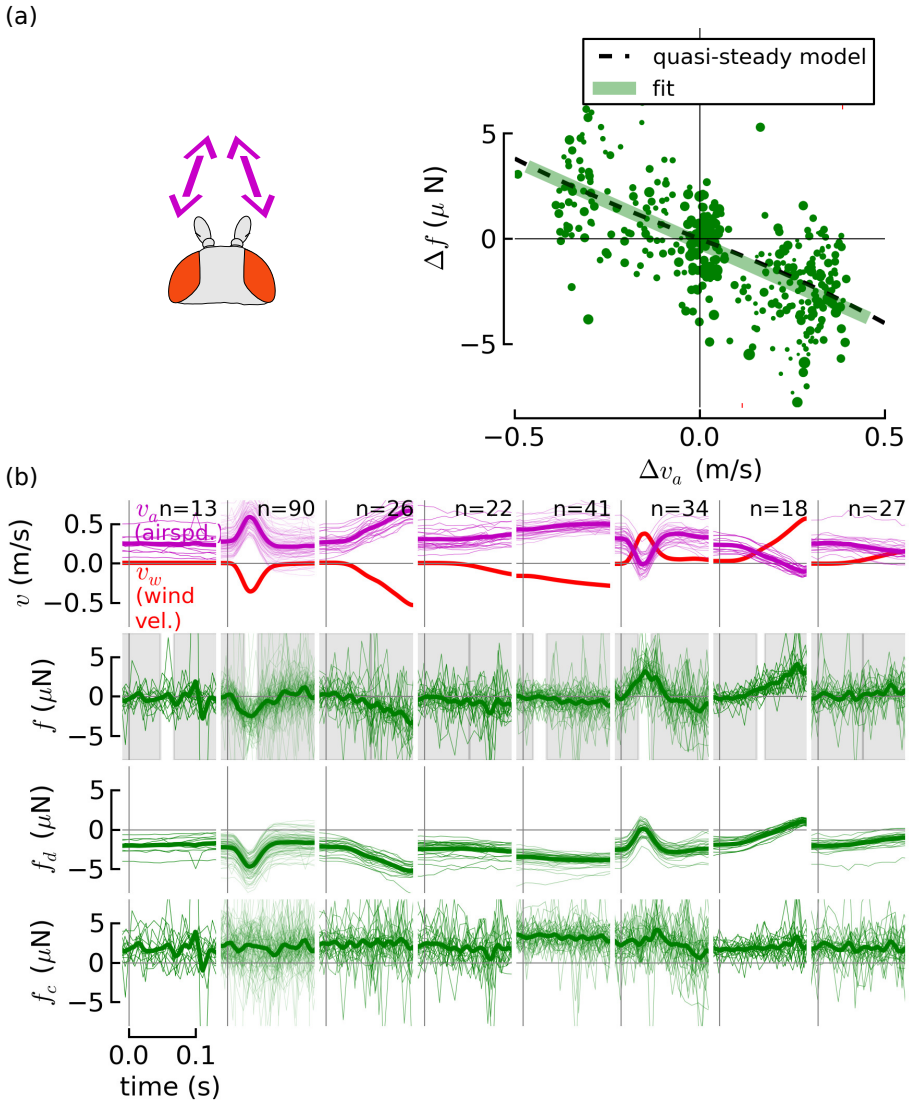

Figure 3.2: Aerodynamic wing drag is roughly proportional to airspeed. To estimate the drag coefficient b of $f_d = -bv_a$ where f_d is drag force and v_a is airspeed, the accelerations of flies were fit to their airspeeds as they were perturbed by gusts of wind (a). The fit is in remarkable agreement with a prediction from a quasi-steady aerodynamic model of the fly (dashed line). To eliminate feedback, the aristae were removed and only data collected during the early, pre-visual phase of the gust were considered. Tunnel-frae wind inputs v_w (red) and corresponding per-fly airspeed inputs v_a (magenta) are shown in the first row of (b), with the mean shown with a thicker line. The second row shows total force measured from fly accelerations $f = m\dot{v}_g$, with the mean in a thicker line. Only data in the unshaded time periods were used, when the windspeed $v_w > 0.1$ m/s and before the visual response. The resultant calculated drag force $f_d = -bv_a$ is shown in the third row. The corresponding active force $f_c = f - f_d$ (bottom row) is essentially unchanged during the different gusts, suggesting the fit is good.

2. Calcuate the response of the candidate controller to the input $\tilde{v}_a$ without the gain factor K_a. For delays, shift the input vector to the nearest whole-number frame, and for the lag controller, calcuate the unity-gain controller response using the `lsim` command in python-control (equivalent to the command of the same name in MATLAB). This gives $C'_a(\tilde{v}_a)$, response of the delay or lag component of the controller (but not the gain).

3. Find the value K_a that minimizes the squared residual error $E = (\tilde{f}_c - K_a C'_a(\tilde{v}_a))^2$ by taking the pseudoinverse.

4. With an estimated K_a for each trajectory in a given condition, take the median K_a from the sample distribution as the estimate to minimize the effect of outliers. This was necessary because of the tendency of flies to occasionally accelerate or accelerate, likely an adaptive behavior so that their paths cannot be easily predicted by predators.

5. Take the mean K_a for all conditions so that each condition is weighted equally, so that conditions with more data are not weighted more heavily. Because weaker stimuli had a low signal-to-noise ratio, only a subset of the conditions, conditions 2, 3, 6, and 7 (counting from the left in Figure 3.3), were used in the fit, but others are shown for validation purposes.

The preceeding calculates the gain for a given delay. To calculate the best delay, the following procedure was used

1. Calculate a gain K_a as described above for each of a range of different delays/lags T_a.

2. Choose the delay (and corresponding gain) with the least residual error for conditions 2 and 5 (the only ones with significant timing information during the pre-visual period).

To then choose the best candidate controller, take the simplest one whose best fit had the lowest overall residual error.

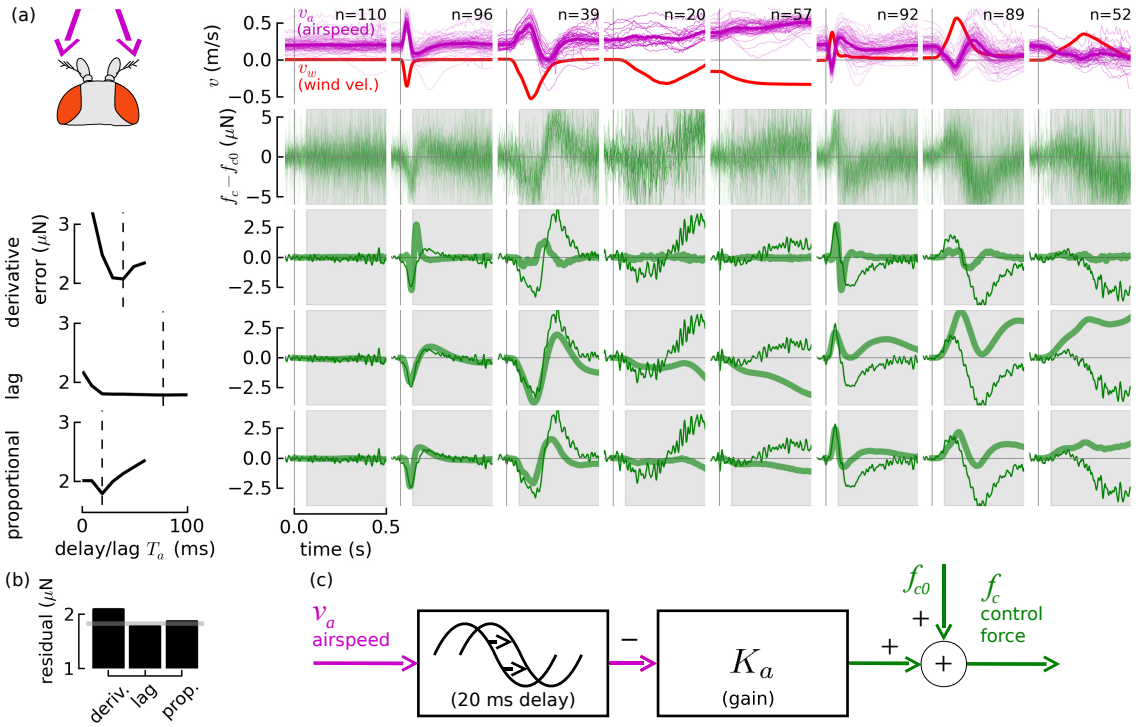

Figure 3.3: System identification supports a proportional controller with a 20 ms delay as the antennae feedback controller. The effect of visual stimuli was eliminated by considering only responses during the early, pre-visual period (unshaded portion of plots). A range of tunnel-frame wind stimuli conditions v_w (a) (top row, red) and corresponding fly-frame airspeed stimuli v_a (magenta) were presented to flies. Their force responses f_c, recentered to zero, are shown in the second row. We proposed three controller models, derivative, lag, and proportional, and fit them to the flies' input-output behavior. To fit the delay or lag time, we performed the fit for a range of delays (left column) and chose the delay or lag (dashed vertical line) with the lowest residual error. We simulated each controller with the mean v_a as input (thick green line, lower three rows) and compared the response to flies' mean force respones (thin green line). The model and the data diverge after 0.1 s because of the effect of the visual response. The residual error for all trajectories and conditions for each controller is shown in (b), compared with the standard deviation for condition 1 (grey line), which had no input. The simplest controller that can explain the observed controller with delay, shown it in block diagram form (c).

The simplest controller that could explain observed behavior with reasonable error was

the proportional controller with a short delay, shown in block diagram form in Figure3.3. The gain for the chosen proportional controller is 7.8×10^{-6}, similar to damping the damping coefficient of the wings at 8×10^{-6}, so the effect of the active antennae response is to essentially double wing damping.

We remark that the model for the airspeed feedback exhibits the behavior observed in the previous chapter, that the feedback response in effect augments the gust. The effect is observed in about equal measure in both directions. The fly appears to be regulating around a desired airspeed, decelerating or accelerating if it is not matched. If subject to a headwind gust, a fly with intact aristae would be expected to quickly decelerate to match the airspeed, as in shown Figure 2.5.

We emphasize that a derivative controller cannot account for the observed behavior. A derivative controller would exhibit a very weak force response in gentle gusts (as shown in the derivative controller's response to gentle gusts in Figure 3.3) because of the small airspeed derivative, but this is not observed.

3.6 Visual model

Next we found a model for visual feedback by subjecting arista-ablated flies to visual stimuli so that their feedback response was due only to vision. We assumed feedback from the antennae became steady-state after arista removal because of the slow adaptation observed in all mechanosensory systems.

We proposed three simple linear models, proportional, integral, and proportional + integral controllers, with corresponding frequency-domain representations $e^{-sT_v}K_{vp}$, $e^{-sT_v}\frac{1}{s}K_{vi}$, $e^{-sT_v}(\frac{1}{s}K_{vi} + K_{vp})$. A derivative model was eliminated out of hand because it would require an extremely long associated delay (taking a derivative adds phase lead—the opposite of a delay). Visual velocity timecourses were chosen by finding flies' mean velocity responses to the wind stimuli chosen in the previous section. The fit procedure was performed as described in the previous section, except with visual velocity v_v as the input and the following changes: (1) Fits were performed over 0.5 s rather than just the 60 ms pre-visual period. Longer fit periods did not yield reasonable fits (fit gains too small), possibly because of visual motion adaptation in the vision system [75, 76]. (2) We used conditions 2, 3, 5, and 6 (from the left) of Figure 3.4 to perform the fit. The other conditions had a much

lower signal-to-noise ratio. (3) Condition 2 was used for finding time delays. For the proportional+integral controller, the least-squares was extended to fit to fit both gains at once, and the median of each gain vector was taken.

The visual feedback model includes a v_d input, which was included to be consistent with the finding that flies maintain a steady-state groundspeed independent of airspeed [54], a behavior likely dependent on vision.

The overall residual error for the three different controllers showed comparable performance (Figure 3.4). The proportional controller was discarded because it could not explain the finding by [54] and reproduced in this lab (data not shown) that flies' groundspeeds are independent of wind velocity, which necessitates a high gain at low frequencies, or a more complicated wind-speed estimator (see Discussion) to eliminate steady-state error. The proportional+integral controller was slightly better, but is more complicated. For the sake of simplicity, and because it required a more neurologically plausible shorter delay (60 ms vs 150 ms), we prefer the integral controller.

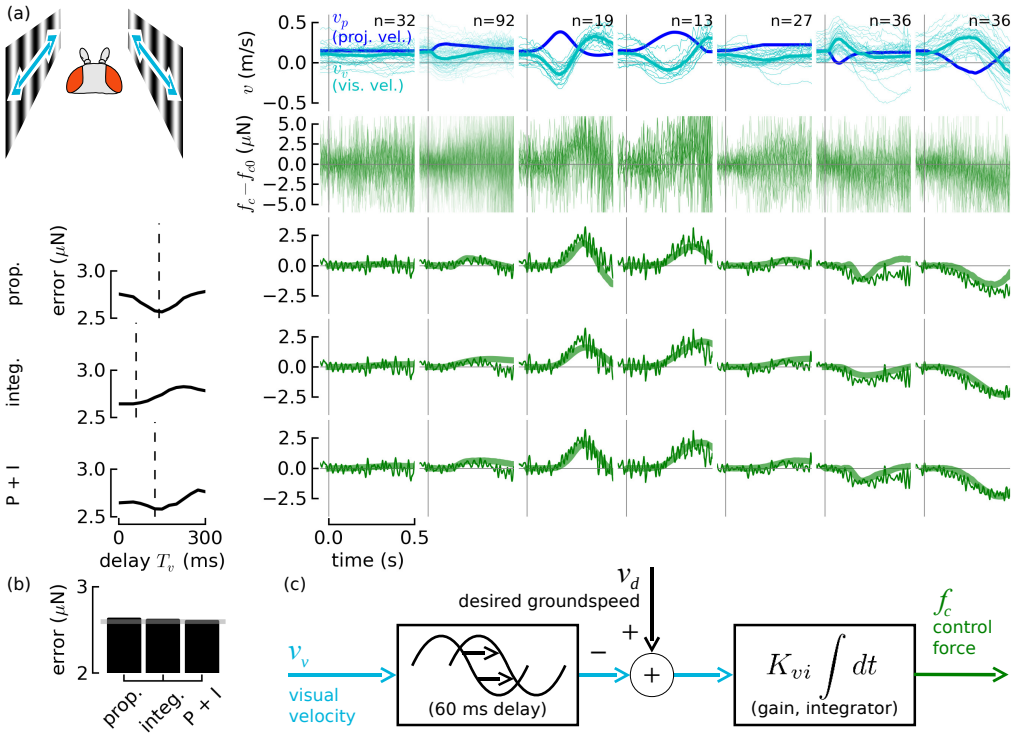

Figure 3.4: A good fit for the visual controller is an integral controller with a 60 ms time delay. The effect of antenna feedback was eliminated by removing the aristae. A range of tunnel-frame visual stimuli conditions v_p (a) (top row, dark blue) and corresponding fly-frame visual stimuli v_v (turquoise) were presented to flies. Their force responses f_c, recentered to zero, are shown in the second row. We proposed three controller models, proportional, integral, and proportional+integral, and fit them to the flies' input-output behavior. To find the correct delay, we performed the fit for a range of delays (left column) and chose the delay (dashed vertical line) with the lowest residual error. For validation we simulated each controller with the mean v_v as input (thick green line, lower three rows) and compared the response to flies' mean force respones (thin green line). The residual error for all trajectories and conditions for each controller is shown in (b), compared with the standard deviation for condition 1 (grey line), which had no input. Based on this data, the simplest controller that can explain observed behavior is an the integral controller, shown it in block diagram form (c).

3.7 Sensory Fusion

With models for the fly's response to pure visual and pure antennae, the next question was how these two senses interacted. Does one supersede the other during a strong gust or do they sum?

We addressed this issue by subjecting intact flies to different gusts with different visual stimuli. Under conditions of an impulsive gust of wind, a fly is normally subject to both an airspeed stimulus as well as a visual stimulus as its groundspeed changes due to drag forces arising from the gust. We subjected the fly to either only the visual stimulus associated with the gust, or only the airspeed stimulus associated with the gust by animating the walls in lockstep to the average velocity response of the flies, thereby reducing the visual slip normally associated with the gust. In an attempt to further eliminate visual stimulus during the gust, we implemented a virtual open loop to move the walls with the fly's actual motion in realtime so that it experienced no visual slip, but because of the ≈ 50 ms tracking latency it was impossible to keep up with the fast dynamics of the antennae response during the faster gust transients, so this approach was abandoned.

When the responses in the two cases were compared to the response of flies subject to a normal gust, that is, with both the visual and antennae stimuli, we found that the behavior of the latter is close to the sum of the two former (Figure 3.5). That is, under the conditions tested, the responses of the two senses simply add. Strengthening this conclusion, we performed an experiment in which we reversed the direction of visual stimulus during the gust, doubling the normal gust visual stimulus. Under these conditions, the fly's behavior is is equivalent to the sum of the regular gust and the pure visual gust responses (Figure 3.5 (e)).

In Figure 3.5 we compare the response of a model that takes sum of the responses of the two individual senses to the data. For the most part, they match, except for after heavy positive accelerations (denoted by grey marks at the bottom of the figures). A possible explanation is that the flies may be decelerating in response to visual looming of the darkened end of the wind tunnel as it approaches them from the front. Dark looming stimuli induce takeoff in walking fruit flies [39] and a landing response in tethered flies [77], even in the absence of a wind stimulus, so it seems reasonable they should elicit some effect on flying flies. The effect is observed about 100 ms after the acceleration, a delay that is approximately

equal to the visual delay. An avoidance behavior would not be expected (and is not observed) in headwind gusts, which induce negative acceleration, because any visual looming would occur in the rearward portion of the visual field is not sampled by the fly's eyes. To verify that there was not a directionality associated with the apparatus or the room, the direction of the experiment was reversed so that the fly was flying toward the air piston rather than away when the gust happened, and the looming response was observed only in tailwind gusts in that direction as well (data not shown).

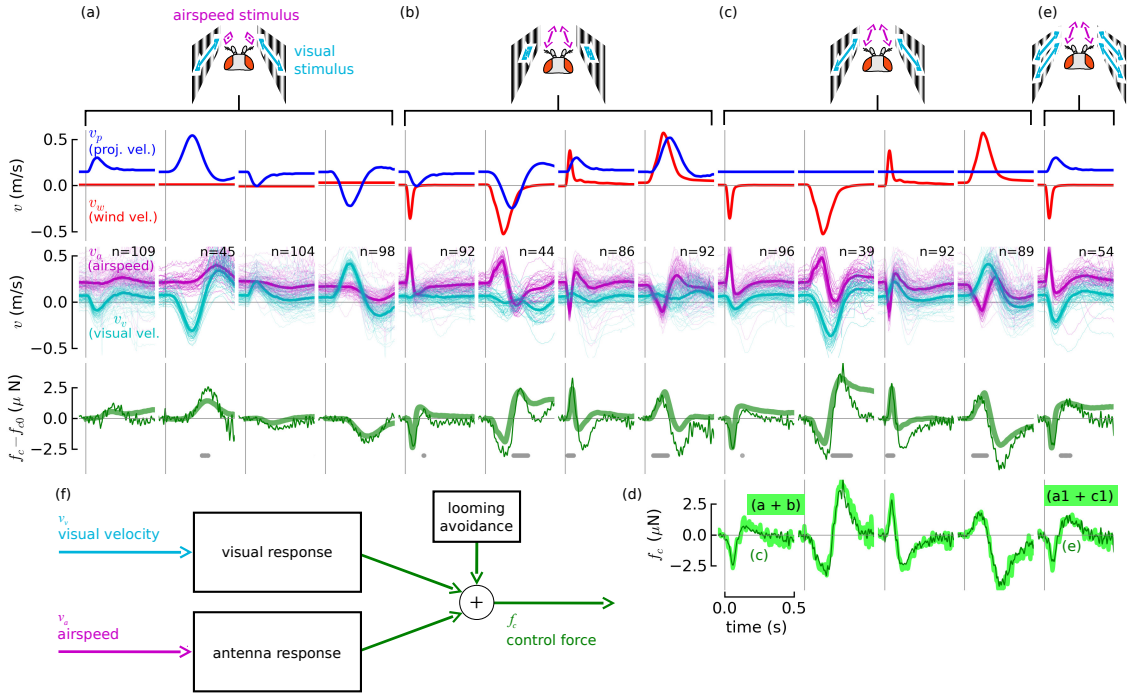

Figure 3.5: Vision and antennae response forces combine by summing. Force responses were measured for flies subjected to either visual gusts with little airspeed input (a) or pure wind gusts in which visual stimulus was substantially removed (b), or regular gusts with no change in v_p during the trial (the naturalistic scenario). In (d), the sum of (a) and (b) is compared to (c), showing a remarkably close correspondence. Strengthening this conclusion, flies' response to doubling the visual input (e) is roughly equal to the sum of a regular gust (c) and a visual gust (a). A block diagram showing this model is shown in (f). Simulating the responses to mean visual and antennal inputs and then summing (thick green line, third row) compares favorably to flies' measured responses (thin line, third row). Where the two diverge seems to be ≈100 ms after tailwind-induced accelerations (denoted by grey marks at the bottom of the figures), consistent with a looming avoidance response.

3.8 Predictions of the model

The findings are collected in block diagram form in Figure 3.6. It was impossible vary v_d experimentally, so we could not know exactly where it entered into the fly's feedback system. Given the ambiguity, we fed v_d into the antenna feedback loop as well, which is equivalent to an efferent copy from the central nervous system [78]. This arrangement would speed up the fly' response to changes in v_d by recruiting the faster antenna feedback system.

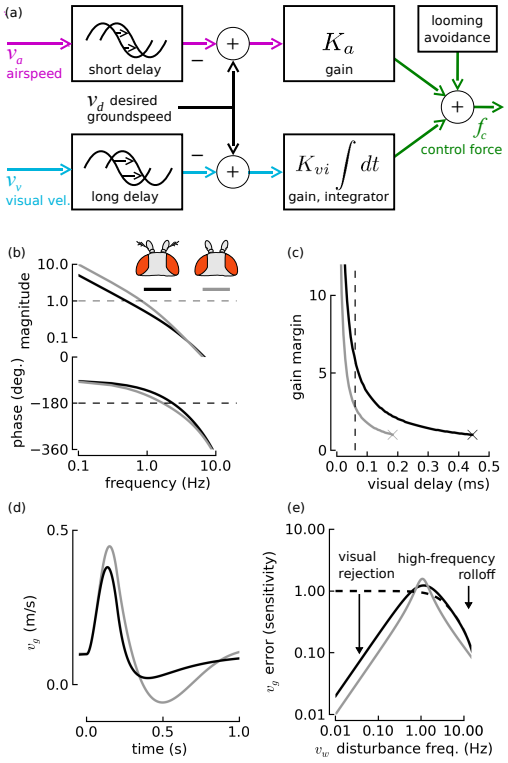

Figure 3.6: The antennae wind sense increases robustness but increases susceptibility to wind gusts at intermediate frequencies. The complete model is shown in (a). A frequency response magnitude-phase Bode plot (b) of the open-loop transfer functions (breaking the loop at the visual controller) shows how the antennae increase both gain and phase margin, reducing overshoot and enhancing robustness. Across a range of possible visual delays, the gain margin of flies with antennae is higher (c), with the × denoting onset of instability. The dashed line shows experimentally-derived time delay. A time-domain simulation on a quasi-steady flapping-wing aerodynamic model [49] subject to a sudden force disturbance (d), shows how the antennae reduce overshoot. Simulation of the closed-loop feedback system to wind disturbances at different frequencies (e) shows that flies with intact aristae show greater relative error between desired and actual groundspeed in gust disturbances from about 0.5 to 0.9 Hz, but acceptable performance elsewhere. The dashed line indicates performance of a hypothetical fly without visual feedback but with antenna feedback, showing that the effect of vision is to reject wind disturbances at low frequencies. The effect of the antennae reduces performance somewhat, by a factor of approximately 2. But this is a negligible effect at low frequencies, where vision-mediated disturbance rejection is the dominant effect with a factor much greater than 2.

Because of the mathematical nature of the model, there are a number of equivalent interpretations. One is that the antenna mediated behavior has a damping effect, counteracting the oscillatory behavior observed in flies with ablated aristae that have to rely only on long-delay visual feedback (Figure 2.4). Equivalently, the antennae increase gain and phase margin of the control loop, raising the gain margin from 2.8 to 5.9 and raising the phase margin from 39° to 70°. Increased margins are equivalent to being farther from the onset of instability, enhancing robustness to variations in these parameters. A controller at the edge of stability is expected to show greater overshoot. The model predicts just that, showing much less overshoot after an impulse force disturbance for flies with intact antennae (Figure 3.6).

For validation, we have plotted the predictions of the model compared to the mean flight velocities of flies in the velocity domain for all conditions tested (Figure 3.7). The model predicts a step response similar to that found in previous work [47] (Figure 3.8). In that figure we show that if the visual gain is increased for flies without their aristae, the feedback system can go unstable.

3.9 Discussion

Earlier work suggested flies could sense wind with their antennae, but exactly what the antennae measured and how it was used to control flight was not known. By tracking the flight behavior of flies in gusts of wind and measuring their acceleration rates to estimate force outputs, and fitting that to a linear model, we conclude the following about how antennae augment vision in forward flight:

1. The antennae response to airspeed changes is much faster than visual response.

2. The antenna-mediated force response is linearly proportional to airspeed. This is functionally equivalent to a "dashpot" or "viscous damper" terms used in control engineering, and doubles the airspeed-proportional damping drag from aerodynamic on the wings alone.

3. The force responses of the two senses simply add. This allows the visual response to override the counteracting force of the antennae if there is a headwind.

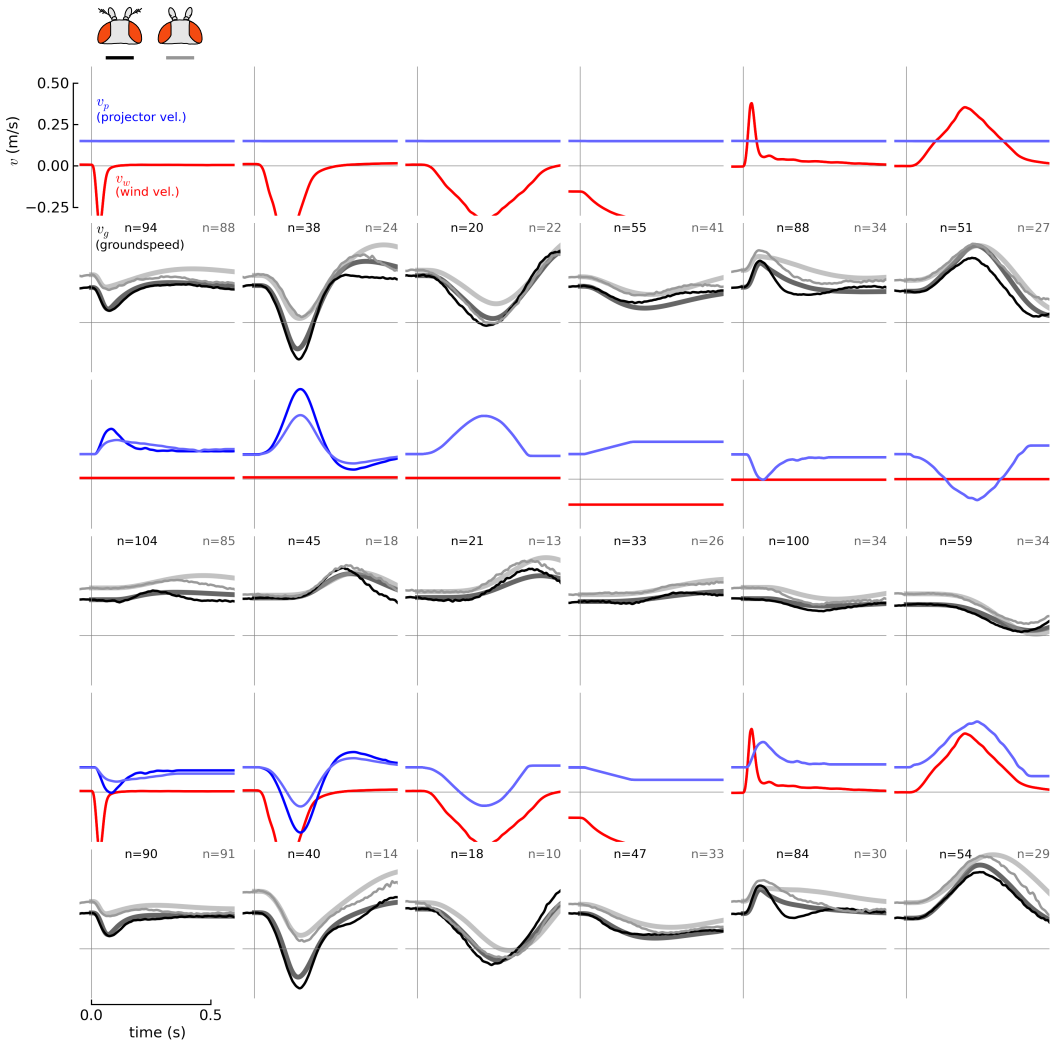

Figure 3.7: Comparison of response of model compared to data in the velocity domain. Thicker translucent lines are model. Initial mean velocity of flies was was added to model response in each condition for easier comparison to data. Light blue lines correspond to visual stimuli given to aristae-less flies. The light blue lines are different because goal was to eliminate visual stimulus during gust, and the gust responses of the aristae-less flies are different in the fast gusts.

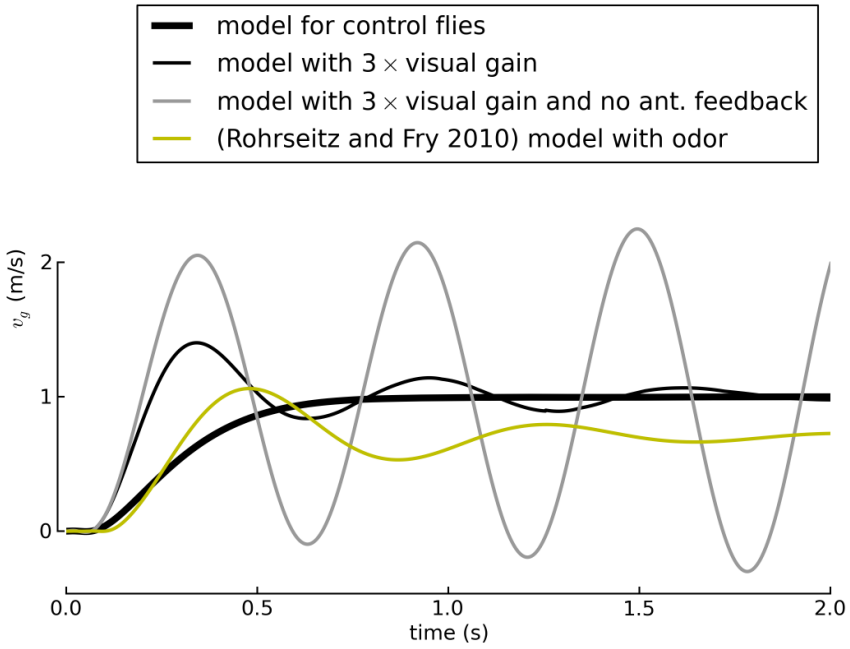

Figure 3.8: Step response of the velocity-domain model reported in [47] to our model. Also shown is a step response with a 3× higher visual gain. This could occur if the sun was brighter or visual environment was smaller, both of which could result in a higher effective gain. But if the gain is increased on flies with no aristae, the controller goes unstable.

4. The effect of the antenna response is to damp out, and thus to slow, the forward flight dynamics so it is easier to control by long-delay visual feedback, giving the fly robustness to parameter variability.

An equivalent damping effect could be achieved by adding a physical damper to the fly, such as large bristles or a tail. With the added drag, however, the fly would have to expend more energy to move upwind. With an active, controllable antennae response, on the other hand, the fly accrues the dynamic benefits of damping but need only override their effects inside the brain, a much more efficient proposition.

3.9.1 The gust response paradox

A paradox arises out of this analysis. The fly's antennae response seems to augment the force of the wind, rather than diminishing it. We remark that this may be indirect evidence against a fast, proprioceptive acceleration sense in fly, because if present it would likely act to counteract the observed gust-force augmentation by the antennae. Flies can sense the direction of gravity in the dark [57], but it may be a slow response, or may not function during the vibration-heavy motion of flight.

One possible explanation lies with the character of air flow encountered by the fly. The spectrum of wind turbulence tends to follow a power law with decreasing strength with increasing frequency (for example, the Kolmogorov spectrum, which goes as $f^{-5/3}$). It may be that the magnitude of turbulence disturbances at high frequencies, where the antenna response is significant, is small enough not to be of concern for the fly. Whereas the larger, low-frequency disturbances can be adequately compensated for by vision and its high gain at low frequencies (Figure 3.8(e)). At low frequencies below about 0.5 Hz, the increase in error due to the antennae remains a constant factor of only 2, which is greatly overridden by the visual sense.

3.9.2 Remarks on feedback architecture

The antenna provide some of the same benefits as the derivative term in a standard proportional-integral-derivative (PID) controller: less phase delay at the gain crossover frequency. We did not expect the antenna feedback behavior to match a proportional regulator, however. A derivative response would have provided more phase lead and would have been in keeping

with usual phasic, derivative-like behavior of mechanosensors [50, 79]. A potential benefit of measuring the absolute airspeed, however, rather than its derivative, is that the information could be used in combination with a visually-derived estimate of groundspeed to estimate wind direction and magnitude for the purpose of odor localization and flight control.

Among the three proposed visual controllers, all had nearly the same performance on the test data, including the controller with proportional feedback. While compelling for its simplicity, we do not think the proportional controller is a reasonable possibility. Unlike the controllers with integral terms, which have high gain at low frequency to eliminate steady-state error arising from headwind drag, the proportional controller has no such effect. In a steady headwind, a fly with a basic proportional visual controller would show significant error, failing entirely to maintain a forward groundspeed if the wind was too strong. But fruit flies are able to maintain their groundspeed independent of the strength of the wind up to 1 m/s [54] and reproduced in our lab (data not shown). This feat could nonetheless still be achieved with a proportional visual controller, but would require a number of elaborations. The fly would need an internal model of the value of b, the drag coefficient on the wings, and would require computing a delayed version of the antenna feedback to compare with visual feedback to avoid compromising performance. This seems complicated, but not impossible. The best-fit proportional controller also has a long, neurophysiologically unlikely 150 ms delay, whereas visual response delays reported elsewhere are on the order of 30 - 100 ms [80, 45, 47]. The integral controller, on the other hand, has a delay of only 60 ms because the integration inherently adds some phase lag. For simplicity and neurophysiological plausibility, we prefer the integral controller.

We have proposed a model in which feedback from vision and the antennae operate in parallel feedback loops, each of which produces a force output. Another possible arrangement would be an inner-loop/outer-loop topology. The typical application is for two interdependent dynamic systems that have significantly different response times. For the purposes of designing the outer loop, the response of the inner loop is assumed to be nearly instantaneous. An example is a helicopter using its pitch angle to regulate forward velocity: the helicopter pitches forward to accelerate. The inner loop controls pitch angle, and an outer loop gives a desired pitch to the inner pitch angle controller to regulate forward velocity. An example of this principle applied to the fly in a simulator is given in [81].

It is a matter of block diagram algebra to rearrange the feedback loop structure reported

here in Figure 5.7 to be inner-outer loop, but with indistinguishable behavior. The output from the visual system could feed into the antennae system by adding to its v_d term as a correction factor δv_d compensating for wind disturbances. If there had been evidence of of similar poles or gains appearing in the feedback dynamics in multiple places, this could be evidence that they were in fact outputs of the same block. For instance, in the PI visual model, the proportional gain fit is 5e-6, not far from the antennae proportional gain of 7.8e-6, so it is possible they could pass through the same gain.

3.9.3 Relation to previous work

Previous work in bio-inspired feedback control has considered the role of vision, but the ramifications of using visual feedback with a significant delay—and compensation with a second sense—have not previously been addressed. Neumann [82] created a simulated helicopter that maintained its attitude relative to the world using horizon detection and could traverse terrain using an omnidirectional visual sensor and correlators. In that work, there was no dynamic feedback for forward velocity regulation: forward velocity was regulated by simply applying a forward force and letting the modeled robot accelerate until the force was counteracted by an equal amount of drag. Thus, the matter of visual time delay in forward flight control is avoided by dispensing with velocity control altogether. Optimistic open-loop feedforward flight control of this form is unlikely to be used by the fly for the reason that disturbances, such as wind or wing damage, cannot be compensated for. And flapping wing flight is a fast, almost violent and complicated mechanical motion that can generate strong aerodynamic forces. It seems unlikely that such a complicated mechanism can produce an arbitrary desired thrust without some feedback.

In [83] a robot controller was considered that used an array of hair sensors to augment vision. But while in that work it is remarked in passing that the different senses may have different bandwidths, no effort is made to address the question and the feedback controller implemented in simulation does not have any visual delay.

Compared to the work of [47], our visuo-mechanosensory model exhibits roughly similar step response behavior, including a similar visual delay (Figure 3.8), despite the different forms of the two controllers. The slower performance of our flies may be attributable to the fact that they were not exposed to an odor plume during the experiments (for why, see Section 2.3.5). To illustrate the instability that could arise with higher visual gain (such

as from closer obstacles) without antennae, included is the step response for a higher-gain visual controller with and without antenna feedback (Figure 3.8).

We remark that feedback control in the presence of proportional airspeed drag has parallels with the study rotational motion. As in linear flight control models [47], early rotation flight control models assumed that the dynamics were essentially viscous [84, 70]. However, rotation rates of body saccades in free-flight are fast enough that their dynamics were found to be primarily inertial, necessitating an active counter-torque to slow and stop rotation at the end of the saccade [85], presumably mediated by the halteres [86]. Other studies, however, argued instead that passive drag is sufficient to account for the dynamics of saccade completion [87, 88]. The matter has been resolved in the case of drosophila using a dynamically-scaled flapping model in a tow tank of oil. The result, as is often the case, was that both sides were partially correct: passive and active forces act in roughly equal proportion during the completion of a saccade [63]. It is interesting to note that the same holds in forward flight: active and passive force components act in roughly equal proportion during a gust of wind. We are left to wonder, could there be an underlying biological feedback principle underlying both, or is it just coincidence?

3.9.4 Mechanisms

We emphasize that the antennae response seems to be proportional to airspeed input, down to the 1 Hz half-sinusoid impulse signals we tested. This is unlike most mechanosensory input, which is usually fast-adapting, or phasic. While [50] found only phasic, derivative-like electrophysiological results from moving the antennae, it is possible a smaller tonic signal was not detected. Yorozu [62], however, using genetic expression of a calcium-sensitive dye, found locations in the brain innervated by projections from the Johnston's Organ that responded to the magnitude of the wind, rather than its derivative, supporting the finding that the antennae are somehow able to measure absolute airspeed.

We remark that the visual controller has the form of an integral. An integral controller could use the same neural hardware as visual odometry, that is, by taking the integral of velocity, which is position. Visual odometry exists elsewhere among insects: is a vital mechanism for bees [43, 22].

As for thrust generation, a number of possible mechanisms can be imagined, from varying the angle of attack between the upstroke and downstroke (fast) [49, 89], to pitching the nose

downward like a helicopter to generate forward thrust (slower, because it involves rotating the body) [90]. Do visual and antenna input lead to the same changes in wing kinematics? Considerations of efficiency suggest that they should, and that summation of the two senses should take place in the brain, for example by excitatory synaptic input to a common neuron, rather than by opposing but distinct perturbations in wing kinematics, which would presumably require more energy expenditure. We leave this question for future work.

3.9.5 Robustness to variability

One of the most beneficial effects of the antenna response, in control terms, is to give the visual feedback loop more robustness to variation in gain and delay. The delay could increase in lower temperatures or perhaps because of ingestion of ethanol. Variable visual gain could arise from the presence of odor [91] (though free-flight studies have indicated odor may not affect the gain of velocity regulation [92]), brightness of visual scenery, fly-to-fly variability, or distance to obstacles as described in the next section.

3.9.6 Robustness to environment geometry

Gain margin provided by the antennae may be particularly important when the fly is near obstacles or in confined spaces. The closer the visual obstacle, the higher the gain in the visual feedback loop. Bees [43] and flies [54] seem to regulate around the angular rate of visual motion across their retinas. Thus for a given thrust impulse, and thus a given change in forward velocity, a fly near an obstacle would receive a much greater visual stimulus, as shown in Figure 3.9. Because the response is proportional to the error, the stronger response propagates feedback loop and is equivalent to a higher visual gain. The higher gain is problematic because the combination of high visual gain and long visual delay can cause the controller to go unstable. Extra information from the antennae could damp this out and allow the fly to fly stably in a much larger range of visual geometries.

3.10 Appendix: Parameters used in simulations

Parameters used in simulations and curve fitting are given in Tables 3.1-3.3. To ease analysis during fitting, we rearranged the blocks of the feedback model, as shown in Figure 3.10. With this arrangement, closed-loop behavior could be more easily analyzed by using Black's

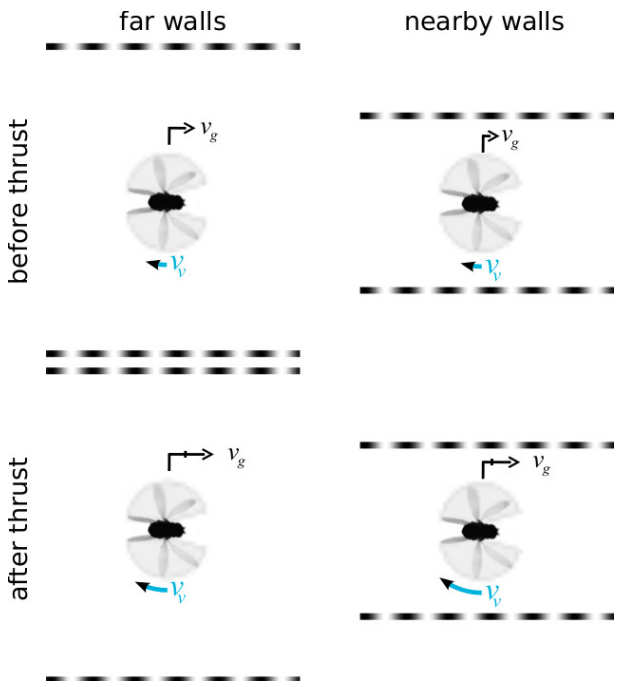

Figure 3.9: How distance to obstacles can affect gain in the visual feedback loop. For a given thrust impulse, and thus a given change in velocity, the fly in the narrow corridor receives much greater visual feedback because the change in angular rate of visual motion flow v_v is much higher. Because the response is proportional to the error, the stronger response passes though the whole feedback loop and is equivalent to a higher gain.

controller	frequency-domain form	gain K_a	delay/lag T_a (ms)
derivative	$e^{-sT_a}sK_a$	1.6×10^{-7}	40
lag	$\frac{1}{T_a s+1}K_a$	2.5×10^{-5}	77
proportional	$e^{-sT_a}K_a$	7.8×10^{-6}	20

Table 3.1: Fit values for candidate antenna controllers

controller	freq. domain form	prop. gain K_{vp}	integ. gain K_{vi}	delay T_v (ms)
proportional	$e^{-sT_v}K_{vp}$	7.5×10^{-6}		150
integral	$e^{-sT_v}\frac{1}{s}K_{vi}$		5.0×10^{-5}	60
prop. + integ.	$e^{-sT_v}(\frac{1}{s}K_{vi} + K_{vp})$	5.0×10^{-6}	3.8×10^{-5}	125

Table 3.2: Fit values for candidate visual controllers

law for negative feedback loop closure. Black's law states that the transfer function of a closed-loop system with forward transfer function $G(s)$ and feedback transfer function $H(s)$ is equal to $G/(1+GH)$ [93].

parameter	symbol	value
fly mass	m	1.0 mg
wing damping coefficient	b	8.0×10^{-6} Ns/m
simulation time step	dt	1 ms

Table 3.3: Parameters used in simulation

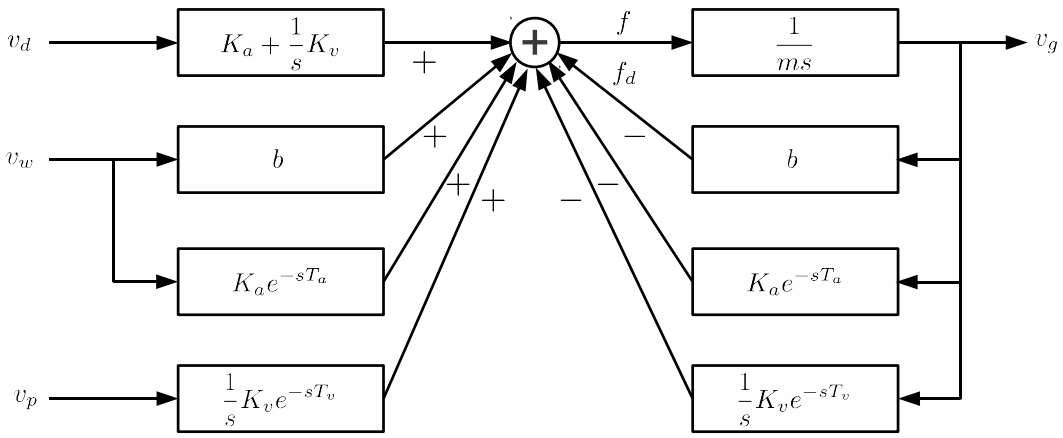

Figure 3.10: Block diagram of feedback system, rearranged for ease of analysis. This diagram shows the best models for antenna and visual feedback, but other models can be accomodated by substituting the corresponding delay and control blocks. Key: v_g ground speed; v_d desired groundspeed; v_w wind velocity; $v_p$projector speed; f total force; f_d drag forcel; subscript a denotes the antenna feedback system; subscript v denotes the visual feedback system; K is a gain and T is a time delay. All velocities in this figure are expressed in tunnel-frame coordinates.

Chapter 4

An Insect-inspired Autocorrelation Model for Visual Flight Control in a Corridor

4.1 Abstract

In this chapter we propose and demonstrate stable robot controllers that use a small omnidirectional array of visual sensors and a fly-like autocorrelation scheme for sensing visual motion using minimal computation. Designing the controllers required deriving a model for the response of an array of correlators observing the motion of a flat moving surface such as the ground or a wall. The model operates in the frequency domain and incorporates the effects of perspective, motion parallax, and spatial blurring. Using it, suitable parameters for inter-sensor spacing and blur width were found that mitigated the effect of incorrect estimates arising from aliasing. Controllers that decomposed the correlator response into harmonics to observe and control the robot's state were implemented on a fan-actuated hovercraft robot. They were able to stabilize it as it moved through a corridor, the first use of correlators to control the motion of a flight-like (non-kinematic) dynamic vehicle.

4.2 Introduction

Autonomous flight by an insect-sized robot will require feats of miniaturization on multiple fronts: actuation, system integration, power use, and even computation. On the actuation side, it may be necessary to replace electric motors and ball bearings with piezo actuators and flexure joints because of scaling considerations [94]. In this work we are concerned with

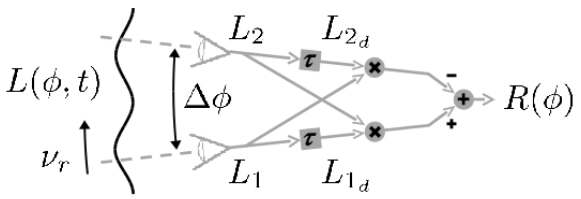

Figure 4.1: The correlator. A pair of visual elements aimed at different angles observes luminance signal L moving across the retina with retinal velocity ν_r. The element τ is a delay or time-lagging low-pass filter and $\times$ is a multiplication or "correlation." For a sinusoid of a given spatial frequency, the correlator response R increases linearly with ν_r for ν_r sufficiently small.

the sensors and feedback control of such a vehicle, aiming to insure it flies stably and keeps collisions to a minimum.

Miniaturization may require dispensing with traditional sensors and looking toward biology for inspiration. Larger unmanned aerial vehicles (UAVs) use radar, lidar, and the global positioning system (GPS) for self-localization. But for a tiny robot that may need to fly in enclosed or cluttered environments, the GPS signal will be compromised. And emissive sensors may use too much power or be too heavy. Nature has found a remarkable solution that functions under these constraints: the fly. Flies are virtuosos among small flying animals, demonstrating superlative agility [38]. What lies behind their agility is a particularly capable suite of sensors. Unlike four-winged bees and dragonflies, flies are dipterans, meaning "two wings," having repurposed their rear set to perform rotation rate sensing. Called halteres, the rear vestigial stubs beat anti-phase to the wings over nearly 180° and can sense rotation around all three axes by sensing strains arising from coriolis forces [95]. Flies carry large omnidirectional eyes and corresponding large brain that both require a significant metabolic investment [96]. They also have wind sensors (antennae and bristles), and possibly use their wings as the emitter of a sonar [38].

Thus the suite of sensors available to our robot may prove to be quite similar to the fly's. There is an ongoing effort to reverse-engineer the fly both for biological understanding and engineering insight [49, 81]. We remark that nature's solution does not compromise: flies are able land on inverted surfaces and right themselves within tens of wingbeats after tumbling surprise take-offs [39], and can navigate complicated enclosed spaces including forests and kitchens with remarkable speed and agility.

In this work we report a controller that is able to use fly-like correlators to stabilize a

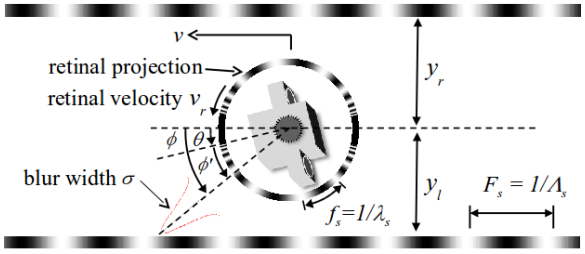

Figure 4.2: A fan-actuated robot with hovercraft dynamics carries an omnidirectional array of correlators and navigates a corridor. The goal is to move along the axis, avoiding the walls. A representation of the projection of the visual environment onto the retina surrounds the robot.

dynamic, force-actuated vehicle with flight-like dynamics. It requires minimum computation, eschewing serial, digital computation (for example, comparing lists of visual features [97]) in favor of simple operations like addition and multiplication that could be implemented in parallel low-power analog on silicon or by neurons. Like the fly, our controller uses an autocorrelation scheme to estimate visual motion using a relatively small number of visual pixels arrayed omnidirectionally (Figures 4.1 and 4.2). This visual motion estimate is then integrated over wide fields of the visual sphere in a clever way, taking sums and differences, to estimate vehicle state in a corridor.

In particular, we report a new model for the response of an array of correlators moving relative to a large flat textured surface. The model is applicable to the common cases of flying over the ground or along walls. A state estimate derived from the flat-surface model is coupled to a controller for a fly-like dynamic (non-kinematic) hovercraft robot that uses fans for thrust and torque generation. This work differs from previous work in that it uses an explicit model for correlators rather than a pure optic flow formulation to control self motion [27], [2]. Other work has reported controlling a kinematic wheeled robot [19] using correlators, but here we tackle the more difficult problem of controlling a dynamic, force-actuated robot. We are concerned with motion in a corridor—particularly centering—because it is an essential behavior required for navigating between obstacles in a cluttered environment. Such a low-level, reflexive controller could operate in real-time, freeing a high-level controller to pursue long range goals such as searching or path planning.

4.3 Insect Flight Control

The fly has a small array of visual elements in its eye, numbering only approximately 1400 in the fruit fly *Drosophila melanogaster* [70]. They are arranged nearly omnidirectionally, except for a small 20° slice in the rear. Evidence from a number of studies [98, 99, 70] suggests these flies use an autocorrelation scheme in which luminance readings between nearest-neighbor visual elements are compared by delaying one and multiplying it by the other (Figure 4.1) to estimate visual motion.

It is not yet well understood the underlying neural mechanism that performs the delay and multiply operations in the correlator, recent clues have emerged [100, 15], and it is not entirely certain why correlation is used rather than a different method such as the gradient method [101, 102]. It may be that autocorrelation-based control is easy to learn [32] or that performing a division or matrix inversion as required by the gradient scheme is too neurally expensive [101]. In any case, correlators are able to sense visual motion because their time-mean response has the same sign and, below saturation, is proportional to the magnitude of the visual motion ν_r. However, they have a number of non-idealities. Their response is strongly dependent on both spatial frequency and contrast in the image [70]. For naturalistic imagery, this leads to a low signal-to-noise ratio that requires averaging over space and time to obtain a lower-noise estimate [69, 76, 102]. Nonetheless, their advantages may outweigh their disadvantages, and in particular they may be sufficient to carry out the feedback controls tasks required by the fly.

Downstream from the correlators, visual imagery is projected retinotopically into deeper layers of neurons in the brain (that is, the topological arrangement image information is preserved). A small number of larger cells, the lobula plate tangential cells (LPTCs) (approximately 30 in *Drosophila*) are found that sum the correlator respose over visual fields as large as a hemisphere [11] [23]. Sums and differences taken of their response may explain insects' tendency to center in a corridor and decelerate in confined environments [101]. Outputs from these cells project into motor centers which then orchestrate changes in wing kinematics to perform flight maneuvers.

4.4 Frequency-Domain Analysis of Correlators

4.4.1 Correlator response to panoramic image motion

To analyze correlators, we first consider the case of luminance readings coming from panoramic image motion as would be induced by self-rotation. Since by Fourier decomposition an arbitrary image can be represented by a sum of sinusoids of different frequencies and amplitudes, we start with an analysis of a single, arbitrary sinusoid.

A single correlator consists of two luminance sensors oriented at slightly different body-centric angles separated by an angle $\Delta\phi$ (Figure 4.1). Suppose a sinusoid luminance signal L with spatial frequency f_s (cycles/rad) moves at ν_r rad/sec in front of the retina. Each sensor reads

$$L(\phi, t) = C_0 \cos(2\pi f_s \phi + 2\pi f_s \nu_r t), \tag{4.1}$$

where t is time. The correlator response is $R = L_2 L_{1_d} - L_1 L_{2_d}$ where the subscript d indicates a delayed or filtered version of the luminance signal. If a pure delay is used, the correlator response can oscillate between positive and negative with increasing ν_r. Accordingly, we preferred using a low-pass filter $\frac{1}{\tau s+1}$ as the delay element because it never goes negative. Assuming zero-mean input, it can be shown that this form of the correlator asymptotically reaches a steady-state (constant in time) response [69]

$$R(t) = \frac{1}{2\pi\tau} C_0^2 \frac{f_t}{f_t^2 + 1/(2\pi\tau)^2} \sin(2\pi f_s \Delta\phi), \tag{4.2}$$

where

$$f_t = f_s \nu_r \tag{4.3}$$

is the temporal frequency of the sinusoid's oscillation as it moves. Initial transients die off with an exponential time constant τ.

4.4.2 Decomposing correlator response

The correlator response to a sinusoid (4.2) can be decomposed into a product of factors

$$R = \frac{1}{2\pi\tau} C^2 T A, \tag{4.4}$$

where

$$C = C_0 \tag{4.5}$$

is a contrast factor that depends only on the amplitude of the luminance input L,

$$T = \frac{f_t}{f_t^2 + 1/(2\pi\tau)^2} \tag{4.6}$$

is a temporal factor that depends only on the temporal frequency f_t of the sinusoid, and

$$A = \sin 2\pi f_s \Delta\phi \tag{4.7}$$

is an aliasing factor that depends only on the product of spatial frequency f_s and angular separation between the pair of luminance sensors $\Delta\phi$.

4.4.3 Incorporating the effect of spatial blurring

We model the luminance sensors as having a Gaussian sensitivity profile, blurring the image. This is effectively a spatial-frequency-dependent attenuation of C. Each luminance sensor has an angle-dependent blurring function

$$G(\phi) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\phi^2}{2\sigma^2}\right),$$

where σ is proportional to the width of the blurring function. If the original luminance signal sinusoid is $L_0(t)$ (Equation 4.1), then by convolving it with G we get the blurred lumincance signal

$$L(t) = G(t) \circ L_0(t).$$

To find the resultant amplitude attenuation, we turn to the frequency domain. Because it contains only one frequency, the Fourier transform $\hat{L}_0$ of L_0 is a pair of delta functions at $\pm f_s$. Using the property that convolution in the spatial domain is equivalent to multiplication in the frequency domain, we take the Fourier transform of G,

$$\hat{G}(f_s) = \exp\left(-\frac{1}{2}(2\pi f_s \sigma)^2\right),$$

where f_s is the spatial frequency, and thus $\hat{L}(f_s) = \hat{L}_0(f_s)\hat{G}(f_s)$. The two delta functions

of $\hat{L}_0$ are scaled by $\hat{G}(f_s)$. When the inverse transform is applied on $\hat{L}(f_s)$, a sinusoid is recovered, but in general with a change in amplitude and phase according to $\hat{G}(f_s)$. Because $\hat{G}$ is a real-valued function, the phase is unchanged and the effective amplitude of the blurred image is thus

$$C(f_s) = \hat{G}(f_s)C_0, \tag{4.8}$$

where C_0 was the amplitude of the original luminance sinusoid (4.1). At low frequencies where $f_s \ll 1/(2\pi\sigma)$, $C \approx C_0$, and at high frequencies $C \ll C_0$.

4.4.4 Incorporating motion parallax and perspective

We would like to extend the equation for the response $R(t)$ (Equation 4.2) to a moving flat surface of infinite extent. To do so, we need only consider how know the spatial and temporal frequencies (f_s and f_t) project onto the retina and change as a function of angle ϕ and state of the vehicle $\mathbf{q}$. The vehicle is moving at a velocity v (m/sec) near the midpoint between two walls pattered with sinusoids with spatial frequency F_s cycles/m (Figure 4.2). The walls are separated by $2y_d$ where y_d is the desired distance the robot wants to keep from the walls. The distance to the left and right walls are y_l and y_r respectively, with $y_l + y_r = 2y_d$. The fly's position $\tilde{y}$ is its distance from the centerline, giving $y_l = y_d + \tilde{y}$ and $y_r = y_d - \tilde{y}$.

The first matter is to find an expression for how the spatial frequency F_s on the walls is projected onto the retina as a (spatially varying) spatial frequency f_s. By multiplying the effect of changing distance to the wall $f_s = \frac{y}{\sin\phi}F_s$ by the effect of changing the angle of the wall $f_s = \frac{1}{\sin\phi}F_s$, the spatial frequency f_s projected onto the retina is

$$f_s(\phi) = \frac{F_s y}{\sin^2\phi}, \tag{4.9}$$

where y is the distance to the wall seen by that sensor and ϕ is the angle relative to the axis of the corridor. (Or alternately, equate frequency $\omega_s = 2\pi f_s$ to rate of phase change $\omega_s = \frac{d\varphi}{d\phi} = \frac{d\varphi}{dx}\frac{dx}{d\phi}$ where $x = y\tan(\phi - \pi/2)$ is the linear distance along the wall from the vehicle and $\frac{d\varphi}{dx} = 2\pi F_s$.)

Second, the retinal velocity or "optic flow" ν_r in the corridor is [27]

$$\nu_r = -\dot{\theta} + \frac{1}{r}\left(v\sin\phi - \dot{y}\cos\phi\right), \tag{4.10}$$

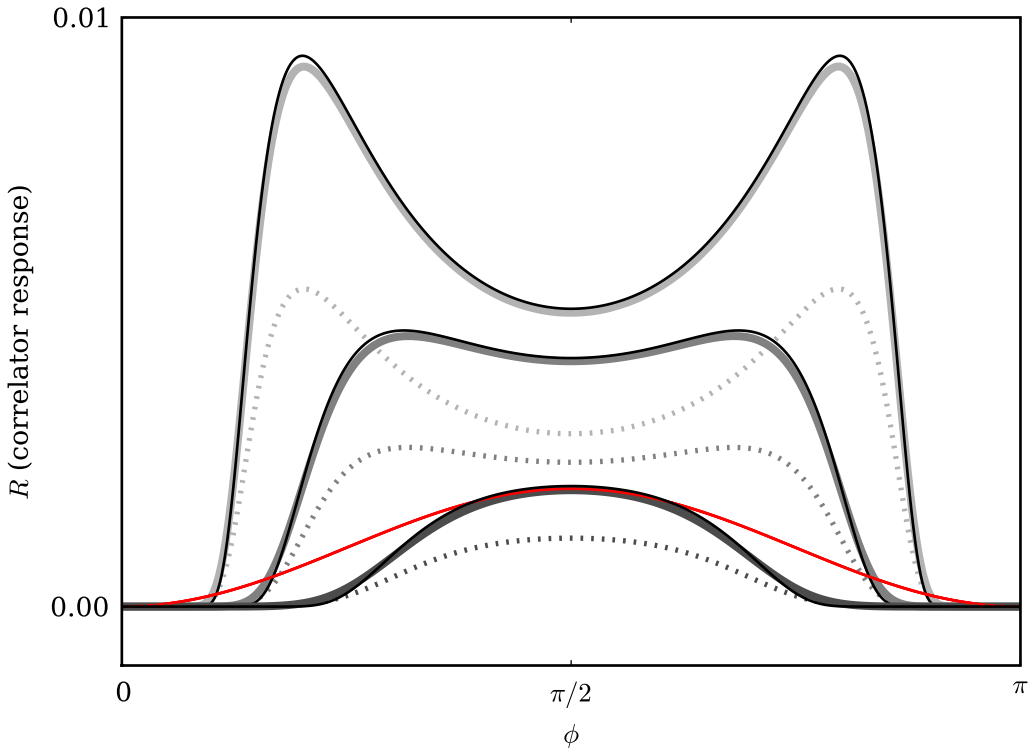

Figure 4.3: Simulated and analytic model correlator response R for different visual sensor blurring widths σ. As σ increases, the response diminishes and simplifies, until it begins to resemble the red line, a scaled version of the retinal velocity or "optic flow" ν_r. Thick lines are simulated at 5 kHz, black lines are from (4.11-4.13), and dashed lines are simulated result at 60 Hz with a zero-order hold and are qualitatively the same. The vehicle is situated 1.5 m from the wall moving at v=0.25 m/s. The red line is scaled for easier comparison.

where $r = \frac{y}{\sin\phi}$ is the distance to the wall at angle ϕ. θ is the angle of the fly in the counter-clockwise direction relative to the axis of the corridor.

To find the correlator response in the corridor, we need only substitute the corresponding f_s, $f_t = \nu_r f_s$, and $C(f_s)$ into (4.5-4.7), to arrive at

$$C = C_0 \exp\left(-\frac{1}{2}\left(\frac{2\pi F_s y \sigma}{\sin^2\phi}\right)^2\right) \tag{4.11}$$

$$T = \frac{F_s\left(v - \frac{\dot{\theta} y}{\sin^2\phi} - \dot{y}\cot\phi\right)}{F_s^2\left(v - \frac{\dot{\theta} y}{\sin^2\phi} - \dot{y}\cot\phi\right)^2 + \frac{1}{(2\pi\tau)^2}} \tag{4.12}$$

$$A = \sin\left(\frac{2\pi F_s y \Delta\phi}{\sin^2\phi}\right) \tag{4.13}$$

for the correlators facing the left wall of the corridor. For the right wall, substitute y_r for instead of y_l and negate any appearance of v. In the rotated frame of the robot, substitute $\phi' - \theta = \phi$.

The analytic form is faithful to the response under full simulation (Figure 4.3). One limitation is that the model assumes both visual sensors of the correlator observe the same local spatial frequency f_s, but in fact f_s is continually varying across the retina. This is not a significant problem because rapid changes in f_s coincide with high f_s, which are blurred out.

4.5 A Controller That Uses Correlators to Approximate Retinal Velocity

The control task is to provide thrust force commands to a pair of fans operating in unison (thrust u_1) or differentially (torque u_2) on a hovercraft robotic testbed. To avoid impacting the walls, the controller uses omnidirectional visual imagery to stabilize its motion along the middle of the corridor (for details of the robotic implementation, see Section 4.6). The vehicle rolls on uni-directional roller balls which allow motion in all directions like a hovercraft with small linear b and rotational c damping coefficients. The dynamics of the vehicle in the moving coordinate frame of the vehicle are thus modelled as

$$\begin{aligned} m\dot{v} &= -bv + u_1 \cos\theta \\ m\ddot{y} &= -b\dot{y} + u_1 \sin\theta \\ J\ddot{\theta} &= -c\dot{\theta} + u_2, \end{aligned}$$

where m is mass and J is the rotational moment of inertia about the center of mass. The vehicle is underactuated because it cannot generate lateral force directly, but lateral dynamics are controllable because they are coupled to forward dynamics by the θ term.

This inertial (non-kinematic) control problem is similar to that encountered by the fly, with its aerial dynamics and pair of independently-controlled wings. An approach for flight control proposed by Humbert et. al. [27] decomposes the retinal velocity function $\nu_r(\phi)$, also known as "optic flow," into sinusoid basis functions (Fourier harmonics), emulating the lobula plate tangential cells (LPTCs) of insects. It can be shown that within a straight

corridor, the first few harmonics correspond implicitly to the state variables of the vehicle, $\{v, y, \theta, \dot{\theta}\}$. Linearizing about an operating point of baseline motion along the center of the corridor and using the first few harmonics as the outputs of the system, the vehicle's state is both observable and controllable by the fan pair.

4.5.1 Tuning σ and $\Delta\phi$ angles for the environment to approximate retinal velocity ν_r

While correlators do not measure pure retinal velocity, their response rises monotonically with it under certain conditions. Figure 4.3 shows that the correlator response have a complicated shape for low σ, but as it is increased (increasing blurring), its response resembles a Gaussian function and also the retinal velocity ν_r. It is possible tune the inter-sensor angle $\Delta\phi$ and the width of the Gaussian blurring kernel σ together to insure that response appears Gaussian and does not alias to a negative response.

There are two criteria to satisfy. The first is to insure the spatial aliasing term A does not go negative in the operating regime. For this we require that A be positive for a correlator facing laterally ($\phi = \pm\pi/2$) no matter where the robot is in the corridor $0 < y < 2y_d$ where y_d is the desired distance to be maintained from each wall. This can be insured if the argument to the sin function $2\pi F_s y \Delta\phi \leq \pi$. This puts an upper limit on $\Delta\phi$, and since R is strongest for larger $\Delta\phi$ (thereby minimizing the effects of noise) we choose the upper limit

$$\Delta\phi = \frac{1}{4F_s y_d}. \tag{4.14}$$

The second criterion is to insure that attenuated contrast factor C^2 falls off sufficiently fast with increasing y that when A does go negative at $y > 2y_d$, the blurring effect in C has attenuated the response to near zero. Under this condition, as the vehicle moves still further from the wall the correlators' response will remain near zero because all of the scenery is blurred away. This is preferable to having aliasing cause the response to go negative. Since the contrast factor dies off as $e^{-(2\pi F_s \sigma y)^2}$, we set the blurring width σ such that at $y = 2y_d$ the signal attenuation has diminished by three standard deviations, or

$$\sigma = \frac{3}{4\sqrt{2}\pi F_s y_d}. \tag{4.15}$$

Lastly, scale the correlator response so that it matches the retinal velocity at $\phi = \pi/2$ by the factor

$$\alpha = \frac{\nu_r}{R} = \frac{2\pi\tau\left(F_s^2 v_d^2 + 1/(2\pi\tau)^2\right)}{C_0 F_s y_d \exp\left(-\frac{1}{2}\left(2\pi F_s y_d \sigma\right)^2\right)\sin 2\pi F_s \Delta\phi},$$

where v_d is the desired forward velocity operating point for the robot.

The foregoing insure that the correlator response will resemble the pure retinal velocity across ϕ near the desired operating point of forward motion along the axis of the corridor, enabling the use of the Humbert controller. Figure 4.3 shows the correlator response function $R(\phi)$ for a few different values of σ, showing that the model is faithful to a full correlator simulation.

4.5.2 Implementation in simulation

A simulation environment was written in MATLAB to render the scene from the perspective of the robot. Visual updates were received at 60 Hz, equal to the servo rate of the physical robot's controller software. The visual environment consisted of 2880 visual rays, and each of the 64 visual sensors was simulated with a local Gaussian pool across a 3-σ neighborhood of the sensor's orientation, blurring the image. The discrete-time low-pass filters in the correlators were also updated at 60 Hz, and dynamics were integrated with a zero-order hold. To replicate the actual robot, the following parameters were used: $F_s = 1$ cycle/m, corridor width $2y_d = 2.5$ m and length 12 m, $v_d = .25$ m/s, $\sigma = .14$ rad from (4.15), and $\Delta\phi \approx .2$ rad from (4.14) (achieved by using second-nearest neighbors when constructing correlators), $m = 6$ kg, $b = 4.45$ N-s/m, $J = .06$ kg-m^2, and $c = .06$ N-m-s.

Because of the difference between correlators and pure retinal velocity, it was necessary to select different gains than those reported in [27]. Following the same naming convention, the lateral gains of $K_{a_0} = -0.1$, $K_{a_1} = 0.8$, and $K_{a_2} = 4$ were found to give stable performance as shown in Figure 4.4. However, because the difficulty of decomposing the response R (4.11-4.13) into sinusoid harmonics analytically, it was impossible to provide a satisfactory stability analysis. This led to the design of the more-tractable controller in Section 4.7.

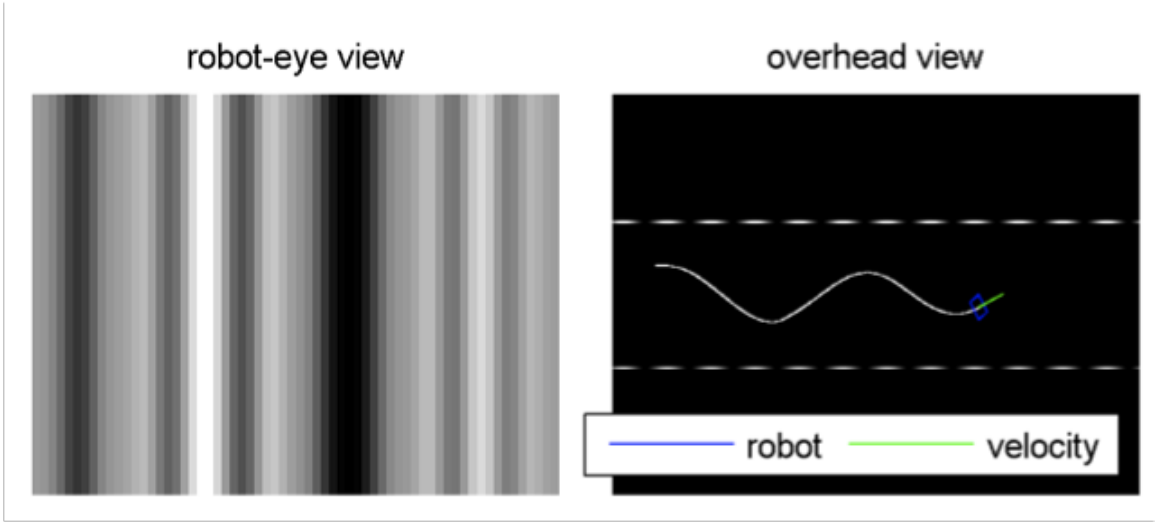

Figure 4.4: Simulation of robot using Humbert controller (sinusoid harmonics) and tuned blur width σ and correlator distance $\Delta\phi$. On the left is the 360° visual environment used for state estimation by the robot. The dark band in the middle corresponds to the end of the corridor. The robot was stable with this controller, but not strongly.

4.6 Robotic Implementation

A robot was constructed with an infrared omnidirectional visual sensor that emulated the eyes of the fly. The robot rolled on low-friction rollerballs and was actuated by fans [103]. Computation was performed by an onboard laptop running the RHexLib library on top of the real-time QNX operating system. Fan forces were calibrated beforehand and interpolated from a look-up table to achieve desired forces. An overhead vision system could track position and pose of the robot for analysis afterward (Figure 4.5). The parameters of the vehicle and dimensions of the arena were what were used in the simulation and are given in Section 4.5.2.

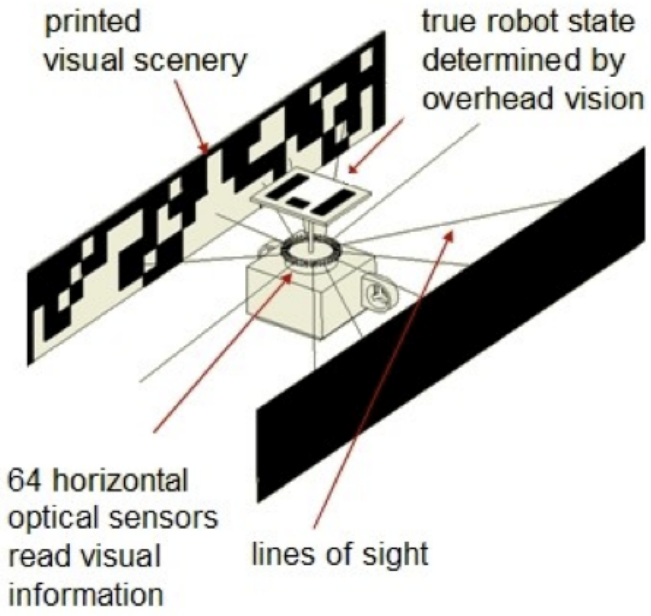

Figure 4.5: Diagram of the fan-actuated hovercraft robot in its environment.

Rather than use a camera, a custom fly-like eye with an array of 64 infrared light-sensitive diodes was constructed so that the view was omnidirectional. Heat-shrink tubing was attached to each sensor to narrow its field of view, giving a sensor blur width of $\sigma \approx .077$ rad or 4.4° (Figure 4.7). The array of light sensors were passed through analog multiplexors, amplified by operational amplifier, and read into the 8-bit analog-to-digital converter of a PIC microcontroller. The microcontroller read the entire array of luminance readings and communicated the result to the host laptop at 60 Hz over its parallel port. A spatial discrete Gaussian blurring was performed in software with $\sigma = .11$ using the property that the convolutions of two Gaussians with standard deviations σ_1 and σ_2 is a Gaussian with standard deviation $\sqrt{\sigma_1^2 + \sigma_2^2}$ to arrive at the desired $\sigma = .14$ rad. Parallel walls were constructed and illuminated by DC incandescent lights powered by a large power supply to minimize 60 Hz line interference. The infrared sensors could detect printed black vs. white on paper if it was printed by laser, but not by other printing technologies such as ink-jet printing [104]. To normalize the luminance to zero, the mean was taken of all sensors at the beginning of each trial and subtracted out. The robot and its environment are shown in Figure 4.6.

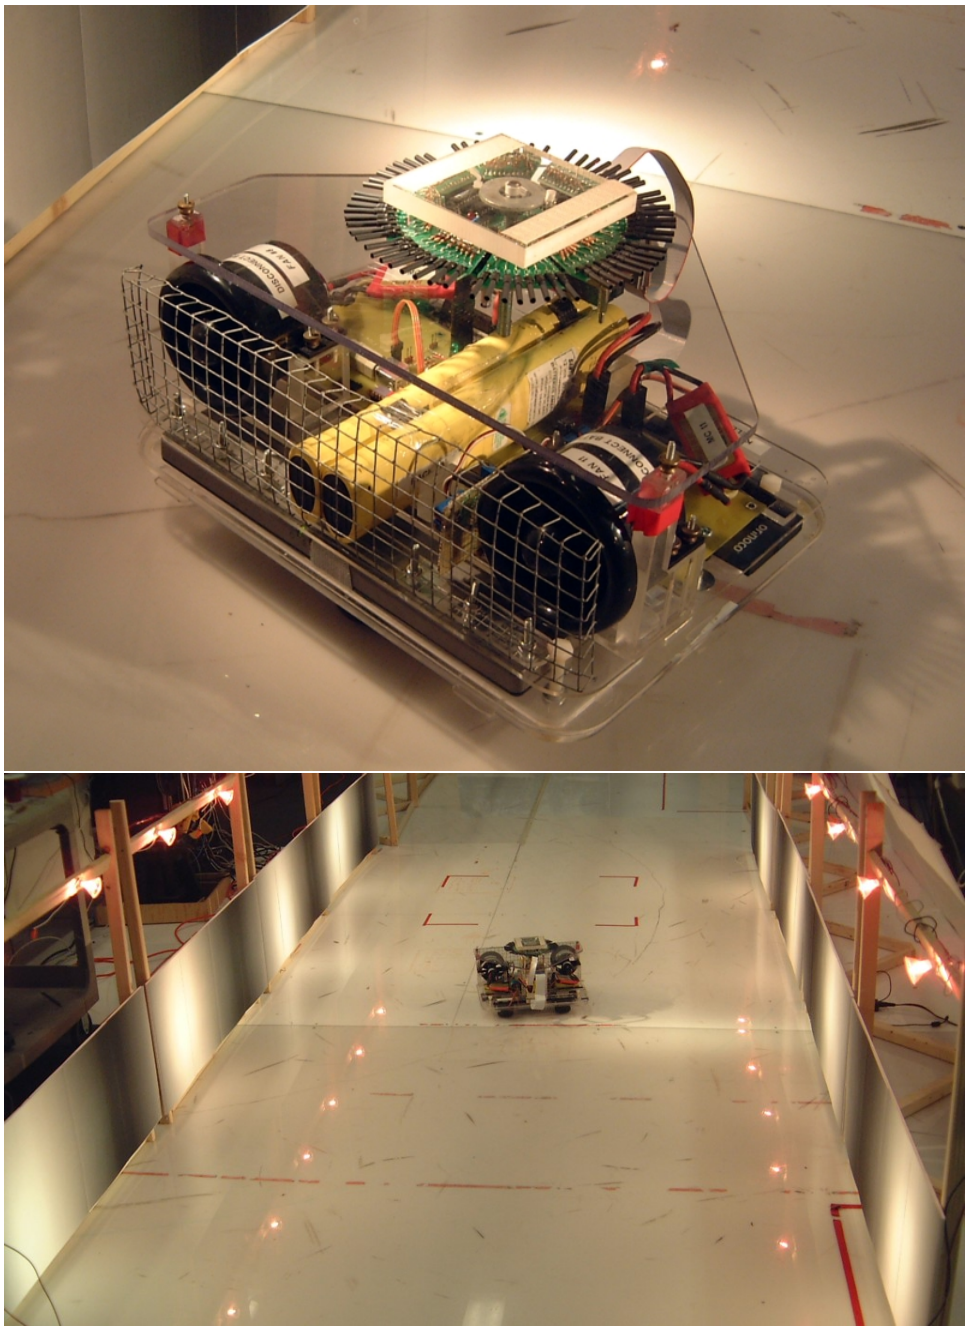

Figure 4.6: Fan-actuated hovercraft robot with a 64-element circular omnidirectional array of infrared light sensors (top). The view of each infrared sensor was constricted by a segment of heat-shrink tubing to increase visual accuity and eliminate spurious sources of light. Computation was performed by a subnotebook laptop at the base of the robot and forces were generated by two ducted fans (black) powered by a NiMH battery (yellow). The "hat" used by the overhead tracking system attached by velcro to the translucent piece of acrylic at the top and was removed so that the visual sensor was visible. The vehicle is shown in its environment, patterned walls lit by DC incandescent lighting and a low-friction floor.

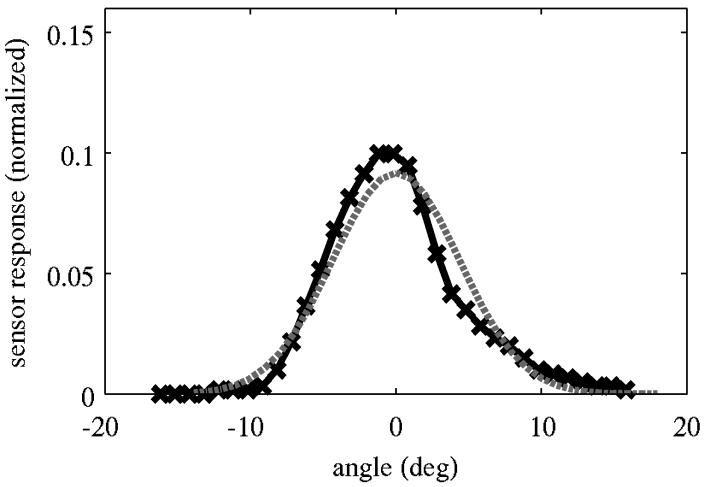

Figure 4.7: The response profile of an individual infrared sensor was measured by moving a point light source laterally across its view in small increments. The response (black) is fit by a Gaussian by calculating the sample standard deviation (dashed grey) giving a width of approximately 4.4°

Representative trajectories captured by the overhead vision system are shown in Figure 4.8. The robot could navigate the corridor, but not always consistently. Clumps of lint and cracks in the floor were large disturbances that were hard to compensate for. A significant problem was that if it got too close to the walls it would turn into them rather than away. This can be explained as follows. The first few sinusoid harmonics of the retinal velocity ν_r give enough information to observe the vehicle state. But the zeroth cosine harmonic (the mean $\overline{\nu_r}$) is $a_0 = -\sqrt{2}\dot{\theta} + \frac{v_d}{\sqrt{2}y_d^2}y$ and is a function of two state variables. This is not a problem using pure optic flow because they can be separated using information from the second cosine harmonic $a_2 = \frac{v_0}{2y_d^2}y$ also has information on y alone. However, using correlators, the two states cannot be disambiguated because there is no response near $\phi = \{\pi, 0\}$ because of blurring of the high spatial frequencies. Hence, being leftward of the centerline of the corridor (stronger correlator response on left hemisphere) is indistinguishable from a rightward rotation $\dot{\theta} < 0$ (larger response on left hemisphere) (see Figure 4.9). These two conditions require opposite torque responses to reach the desired $y = y_d, \dot{\theta} = 0$. As soon as the robot turned away from the wall arising from a large detected error in y, its large rotation rate $\dot{\theta}$ would immediately eliminate the perceived error and, overcompensating, the robot would turn back into the wall.

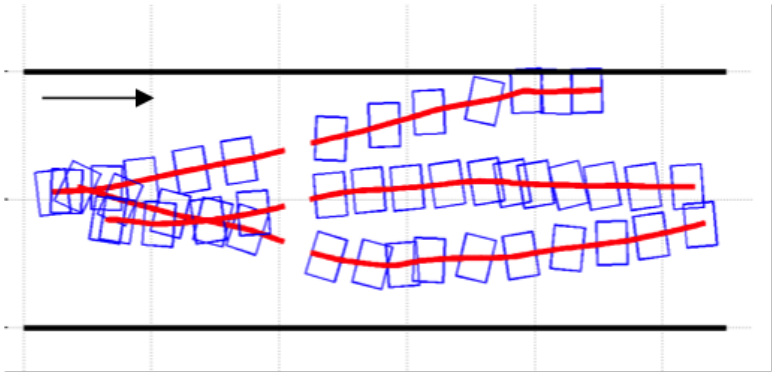

Figure 4.8: Trajectories (left to right) of the robot in the corridor captured by the overhead vision system. Gaps in the trajectories coincide with areas where tracking was not available.

4.7 A Controller Designed for Correlators

The limitation with applying the Humbert controller to the correlator response is that correlators do not behave like the retinal velocity. In areas of low visual contrast (such as where it is blurred out at angles near the axis of the corridor) there is no correlator response, leading to an ambiguity between $\dot{\theta}$ and y (Section 4.6). In addition, because of the complexity of the correlator response R, taking sinusoid harmonics of this function is analytically intractable, making a stability analysis impractical.

Our approach is to instead seek an approximation R' of the correlator response that captures its salient characteristics. With the approximation, it becomes possible to analytically take square harmonics which are also functionally equivalent to the wide-field integration performed by the tangential cells in flies [23], sidestepping the difficulty of taking inner products with sinusoids. To disambiguate y and $\dot{\theta}$, an estimate $\hat{\dot{\theta}}$ must be made from a non-visual source such as a gyroscope. With that information, as well as the centroid of the correlator response to estimate θ, by taking first-order Taylor expansions of the sum and difference of the harmonics, it is possible to explicitly estimate the state.

4.7.1 Approximations to correlator model

We start with an approximation to the correlator model, simplifying (4.11-4.13) by making certain assumptions that hold for the conditions of the robot described in Section 4.6.

4.7.1.1 Aliasing factor (Equation 4.13)

A attains an infinite frequency in the vicinity of $\phi = \{0, \pi\}$. However, the correlator response is also heavily attenuated by Gaussian blurring C factor there, meaning that the correlator response at these angles is nearly zero and can be ignored. In addition, the Gaussian blurring "bump" around $\phi = \pm\pi/2$ in C is much narrower across ϕ than the corresponding "bump" in A, so we may approximate A as being constant across ϕ.

4.7.1.2 Temporal factor (Equation 4.12)

T rises to a "saturated" maximum at $f_t = 1/2\pi\tau$ and then drops off, mimicking the temporal frequency peak observed in studies on insect behavior. If, however, the temporal frequency is far below saturation, $f_t \ll 1/2\pi\tau$, then the f_t^2 term in the denominator can be neglected. We also assume assume $\dot{y}$ is small and neglect it. Though it is amplified by the large factor $\cot\phi$ near $\phi = \{0, \pi\}$, temporal frequency saturation negates the effect of this term as $\cot\phi$ grows large, as does the Gaussian blurring. In simulation, it was found to have little effect on the correlator response (Figure 4.9).

4.7.1.3 Perspective approximation

The quantity $1/\sin^2(\phi)$ appears repeatedly. To approximate it, substitute $\alpha = \phi \pm \frac{\pi}{2}$ and perform a Taylor expansion around $\alpha = 0$ (the same approximation holds in both cases) to find

$$\left.\frac{1}{\sin^2\phi}\right|_{\phi=\pm\frac{\pi}{2}} = \frac{1}{\sin^2\left(\alpha+\frac{\pi}{2}\right)} = \frac{1}{\cos^2\alpha} \approx 1 + \alpha^2 + \frac{3}{2}\alpha^4 + \dots$$

Applying the approximations to equations (4.11-4.13) and truncating the Taylor expansion to 2nd order terms, we arrive at the tractable approximation R' whose factors are

$$C' = C_0 \exp\left[-\tfrac{1}{2}\left(2\pi F_s y \sigma\right)^2\left(1+2\alpha^2\right)\right] \approx C \tag{4.16}$$

$$T' = (2\pi\tau)^2 F_s\left(v - \dot{\theta} y\left(1+\alpha^2\right)\right) \approx T \tag{4.17}$$

$$A' = \sin(2\pi F_s y \Delta\phi) \approx A. \tag{4.18}$$

Full analytic and approximate responses under different perturbations are shown in Figure 4.9, showing reasonable agreement across the perturbations of interest.

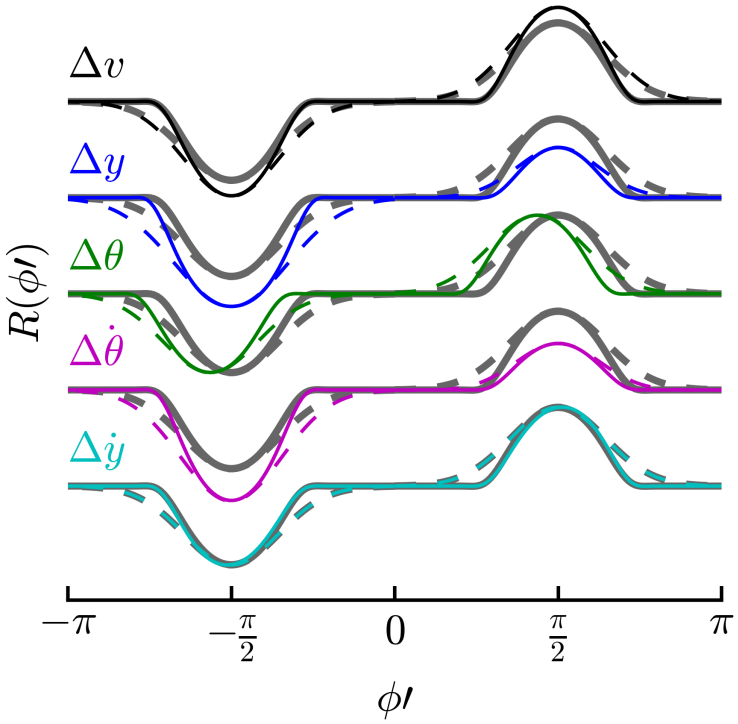

Figure 4.9: Correlator response to state perturbations as a function of body-frame angle ϕ'. The baseline response $R(\mathbf{q}_d, \phi')$ is shown in grey and perturbation responses are shown as solid ($R(\mathbf{q}, \phi')$) or dashed (approximation $R'(\mathbf{q}, \phi')$) lines. Changes in $\dot{\theta}$ are essentially indistinguishable from changes in y, necessitating an externally-derived estimate $\hat{\dot{\theta}}$. Changes in $\dot{y}$ have essentially no effect on R. The perturbation magnitudes are $\Delta v = 0.05$ m/s, $\Delta y = 0.2$ m, $\Delta\theta = .2$ rad, $\Delta\dot{\theta} = 0.08$ rad/s, and $\Delta\dot{y} = 0.05$ m/s.

4.7.2 Finding state variables by square harmonics

There is a finite minimum number of correlator pairs at different directions ϕ needed to make the system observable. But we seek an explicit inverse relation between correlator response and state variables with a significant effect on the correlator response for ease of constructing a controller. In analogy to the fly, suppose we have two lobula plate tangential cells, call them L_L and L_R that take the mean response over the left or right hemisphere, respectively. Because this integration range is much larger than the width of each $R(\phi)$ Gaussian, the integration range may be extended out to infinity without significantly affecting the result, assuming the sensor blur width $\sigma \ll \pi/2$. Thus,

$$L_L = \int_0^{\pi} R(\phi) d\phi \approx \int_{-\infty}^{\infty} R'(\phi) d\phi = \int_{-\infty}^{\infty} R'(\alpha) d\alpha = L'_L.$$

By approximating this π-width square harmonic as an infinite-domain integral, we need only evaluate integrals of the form $\int x^2 e^{-ax^2} dx$ and $\int e^{-ax^2}$ for which closed-form limits exist. This avoids taking analytically-intractable inner products with sinusoids. And because the limits of integration are much wider than the width of these Gaussian-like functions, for small θ the orientation of the robot may be ignored here. Integrating,

$$L'_L = \gamma e^{-\sigma_s^2 y_l^2} \left(\frac{v}{\sigma_s y_l} - \frac{\dot{\theta}}{\sigma_s} - \frac{\dot{\theta}}{4\sigma_s^3 y_l^2} \right) \sin(2\pi F_s y_l \Delta\phi), \tag{4.19}$$

where we have defined $\sigma_s = 2\pi F_s \sigma$ and the constant $\gamma = 2\pi\tau C_0^2 F_s \sqrt{\pi/2}$ for compactness. For L'_R, negate v and substitute y_r for y_l.

By taking the sum and difference of the two L's (compare to the zeroth cosine and first sine harmonic),

$$\Sigma = L_L + L_R \tag{4.20}$$

$$\Delta = L_L - L_R \tag{4.21}$$

as well as the centroid of the bumps, and using internal knowledge of $\dot{\theta}$ gathered from e.g. gyroscopes, the full state can be extracted as follows.

For a chosen desired operating point $\mathbf{q}_d = \{y_d, v_d, \theta_d = 0, \dot{\theta}_d = 0\}$ there is a corresponding Σ_d and Δ_d. Shifting the origin to the desired operating point, $\tilde{v} = v - v_d$, $\tilde{y} = y_l - y_d = y_d - y_r$, $\tilde{\Sigma} = \Sigma - \Sigma_d$, $\tilde{\Delta} = \Delta - \Delta_d$, the shifted coordinates can be interpreted as error to be driven to zero.

First, $\hat{\dot{\theta}}$ is estimated from gyroscopes, a reasonable proposition for both flies and robots. Then, $\hat{\theta}$ is found by taking the centroid of the two opposite-hemisphere "bumps" from the perspective of the vehicle, that is, using ϕ' instead of ϕ,

$$\hat{\theta} = \frac{\langle R(\phi'), |\phi'| - \pi/2 \rangle}{\overline{R(\phi')}} \approx \frac{\langle \hat{R}(\phi'), |\phi'| - \pi/2 \rangle}{\Sigma'/2\pi}. \tag{4.22}$$

where R is the true reading and $\hat{R}$ is the response read in from sensors. The divisor Σ' is known from the model and thus can be inverted beforehand to be a multiplicative scaling factor, avoiding a division operation. This inner product is similar to the first cosine harmonic, but has a larger domain.

The other states are found by taking Taylor expansions of $\tilde{\Sigma}$ and $\tilde{\Delta}$ about the origin. For Σ,

$$\tilde{\Sigma} = \left.\frac{\partial \Sigma}{\partial \tilde{y}}\right|_{\tilde{y}=0} \tilde{y} + \left.\frac{\partial \Sigma}{\partial \tilde{v}}\right|_{\tilde{v}=0} \tilde{v} + \left.\frac{\partial \Sigma}{\partial \tilde{\theta}}\right|_{\tilde{\theta}=0} \tilde{\theta} + \left.\frac{\partial \Sigma}{\partial \dot{\tilde{\theta}}}\right|_{\dot{\tilde{\theta}}=0} \dot{\tilde{\theta}} + \ldots \tag{4.23}$$

where the ellipsis denotes higher-order terms. Using the approximation Σ', an analytic form can be found for each of the derivatives. Both $\frac{\partial \Sigma'}{\partial \tilde{v}} = 0$ because the v term changes sign between L_R and L_L (4.19) and $\frac{\partial \Sigma'}{\partial \tilde{\theta}} = 0$ because by construction Σ does not depend on θ. Because we have an estimate $\hat{\dot{\theta}}$ from gyros, we can rearrange (4.23) to get the estimate

$$\hat{\tilde{y}} = \left(\hat{\tilde{\Sigma}} - \hat{\dot{\theta}} \frac{\partial \Sigma'}{\partial \dot{\theta}}\right) \Big/ \frac{\partial \Sigma'}{\partial y}, \tag{4.24}$$

where $\hat{\tilde{\Sigma}}$ is read from the sensors. Both derivatives of Σ' are evaluated at $\mathbf{q}_d$ and can be found from (4.19) and (4.20) and can be easily calculated with symbolic software such as Sympy or Mathematica. As in (4.22), the divisor can be inverted beforehand. We can then calculate

$$\hat{\tilde{v}} = \left(\hat{\tilde{\Delta}} - \hat{\tilde{y}} \frac{\partial \Delta'}{\partial y} - \hat{\dot{\theta}} \frac{\partial \Delta'}{\partial \dot{\theta}}\right) \Big/ \frac{\partial \Delta'}{\partial v}. \tag{4.25}$$

The controller was implemented by using an outer loop to set a desired orientation θ_d according to the estimate of the lateral distance from the center, $\theta_d = 0.25\hat{\tilde{y}}$ to steer the vehicle to the center. A slightly-damped inner loop regulated $\hat{\theta}$ with the torque command $\tau = K_\theta(s + 10)\hat{\theta}_e$ where $\hat{\theta}_e = \theta_d - \hat{\theta}$ is the θ error and $K_\theta/J = 3$. This controller was implemented in simulation and was much more stable. Though we can offer no proof of stability in this report, we note that this controller (with $2\times$ larger blur width) is able to stabilize in a corridor patterned with a texture taken from a real photograph (Figure 4.10), suggesting its ability to control motion under more real-world circumstances.

4.8 Conclusions and Future Work

This paper reports progress on the visual control of micro-aerial vehicles using a small number of visual sensors. The controllers reported rely only on simple-to implement neurologically plausible multiplies and adds. Using a model of the response of an array of visual correlators to a moving sinusoid-striped wall of infinite extent, we were able to derive relations to determine a blur width and inter-sensor spacing that eliminated the possibility of

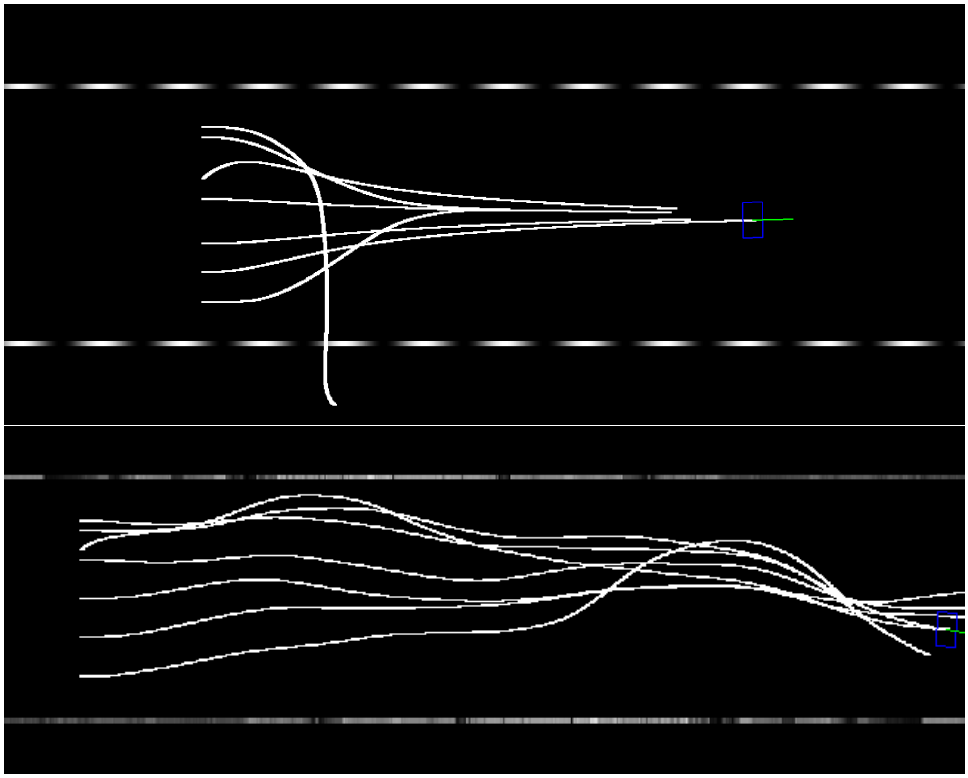

Figure 4.10: The controller introduced in this work, designed explicitly for the correlators, has a larger basin of attraction, extending nearly to the walls of the corridor (too close and it overcompensates) (top). It is also able to stabilize motion in a corridor patterned with imagery taken from a real photograph (bottom).

driving the response negative due to aliasing. Using these relations, we reported navigation in a corridor using a controller that decomposed the correlator response array into sinusoids or square harmonics. The controller is computationally efficient, requiring less than $6n$ multiply-accumulate operations per step, where n is the number of omnidirectional pixels.

The analysis is in the frequency domain so it may extend to naturalistic scenery with known power spectra (for example, $1/f$) following [69]. It may be extended to a 3-dimensional hemisphere by incorporating a longitude-like lateral angle β between the poles of expansion and contraction, substituting $y/\cos\beta$ for y. The sensitivity to frequency or contrast exhibited by correlators may be mitigated by compressive nonlinearities [69] or by local contrast estimation [32]. Another possible direction would be to find a relation for the number of sensors needed to reach a desired variance in the state estimate.

This work is related to [32] and [33] in that both works aim to stabilize around "snapshot" desired visual response, but differs in that this work seeks a desired correlator snapshot corresponding to a continuous state of motion rather than a fixed pose. It may be possible to combine with that work, which demonstrates learning an optimal bilinear state estimator using only the delayed (L_d) and raw (L_d) luminance readings as inputs, to instead learn and stabilize to a correlator response "snapshot."

Chapter 5

How Flies Use Correlators to Control Forward Velocity

5.1 Introduction

Soon after evidence emerged that insects use visual motion to guide their behaviors, attention turned to the mechanism by which they computed visual motion. Two primary competing hypotheses were gradient-based and correlation based [70]. We provide a brief comparison of the two here. Suppose the visual world is divided into picture elements, or pixels, each of which is reading a single luminance value L. These may be thought of as each ommatidium in an insect's eye (though neural superposition in higher insects like dipterous flies makes the picture a little more complicated, but the principle is the same). In the gradient method, the time derivative $\frac{dL}{dt}$ is divided by the spatial derivative $\frac{dL}{d\phi}$ with respect to the angle ϕ across the retina (the one-dimensional gradient) to estimate the rate of motion of the image $\frac{d\phi}{dt}$. Generalizing to two-dimensional image motion requires an assumption of local luminance constancy and requires a matrix inverse [105], but we need only consider one-dimensional motion here. The correlation method, by contrast, multiplies these two derivatives rather than dividing them [106]. In both cases, the spatial derivative could be approximated by taking the difference between neigboring sensors, and the time derivative by subtracting the current sensor reading from a delayed version. A schematic of an implementation of the correlator is shown in Figure 4.1. Both methods rely on a nonlinearity to estimate image motion [107].

Early studies on insect vision were constrained by the technology available at the time. In an influential study, a beetle was tethered at the center of a drum rotating about a vertical

axis. The beetle was made to repeatedly choose left or right as it walked around a small ball it held with its feet. The direction chosen by the beetle indicated that it was applying feedback to limit rotational visual "slip". The strength of this "optomotor response," however, did not vary with with rotational velocity, but varied instead with the temporal frequency of stripes painted on the inside of the drum [108, 106, 98]. Because it was impossible at the time to measure internal activity of the brain, the insect's motor response was assumed to be proportinal to the visual stimulus received by the insect. Based on this study, it was concluded that the elementary motion detectors (EMD) of the beetle were of the correlation type. A diagram of the correlator model is shown in Section 4.3 and with it is included a mathematical analysis.

Later, electrophysiological studies began to illuminate the pathways in the fly brain. At the periphery, it was found that visual motion detection was essentially monochromatic, relying primarly on the blue-green sensitive R1-6 rhabdomeres in the fly's eye [109]. Recent work has begun to uncover the neural mechanism by which the sum and difference operations in the correlator are performed, but an explanation for the multiplication and delay stages has not yet emerged [100, 15]. Downstream, a class of large and accessible cells, the lobula plate tangential cells, were identified that sum input over large fields of the visual sphere [11]. At least part of their function appears to be to encode self-motion of the insect: local preferred directions (the direction of visual motion across the retina that elicits the strongest response) patterns for different cells correspond to rotation around different axes (vertical or VS cells) or linear motion (horizontal or H cells) [23, 24]. Suggesting that they received input from correlators, their responses also exhibited the temporal frequency peak response to image motion [11]. In addition, HS cells from the horizontal system exhibited the higher harmonics expected to arise from the multiplication step of the correlators [99]. Further emphasizing their role in motion control, if these tangential cells were ablated, the optomotor response of flies was diminished [11]. A review of results supporting the correlation method over the gradient method as the correct hypothesis is given in [70].

The correlator model has a number of non-idealities that limit its ability to faithtfully estimate visual motion, such as being both contrast-dependent and spatial-frequency dependent [107, 69]. When it comes to natural images taken by a camera, however, insect visual motion performance is much better [76, 110]. Elaborated correlator models have been proposed that exhibit correlator-like responses to sinusoids but yet are still capable of giving

accurate velocity estimates on varying natural imagery, even in spite of different contrast levels [69, 102], or to the control of a model fly in simulation [111]. Thus, while they exhibit non-idealities when observing single-frequency input, it may be that their behavior for natural imagery is what redeems correlators. Or it may be that computation is more efficient for neurons or less sensitive to noise. In this work we take the viewpoint that the metric for performance should not ultimately be its fidelity at reproducing the velocity of image motion. Instead, it should be measured against its performance in feedback for motor control of flight.

While the peculiar behavior of correlators to rotatory sinusoid imagery is well-documented, the picture appears more complicated for free-flight control of forward velocity. Bees flying in corridors with vertically-striped walls do not significantly change their forward velocity when the spatial frequency on the walls is changed [43, 112, 113, 22], and neither do flies [54] [48]. Hypotheses have emerged to explain this phenomenon, from separate, individually-tuned correlators [114], to the use of a gradient method, which is not dependent on spatial frequency [70].

The study in [48] is perhaps the most detailed study undertaken so far. In that work, the visual stimulus was sinusoid grating patterns back-projected onto the walls of a wind tunnel by computer while fruit flies' forward flight was tracked by cameras. Instead of just measuring their speed, however, the flies were placed in "virtual open loop." In this experimental paradigm, the visual stimulus is moved by the computer in real-time in response to the motion of the fly. In this way the visual stimulus can be closely controlled, combining the advantages of tethered flight with the more realistic stimuli of free-flight. Flies were presented a range of spatial and temporal frequencies and their accelerations were measured as the output and plotted on a 2D response map. They found that over a certain range of spatial frequencies, flies' acceleration responses were relatively independent of spatial frequency. This suggested that a different visual pathway than the basic correlator was at work estimating forward velocity.

In this chapter we consider whether it is necessary to invoke a different or complicated model than the basic correlator to explain the findings of [48] and [54]. In Chapter 4 I demonstrated that it was possible in principle to control lateral and forward motion in a corridor using correlators. Based on the model of correlators observing a flat moving wall derived in that chapter (Equations 4.11-4.13), I observed that the sum of all wall-facing

correlators (the value of Δ from Equation 4.20), plotted across a spatio-temporal frequency map as in [48], could have a similar appearance. In other words, the effects of perspective and motion parallax on an array of correlators, which had not previously been modeled, could be enough to account for observed behavior of the fly. It would not be necessary to invoke a more complicated correlator model to account for independence of spatial frequency.

One of the features that characterizes correlators is their temporal frequency tuning peak. That is, though the correlator responds differently to different spatial frequencies, its peak response is always at the same temporal frequency, which we will call f_t^p, the peak temporal frequency. If flies exhibited a temporal frequency tuning peak, this could be construed as evidence that the correlator model was correct.

To probe the matter further, we performed experiments on flies in an attempt to reproduce the results of [48], but with greater emphasis on whether a temporal frequency peak could be detected. The results suggest that the fly's behavior matches predictions for a array of identical correlators. The correlators exhibited a temporal frequency peak of 5-10 Hz at all spatial frequencies tested in both the drum and tunnel conditions. This corresponds to a time constant in the delay block of the correlator of 21 ms. The model and the fly's behavior was also consistent with spatial frequency independence for the spatial frequencies reported in [54, 48]. Because a visual motion pathway using correlators already exists for the fly, the simplest possible explanation is that they are using a similar pathway in forward flight regulation.

5.2 Methods

5.2.1 Model

A derivation of a frequency-domain model of the correlator response is given in Section 4.4 that incorporates the effects of perspective and motion parallax. We briefly review the relevant results.

5.2.2 Drum case

It has been shown in [69] that a correlator with a low-pass filter as its delay element has the steady-state response to a sinusoid of

$$R(t) = \frac{C_0^2}{2\pi\tau} \frac{f_t}{f_t^2 + 1/(2\pi\tau)^2} \sin(2\pi f_s \Delta\phi), \tag{5.1}$$

where C_0 is the amplitude of the sinusoid luminance pattern,

$$f_t = f_s \nu_r$$

is the temporal frequency (cycles/s) associated with rate of visual motion ν_r (rad/s) of the sinusoid across the retina of the eye, τ is the time constant of the low-pass filter, f_s is the spatial frequency of the sinusoid (in cycles/rad around the retina of the fly) and $\Delta\phi$ is the angular separation between the pair of luminance sensors observing the sinusoid. This model predicts the behavior of insects in a rotating drum in which the spatial frequency is constant across the visual field. The temporal factor (4.6) is

$$T = \frac{f_t}{f_t^2 + 1/(2\pi\tau)^2}.$$

Because of the f_t^2 term in the denominator, T first rises and then falls with f_t, exhibiting a temporal frequency tuning peak at

$$f_t^p = \frac{1}{2\pi\tau}. \tag{5.2}$$

For an example plot, see Figure 5.7, left.

5.2.3 Tunnel case

For the special case of a correlator observing a flat moving wall with no lateral motion and no rotation ($\dot{y} = \dot{\theta} = \theta = 0$), Equations (4.11-4.13) simplify to

$$R(t,\phi) = \frac{1}{2\pi\tau} C_0^2 \exp\left(-\frac{1}{2}\left(\frac{2\pi F_s y \sigma}{\sin^2\phi}\right)^2\right) \frac{F_s v_g}{(F_s v_g)^2 + 1/(2\pi\tau)^2} \sin\left(\frac{2\pi F_s y}{\sin^2\phi}\Delta\phi\right), \tag{5.3}$$

where v_g is the forward velocity, F_s is the spatial frequency on the wall in cycles/m, y is the distance to the wall, and σ is the standard deviation of the Gaussian blur. Some examples for different forward velocities and spatial frequencies are shown in Figure 5.2.5. Note that the temporal factor in this case is simply

$$T = \frac{F_s v_g}{(F_s v_g)^2 + 1/(2\pi\tau)^2}$$

and is the only factor that depends on the forward velocity v_g. The temporal frequency

$$f_t = F_s v_g \tag{5.4}$$

does not depend on the distance to the wall. This can be understood by considering a luminance sensor taking readings at a certain angle ϕ. This sensor will experience the same temporal frequency regardless of the distance y to the wall or ϕ if the fly is moving straight and parallel to the wall. Hence, an array of correlators observing a moving wall patterned with a constant spatial frequency would all experience the same temporal frequency as in the drum condition.

The model does not take into account second-nearest neighbors as reported in [115, 70], but their effect could in principle be added in using the appropriate choice of $\Delta\phi$.

It is assumed that there is some neural computation in the fly that takes the sum of all correlators on the left or right hemisphere to calculate the quantity around which it regulates its forward velocity. This is likely performed by neurons of the horizontally-sensitive system in the lobula plate [11, 23]. This quantity is Δ (Equation 4.20), that is, the difference between the sum of left-facing correlators and right-facing correlators, shown as the area under the curves in Figure 5.2.5. (The difference is needed rather than the sum because left-facing and right-facing correlators would have opposite sign when flying forward as shown in Figure 4.9.)

While it is impossible for a small fly to measure correlator response in the fly brain in a flying *Drosophilia* (though the feat has recently been accomplished on tethered *Drosophila* [100]), the magnitude of acceleration is expected correspond to responses to changes in this value, as described in Section 5.2.4.

Figure 5.1: Photograph of corridor with short-wavelength gratings back-projected onto the walls. Because of the sensitivity of the CMOS sensor used to take this picture to infrared, the infrared backlights can be seen on the walls as well, but these would not be visible to the fly.

5.2.4 Apparatus

The system for tracking flies and providing visual stimulus is described in Section 2.2, but for the results reported in this chapter the wind gust apparatus was removed. As in that section, 12 female flies were water but no food for 2-4 hour before the experiment and the experiment continued for 24 hours.

The visual stimulus consisted of vertically-aligned sinusoid gratings. One alternate hypothesis was that there could be some unknown peculiarity of drum-like stimulus that lead to the correlator-like behavior that was not induced in a tunnel geometry. Thus, a "drum" condition was introduced, rendered in OpenGL using VisionEgg [67] onto the walls of the wind tunnel so that the fly was exposed to a stimulus with constant spatial frequency f_s across the retina. The tunnel case was simply a sinusoid grating of constant linear spatial frequency on the walls F_s (a photograph is shown in Figure 5.1). Spatial frequencies can also be represented by wavelengths, with $\lambda = 1/f_s$ expressed in radians or degrees. In contrast to the work in Chapter 2, a range of different spatial frequencies were presented, rather than just one.

Both of these stimuli required updating in real-time with the motion of the fly in "virtual open loop" [48]. For the drum case, the virtual drum moved so that its center coincided with the fly. Trials were started when the fly passed through the plane transverse to the axis of the tunnel at the middle of the length of the tunnel, in either direction, and continued for 1.0 seconds. A new trial could not start for a 3-second refractory period. The trial consisted

of putting the fly in a visual clamp of a desired temporal frequency so that regardless of its own motion it was subject to this constant stimulus. A fly's response was the mean acceleration rate during a 100 ms period starting 200 ms after trial start. To calculate the average response over all trials, the median was taken instead of the mean to insure that strong outliers, which occur when the fly makes sudden body-saccades, did not influence the results. While the visual motion is known to exhibit motion adaptation, that is, a gradually diminishing response over time to a constant stimulus [76, 75], (thought to increase the dynamic range of neurons), because responses were measured within 300 ms of the trial start, it was assumed this effect would be negligible.

The result of Chapter 3 suggest that an integral or a proportional+integral controller are accurate models for the fly's response to a given visual stimulus. But because of linearity, regardless of the model, the rate of acceleration during that short period is sufficient information to measure the relative strengths of the different visual stimuli.

To subject the fly to a specific temporal frequency, we set the velocity of the stimulus according to Equation 5.4 (tunnel case) or Equation 4.3 (drum case). The frame rate of the projector, 120 Hz, limited the top temporal frequency to below 1/4 of this value to insure unambiguous directionality to the stimulus. We verified that all frames were being projected correctly by taking a video of the stimulus with a high-speed 500 fps camera.

5.2.5 Simulation

I wrote a simulation to model correlators in python to verify the validity of the analytic model of Equation 5.3. The sinusoid grating was simulated for 200 different angles (rays) and a smaller number of visual pixels arranged with angular separations roughly equal to the fruit fly's average ommatidia spacing [115], each of which sampled a small area of the image using a Gaussian blur (Figure 5.2). The simulation step time was 1 kHz. The simulation and Equation 5.3 are closely matched, as shown in Figure 5.2.5.

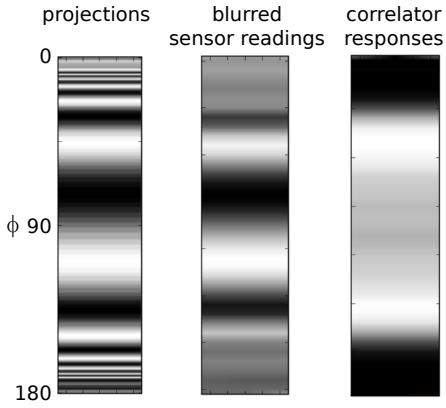

Figure 5.2: Visualization of correlator simulation. Rays are cast out to sample the sinusoidal luminance pattern (left) and a local Gaussian blur is applied (middle). The correlator response is computed (right) and shown with with white indicating a strong response and black to a weak response. See also Figure 4.2.

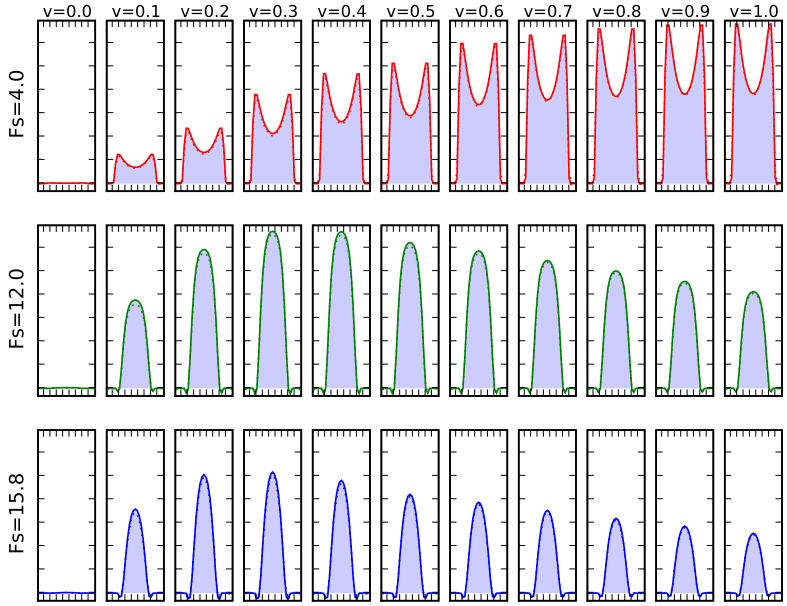

Figure 5.3: Simulated and analytic correlator responses for different spatial frequencies closely match. Because they are nearly coincident, the simulated (solid) and analytic (dashed) are hard to distinguish, indicating the fidelity of the model. The x-axis on each plot is ϕ, the angle across the retina of the fly (see Figure 4.2), and ranges from 0° to 180°. The distance to the wall is $y = 0.15$ m, corresponding to a fly at the mid-line of our wind tunnel arena. F_s is the spatial frequency (cycles/m) projected onto the wall; v is the ground speed (m/s).

5.3 Results

5.3.1 Data

Flies accelerated in response to the visual stimulus. Responses to the two cases of drum and tunnel stimuli are shown in Figures 5.4 and 5.5, respectively. The pattern of acceleration response in both cases showed a temporal frequency peak of about 7 Hz for all spatial frequencies tested.

In [48], the finding that there is spatial frequency independence amounts to that in a certain area of the 2D acceleration strength map, the contours are straight. In Figure 5.10

(top), in the range F_s =4–16 cycles/m, the viewer can observe that though the contours are not straight, they are nearly so. The difference is within the measurement noise of the corresponding figure in [48].

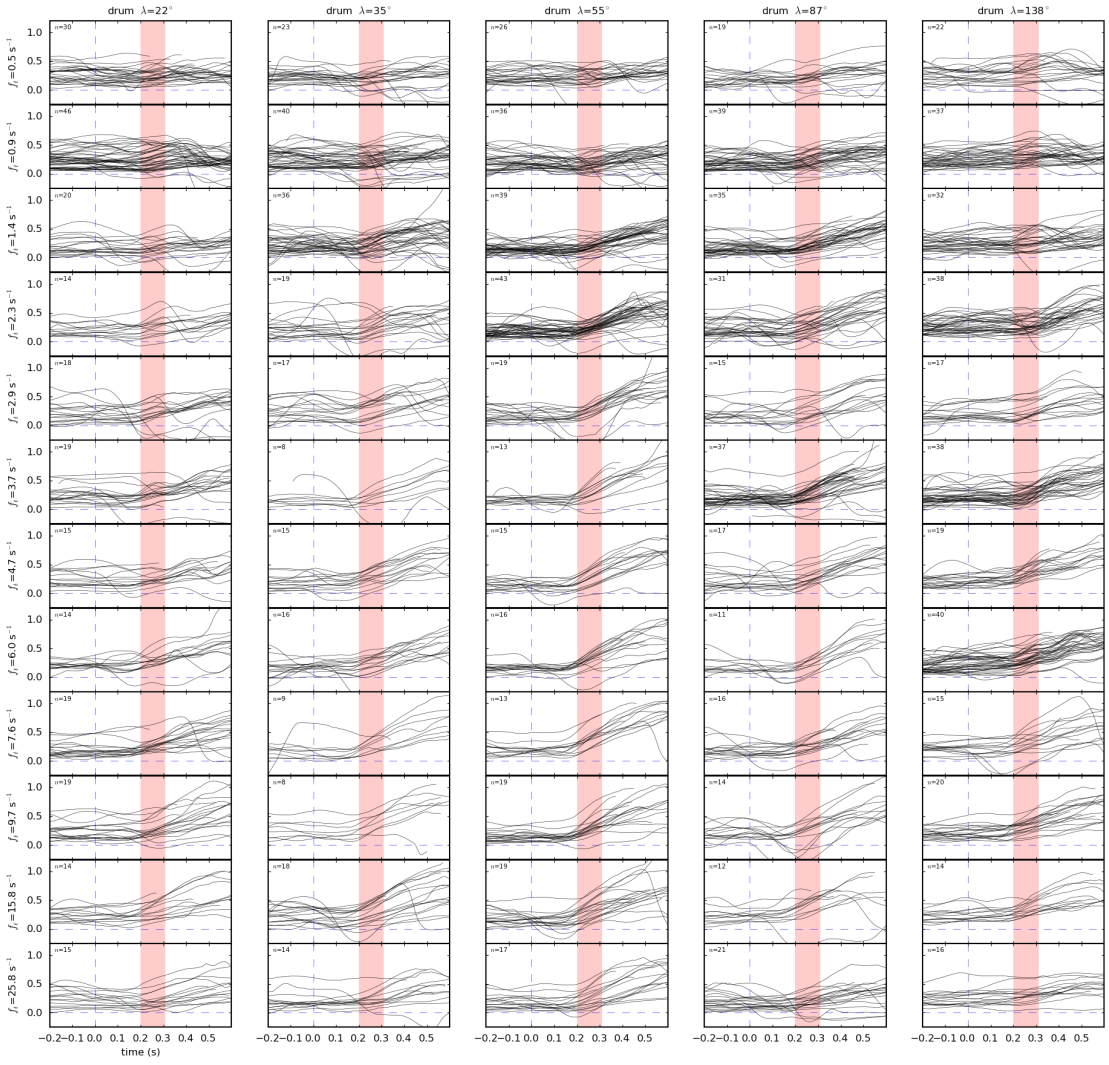

Figure 5.4: Drum ground speed trajectories. The shaded portion of the figure denotes where the acceleration of flies was measured. Not all temporal frequency conditions are shown because of space limitations.

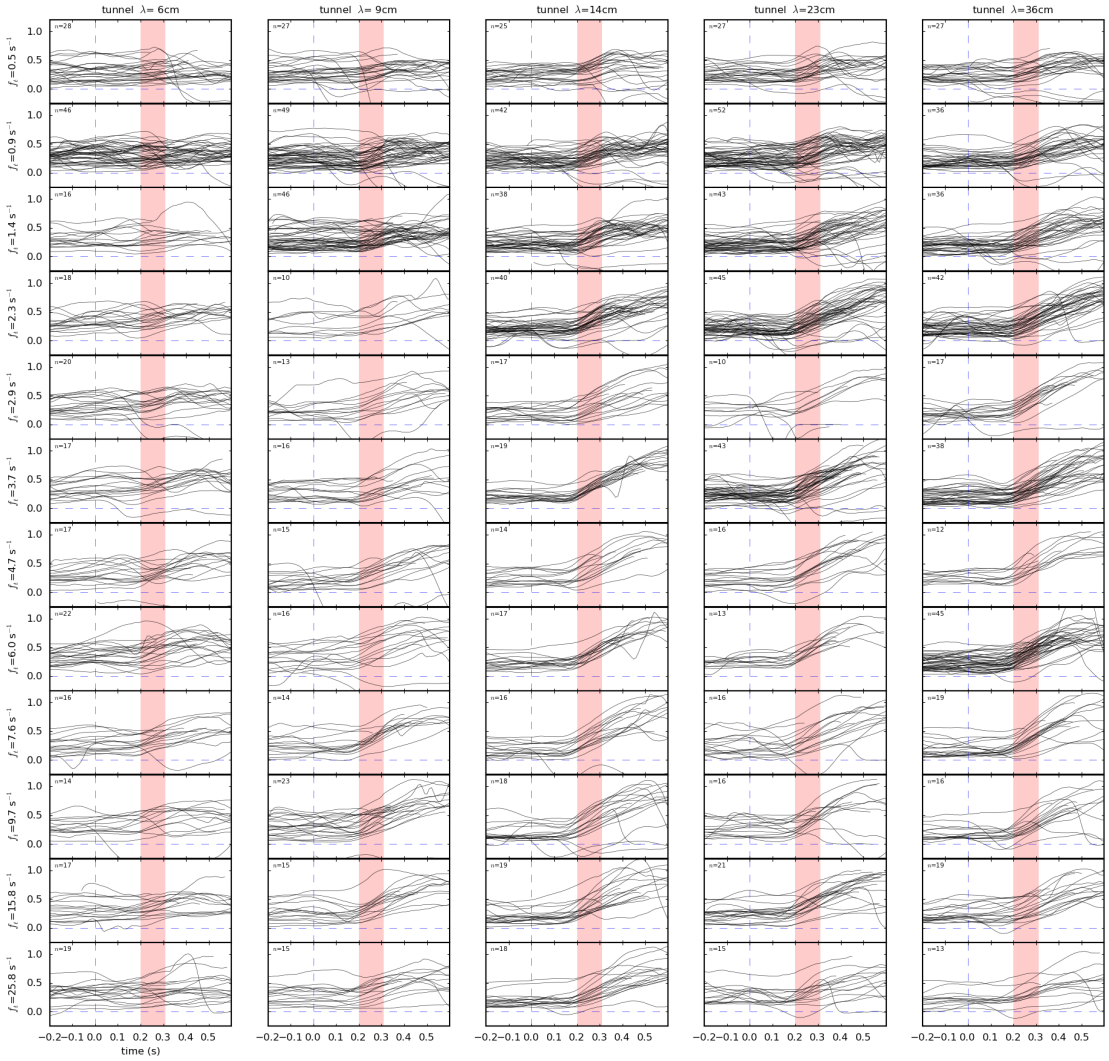

Figure 5.5: Tunnel ground speed trajectories. The shaded portion denotes where the acceleration of the flies was measured. Not all temporal frequency conditions are shown because of space limitations.

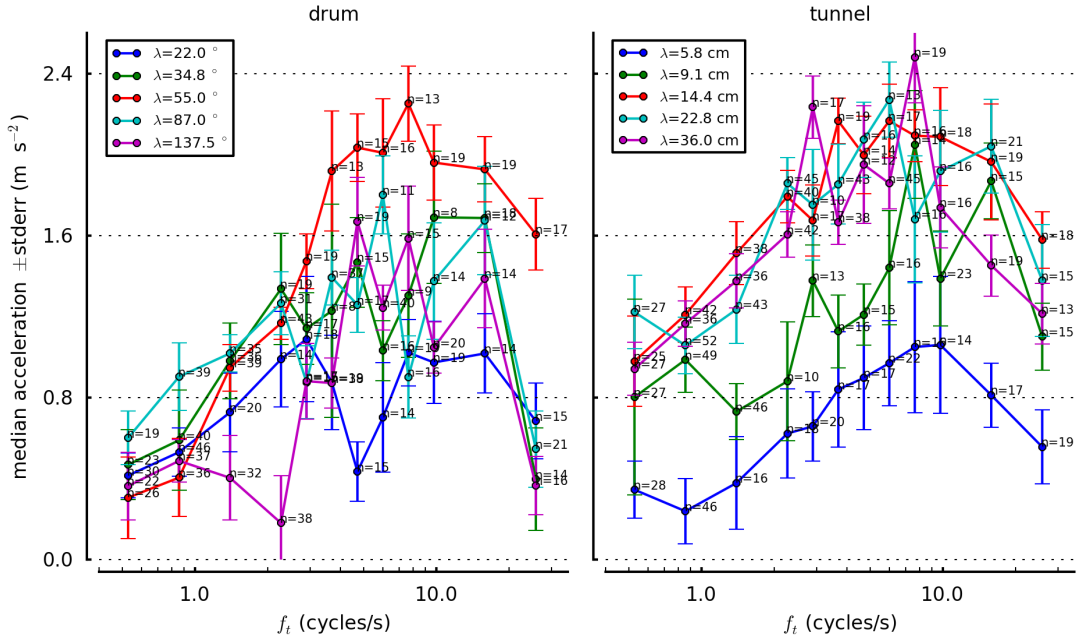

Figure 5.6: Fly accelerations in response to visual stimuli exhibit a temporal frequency tuning peak. Stimuli were presented in virtual open loop in both the drum and tunnel cases. The shortest wavelength (blue) exhibits the weakest response likely because of its amplitude is diminished by spatial blurring.

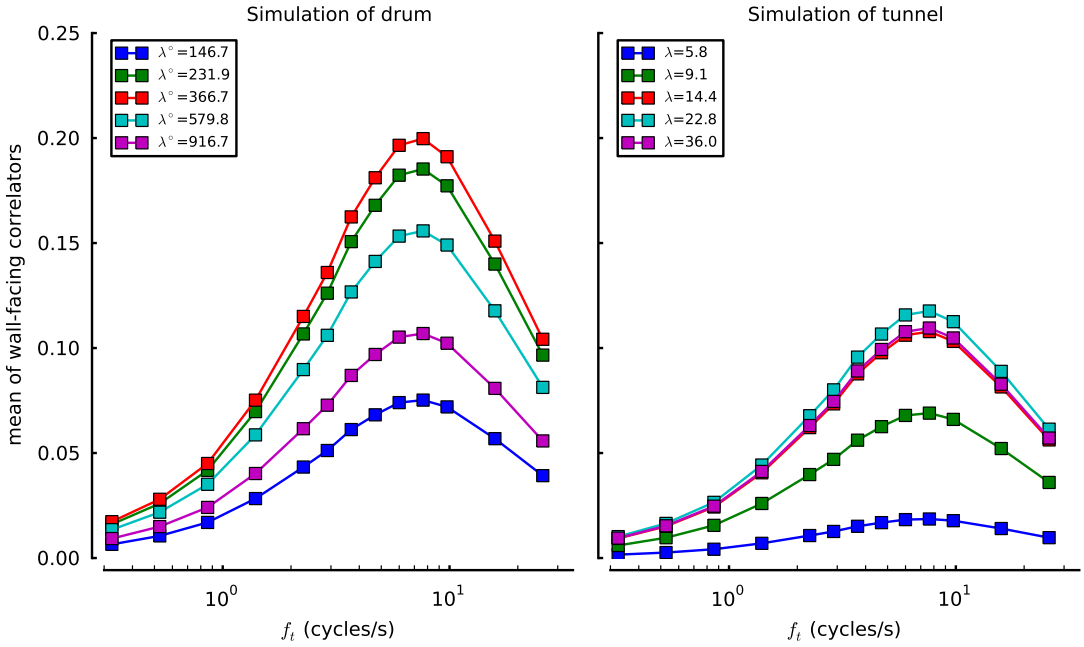

Figure 5.7: Model correlator responses to drum and tunnel stimuli

name		value		
inter-ommatidia angle	$\Delta\phi$	4.8°		
Gaussian blur width	σ	4.7°	$\Delta\rho$	11°
correlator time constant	τ	21 ms		

Table 5.1: Simulation parameters for model of fly's forward-velocity correlator-based velocity controller. The term σ is the standard deviation of the Gaussian function. The term $\Delta\rho = 2\sqrt{2\log 2}\sigma$ measures the width at half-maximum of the Gaussian function and is commonly reported in the biological literature.

5.3.2 Model fitting

Two free parameters were available in the model to fit the observed behavior in flies (Figure 5.7). The temporal frequency peak was was set by adjusting time constant τ and the spatial frequency peak was set by adjusting the Gaussian blur width σ. The value of $\Delta\phi$ was taken from [109] and reflects the effective distance between neighboring ommatidia for horizontal motion. Correlator parameters corresponding to good fits are given in Table 5.1).

Figure 5.8 shows the relative strengths across spatial frequency of the peak temporal frequency responses for simulation and data.

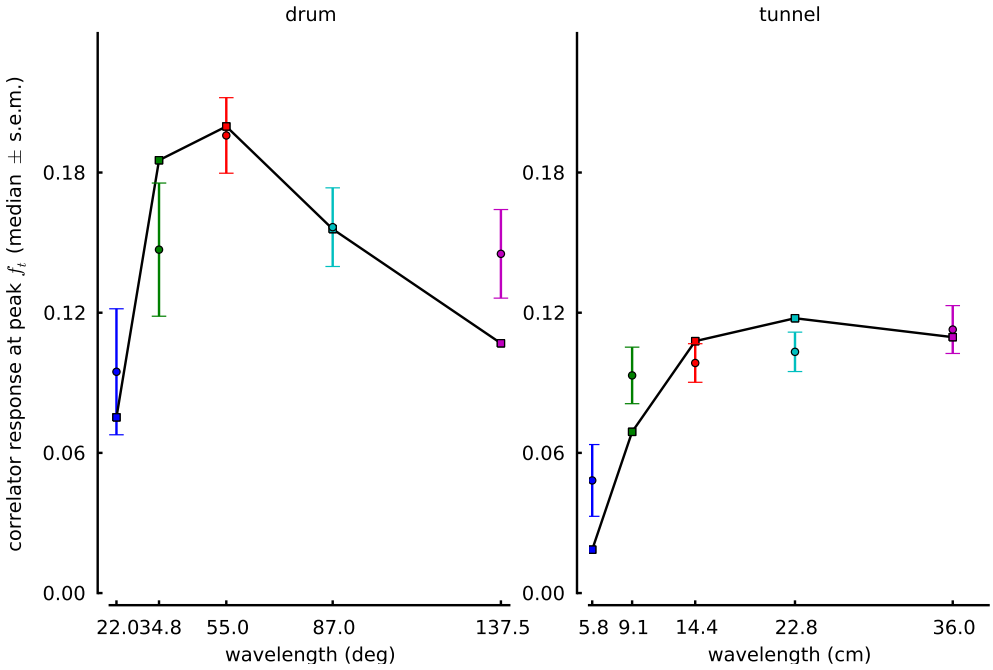

Figure 5.8: By tuning the blur width in the correlator, it is possible to achieve an agreement between the relative responses at different spatial frequencies. The peak responses in simulation (correlator response) and in data collected from flies (normalized with error bars showing median acceleration ±std. err.) have similar relative response characteristics.

5.3.3 Velocity sensing

The preceding section showed how fruit flies' forward acceleration behavior exhibits a temporal frequency tuning peak. Suppose the results are instead plotted versus linear velocity rather than temporal frequency. Can they still agree generally with the findings of [48] that the response is roughly independent of spatial frequency over a certain range of spatial frequencies?

Results are shown in Figure 5.9. For the tunnel condition, the velocity is $v_g = f_t/F_s$ from Equation 5.4 and for the drum case, the velocity is the retinal velocity $\nu_r = f_t/f_s$ from Equation 4.3. The model response is shown in Figure 5.10, showing close agreement. For all spatial frequencies but the highest tested, the correlator response is roughly proportional to ground speed up to 1 m/s.

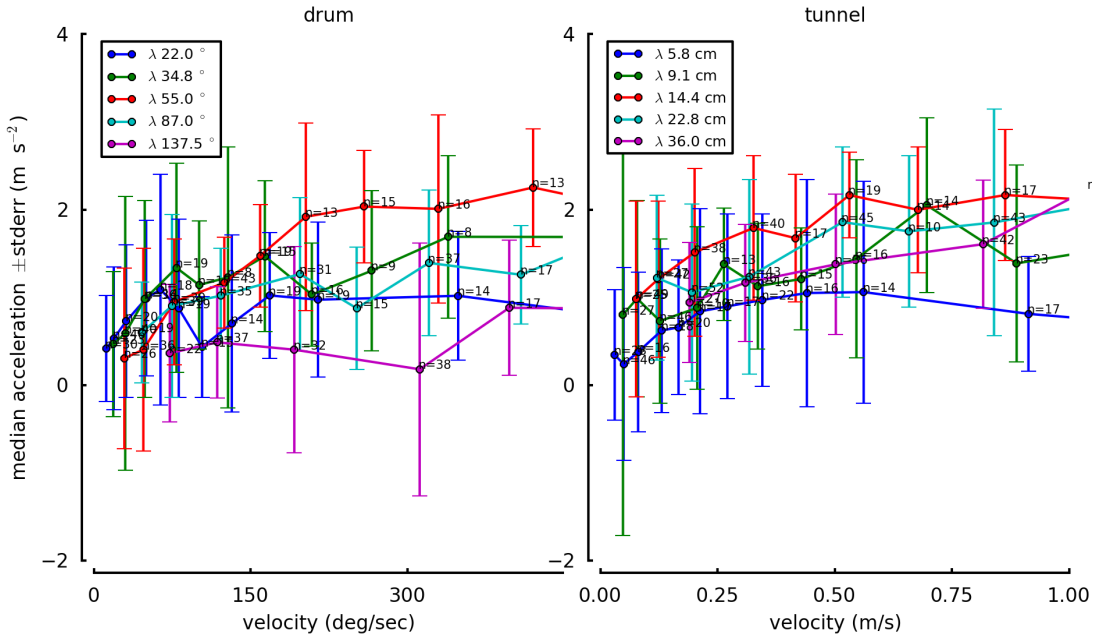

Figure 5.9: Fly acceleration response plotted as a function of velocity shows that its system is a good estimator of forward velocity for most of the tested spatial frequencies in the range of 0-1 m/s for the tunnel condition (right). The data for the drum condition is not as interpretable, as expected. Data shown in this plot is the same as in Figure 5.6, but replotted with the transformation $v_g = f_t/F_s$. Only the high-frequency blue trace appears to provides an inaccurate estimate of forward velocity for the fly.

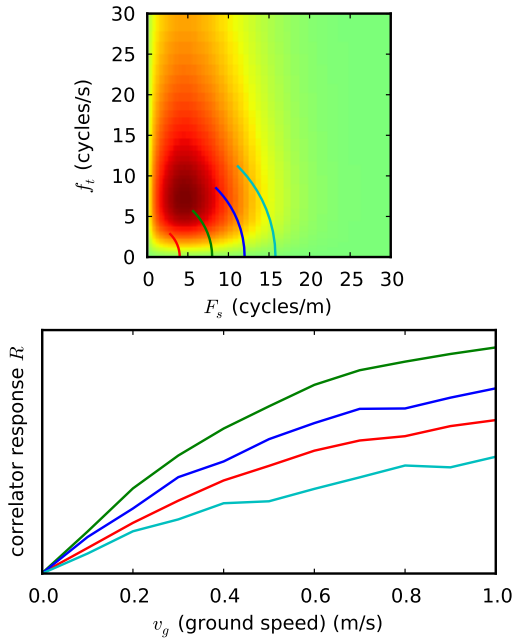

Figure 5.10: The model predicts that correlators could perform good-enough velocity estimation in the range of $F_s = 4$ to 16 cycles/m in a tunnel geometry. The response map (top) shows the strength of the sum Δ (4.20) of wall-facing correlators as a function of spatial and temporal frequencies (red=strong; green=weak). The behavior of the model correlators appears roughly linear with forward velocity up to 1 m/s (below). Compare the data of Figure 5.9, right, which plots fly responses for slightly different spatial frequencies. The response map and velocity responses resemble the plots of Figure 4 of [48], indicating that correlators could account for spatial frequency independence in the noise of the measurements of that work.

5.4 Discussion

In this work we asked whether flies' forward-flight visual velocity regulation in response to sinusoid gratings could be explained by a uniform array of correlators. We found that flies' acceleration responses have a temporal frequency tuning peak that is independent of spatial frequency, consistent with a basic correlator. We found that a model for the response of an array of correlators to a flat moving wall derived in Chapter 4 could account

for both the temporal frequency peak and the spatial frequency independence. By setting the time constant τ in the low-pass filters of the correlator model, the temporal frequency peak coincided with the measured value. And by setting the width of the blur σ of the luminance sensors, we could adjust the location of the spatial frequency peak so that the relative strengths of the responses at different spatial frequencies matched the data. The fit resembled measured data for both the drum and tunnel geometries simultaneously.

This model and the observed data both predicted a nearly spatial-frequency independent velocity response for spatial frequencies previously reported. While not a perfect velocity estimator, the response values were within the uncertainty envelope of previously-reported data in [48]. Here we propose that, in essence, the reason for the observed spatial frequency independence was that the tested spatial frequencies were all clustered around the peak of a bell-shaped response curve (Figure 5.9). For [48], they ranged from 4-16 cycles/m and for [54] they were $\approx$5 and 10 cycles/m after being scaled according to the width of that apparatus. That correlators could exhibit this spatial frequency independence had not previously been observed because an adequate model for correlators in conditions of tunnel-like geometry had not been realized. Given the strong evidence for correlators in rotatory optomotor response pathway [108, 98, 72, 70, 11, 99], then by parsimony, strong weight should be given to any explanation in which they they are used in forward flight control, as we have done here. Hence, it may not be necessary to invoke a separate pathway with a different model, such as multiply-tuned correlators [114] or the gradient method [116] to explain forward flight behavior.

In [117, 109, 70, 41], estimates of the width of the blur kernel $\Delta\rho$ have ranged from 3° to 5.5°, relatively distant from the 11° value reported here. However, a recent electrophysiological study on *Drosophila* photoreceptors found a much closer value of 8° [42].

The width of the temporal frequency peaks was much wider in the data than in the models. A number of factors could account for this. First, we selected for trajectories starting in the middle third of the wind tunnel, but even with that restriction different flies were subject to different stimuli because of different distances from the walls. The peak broadened if trajectories closer to the walls were included. Second, if the fly was rotating at all that could increase its visual stimulus. Hence, the the measured peak would be expected to be wider than the ideal peak because of variable stimulus conditions for the fly.

Future work could attempt to repeat the analysis for bees. Papers indicating spatial

frequency independence for bees have also only considered a small range of frequencies varying by at most a factor of 4 [113, 43, 22].

While different spatial frequencies exhibit small differences in velocity responses, this is likely to have a small effect on the stability of the fly. Because of the effect of the antennae described in Chapter 3, the fly should be robust to these variations.

Chapter 6

Conclusion

This thesis is concerned with dynamic control of small flying vehicles. The approach includes both scientific analysis of free-flight fruit fly behavior and engineering synthesis of fly-inspired control algorithms on a robot. The main results concern the dynamic control of forward flight motions. I found that flies use their antennae as fast sensors in an active feedback loop that damps out and thus slows their forward flight dynamics, giving the fly robustness. I demonstrated that it was possible to control the planar motion of fly-like hovercraft dynamics in a corridor using minimal computation and viual autocorrelation by controller synthesis. And I found that by invoking a model for correlators observing a flat moving surface developed for the robot controller synthesis, I found that flies' visual forward flight-speed regulation could be explained by correlators alone, rather than necessating a second distinct pathway for forward-flight visual motion detection.

The picture that emerges is an insect that uses visual correlators as a powerful sensor to measure the fly's motion relative to its fixed surroundings. But correlator-based visual feedback also comes with limitations: correlators can give variable results depending on contrast and frequency (though they may perform better on naturalistic scenery with typical power spectra [76]). Visual feedback performance is limited by a significant delay [48]. And because rate of visual flow of an obstacle motion is proportional to the inverse distance to the obstacle [43, 71, 118], the effective visual gain propagated through the feedback loop is highly variable. The fly's flight regulator must somehow function despite these sensor limitations. Wind-sensing antenna provide a valuable secondary piece of information. By using antenna feedback as a dynamic damper, the fly slows its dynamics so that visual feedback control becomes more stable. The cost is lower performance in the form of greater susceptibiliy to wind gusts at intermediate frequencies (0.1 to 0.9 Hz), but because of the

nature of turbulent flow, which follows a power spectrum with stronger components at lower frequencies, the magnitude of disturbances at these intermediate frequencies may not be of significant concern to the fly. The integral visual controller, with its high gain at low frequencies, can then act to override the damping effect of the antenna and reject strong, low-frequency disturbances including steady wind.

The results illustrate the value of approaching the problem by simultaneously performing scientific analysis and engineering synthesis. Engineering synthesis often leads to a problem-solving approach that uncovers issues that could not be foreseen by scientific study alone. The synthesis of the robot controller necessitated a more sophisticated model of correlators observing geometry wall geometry that had not previously been proposed. This in turn led to a new study on flies and the finding that their forward flight could be explained by correlators alone. On the biological side, by using quantitative tools of engineering, I uncovered how flies use antenna feedback to compensate for limitations in visual feedback.

Future work can be divided into analysis and synthesis. On the analysis side, the question remains as to how the antennae are performing wind sensing, and how feedback is translated into wing kinematic changes, including pitch angle changes. Circumstantial evidence from a small number of videos we took (n=3) seems to suggest that flies with intact antennae pitch back much farther in headwind gusts than do arista-ablated flies. What is the neuron or neurons that perform the sensory summation? What role, if any, do the antennae play in the control of unstable pitch dynamics [81, 49, 119]? What is the relationship between wind sensing and haltere feedback and is it anything like the visual efferents to halteres [120]? Does the fly use antenna feedback in a similar way in lateral wind gusts?

On the the side of robotic application, correlator-based motion control has only been demonstrated for two-dimensional robot motion, but could be extended to 6-degree-of-freedom robot. In either case, inherent noisiness of correlators suggest that an optimal Kalman filter estimator might low-pass filter their output before feeding into the controller, adding phase lag to the system. Wind sensing might thus be faster than vision for the robot as well, and provide a similar benefit. A future correlator analysis could combine the refinements of [69] and [102] with the analysis of chapter 4. A navigation or search planner [121] that could modulate the behavior of this reflexive controller as in [122] could be added on top. The reflexive controller could avoid collisions, freeing the planner from the demands of realtime responses while it worked in parallel.

Longer term, once there is a quantitative 6 degree of freedom description of the reflexive stability control in flies, inspiration will continue to be taken from their higher-level behaviors. Airborne search for odor sources [51] is a valuable robotic capability for tasks such as search and rescue, as is visual odometry and mapping performed by bees[22]. Ultimately, the powerful genetic tools available for the fruit fly will make it a particularly valuable organism for understanding complex mental processes (Figure 6.1), including the nature of mind.

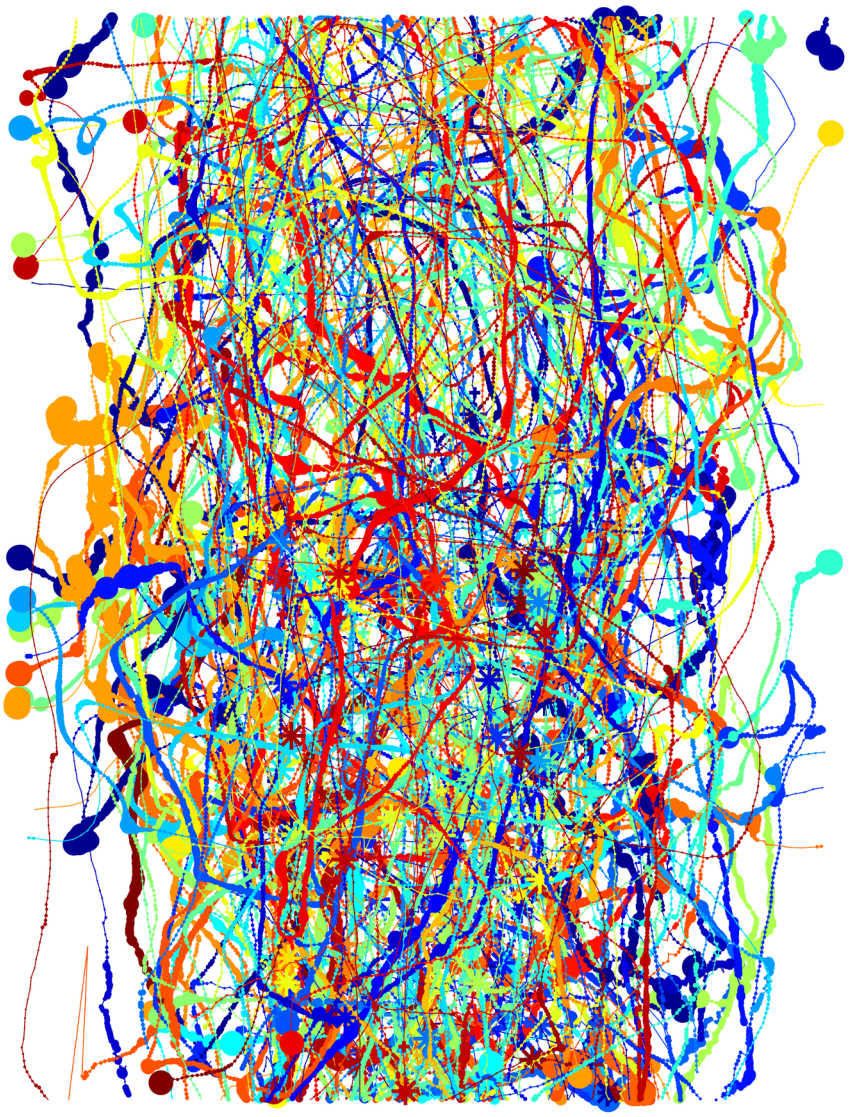

Figure 6.1: The flight trajectories of the fruit fly as they explore in the flight arena, as seen from above. Each dot is scaled according to flight speed, as if the animal was dribbling paint as it was flying. The stars indicate when the flies were subjected to a rapid gust of wind. Can we explain this complicated behavior?

Bibliography

[1] K. P. Valavanis. *Advances in Unmanned Aerial Vehicles*. Springer, 2007.

[2] James Sean Humbert and Andrew Maxwell Hyslop. Bioinspired visuomotor convergence. *IEEE Transactions on Robotics*, 26:121–130, 2010.

[3] Antoine Beyeler, Jean-Christophe Zufferey, and Dario Floreano. Vision-based control of near-obstacle flight. *Autonomous Robots*, 27:201–219, 2009. 10.1007/s10514-009-9139-6.

[4] S. Shen, N. Michael, and V. Kumar. 3D estimation and control for autonomous flight with constrained computation. In *IEEE Conference on Robotics and Automation*, Shanghai, China, May 2011.

[5] Larry Matthies, Mark Maimone, Andrew Johnson, Yang Cheng, Reg Willson, Carolos Villalpando, Steve Goldberg, Andres Huertas, Andrew Stein, and Anelia Angelova. Computer vision on mars. *International Journal of Computer Vision*, 74:67–92, 2007.

[6] Dario Floreano and Claudio Mattiussi. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. The MIT Press, 2008.

[7] Michael H. Dickinson, Claire T. Farley, Robert J. Full, M. A. R. Koehl, Rodger Kram, and Steven Lehman. How animals move: An integrative view. *Science*, 288:100–106, 2000.

[8] Werner Nachtigall. *Insects in flight*. HarperCollins Publishers Ltd., 1974.

[9] Michael H. Dickinson, Fritz-Olaf Lehmann, and Sanjay P. Sane. Wing rotation and the aerodynamic basis of insect flight. *Science*, 204:1954–1960, 1999.

[10] J. L. D. Williams. Anatomical studies of the insect central nervous system: A ground-plan of the midbrain and an introduction to the central complex in the locust *Schistocerca Gregaria*. *Journal of Zoology*, 176:67–86, 1975.

[11] Klaus Hausen. *Photoreception and vision in invertebrates*, chapter The lobula-complex of the fly: Structure, function, and significance in visual behavior, pages 523–559. Plenum, 1984.

[12] K.-F. Fischbach and A. P. M. Dittrich. The optic lobe of *Drosophila melanogaster* i: A golgi analysis of wild-type structure. *Cell Tissue Research*, 258:441–475, 1989.

[13] Roland Strauss. The central complex and the genetic dissection of locomotor behavior. *Current Opinion in Neurobiology*, 12:633–638, 2002.

[14] Mark A. Frye and Michael H. Dickinson. Fly flight: A model for the neural control of complex behavior. *Neuron*, 32:385–388, 2001.

[15] Alexander Borst, Juergen Haag, and Dierk F. Reiff. Fly motion vision. *Annu Rev Neurosci*, 33:49–70, 2010.

[16] Norbert Wiener. *Cybernetics: or control and communication in the animal and the machine*. The MIT Press, 1965.

[17] Rodney A. Brooks. New approaches to robotics. *Science*, 253:1227–1232, September 1991.

[18] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, March 1986.

[19] Nicolas Franceschini, J. M Pichon, C Blanes, and J. M. Brady. From insect vision to robot vision. *Philosophical Transactions: Biological Sciences*, 337:283–294, 1991.

[20] Matthias O. Franz, Titus R. Neumann, Michael Plagge, Hanspeter A. Mallot, and Andreas Zell. Can fly tangential neurons be used to estimate self-motion? In *Proceedings of 9th International Conference on Artificial Neural Networks*, pages 994–999, 1999.

[21] K. Weber, S. Venkatesh, and M. V. Srinivasan. *From Living Eyes to Seeing Machines*, chapter Insect inspired behaviors for the autonomous control of mobile robots, pages 226–248. Oxford University Press, 1997.

[22] Mandyam V. Srinivasan. Honeybees as a model for the study of visually guided flight, navigation, and biologically inspired robotics. *Physiological Reviews*, 91(2):413–460, 2011.

[23] Holger G. Krapp and Roland Hengstenberg. Estimation of self-motion by optic flow processing in single visual interneurons. *Nature*, 384:463–466, 1996.

[24] Holger G. Krapp, Barbel Hengstenberg, and Roland Hengstenberg. Dendritic structure and receptive-field organization of optic flow processing interneurons in the fly. *American Physiological Society*, pages 1902–1917, 1998.

[25] Rüdiger Wehner. Matched filters: Neural models of the external world. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 161:511–531, 1987. 10.1007/BF00603659.

[26] Hans-Jurgen Dahmen, Matthias O. Franz, and Holger G. Krapp. *Motion Vision - Computational, Neural, and Ecological Constraints*, chapter Extracting Egomotion from Optic Flow: Limits of Accuracy and Neural Matched Filters, pages 143–168. Springer Verlag, Berlin, 2001.

[27] J. Sean Humbert, Richard M. Murray, and Michael H. Dickinson. Sensorimotor convergence in visual navigation and flight control systems. In *16th International Federation of Automatic Control World Congress, Prague, Czech Republic*, 2005.

[28] 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*, 205:327–343, 2002.

[29] J.-C. Zufferey and D. Floreano. Fly-inspired visual steering of an ultralight indoor aircraft. *IEEE Transactions on Robotics*, 22(1):137–146, Feb. 2006.

[30] Nicolas Franceschini, Franck Ruffier, and Julien Serres. A bio-inspired flying robot sheds light on insect piloting abilities. *Current Biology*, 17:1–7, 2007.

[31] Andrew D Straw, Serin Lee, and Michael H Dickinson. Visual control of altitude in flying *Drosophila*. *Current Biology*, 20(17):1550–1556, Sep 2010.

[32] Andrea Censi, Shuo Han, Sawyer B. Fuller, and Richard M. Murray. A bio-plausible design for visual attitude stabilization. In *IEEE Conference on Decision and Control*, 2009.

[33] Shuo Han, Andrea Censi, Andrew D. Straw, and Richard M. Murray. A bio-plausible design for visual pose stabilization. In *International Conference on Intelligent Robots and Systems*, 2010.

[34] Valentino Braitenburg. *Vehicles: Experiments in synthetic psychology.* The MIT Press, 1986.

[35] Auke Jan Ijspeert, Alessandro Crespi, Dimitri Ryczko, and Jean-Marie Cabelguen. From swimming to walking with a salamander robot driven by a spinal cord model. *Science*, 315:1416–1420, March 2007.

[36] P. S. Sreetharan and R. J Wood. Passive torque regulation in an underactuated flapping wing robotic insect. In *Robotics: Science and Systems*, Zaragoza, Spain, June 2010.

[37] Michael H. Dickinson and Michael S. Tu. The function of the dipteran flight muscle. *Journal of Comparative Physiology A*, 116A:230–238, 1997.

[38] Graham K. Taylor and Holger G Krapp. Sensory systems and flight stability: What do insects measure and why? *Advances in Insect Physiology*, 34:231–316, 2008.

[39] Gwyneth Card and Michael H. Dickinson. Performance trade-offs in the flight initiation of *Drosophila*. *The Journal of Experimental Biology*, 211:341–353, 2008.

[40] Nicholas J. Strausfeld. *Atlas of an Insect Brain.* Springer-Verlag, 1976.

[41] M. F. Land. Visual acuity in insects. *Annu Rev Entomol*, 42:147–177, 1997.

[42] Paloma T. Gonzalez-Bellido, Trevor J. Wardill, and Mikko Juusola. Compound eyes and retinal information processing in miniature dipteran species match their specific ecological demands. *Proceedings of the National Academy of Sciences*, 108(10):4224–4229, 2011.

[43] Srinivasan, Zhang, Lehrer, and Collett. Honeybee navigation en route to the goal: visual flight control and odometry. *J Exp Biol*, 199(Pt 1):237–244, 1996.

[44] M.B. Reiser, J.S. Humbert, M.J. Dunlop, D. Del Vecchio, R.M. Murray, and M.H. Dickinson. Vision as a compensatory mechanism for disturbance rejection in upwind flight. In *American Control Conference*, volume 1, pages 311–316, July 2004.

[45] Michael B Reiser and Michael H Dickinson. *Drosophila* fly straight by fixating objects in the face of expanding optic flow. *Journal of Experimental Biology*, 213(Pt 10):1771–1781, May 2010.

[46] Schilstra and Hateren. Blowfly flight and optic flow. i: Thorax kinematics and flight dynamics. *The Journal of Experimental Biology*, 202:1481–1490, Jun 1999.

[47] Nicola Rohrseitz and Steven N. Fry. Behavioural system identification of visual flight speed control in *Drosophila melanogaster*. *Journal of the Royal Society Interface*, 2010.

[48] Steven N. Fry, Nicola Rohrseitz, Andrew D. Straw, and Michael H. Dickinson. Visual flight speed control in *Drosophila melanogaster*. *The Journal of Experimental Biology*, 212:1120–1130, 2009. In preparation.

[49] William B. Dickson, Andrew D. Straw, and Michael H. Dickinson. Integrative model of drosophila flight. *Journal of the American Institute of Aeronautics and Astronautics*, 46(9):2150–2164, September 2008.

[50] Dietrich Burkhardt and Michael Gewecke. Mechanoreception in]textitArthropoda: the chain from stimulus to behavioral pattern. *Cold Spring Harbor Symposium on Quantitative Biology*, 30:601–614, 1965.

[51] Seth A Budick and Michael H Dickinson. Free-flight responses of *Drosophila melanogaster* to attractive odors. *The Journal of Experimental Biology*, 209(Pt 15):3001–3017, Aug 2006.

[52] Seth A. Budick, Michael B. Reiser, and Michael H. Dickinson. The role of visual and mechanosensory cues in structuring forward flight in *Drosophila melanogaster*. *The Journal of Experimental Biology*, 210:4092–4103, 2007.

[53] John S. Kennedy. The visual responses of flying mosquitos. *Proceedings of the Zoological Society of London*, page 22, 1940.

[54] Charles T. David. Compensation for height in the control of groundspeed by *Drosophila* in a new, 'barber's pole' wind tunnel. *Journal of Comparative Physiology A*, 147:485–493, 1982.

[55] F. S. J. Hollick. The flight of the dipterous fly muscina stabulans fallen. *Philosophical Transactions of the Royal Society of London B*, 230:357–390, 1940.

[56] Martin C Göpfert and Daniel Robert. The mechanical basis of *Drosophila audition*. *The Journal of Experimental Biology*, 205(Pt 9):1199–1208, May 2002.

[57] Alice A. Robie, Andrew D. Straw, and Michael H. Dickinson. Object preference by walking fruit flies, *Drosophila melanogaster*, is mediated by vision and graviperception. *The Journal of Experimental Biology*, 213(Pt 14):2494–2506, Jul 2010.

[58] Azusa Kamikouchi, Hidehiko K Inagaki, Thomas Effertz, Oliver Hendrich, André Fiala, Martin C Göpfert, and Kei Ito. The neural basis of drosophila gravity-sensing and hearing. *Nature*, 458(7235):165–171, Mar 2009.

[59] Akira Mamiya, Andrew D. Straw, Egill Tómasson, and Michael H. Dickinson. Active and passive antennal movements during visually guided steering in flying drosophila. *The Journal of Neuroscience*, 31(18):6900–6914, 2011.

[60] M. Gewecke. Antennae: A nother wind-sensitive receptor in locusts. *Nature*, 225(5239):1263–1264, Mar 1970.

[61] M. Gewecke, H.-G. Heinzel, and J. Philippen. Role of antennae of the dragonfly *Orthetrum cancellatum* in flight control. *Nature*, 249:584–585, 1974.

[62] Suzuko Yorozu, Allan Wong, Brian J. Fischer, Heiko Dankert, Maurice J. Kernan, Azusa Kamikouchi, Kei Ito, and David J. Anderson. Distinct sensory representations of wind and near-field sound in the *Drosophila* brain. *Nature*, 458(7235):201–205, Mar 2009.

[63] William B. Dickson, Peter Polidoro, Melissa M. Tanner, and Michael H. Dickinson. A linear systems analysis of the yaw dynamics of a dynamically scaled insect model. *The Journal of Experimental Biology*, 213(Pt 17):3047–3061, Sep 2010.

[64] Steven N. Fry, Rosalyn Sayaman, and Michael H. Dickinson. The aerodynamics of hovering flight in *Drosophila*. *The Journal of Experimental Biology*, 208:2303–2318, 2005.

[65] A. D. Straw, K. Branson, T. R. Neumann, and M. H. Dickinson. Multi-camera realtime 3d tracking of multiple flying animals. *Computer Vision and Image Understanding*, 2010.

[66] Erich Buchner. *Dunkelanregung des stationaeren Flugs der Fruchtfliege Drosophila*. PhD thesis, University Tuebingen, Tuebingen, Germany, 1971.

[67] Andrew D. Straw. Vision egg: an open-source library for realtime visual stimulus generation. *Frontiers in Neuroinformatics*, 2008.

[68] 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, Aug 2009.

[69] Ron O. Dror, David C. O'Carroll, and Simon B. Laughlin. Accuracy of velocity estimation by Reichardt Correlators. *Journal of the Optical Society of America A*, 18(2):241–252, February 2001.

[70] E. Buchner. *Photoreception and Vision in Invertebrates*, chapter Behavioral Analysis of Spatial Vision in Insects, pages 561–621. Plenum, 1984.

[71] Jan J. Koenderink. Optic flow. *Vision Research*, 26(1):161–180, 1986.

[72] W. Reichardt. *Movement Perception in Insects*, chapter Movement perception in Insects, pages 465–493. Academic Press, New York, 1969.

[73] Klaus Hausen. *Visual Motion and its Role in the Stabilization of Gaze*, chapter Decoding of retinal image flow in insects, pages 203–235. Elsevier, 1993.

[74] Steven Vogel. Flight in *Drosophila*: I. Flight performance of tethered flies. *Journal of Experimental Biology*, 44(1):567–578, 1966.

[75] Robert A. Harris, David C. O'Carroll, and Simon B. Laughlin. Contrast gain reduction in fly motion adaptation. *Neuron*, 28:595–606, November 2000.

[76] Andrew D Straw, Tamath Rainsford, and David C O'Carroll. Contrast sensitivity of insect motion detectors to natural images. *Jornal of Vision Research*, 8(3):32.1–32.9, 2008.

[77] Lance F Tammero and Michael H Dickinson. Collision-avoidance and landing responses are mediated by separate pathways in the fruit fly, *Drosophila melanogaster*. *The Journal of Experimental Biology*, 205(Pt 18):2785–2798, Sep 2002.

[78] Barbara Webb. Neural mechanisms for prediction: do insects have forward models? *Trends in Neuroscience*, 27(5):278–282, May 2004.

[79] David J. Randall, Warren W. Burggren, Kathleen French, and Roger Eckert. *Eckert Animal Physiology: Mechanisms and Adaptations*. W.H. Freeman and Co., 2002.

[80] M. F. Land and T. S Collett. Chasing behavior of houseflies. *Journal of Comparative Physiology A*, 89:331–357, 1974.

[81] Michael Epstein, Stephen Waydo, Sawyer Fuller, Andrew D. Straw, William B. Dickson, Michael H. Dickinson, and Richard M. Murray. Biologically inspired feedback design for *Drosophila* flight. In *American Control Conference*, pages 3395–3401, July 2007.

[82] Titus R. Neumann and Heinrich H. Bulthoff. Behavior-oriented vision for biomimetic flight control. In *Proceedings of EPSRC/BBSRC International Workshop on Biologically Inspired Robotics*, pages 196–203, August 2002.

[83] J. Keshavan and J.S. Humbert. Mav stability augmentation using weighted outputs from distributed hair sensor arrays. In *American Control Conference*, pages 4445 –4450, 30 2010-july 2 2010.

[84] Werner Reichardt and Tomaso Poggio. Visual control of orientation behavior in the fly. *Quarterly Reviews in Biophysics*, 9(3):311–375, 1976.

[85] Steven N. Fry, Rosalyn Sayaman, and Michael H. Dickinson. The aerodynamics of free-flight maneuvers in *Drosophila*. *Science*, 300:495–498, 2003.

[86] John A. Bender and Michael H. Dickinson. A comparison of visual and haltere-mediated feedback in the control of body saccades in *Drosophila melanogaster*. *J Exp Biol*, 209(Pt 23):4597–4606, Dec 2006.

[87] Tyson L Hedrick, Bo Cheng, and Xinyan Deng. Wingbeat time and the scaling of passive rotational damping in flapping flight. *Science*, 324(5924):252–255, Apr 2009.

[88] Thomas Hesselberg and Fritz-Olaf Lehmann. Turning behaviour depends on frictional damping in the fruit fly drosophila. *The Journal of Experimental Biology*, 210(Pt 24):4319–4334, Dec 2007.

[89] Leif Ristroph, Attila J. Bergou, Gunnar Ristroph, Katherine Coumes, Gordon J. Berman, John Guckenheimer, Z. Jane Wang, and Itai Cohen. Discovering the flight autostabilizer of fruit flies by inducing aerial stumbles. *Proceedings of the National Academy of Sciences*, 107(11):4820–4824, Mar 2010.

[90] Charles T. David. The relationship between body angle and flight speed in free-flying *Drosophila*. *Physiological Entomology*, 3:191–195, 1978.

[91] Dawnis M. Chow and Mark A. Frye. Context-dependent olfactory enhancement of optomotor flight control in *Drosophila*. *The Journal of Experimental Biology*, 211(Pt 15):2478–2485, Aug 2008.

[92] Finlay J Stewart, Dean A Baker, and Barbara Webb. A model of visual-olfactory integration for odour localisation in free-flying fruit flies. *The Journal of Experimental Biology*, 213(11):1886–1900, Jun 2010.

[93] Karl Johan Astrom and Richard M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 41 William Street, Princeton, NJ 08540, 2008.

[94] Michael Karpelson, John P. Whitney, Gu-Yeon Wei, and Robert J. Wood. Energetics of flapping-wing robotic insects: Towards autonomous hovering flight. In *International Conference on Intelligent Robots and Systems*, 2010.

[95] Michael H. Dickinson. Haltere-mediated equilibrium reflexes of the fruit fly, *Drosophila melanogaster*. *Philosophical Transactions of the Royal Society of London B*, 354:903–916, 1999.

[96] Simon B. Laughlin, Rob R. de Ruyter van Steveninck, and John C. Anderson. The metabolic cost of neural information. *Nature Neuroscience*, 1:36–41, 1998.

[97] Yang Cheng, Mark Maimone, and Larry Matthies. Visual odometry on the mars exploration rovers - a tool to ensure accurate driving and science imaging. *Robotics and Automation Magazine*, 13(2):54 – 62, June 2006.

[98] W. Reichardt. *Principles of Sensory Communication*, chapter Autocorrelation, a principle for the evaluation of sensory information by the central nervous system., pages 303–317. John Wiley & Sons, New York, 1961.

[99] Martin Egelhaaf, Alexander Borst, and Werner 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 A*, 6(7):1070–1087, Jul 1989.

[100] Maximilian Joesch, Johannes Plett, Alexander Borst, and Dierk F. Reiff. Response properties of motion-sensitive visual interneurons in the lobula plate of *Drosophila melanogaster*. *Current Biology*, 18(5):368–374, Mar 2008.

[101] Mandyam V Srinivasan and Shaowu Zhang. Visual motor computations in insects. *Annual Reviews of Neuroscience*, 27:679–696, 2004.

[102] 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, 11 2009.

[103] Timothy Chung, Lars Cremean, William B. Dunbar, Zhipu Jin, Eric Klavins, David Moore, Abhishek Tiwari, Dave van Gogh, and Stephen Waydo. A platform for cooperative and control of multiple vehicles: the caltech multi-vehicle wireless testbed. In *Conference on Cooperative Control and Optimization*, 2002.

[104] Sawyer B. Fuller, Eric J. Wilhelm, and Joseph M. Jacobson. Ink-jet printed nanoparticle microelectromechanical systems. *IEEE Journal of Microelectromechanical Systems*, 11(1):54–60, Feb. 2002.

[105] Berthold K. P. Horn and Brian G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

[106] Werner Reichardt. *Sensory Communication: Contributions to the Symposium on Principles of Sensory Communication*, chapter Autocorrelation, a principle for the Evaluation of Sensory Information by the Central Nervous System, pages 303–317. MIT Press, 1959.

[107] Alexander Borst and Martin Egelhaaf. Principles of visual motion detection. *Trends in Neuroscience*, 12:297–306, 1989.

[108] B. Hassenstein and W. Reichardt. Systemtheoretische analyse der zeit-, reihenfolgen- und vorzeichenauswertung bei der bewegungsperzeption des russelkafers chlorophanus. *Zeitschrift Fur Naturforschung*, 11b:513–524, 1956.

[109] Martin Heisenberg and Erich Buchner. The role of retinula cell types in visual behavior in *Drosophila melanogaster*. *Journal of Comparative Physiology A*, 117:127–162, 1977.

[110] Pat A Shoemaker, David C. O'Carroll, and Andrew D. Straw. Velocity constancy and models for wide-field visual motion detection in insects. *Biological Cybernetics*, 93(4):275–287, October 2005.

[111] J. P. Lindemann, R. Kern, J. H. van Hateren, H. Ritter, and M. Egelhaaf. On the computations analyzing natural optic flow: Quantitative model analysis of the blowfly motion vision pathway. *Journal of Neuroscience*, 25:6435–6448, 2005.

[112] Martin Egelhaaf and Roland Kern. Vision in flying insects. *Current Opinion in Neurobiology*, 12:699–706, 2002.

[113] Emily Baird, Mandyam V. Srinivasan, Shaowu Zhang, and Ann Cowling. Visual control of flight speed in honeybees. *Journal of Experimental Biology*, 208(20):2895–3905, 2005.

[114] Mandyam V. Srinivasan, Michael Poteser, and Karll Kral. Motion detection in insect orientation and navigation. *Vision Research*, 39:2749–2766, 1999.

[115] Eric Buchner. Elementary movement detectors in an insect visual system. *Biological Cybernetics*, 24:85–101, 1976.

[116] M. V. Srinivasan. Generalized gradient schemes for the measurement of two-dimensional image motion. *Biological Cybernetics*, 63:421–431, 1990.

[117] K. G. Götz. [optomotor studies of the visual system of several eye mutants of the fruit fly drosophila]. *Kybernetik*, 2(2):77–92, Jun 1964.

[118] Jan J. Koenderink and Andrea J. van Doorn. Facts on optic flow. *Biological Cybernetics*, 56:247–254, 1987.

[119] Imraan Faruque and J. Sean Humbert. Dipteran insect flight dynamics. Part 1: Longitudinal motion about hover. *Journal of Theoretical Biology*, 264:538–552, 2010.

[120] Wai Pang Chan, Frederick Prete, and Michael H Dickinson. Visual input to the efferent control system of of a fly's "gyroscope". *Science*, 280:289–292, 1998.

[121] Matthias O. Franz and Hanspeter A. Mallot. Biomimetic robot navigation. *Robotics and Autonomous Systems*, 30:133–153, 2000.

[122] Rodney A. Brooks. A robot that walks: Emergent behaviors from a carefully evolved network. *IEEE International Conference on Robotics and Automation*, pages 292–296, May 1989.