

**TASK-LEVEL MODELS FOR  
IMAGE-STABILIZATION BEHAVIORS IN ANIMALS**

by

Eatai Roth

A dissertation submitted to The Johns Hopkins University in conformity with the  
requirements for the degree of Doctor of Philosophy.

Baltimore, Maryland

June, 2012

© Eatai Roth 2012

All rights reserved

# Abstract

This research addresses a fundamental question in biology and neuroscience: how do animals process sensory information for the control of locomotor behaviors? Behaviors can be described as a sensorimotor loop: sensing (sensori-) governs action (-motor), action changes the environment, and these changes are perceived via sensing. Animal behavior arises from a concert of sensory, computational, and mechanical systems. Often, these mechanisms are studied independently (and often isolated from the context of the behavior) and the behavioral model is constructed from knowledge of the constituents, a bottom-up synthesis. Complementary to this approach, we model behavior at the level of the sensorimotor loop (the task-level) and subsequently generate hypotheses as to the mechanistic constituents. These top-down models serve to constrain permissible mechanisms and identify necessary neural computations.

We design an assay of experiments and frequency-domain analyses to identify task-level behavioral models, specifically for image-stabilization behaviors. Image-stabilization describes a broad class of behaviors in which animals modulate movement to fixate a sensory signal. In this dissertation, we study analogous behaviors in two

## ABSTRACT

species: refuge-tracking in weakly electric knifefish and stripe-fixation in fruit flies.

Glass knifefish swim forward and backward to maintain their position relative to a moving refuge. Fish were recorded performing refuge-tracking behavior for sinusoidal (predictable) and sum-of-sines (pseudo-random) refuge trajectories. System identification reveals a notable nonlinearity in the behavior; the frequency response functions (FRFs) generated from predictable and pseudo-random experiments are categorically different. The data support the hypothesis that fish generate an internal dynamical model of the stimulus motion, hence enabling improved tracking of predictable trajectories (relative to unpredictable ones) despite similar or reduced motor cost.

Fruit flies adeptly coordinate flight maneuvers to seek, avoid, or otherwise interact with salient objects in their environment. In the laboratory, tethered flies modulate yaw torque to steer towards a dark vertical visual stimulus. This stripe-fixation behavior is robust and repeatable; in series of experiments, flies stabilize moving stripes oscillating over a range of frequencies. We parameterize this FRF description to hypothesize a Proportional-Integral-Derivative (PID) control model for the fixation behavior. We demonstrate that our hypothesized PID model provides a parsimonious explanation for several previously reported phenomena.

Primary Reader: Noah J. Cowan

Secondary Readers: Thomas L. Daniel, Eric S. Fortune, and Sridevi V. Sarma

# Acknowledgments

Above all, family. The family you're born into is the greatest lottery in life, and I won. I have an amazing family who have encouraged and supported me my whole life. My parents are my role models: smart and selfless, they exemplify altruism, idealism, and pragmatism. I hope the person I am reflects well on my parents. My brother and sister, Amir and Nurit, are diametric opposites in a beautiful way. Nurit showed me what it was to be rebellious and an individual and she still marches to the beat of her own drum. In my youth, Amir was my tormentor and, as an adult, he is my best friend. He told me time and again not to get a PhD, to turn back, that it wasn't worth it; that was reason enough for me to do it. And to Marci and Irah, my two smart siblings-in-law, thanks for being supportive and engaged and interested even when I was down and standoffish and boring (hardly ever). Daniel, Elliot, Allison, and Leela, thanks for being kids, my favorite kids; you keep me young, so don't grow up. Also, thanks Ramona. I love you all.

Noah Cowan showed me what it is to be a scientist. He is everything a scientist and an advisor should be: brilliant and continuously curious, demanding and patient,

## ACKNOWLEDGMENTS

pragmatic and idealistic, fair and compassionate. He was my boss when he needed to be my boss and he was my friend always. My thesis committee —Tom Daniel, Eric Fortune and Sri Sarma and informally David Solomon and Louis Whitcomb—have helped me gain new perspective on my work. Our conversations, their comments on this dissertation and their questions and discussion at the thesis defense have provided valuable insights that will shape the future direction of my research. Where I was narrowly focused on the questions addressed within this thesis, they showed me the forest. Thank you, professors. And while I’m thanking those doctors who helped so much in shaping my thesis, I thank Larry David for giving me perspective on everything else.

Thank you to the LIMBS lab present and past, my academic family: Jusuk Lee, Vinutha Kallem, John Swensen, Manu Madhav, Mert Ankarali, Aliçan Demir, Sarah Stamper, Sean Carver, Eric Tytell, Elyse Edwards, Rob Grande, Terrence Jao, Katie Zhuang, Avik De, Rachel Jackson, and Shahin Sefati. I’ve made many great friends in the LCSR and ME department: Mike Kutzer, Matt Moses, Kevin Wolfe, Yan Yan, Tom Wedlick, Tricia Gibo, David Grow, Paul Griffiths, Sarah Webster, Amy Blank, Giancarlo Troni, Topher McFarland, Yixin Gao, Ehsan Basafa, Sue Sue U-Thainual, Garrick Orchard, Alex Russell, Cindy Byer, Rika Wright Carlsen, Jessica Meulbroek, Emily Huskins, Andy Tonge and Tian Xia. I’ve learned a lot from this group. Thanks for the help. Thanks for discussing research. Thanks for that crucial citation. Thanks for finding the error in my code (that bug was annoying). Thanks

## ACKNOWLEDGMENTS

for lunch and coffee and tiramisu. Thanks for abetting in my procrastination. Thanks for keeping me sane. I look forward to our paths crossing in the future, my esteemed colleagues.

Then there are the people behind the curtain that make everything run smoothly, the administrative staff at LCSR and the ME department: Alison Morrow, Jamie Meehan, Mike Bernard, Megan Herbert, Deana Santoni, Lorrie Dodd, Katy Sanderson and Christine Parks. Admittedly, I'd sometimes look forward to getting a package or having a mistake on my payroll just for an excuse to gab with the staff. Thank you all for all your help over the years and I'll miss shooting the breeze. And thanks to the original Café Q gang for keeping me caffeinated and keeping me company: Kamilah, B. Josh, Angela, Toni, Amy, Kim, Skulls, and Rome.

Very sincerely, thank you all.

# Dedication

This thesis is dedicated to my parents, Zipora and Itzhak Roth, who raised me to be a mensch.

# Contents

|  |           |
|--|-----------|
| <b>Abstract</b>  | <b>ii</b> |
| <b>Acknowledgments</b>   | <b>iv</b> |
| <b>1 Introduction</b>  | <b>1</b>  |
| 1.1 The Allure of Animal Behavior . . . . .                            | 2         |
| 1.2 Image-Stabilization Behaviors . . . . .                            | 4         |
| 1.3 Contrasting Approaches to Ethology . . . . .                       | 5         |
| 1.3.1 Bottom-up: a place for reductionism . . . . .                    | 7         |
| 1.3.2 Top-down: why complex systems should be studied intact . . . . . | 10        |
| 1.3.3 Reconciling the top-down and bottom-up approaches . . . . .      | 15        |
| 1.4 Dissertation Organization . . . . .                                | 18        |
| 1.5 Contribution and Dissemination . . . . .                           | 19        |
| <b>2 Frequency-domain System Identification</b>                        | <b>21</b> |
| 2.1 Motivating the Frequency-domain Approach . . . . .                 | 22        |



# CONTENTS

|          |  |           |
|----------|--|-----------|
| 2.1.1    | Why a non-parametric approach? . . . . .                             | 24        |
| 2.1.2    | Why the frequency domain? . . . . .                                  | 26        |
| 2.1.3    | A role for linear models in describing image-stabilization behaviors | 30        |
| 2.1.4    | Extending frequency analyses to other image-stabilization tasks      | 32        |
| 2.2      | Mathematical Tools . . . . .   | 34        |
| 2.2.1    | Coherence . . . . .  | 34        |
| 2.2.2    | Frequency response functions (FRFs) . . . . .                        | 36        |
| <b>3</b> | <b>Refuge-Tracking in <i>Eigenmannia virescens</i></b>               | <b>42</b> |
| 3.1      | Why Weakly Electric Knifefish? . . . . .                             | 43        |
| 3.1.1    | Mechanisms . . . . .   | 45        |
| 3.1.2    | Behaviors . . . . .  | 46        |
| 3.1.3    | Our approach . . . . .   | 48        |
| 3.2      | Materials and Methods . . . . .                                      | 49        |
| 3.2.1    | Animal husbandry and preparation . . . . .                           | 49        |
| 3.2.2    | Experimental apparatus . . . . .                                     | 49        |
| 3.2.3    | Experiment Design . . . . .  | 51        |
| 3.2.4    | Analysis . . . . .   | 53        |
| 3.2.4.1  | Frequency response functions . . . . .                               | 53        |
| 3.2.4.2  | Continuous phase estimation . . . . .                                | 54        |
| 3.3      | Results . . . . .  | 57        |
| 3.3.1    | Responses to motion stimuli are coherent . . . . .                   | 57        |

## CONTENTS

|          |  |           |
|----------|--|-----------|
| 3.3.2    | Linear models do not generalize across stimulus classes . . . . .                      | 59        |
| 3.3.3    | Fish adapt to changes in stimulus . . . . .  | 63        |
| 3.3.4    | Adaptation to single-sine stimuli reduces tracking error . . . . .                     | 66        |
| 3.4      | Discussion . . . . .   | 69        |
| 3.4.1    | Responses to single-sine and sum-of-sine stimuli . . . . .                             | 69        |
| 3.4.2    | An internal model predicting refuge movement explains phase<br>discrepancies . . . . . | 71        |
| 3.4.3    | Gain discrepancies indicate improved tracking for predictable<br>stimuli . . . . .     | 74        |
| <b>4</b> | <b>Optomotor Yaw Regulation in <i>Drosophila melanogaster</i></b>                      | <b>76</b> |
| 4.1      | Why the Fruit Fly? . . . . .   | 77        |
| 4.1.1    | Mechanisms . . . . .   | 78        |
| 4.1.2    | Behaviors . . . . .  | 81        |
| 4.1.3    | Our approach . . . . .   | 83        |
| 4.2      | Materials and Methods . . . . .  | 84        |
| 4.2.1    | Animal husbandry and preparation . . . . .   | 85        |
| 4.2.2    | Experiment design . . . . .  | 87        |
| 4.2.3    | Analysis . . . . .   | 90        |
| 4.3      | Results . . . . .  | 93        |
| 4.3.1    | Empirical frequency responses . . . . .  | 93        |
| 4.3.2    | A linear transfer function model . . . . .   | 95        |

## CONTENTS

|          |  |            |
|----------|--|------------|
| 4.4      | Discussion . . . . .   | 97         |
| 4.4.1    | Progressive-regressive response concurs with PID control . . .                                 | 99         |
| 4.4.2    | Positional cues explain differences between wide-field and small-<br>field responses . . . . . | 101        |
| 4.4.3    | Replay experiments yield predicted results . . . . .   | 102        |
| <b>5</b> | <b>Conclusion</b>  | <b>107</b> |
| 5.1      | Future work: A Systems Approach to Modularity . . . . .  | 109        |
|          | <b>Bibliography</b>  | <b>113</b> |
|          | <b>Vita</b>  | <b>132</b> |

# Chapter 1

## Introduction

*Si Dieu nous a faits à son image,  
nous le lui avons bien rendu.*

If God has made us in His image,  
we have returned Him the favor.

---

Voltaire

## 1.1 The Allure of Animal Behavior

Animal behavior has long inspired scientists and engineers to discover and recreate Nature's implementation of sensory processing, motor control, learning and navigation. Animals interact with the environment through a diverse set of behaviors, requiring them to draw relevant information from their surroundings from numerous sensors (eyes, antennae, ears, etc.) to produce motion via similarly diverse effectors (legs, wings, fins, etc.). And though these behaviors manifest from evolved mechanisms and an economy of neural architecture (merely 100,000 neurons comprise the fruit fly nervous system), the performance achieved by animals far exceeds that of even the most impressive robots they've inspired (*e.g.* terrestrial walkers [Saranli et al.(2001), Buehler et al.(2005), Clark et al.(2001)], wing-flapping fliers [Fearing et al.(2000), Wood(2007)], undulating swimmers [Mason and Burdick(2000), Ijspeert et al.(2007)] and jumping robots [Bergbreiter and Pister(2007)]). Much of the performance disparity between these state-of-the-art robots and their natural counterparts can be attributed to shortcomings in technology (*e.g.* a lack of amenable materials and fabrication processes, insufficient batteries or actuators, etc.). But in terms of the hardware required for computation, we have sufficient technology; modern microprocessors and GPUs boast several billion transistors, dwarfing the number of neurons in the fruit fly or the knifefish or even mice, dogs and cats (yet still an order of magnitude shy of the number of neurons comprising the human nervous system). So

## CHAPTER 1. INTRODUCTION

what are we missing? As a mechanical engineer and roboticist, I marvel at the diversity and robustness of behaviors exhibited by even some of the “simplest” animals. This research addresses a fundamental question in biology and neuroscience: how do animals process sensory information for the control of locomotor behaviors?

In this work, we investigate fixation behaviors in two different species: refuge tracking in the weakly electric knifefish *Eigenmannia virescens* [Rose and Canfield(1993a), Cowan and Fortune(2007)] and optomotor yaw regulation in the fruit fly *Drosophila melanogaster* [Götz(1968), Götz et al.(1979), Reichardt and Poggio(1976), Poggio and Reichardt(1976), Poggio and Reichardt(1981), Heisenberg and Wolf(1988), Wolf and Heisenberg(1990), Tammero and Dickinson(2002), Reiser and Dickinson(2008), Duistermars et al.(2007)]. In the refuge-tracking behavior, knifefish swim forwards and backwards to maintain a constant relative position with moving objects in their environment, relying on the integration of visual, electrosensory and mechanosensory information; for our experiments, fish track an actuated refuge. *Drosophila* exhibit an analogous behavior, flying towards salient vertical features; in the laboratory, tethered flies modulate yaw turning to frontally fixate a moving stripe displayed on an LED array. Both behaviors occur naturally (without training) and robustly. But most importantly, these behaviors are representative of a class of animal reference-tracking behaviors regulated about an equilibrium (zero sensory slip).

In laboratory preparations, we can constrain reference trajectories and locomotor output each to a single degree of freedom, resulting in a single-input-single-output

(SISO) system. Both the sensory input (refuge velocity or stripe position) and locomotor output (fish velocity or fly yaw torque) for these behaviors are described by purely kinematic measures, unobtrusive to the behaving animal. Through an assay of carefully constructed sensory perturbations and a frequency-domain system identification analysis, we generate task-level dynamical models describing the behavior for a neighborhood about the fixation equilibrium. These descriptive and predictive task-level models are subsequently used to generate hypotheses about the neuromechanical systems (mechanistic models) underlying the behaviors.

## 1.2 Image-Stabilization Behaviors

The analyses presented in this dissertation can be extended to a broad category of behaviors we will refer to as *image-stabilization behaviors*, those in which animals perform motor actions to stabilize some sensory image about an equilibrium. The usage of “image” in this moniker is not limited to visual images but rather suggests the output of the sensory transform. So in this sense, the image can refer to the output from a single sensory modality (*e.g.* an optical flow estimate extracted from the image projected onto the retina) or a signal resulting from the convergence of several sensory modalities (*e.g.* for human posture control, a representation of “uprightness” may rely on visual, vestibular and proprioceptive cues [Oie et al.(2002), Jeka et al.(2004)]).

But by no coincidence, image-stabilization makes allusion to the human oculomotor system. We observe a lot of kinematic similarity between the locomotor behavior

## CHAPTER 1. INTRODUCTION

of knifefish and flies and human eye movements in target pursuit and fixation. In the later chapters, we borrow vocabulary from the study of the oculomotor system to describe observed locomotor phenomena.

The term “smooth pursuit” typically refers to visual target tracking behaviors in the oculomotor system of foveate animals, primates in particular [Fuchs(1967), Lisberger et al.(1987)]. In tracking visual targets, eye motions serve to stabilize the target on the fovea, the area of the retina most densely populated with photoreceptor cells. Visual tracking involves the cooperation of two distinct categories of eye movements, *smooth pursuit eye movements* and catch-up *saccades*, the distinction between which is typically made in kinematic terms: smooth pursuit eye movements are composed of continuous eye trajectories with low limits on velocity and acceleration, while saccades are ballistic, short-duration motions thought to correct for discontinuous positional errors which might accumulate during smooth pursuit [Rashbass(1961), Becker and Fuchs(1969), de Brouwer et al.(2001)]. Both fruit flies and fish exhibit smooth and ballistic locomotor outputs, though for the purpose of reference tracking, we focus on smooth pursuit.

### 1.3 Contrasting Approaches to Ethology

Ethology, the study of behavior, has been propelled by contributions from the fields of biology, neuroscience, psychology, applied mathematics, and engineering (to name a few). Yet, the analyses used in this research are not typical to the fields of biology or



## CHAPTER 1. INTRODUCTION

neuroscience and for the most part, the proposed models avoid explanation in terms of biological mechanisms. The contribution of this research is best explained in the context of two opposing scientific philosophies, *reductionism* versus *complex system science*, and corresponding methodologies, bottom-up and top-down respectively.

Reductionism comes in many flavors, some more established and others still philosophically controversial. But a commonly accepted notion from reductionist theory is that for a given system, any process arises as a consequence of processes occurring in the constituent components of the system; this notion is a weak interpretation of *ontological reductionism* [Brigandt and Love(2008)] and is rather widely accepted. New properties at one hierarchical level are *epiphenomena* of the constituent systems; the whole can be described as the sum of the parts. *Methodological reductionism* extends this notion, prescribing that systems *should* be studied as the composition of subsystems at “the lowest possible level” [Brigandt and Love(2008)]. In some cases which I will illustrate, methodological reduction yields misleading or erroneous observations. Additionally, from a pragmatic standpoint, as systems increase in the number of components and the complexity of interaction, reductionism becomes less constrained (due to combinatorics) and less informative. We cannot measure every component and connection within a complex system and the problem of extrapolating our understanding of a few parts into a full system description is underconstrained.

In contrast to reductionism, the *complex system* view posits that at each level of added complexity, *emergent phenomena* require a new set of general laws. For

## CHAPTER 1. INTRODUCTION

designing a research plan, the complex systems approach suggests that experiments should aim to study mechanisms in the context of behavior, leaving the animal and the behavior as intact as is feasible being that emergent properties which exist at a high level may not be deduced from knowledge of the parts. Free behavior, however, is rarely feasible for neurophysiological experiments. So methodological reduction, despite the noted drawbacks, is necessary in constructing mechanistic models for behavior.

The approach in this research primarily adopts the complex system approach, utilizing a top-down analysis of data collected from (mostly) freely behaving animals to generate holistic models of behavior. These empirically derived task-level models provide constraints for proposed mechanistic models (hypotheses). In the following sections, we further explain the application of reductionist methodology and subsequently motivate the complementary role of top-down methods.

### **1.3.1 Bottom-up: a place for reductionism**

Towards understanding animal behavior, we ultimately desire mechanistic models. These models depict behaviors as built out of subsystems (just as we think of a machine as an assemblage of parts). In the identification of biological behaviors, the decomposition into parts is of particular interest because many neuromuscular mechanisms contribute to multiple behaviors. Moreover, many of these mechanisms are preserved in part across different taxa. So identification of a subsystem may

## CHAPTER 1. INTRODUCTION

inform our understanding of many behaviors in different animals. The reductionist philosophy implies a constructionist approach towards generating mechanistic models: a system is studied at the lowest feasible level, acquiring a vocabulary of elementary parts and connections, and subsequently, a mechanistic model is composed *bottom-up* from the vocabulary to give rise to a behavior. The goal towards mechanistic models is driven by a reductionist ideology, but there are many challenges and pitfalls (which will be discussed shortly) involved in empirically identifying the subsystems and constraining the designed models.

Reductionist methodologies arise partially out of necessity. The complex systems approach prescribes that mechanisms be studied in the context of behavior, but there are limited preparations which allow for recording neural signals from awake and behaving animals; several of the successes in this area will be acknowledged in the following subsection. For the most part, neural mechanisms are identified in highly constrained or *ex vivo* preparations due to physical limitations imposed by the compatibility of the instrumentation and the behavior.

The bottom-up approach is a useful starting point for generating mechanistic models. Particularly for behaviors for which no models exist (or for which models exist for analogous behaviors), the bottom-up approach can rapidly furnish competing hypothetical models which can be supported or refuted empirically.

Bottom-up connotes design, the synthesis of a new system from components. The reductionist approach is the cornerstone of bio-inspired robotics. In his book *Vehicles:*

## CHAPTER 1. INTRODUCTION

*Experiments in Synthetic Psychology*, Valentino Braitenberg neatly demonstrates a reductionist approach to bio-inspired robotics [Braitenberg(1986)]. Using a limited library of sensors mounted on small wheeled vehicles, Braitenberg illustrates how different wiring schematics can generate behaviors which convey animal-like intention and emotions: aversion, pursuit, aggression, cooperation, and love. These behaviors are designed to capture some qualities of animal behaviors. Many areas of robotic study have been spawned by biological principles: neural networks, genetic algorithms, haptics, etc. But in implementation, these bio-inspired mechanisms diverge from their biological inspirations. And rightfully so. In the analysis of animal behavior, we are constrained by what is actually implemented in animals. In the synthesis counterpart, we are free to optimize the assembly of parts to produce a system which best fits our needs. The bottom-up synthesis is not constrained to corroborate any top-down empirical analysis.

While useful in design where the goal is to maximize performance with respect to some task, the over-extension of optimization in the analysis of animal behavior can be dangerous. The epilogue to this chapter quotes Voltaire:

If God has made us in His image, we have returned Him the favor.

There is a tendency, in the attempt to understand the machinery behind Nature, to approach the mystery as a synthesis problem. *How would I go about designing a system to perform this behavior?* This notion, that scientists should impose “good design” in studying animal behavior is flawed. An animal produces myriad behaviors

## CHAPTER 1. INTRODUCTION

using a limited set of mechanisms; so any particular mechanism may contribute to numerous behaviors and physiological functions. It is rash to assume optimality of a mechanism with respect to a singular task in spite of the diversity of robust behaviors to which the mechanism contributes; suboptimality might be a safer presumption [Fortune and Cowan(2010)]. Moreover, the optimization in animals is constrained by the process of evolution via natural selection. That means that some traits may improve fitness (either for survival or attraction of a mate, but always with the end goal of reproduction) and others may merely be features genetically linked to these preferable traits. For these reasons, it is important to observe and quantify behaviors and not to presuppose optimality in design. However, when similar mechanisms are observed to have evolved independently in dissimilar taxa (convergent evolution), we may search for the overarching factors and constraints which yielded such narrowly similar solutions [Fortune and Cowan(2010)].

### **1.3.2 Top-down: why complex systems should be studied**

#### **intact**

The notion that new properties emerge as the consequence of the interactions between subsystems is central to control theory. In control design, systems are designed as blocks (subsystems) connected through a variety of topologies (feedback loops, serial cascades or parallel summations) to achieve specified behaviors (*e.g.* reference tracking, noise rejection, model identification, etc.) or novel properties (*e.g.* stability,

## CHAPTER 1. INTRODUCTION

robustness, sensitivity, etc.). When a system is designed according to rules, the blocks and interconnections have been selected to achieve specific behavioral characteristics, there is hardly the sense of “emergence” when these characteristics surface. There is little mystery in design.

So from a design perspective, it is no surprise that the whole may exhibit properties not found in its parts. But in the identification problem, the converse also becomes problematic; a constituent subsystem may behave differently removed from the context of the rest of the system or behavior. As a simple illustration, consider the simple op-amp circuit in Figure 1.1.

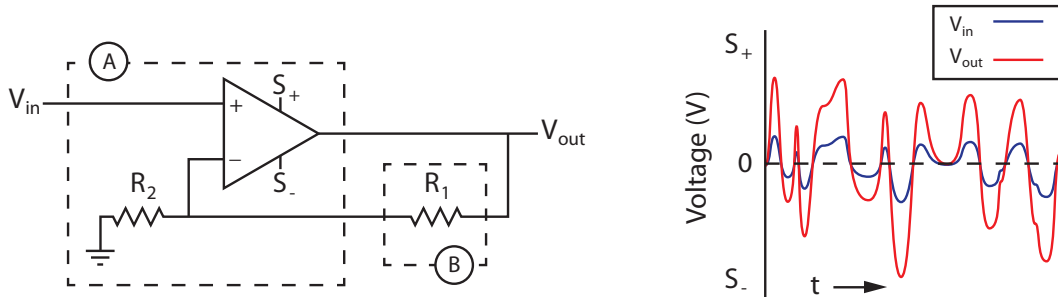


Figure 1.1: A non-inverting op-amp circuit with a closed-loop gain of  $1 + \frac{R_1}{R_2}$ .

Observations of input–output data readily reveals a task-level model; this simple circuit scales the input signal (by a gain of  $1 + \frac{R_1}{R_2}$ ). Now, feigning ignorance of the circuit design, we attempt to identify the contribution of the subsystems using a reductionist approach, decomposing the circuit into subsystems  $A$  from  $B$ . If we replace  $B$  with a short circuit,  $A$  does nothing at all, passing the input to the output unaltered (provided the input doesn’t exceed the supply voltages); if instead we substitute

CHAPTER 1. INTRODUCTION

an open circuit,  $A$  acts as a comparator, saturating to  $S_+$  when  $V_{in}$  is positive and to  $S_-$  when  $V_{in}$  is negative;  $B$  relays the input to the output unaffected (Figure 1.2).

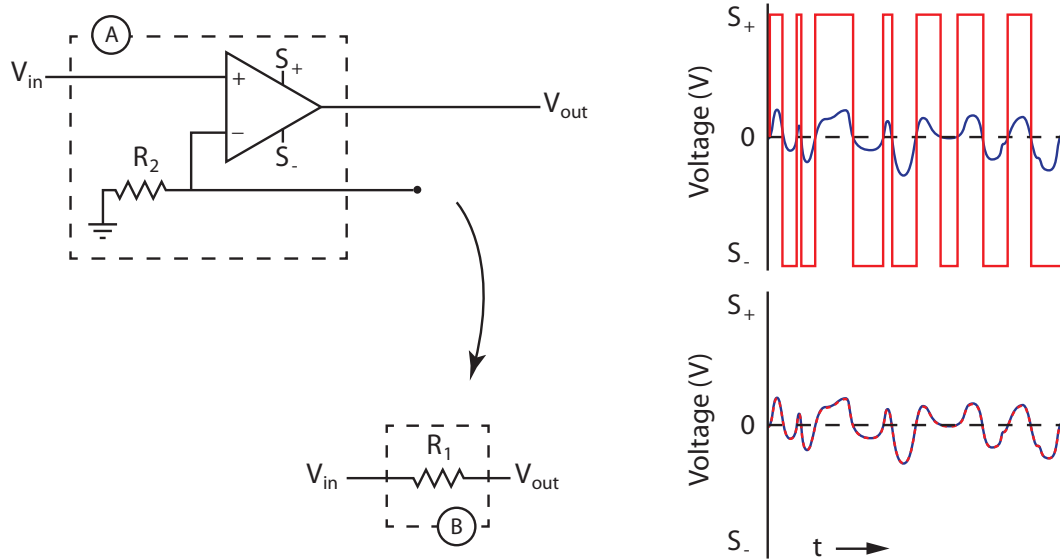


Figure 1.2: The responses of the constituent subsystems from Figure 1.1.

We may interpret this “preparation” in the context of two common neurophysiological methodologies. In the first interpretation, we consider observations of  $A$  in open loop as analogous to *ex vivo* preparations. In these preparations, a subsystem (a physical neuroanatomical structure) of interest is observed excised from the animal and typically stimulated electrically or pharmacologically. For the op-amp example in open loop Figure 1.2, we may interpret the binary output of the op-amp as a polarity indicator. If the combined circuit were believed to participate in motion processing, we might suspect that  $A$  is a direction selective circuit. Under the second interpretation, we posit that  $B$  is a physical structure ablated from the system and we infer its contribution based on any observed deficits in the resultant behavior. Depending on

## CHAPTER 1. INTRODUCTION

how  $B$  is ablated (replaced with an open circuit) we observe the described binary output; from these observations we might conclude that  $B$  contributes the proportional response.

In this analysis, the contribution of the op-amp is severely misrepresented. With a small enough input, it is possible to observe a proportional response, but due to the high gain and the saturation non-linearity imposed by the finite supply voltages, an open-loop analysis mischaracterizes the behavior of this component. And why would we ever suspect that it behave so differently in the context of the full circuit? Compared to the intricacy of animal behavioral systems this circuit is remarkably simple, yet it illustrates one peril of decontextualization as part of the reductionist approach. For image-stabilization behaviors, closed-loop stability will play an integral role in motivating our use of linear system identification analyses.

In studying mechanistic components in the context of behavior, we must be aware that the system is not merely the animal, but an interaction between the animal and the environment. Many behavioral paradigms maintain the animal system as a whole but limit interaction with the environment, performing experiments on awake but highly constrained animals (*e.g.* the direction selectivity studies on *Eigenmannia* discussed in Chapter 3). And there are neurobiological phenomena which simply can not be observed without this interaction (*i.e.* the behavior cannot be elicited with fictive stimulation).

For example, the discoveries of place cells [O'Keefe and Dostrovsky(1971)] and



## CHAPTER 1. INTRODUCTION

more recently grid cells and border cells [Hafting et al.(2005), Solstad et al.(2008)] exemplify the importance of the animal-environment interaction in the role of behavior. As their names suggest, place cells (hippocampus), grid cells (entorhinal cortex), and border cells (entorhinal cortex) are neurons which encode an animal's spatial orientation within an environment. In the experiments which lead to these discoveries, rats were surgically fitted with cortical implants comprising an array of electrodes. This device allowed the rat to freely explore an environment while neural signals are measured and transmitted over a flexible tether. The roles of these specialized cells could not have been observed in any constrained or *ex vivo* preparation.

In the past several decades, numerous technologies have been developed to measure neural activities in behaving animals: electromyography (EMG, both surface and implanted), electrocorticography (ECoG), electroencephalography (EEG), and functional magnetic resonance imaging (fMRI). In the scope of insect electrophysiology, tethered electrodes have been used to study locomotion in freely behaving cockroaches [Ye et al.(1995)]. More recent developments facilitate neurophysiological recordings in flying insects. Maimon *et al.* developed a method for whole-cell patch-clamp recording for in tethered *Drosophila* [Maimon et al.(2010)]. Wireless neural recording (and stimulation) has been demonstrated in hawkmoths [Tsang et al.(2011)] and locusts [Harrison et al.(2010)].

These technologies provide increasingly unobtrusive means for recording neural signals in behaving animals, measuring the outputs of neuromuscular mechanisms in

## CHAPTER 1. INTRODUCTION

the context of behavior. For the scope of this research, however, we develop input–output task-level models based strictly on kinematic data.

### 1.3.3 Reconciling the top-down and bottom-up approaches

Neither reductionist nor systems-theoretic approaches alone can generate a holistic *and* mechanistic understanding of a behavior. Bottom-up and top-down approaches must complement each other in order to achieve an understanding of behavior at different scales of organization. In this dissertation, we focus primarily on *descriptive and predictive* (phenomenological) models in contrast to *mechanistic* models. A descriptive model generalizes *what* is happening; a mechanistic model explains—at a lower-level—*how* the observed phenomenon arises.

Often a single model may be both descriptive (of a lower-level phenomena) and mechanistic (for a higher-level phenomena) (Figure 1.3). This is analogous to the templates and anchor spectrum [Full and Koditschek(1999)] for describing animal locomotion. At the high level, templates are low-dimensional descriptive models which capture salient features of the locomotor dynamics. At the low level, anchors are high-dimensional models which represent the physical system with greater fidelity. In this work, our descriptive models typically take the form of an empirical frequency response function or transfer function. In the template and anchors paradigm, templates are often described as a simple mechanical system. For example, the spring-loaded inverted pendulum (SLIP) model [Blickhan(1989),Holmes et al.(2006)] is used

CHAPTER 1. INTRODUCTION




Figure 1.3: A cartoon decomposition of the sensorimotor transform for refuge-tracking in weakly electric knifefish. The task-level model (A) is decomposed into candidate mechanistic decompositions with increasing refinement, (B) and (C).

to describe the sagittal plane dynamics of walking, primarily the trade-off of kinetic and potential energy. The SLIP model comprises a point mass (a surrogate for the body, co-located with the center of mass) and a spring loaded leg (which can change length during swing phase to inject energy into the system). Though this description sounds mechanistic, the template is an analogy: the SLIP model does not imply that a walking body manifests single-mass and a spring (in a lumped-parameter sense),

## CHAPTER 1. INTRODUCTION

but only that some aspects of the center-of-mass dynamics behave *as if* the system were an spring-loaded inverted pendulum. Similarly, any second-order differential equation which serves as a descriptive model could be framed in physical terms as a spring-mass-damper system or an resistive-inductive-capacitive (RLC) circuit, yet it would not imply that the physical instantiation was such.

Towards developing mechanistic (reductionist) models, Koditschek and Full emphasize the need for high-level descriptive models [Full and Koditschek(1999)]:

In the absence of a principled understanding of a simple model, no firm grasp of any more detailed model is likely. We will do well to view with due curiosity, but deep suspicion, the sort of complex and high-dimensional computational models that computing power allows. . . . We argue that higher-dimensional anchors are indeed required to reveal how mechanisms work, but can best advance our understanding of neuromechanical integration when informed by an underlying template.

To reconcile the systems-theoretic and reductionist approaches, both top-down and bottom-up approaches are necessary: refined descriptive models constrain the formulation of mechanistic models, while mechanistic models furnish predictions that require increasingly refined descriptive models, in a converging dialogue. This work presents a data-driven control theoretic approach towards identifying task-level *descriptive* models for animal behaviors. For both the refuge-tracking behavior in knife-fish and the stripe-fixation behavior in fruit flies, we empirically characterize the frequency response over a range of frequencies salient to the behaviors. The models generated from this study are then used to posit mechanistic decompositions at a lower level.

## 1.4 Dissertation Organization

This dissertation focuses primarily on task-level system identification of image-stabilization behaviors in animals. Chapter 2 provides the mathematical background for system identification, motivating the choice of a frequency domain approach and defining the tools used to assess behavioral performance and generate model descriptions. Chapter 3 and Chapter 4 describe the research contributions towards identifying the stabilization behaviors in knifefish and fruit flies, respectively.

The organization of these chapters reflects the organization of this introduction. Each chapter is introduced with a description of the animal system and the model behavior followed by a literature review of relevant research, roughly categorized as either research at the mechanism or behavior level. For *Eigenmannia*, we present a frequency-domain analysis (top-down) of the refuge-tracking behavior. In the following chapter, we first present the system identification of the stripe-fixation behavior in *Drosophila*. Subsequently, we extend the derived model to explain previous observations in other groups, consolidating disparate phenomena as the consequence of the task-level model. In Chapter 5, we discuss extensions of the task-level identification to include neurophysiological recordings, outlining a multi-input multi-output approach towards model decomposition.

## 1.5 Contribution and Dissemination

In this dissertation, we present a frequency-domain system identification framework for a broad category of animal behaviors, image-stabilization tasks. The framework details the experiments used to elicit smooth pursuit stabilization and the analytical tools developed for recognizing behavior, differentiating stimulus-mediated responses from other behaviors, and fitting parametrized models to frequency data.

We demonstrate this approach for two animal behaviors, refuge-tracking in the weakly electric knifefish *Eigenmannia* and optomotor stripe-fixation in the fruit fly *Drosophila*. It has been observed repeatedly and reaffirmed in this work that these behaviors are low-pass; performance in reference-tracking tasks is strong for slow oscillations and diminishes rapidly for higher frequencies. For this reason, the overwhelming majority of the literature focuses on the low-frequency response, where behavior is robust and easily discernible during experiments. In respect to the frequency band for which these models describe the behaviors, the experiments and resultant models represent the most comprehensive descriptions of these behaviors.

In both animal models, we reveal features of the behavior that were previously unknown or unreported. In knifefish refuge-tracking, the system identification analysis reveals a nonlinearity the nature of which suggests that the behavior utilizes mechanisms for prediction and optimal control. The modeling approach described in this dissertation has since been extended at decoding the sensory fusion problem in

## CHAPTER 1. INTRODUCTION

the same animal. In fruit flies, the empirically derived model (formulated as a PID controller) indicates that position-dependent cues in the visual scene contribute to the motor control policy; prevailing models rely primarily (or exclusively) on optic flow to determine torque. The proposed PID model generalizes, consolidating several previously observed phenomena as consequences of the same underlying model.

The material in Chapters 3 and 4 has in large part been published in the following:

- E. Roth, K. Zhuang, S. A. Stamper, E. S. Fortune, and N. J. Cowan. Stimulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmannia virescens*. *J. Exp. Biol.*, 214(7):1170–1180, Apr 2011.
- E. Roth, M. B. Reiser, M. H. Dickinson, and N. J. Cowan. A task-level model for optomotor yaw regulation in *Drosophila melanogaster*: a frequency-domain approach. In *Proc. Conf. on Decision and Control*, Maui, HI, December 2012. (submitted)

The following publications by the author are relevant to but do not prominently appear in this dissertation:

- S. G. Carver, E. Roth, N. J. Cowan, and E. S. Fortune. Synaptic plasticity can produce and enhance direction selectivity. *PLoS Comp. Biol.*, 4(2), 2008.
- E. Roth, S. Carver, E. S. Fortune, and N. J. Cowan. Mechanisms for encoding velocity in the electrosensory system of weakly electric fish. In *Proc. Int. Symp. Adaptive Motion of Animals and Machines*, Cleveland, OH, June 2008.
- S. Stamper, E. Roth, N. Cowan, and E. Fortune. Active sensing via movement shapes spatiotemporal patterns of sensory feedback. *J. Exp. Biol.*, 215(9):1567–1574, 2012.

## Chapter 2

# Frequency-domain System

## Identification

Look and you will find it—what is  
unsought will go undetected.

---

Sophocles



## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

Within the broad scope of control theory research, system identification focuses on extracting descriptive and predictive dynamical models from empirical data. Though the goal of system identification is narrow, the field encompasses a variety of applications and approaches. For the purpose of controller design, system identification can be used to generate a plant model when one is (partially) unknown. In fact, many developments in system identification have come from the adaptive control community (*e.g.* model identification adaptive control in which a plant model is generated online). And though system identification techniques have been extended to nonlinear and hybrid dynamical systems, the canon of literature in linear modeling is by far the most developed [Ljung(1999), Söderström and Stoica(1988), Pintelon and Schoukens(2001)]. In this chapter, we explore applications of linear system identification to the study of image-stabilization behaviors in animals.

### 2.1 Motivating the Frequency-domain Approach

System identification methods can be broadly separated into two categories: parametric, which includes prediction error methods (PEM), and non-parametric, which includes time-domain correlation methods and frequency analyses. Parametric methods assume a model structure *a priori* and optimize the model representation with respect to the parameters, abstracting the data into a model of substantially lower dimension. Typically, parametric models represent the system as a closed-form mathematical expression (*e.g.* difference or differential equations). These models can furnish

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

hypotheses about the underlying system and generate predictions for novel stimuli (even extrapolating predictions to stimuli outside the support of the initial experimental data). However, if the assumed model structure does not accurately reflect the underlying system, informative features may be lost in the representation. Non-parametric models (*e.g.* lookup tables, graphs, etc) do not presuppose a structure and hence more faithfully retain nuances in the data. However, these representations do not afford the same generalizability or reduction of dimensionality as do parametric models.

The inherent diversity in animals (in many senses) poses a unique challenge for system identification techniques. The goal is to estimate a mean behavioral model from a population. In system identification, we commonly assume that each data set used to infer the underlying model is actually generated by the same underlying system. This assumption is not necessarily valid across a sample set of animals (or across trials for a single individual). So for animal systems, the goal is to identify the behavior within the data and estimate a *mean* behavioral model from the population. We approach the system identification problem with a non-parametric frequency analysis, generating empirical frequency response functions (or Bode plots) as an input–output description of the behavior.

### 2.1.1 Why a non-parametric approach?

Typically, a parametric system identification method consists of three stages: hypothesizing model structure, optimizing model parameters, and selecting from competing hypotheses. In the first stage, we enumerate a number of candidate model structures. Often it's the case that some prior knowledge of the system informs the candidate structures; the hypothetical model structures may be known from first principles (white-box models) or informed by some known physical properties (grey-box). Without such knowledge, we must cast a wide net, canvassing a large assortment of (black-box) candidate structures in order to safely assume that some member of the candidate set closely emulates the true system. Subsequently, we optimize the parameters for each model structure according to some error metric (*e.g.* prediction error). Finally, we select a model according to an information criteria (*e.g.* Akaike or Bayesian information criteria), balancing the prediction error and model order.

Often in collecting behavioral data from animals, it is not clear whether a single model represents each and every data set. Individuals behave differently. Animals fatigue or lose motivation over the course of an experiment regiment. These differences may not only yield a greater variance (uncertainty) in model parameters, but different individuals (or trials) may indicate different model structures. Suppose for example, that the underlying system for a behavior is best fit by a transfer function of relative degree two. At high frequencies, the behavior is analogous to that of a spring-mass-

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

damper system. Now suppose across the sample population, variability between the animals furnishes variability in the natural frequency and damping coefficient in the underlying system. Those individuals in the population who are highly damped may exhibit behavior more suitably modeled by a first-order system (inertial terms are negligible); underdamped individuals may require second-order models to capture the dynamics of behavior. In a non-parametric representation, such as the frequency response functions used in this work, the scope of behaviors is more evident as a continuum of responses. In the parametric approach, this variance could result in a poor model fit or bimodality at the model selection stage (two competing models with similar information criteria).

To complicate matters further, animals may exhibit multiple behaviors over the course of an experiment, and at times, they may exhibit multiple behaviors concurrently. For example, when performing refuge tracking in the dark, *Eigenmannia* perform rapid saccades superimposed over the smooth pursuit tracking behavior [Stamper et al.(2012)]. These high velocity saccades introduce significant tracking error to the behavior. Though the smooth pursuit behavior persists, a parametric model may favor fitting these saccades in lieu of the low-frequency, low-velocity, smooth pursuit dynamics. In this instance, the oscillations which emerge in the dark can be easily observed and filtered prior to a parametric analysis. When these superimposed dynamics are not obvious, the parametric approach may result in a model of the ancillary behaviors.

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

In short, beyond the initial stage in which we select candidate structures and error metrics, parametric approaches leave little room to inject intuition or decision-making into the identification process. In our non-parametric frequency-domain approach, each trial is represented as an empirical frequency response function. This allows us to see the diversity of the responses, cull outliers, and focus on fitting parametrizations that capture the most salient features of the behavior while ignoring any secondary phenomena.

### 2.1.2 Why the frequency domain?

Many image-stabilization behaviors, including the two studied in this dissertation, exhibit low-pass dynamics. These behaviors are typically studied for the regime of stimuli for which the behavior is observed to be robust. In the time domain, attenuated gain in response to a motion stimulus can be masked by distracting volitional movement, causing the persistence of behavior at high frequencies to be overlooked. And even when these reduced responses are observed, they are often dismissed as anemic or unmotivated. For the purpose of system identification, the more vigorous low-frequency responses are less revealing of the underlying dynamics. Excepting differences in gain, for low frequencies (relative to the cut-off frequency), the majority of linear low-pass systems exhibit qualitatively similar responses. To differentiate between competing models—to determine attributes such as model order, cut-off frequency, damping, and delay—requires observation beyond the cut-off frequency,

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

where performance wanes. The frequency domain representation exposes these meager yet informative high-frequency responses.

To further advocate this point, consider the sensitivity of the open-loop fly optomotor response  $F(\omega)$  (the transfer function from perceived sensory error to fly motion) with respect to tracking performance. The tracking behavior can be described by the closed-loop transfer function from reference to output,  $G(\omega)$ , where error is filtered by  $F(\omega)$  and stabilizes the stimulus trajectory through negative feedback:

$$G = \frac{F(\omega)}{1 + F(\omega)}, \quad F(\omega) = \frac{G(\omega)}{1 - G(\omega)},$$

with sensitivity of the optomotor plant with respect to performance calculated as

$$\frac{\partial F}{\partial G} = \frac{1}{(1 - G(\omega))^2}. \quad (2.1)$$

We see in (2.1), that the open-loop optomotor response is most sensitive to performance near  $G(\omega) = 1$  (perfect tracking performance) which occurs when  $F(\omega) \gg 1$ . The heightened sensitivity in this regime has two important implications: small measurement errors will yield highly variable open-loop model predictions, and, conversely, a diversity of open-loop models could have generated the observed closed-loop performance. For this common feedback topology, sensitivity improves to unity as performance degrades. This trade-off must be a critical consideration for system identification, to find the regime of stimuli which elicit observable behavior and

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

informative model constraints, performance and sensitivity near unity.

For both stripe fixation in *Drosophila* and refuge-tracking in *Eigenmannia*, we identify dynamical models for the behaviors through an assay of perturbation experiments. The reference trajectories used are all variations on sinusoidal functions representing a broad range (with respect to the bandwidth of the behavior) of frequencies: individual sinusoids, sums of sinusoids at different frequencies, and chirps (sinusoids with time-varying frequency). The choice of this class of stimuli is motivated by behavioral, analytical and historical factors, as described below.

### **Smooth trajectories elicit smooth pursuit**

Animals interact with their environments through a repertoire of locomotor behaviors and it is important to design stimuli which evoke the desired behavior. It is common practice in system identification experiments to use white noise or pseudo-random noise sequences to generate spectrally inclusive stimuli. In our experience, such signals do not elicit smooth pursuit; in *Eigenmannia*, any abrupt refuge motion induces an escape response. Smooth pursuit requires smooth trajectories. This class of sinusoidal stimuli provides a flexible basis for generating stimuli with desired spectral content and reliably elicits the smooth pursuit response.

### **Spectral composition differentiates behaviors**

Behaviors derived from different dynamical systems have distinct spectral signatures. For the case of image-stabilization behavior, we are most interested in isolating the smooth pursuit response from ancillary behaviors (*e.g.* saccadic corrections, exploratory behaviors, etc.). The stimuli trajectories used in this research share a common and important attribute: the domain is comprised of discrete frequencies (as in the case of sinusoids and sums-of-sines) or of compact intervals of frequency (the case for any time window of a chirp stimulus). The pursuit response will be strongly coherent with the stimulus (as we see in both fish and flies). Hence, output-input coherence serves as an indicator for whether the behavior is present and as a tool for disambiguating pursuit from auxiliary motion.

### **Sinusoidal stimuli are prevalent in the literature**

Though there is remarkably little prior art in this broad-spectrum frequency-domain approach towards behavioral identification [Kiemel et al.(2006)], sinusoids have long been employed as stimuli for reference tracking behaviors. A significant body of literature explores optomotor regulation in fruit flies using similar visual perturbation experiments [Heisenberg and Wolf(1988), Duistermars et al.(2007)]; the knife-fish tracking behavior, on the other hand, was discovered much more recently and has not yet been widely studied [Rose and Canfield(1993a), Cowan and Fortune(2007)]. So while there are few dynamical models against which to compare, past literature



provides results (discrete data points on the frequency response function) against which to compare such models. One critical aim of this research is to consolidate past observations as the consequences of a parsimonious model.

### 2.1.3 A role for linear models in describing image-stabilization behaviors

The frequency response analyses used in previous studies on image stabilization behaviors [Cowan and Fortune(2007), Reiser and Dickinson(2008), Heisenberg and Wolf(1988), Götz(1968), Kiemel et al.(2006), Carver et al.(2005), Jeka et al.(2004), Sprayberry and Daniel(2007), Gilbert(1997)] are predicated on an assumption of linearity. Without this linearity assumption, a frequency response function (FRF) generated from one set of stimuli would not predict the system's response to spectrally distinct stimuli. And, it would be impractical to test the entire range of possible stimuli for any system.

The linearity assumption underlies the predictive and generative power of frequency analyses. But why should we expect any animal behavior be described by such a seemingly restrictive set of models? Admittedly, nonlinearities manifest in many of the biological subsystems which give rise to behaviors, from low-level mechanisms (*e.g.* sensory tuning curves, saturation and hysteresis in muscle force production) to high-level neural processes (*e.g.* long time-scale adaptation, volitional changes between different behaviors). But, linearity at the task-level does not preclude nonlinear

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

constituent subsystems. In this class of closed-loop behaviors, the system is stabilized at a task level to an equilibrium state corresponding to the sensory goal. Local to an equilibrium, many nonlinear systems (and, in fact, *almost all* in a certain mathematical sense [Sastry(1999)]) can be closely approximated by (oftentimes low-order) linear models. Hence, cockroach wall-following, for example, could be faithfully captured by a linear model [Cowan et al.(2006), Lee et al.(2008)].

However, when linear models fail to adequately represent a behavior *i.e.* the behavior does not appear linear for any neighborhood of the equilibrium, the discrepancies in frequency responses to different stimuli can illuminate the underlying nonlinearities. In our analysis of the refuge tracking behavior of *Eigenmannia virescens*, we ascribe the differences in frequency response functions between stimulus types to a model-based prediction mechanism and optimal control. For the proposed model (Figure 3.12) and a fixed stimulus, the Kalman filter and optimal controller are linear; the nonlinearity observed in our experiments is introduced as the Kalman filter adapts to new stimuli, updating an internalized model of the system and external dynamics. The linear analyses we present provide snapshots of an adapting behavior—waypoints which constrain future nonlinear models for the full behavior. Future work can address the mechanisms responsible for these adaptations.

### 2.1.4 Extending frequency analyses to other image-stabilization tasks

Similar assays to those described in this thesis could be used towards identifying control strategies for other animal image-stabilization tasks. The approach outlined in this work is applied to task-level dynamics. For many biological systems, identifying the task-level goal and subsequently measuring a suitable task-level state is not trivial.

Locomotor dynamics often mask the task-level states of interest. For most animal behaviors of interest, the motor dynamics are cyclic (*e.g.* walking strides, flapping wings). The periodicity of locomotor dynamics may or may not manifest in the task-level states. For example, in the case of *Eigenmannia*, the individual undulations of the ribbon fin (which occur at a frequency of about 10 Hz) do not introduce significant variance into the task-level states (longitudinal position and velocity of the body). In contrast, for the control of walking or running in humans, the within stride phase significantly affects the task-level state (often the vertical position and velocity of the center of mass). Walking dynamics are often modeled as an inverted pendulum or some variant on the theme [Alexander(1995)] while running is often represented as a spring-mass hopping system [Blickhan(1989)]; both models clearly illustrate how the task-level state changes periodically, in synchrony with the gait. Similarly, cyclical motor dynamics can manifest in task-level states for flying—particularly in slow-flapping (wingbeat frequencies within the band salient to task-level behavior) animals

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

such as moths and butterflies, bats, and birds—and swimming modalities such as carangiform swimming in which thrust is generated by the caudal fin through body bending.

Many locomotor behaviors are described in terms of stable limit cycles, attracting periodic trajectories in the state space; at a task-level, the goal of an image-stabilization behavior is described as an equilibrium point. We have presented a small sampling of behaviors in which cyclic motor plants are controlled to achieve stationary sensory goals. But in the interest of identifying neural control policies, it is useful for some questions to divest the task-level states from the “artifacts” introduced by the motor plant or external dynamics.

For the cases above, systems theory of cyclic dynamics provides tools for stripping task-level states from the kinematics. Floquet analysis allows the task-level states to be recoordinated according to the phase of the cyclic dynamics, in essence transforming an equilibrium cycle (or limit cycle) into an equilibrium point. Once the kinematic data is transformed to align these Floquet coordinates, data captured from different phases of a “stride” can be compared using the techniques such as those described above [Revzen and Guckenheimer(2008), Revzen(2009)]. In a similar approach, cyclical systems can be discretized through Poincaré analysis. Rather than aligning a cyclically changing coordinate system as in Floquet analysis, Poincaré analysis considers the state of a system at only one phase of a cycle, generating a discrete data point for each cycle. In this way, the task-level states are captured at the same

phase of every stride, fixing the equilibrium point to the state of the limit cycle at that phase [Lee et al.(2008)].

## 2.2 Mathematical Tools

### 2.2.1 Coherence

Coherence, the ratio of the squared cross-spectral density of two signals  $v(t)$  and  $z(t)$  and the product of the respective power spectral densities,

$$C_{vz}(\omega) = \frac{|R_{vz}(\omega)|^2}{R_{vv}(\omega)R_{zz}(\omega)}, \quad (2.2)$$

describes the degree to which two signals are linearly related (correlated) at different frequencies. Unity coherence implies that two signals can be perfectly represented as the input and output of a linear dynamical system; lower coherence may result from the presence of nonlinearities, noisy measurements, or additional unaccounted inputs which contribute to the measured output. In this paper, we perform coherence analysis for sum-of-sines trials to establish that, for a given trial, the input–output relationship is linear.

Consider paired input–output measurements of a linear system. We assume process noise (*e.g.* due to variability of the motor output [Harris and Wolpert(1998)]) corrupts the motor behavior itself. Specifically, letting  $*$  denote the convolution op-

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

erator, suppose the input–output pair  $(u(t), y(t))$  is related by

$$y(t) = f(t) * u(t) + h(t) * m(t) \xrightarrow{\mathcal{F}} Y(\omega) = F(\omega)U(\omega) + H(\omega)M(\omega), \quad (2.3)$$

where the system  $f(t)$  filters the input  $u(t)$  and  $h(t)$  filters a process noise  $m \sim N(0, \nu^2)$ . Here,  $\mathcal{F}$  denotes the Fourier transform. Observations of this pair  $(v(t), z(t))$  are corrupted by measurement noise (which can be minimized to some extent through careful experimentation)  $n_{v,z} \sim N(0, \sigma_{v,z}^2)$ :

$$\begin{aligned} v(t) = u(t) + n_v(t) & \xrightarrow{\mathcal{F}} V(\omega) = U(\omega) + N_v(\omega) , \\ z(t) = y(t) + n_z(t) & \xrightarrow{\mathcal{F}} Z(\omega) = Y(\omega) + N_z(\omega) . \end{aligned} \quad (2.4)$$

In the absence of noise ( $\sigma_v, \sigma_z, \nu = 0$ ), a linear dynamical system yields input–output pairs with unity coherence. Since noise variances appear only as additive terms in the denominator of the coherence function (2.2) any noise introduced to the system or measurements diminish coherence, as shown:

$$C_{vz}(\omega) = \frac{|F(\omega)|^2 |U(\omega)|^4}{|F(\omega)|^2 |U(\omega)|^4 + |H(\omega)|^2 \nu^2 \sigma_v^2 + |H(\omega)|^2 |U(\omega)|^2 \nu^2 + |U(\omega)|^2 \sigma_z^2 + \sigma_v^2 \sigma_z^2} . \quad (2.5)$$

Even a linear system, which in a noiseless case should produce unity coherence, would fail to do so in the presence of noise. Deficiencies in coherence may indicate

either systematic nonlinearities or corruption by noise or both. Moreover, the output of a system may be coherent with the input for a particular choice of stimuli despite nonlinearities in system. Despite being neither a necessary nor sufficient condition, coherence is a useful indicator of linearity, given the above caveats.

### 2.2.2 Frequency response functions (FRFs)

The frequency response function (FRF) of a linear dynamical system is the Fourier transform of the impulse response represented as a complex function of frequency. Given an input–output time series, the FRF can be calculated as the output:input ratio of the Fourier transform of the respective signals. Alternatively, the FRF can be calculated as the ratio of the cross spectral density of the output to the input,  $R_{yu}(\omega)$ , and the power spectral density of the input,  $R_u(\omega)$ . A Bode plot is a graphical representation of the FRF in cylindrical coordinates. In a Bode plot, the FRF is described using both the gain (scaling describing the level of amplification or attenuation of the output with respect to the input) and relative phase imparted by a system (the lead or lag of the output with respect to the input).

For our purposes, we consider linear systems which generate similar Bode plots to be similar linear systems. Admittedly, this measure is subjective. What features of the frequency response are of interest? What are acceptable deviations from the data? Under this interpretation, different parameterizations of a single model structure and even models of different structures may be deemed similar. It is important

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

to note that the FRF is a representation of the input–output behavior, and thus, only portrays the behavior of the observable and controllable (specifically excited) modes. The limitations of observability and controllability are inescapable; this drawback is common to *all* system identification methods<sup>1</sup>. Additionally, FRFs and Bode plots are predicated on the assumptions of linearity; for linear systems, FRFs are generative descriptions in that we can predict the output of the system subject to an arbitrary input. Though we can generate empirical FRFs from input–output data generated by a nonlinear system, the resultant models are not generative. Hence, our notion of similarity does not extend to nonlinear (rather nonlinearizable) systems. The FRF and Bode plot are the primary tools we use to describe image-stabilization behaviors.

At times, we may use the terms “frequency response function” (a mathematical function) and “Bode plot” (a graphical representation of the FRF) interchangeably; a simple (nonlinear) transformation converts one to the other, so in most respects, the two are equivalent representations of a dynamical system. However, the difference between these representations has implications for calculating statistics of the response or goodness-of-fit for a proposed model. Recall that the logarithm of a complex number yields another complex number for which the real and imaginary parts

---

<sup>1</sup>In parametric methods, nonminimal realizations (or realizations with near pole-zero cancellations) are typically rejected by the model selection criteria, since they incur a cost for additional parameters and afford little or no improvement in terms of prediction. In fitting parametrized models to FRFs, we similarly decline nonminimal transfer functions.



## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

represent the log magnitude and phase respectively:

$$\underbrace{\ln(a + i \cdot b)}_{\text{cartesian}} = \underbrace{\ln(m \exp(i \cdot \phi))}_{\text{polar}} = \underbrace{\ln m + i \cdot \phi}_{\text{cylindrical}}, \quad (2.6)$$

where  $m = \sqrt{a^2 + b^2}$  and  $\phi = \arctan_2(b, a)$ . The the dual plots in the Bode representation (gain and phase) are proportional to the real and imaginary parts of the logarithm of the frequency response; magnitude can be scaled by  $\frac{20}{\ln 10}$  to convert absolute gain to decibels. A point in the FRF lives in a cartesian space of complex numbers ( $\mathbb{C} \sim \mathbb{R}^2$ ), in the polar representation lives in the space  $\mathbb{R}^+ \times \mathbb{S}^1$  (which has a singularity at  $m = 0$ ), while in the Bode representation (log-space) lives on the cylinder ( $\mathbb{R} \times \mathbb{S}^1$  with no representation for  $m = 0$ ). We'll consider only the cartesian (as used in Nyquist and pole-zero plots) and log-space representations (as used in Bode plots) since they are ubiquitous in systems and control theory. Throughout this work, we use both representations for our calculations and both have their merits and shortcomings. The selection of representation may be motivated by ease of calculation, properties of the data, convenience or convention.

### Error metrics

In the Cartesian FRF representation, the Euclidean norm is a convenient distance metric. Under this interpretation, the distance between two responses with gains of 1.1 and 1 (assuming equal phase) are equally far as two responses with gain 0.1 and 0,

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

though the difference between the first pair seems minimal while in the second pair, the distance between a little and no response is categorical; an error of fixed magnitude has greater behavioral consequence for responses with low gain. For fitting models, the uniform distance metric seems to preferentially weight closeness for responses with large gain. This shortcoming in the Cartesian representation suggests an inclination for logarithmic (normalized) error.

But the log-space representation does not furnish an immediate or agreeable distance metric. Supposing a similar distance metric, the  $\ell_2$  - norm with a weighting function such that  $\|c\|_W = \sqrt{c^T W c}$  where  $c = [m, \phi]^T$  and  $W$  is a diagonal matrix of positive weights. Again consider a response pair,  $\{\ln m + i \cdot -\pi/2, \ln m + i \cdot \pi/2\}$ ; the  $\|\cdot\|_w$  error is the same regardless of the magnitude  $m$ . But the error signal in the time domain depends significantly on the magnitude; when the magnitude is near zero ( $\ln m \ll 0$  in the log-space), the error signal in the time domain approaches zero despite the phase difference. At low gain, even small tracking errors (in the time-domain) can result in large response errors in the log-space. So while the log-space provides a scale for gain which more intuitively captures our qualitative sense of behavioral similarity, converse to the Cartesian representation, the log-space representation exaggerates errors at low gain.

**Probability distributions**

The Cartesian space accommodates Gaussian distributions and, for refuge-tracking behavior in *Eigenmannia*, Gaussian distributions seem to qualitatively capture the frequency response data (grouped by frequency) on the complex plane. The distribution is calculated in this space and then propagated forward to generate confidence intervals in the Bode plots (described in Section 3.2.4.1). In the log-space domain, there is no common distribution native to the cylinder; if we assume the two dimensions to be separable, we may fit Gaussian and von Mises (the circular analogue to the normal distribution) distributions for the gain and phase respectively.

Alternatively, we could unravel the quotient space  $\mathbb{S}^1$  into the interval  $[-\pi, \pi)$  and stitch together these intervals to form a representation of phase on the real line. In essence, this removes the  $2\pi$  periodicity of phase, allowing the log-space to be represented in  $\mathbb{R}^2$ . From observations of the frequency response at a single frequency, the phase can only be identified within the  $[-\pi, \pi)$ . However, when we construct the FRF densely over an interval of frequencies, we can infer phase lags (or leads) greater than  $\pi$  rad by the continuity of the response. For example, a delay of 1 sec imparts a phase lag of  $4\pi$  rad for a signal at 2 Hz, though by inspection of that frequency alone, we perceive no phase lag at all. Or similarly, at high frequencies, a low-pass system with relative order of three is described as having a phase lag of  $\frac{3\pi}{2}$  rad rather than a phase advance of  $\frac{\pi}{2}$  rad by maintaining the continuity of the phase roll-off. For the

## CHAPTER 2. FREQUENCY-DOMAIN SYSTEM IDENTIFICATION

Bode plot of the stripe-fixation behavior (generated using chirp stimuli), we use this method for calculating statistics, unwrapping phase for each trial.

## Chapter 3

### Refuge-Tracking in

### *Eigenmannia virescens*

Wouldn't the sentence 'I want to put a hyphen between the words Fish and And and And and Chips in my Fish-And-Chips sign' have been clearer if quotation marks had been placed before Fish, and between Fish and and, and and and And, and And and and, and and and And, and And and and, and and and Chips, as well as after Chips?

---

Martin Gardner

### 3.1 Why Weakly Electric Knifefish?

The refuge-tracking behavior in glass knifefish *Eigenmannia virescens* is an untrained and robust behavior in which the fish modulate swimming to maintain a relative position with objects moving in their environment. Native to Central and South America, these fresh-water fish reside in Amazonian tributaries where they may employ this behavior to hide among objects moving in the current, a means of eluding predators during the day. In the laboratory, fish track an artificial refuge (machined from a segment of PVC pipe) moved along prescribed trajectories.

Both the sensory and locomotor mechanisms involved in this behavior are atypical and noteworthy. In addition to vision and mechanoreception, these fish use active electrosensation to detect objects near the body. An electric organ (EO) in the tail emits a weak oscillating electric charge, the electric organ discharge (EOD). Voltage sensitive receptors in the skin, measuring small fluctuations caused by objects in the near electric field. Distributed over the entire body, the population of electroreceptors form a two-dimensional electrosensory image, not unlike photoreceptors in the retina.

Knifefish locomotion utilizes a long actuated ribbon fin along the bottom of the body which undulates to generate thrust. Unlike most other fish which rely on body bending and a sizable caudal fin to propel themselves forward, the knifefish ribbon fin mechanism does not require body bending. More importantly, ribbon fin locomotion allows fish to swim capably both forwards *and backwards* [Rose and Canfield(1993a)]



Figure 3.1: The weakly electric glass knifefish, *Eigenmannia virescens*, relies on vision and electroreception to track objects in its environment. An anal ribbon fin, extending most of the length of the body, allows the fish to swim forwards and backwards.

without changing the orientation of their body. By constraining the motion of the refuge to back-and-forth trajectories along a single axis, we leverage the motor capabilities of these fish to reduce the tracking behavior to a single spatial dimension, and in turn, simplifying the system identification problem to that of fitting a single-input single-output (SISO) system.

Though the tracking behavior has only recently been studied [Rose and Canfield(1993b), Rose and Canfield(1993a), Cowan and Fortune(2007), Roth et al.(2011), Stamper et al.(2012)], the sensory and locomotor mechanisms involved have been studied either in simulation, constrained experiments, or in the context of other behaviors.

### 3.1.1 Mechanisms

The electrosensory system in knifefish contributes to many behaviors, both sensorimotor and social. The electrosensory image and its role in object detection has been explored extensively through simulation [Rasnow et al.(1989), Rasnow(1996), Nelson et al.(2002), von der Emde(2006), Babineau et al.(2007), Sim and Kim(2012)]. Signals from the primary afferents (tuberous electroreceptors) converge at the electrosensory lateral line lobe (ELL) in the medulla. Cells in the ELL lobe receive information from numerous receptors to form spatially localized receptive fields (RFs) of different sizes distributed over the body; the ELL is divided into three parallel segments each a topographic map of the receptors on the body (centromedial, centrolateral, and lateral) and roughly corresponding to RFs of different sizes (smallest to largest respectively) and temporal responses (slowest to fastest respectively) [Shumway(1989)]. The spatio-temporal dynamics of these receptive fields have been investigated in constrained and curarized animals [Shumway(1989), Berman and Maler(1999), Bastian et al.(2002), Chacron et al.(2003)]. Towards understanding how motion in the environment is encoded, direction selective cells in the torus semicircularis (downstream of the ELL, in the midbrain) have been identified, also in curarized animals [Chacron et al.(2009), Chacron and Fortune(2010)]. Integrating empirical descriptions of the spatio-dynamic receptive fields, we have proposed mechanistic models (not discussed in this dissertation) describing how direction selectivity and velocity encoding may



be realized in the nervous system [Carver et al.(2008), Roth et al.(2008)].

Only recently have the mechanics of ribbon fin swimming have become the focus of research. Force production by the ribbon fin locomotor modality has been studied in physical simulation using bio-inspired robotic ribbon fins [Epstein et al.(2006), Curet et al.(2011)], relating the production of thrust and heave to the kinematic parameters of a single undulating ribbon-fin wave passing through water. However, in the tracking behavior, we notice that fish exhibit a different kinematic pattern; knifefish generate inwardly counter-propagating waves, that is, the rostral and caudal sections of the fin generate antagonistic forward and backwards thrust, respectively. Knifefish modulate net thrust by adjusting the proportion of the ribbon fin recruited for forward and backwards swimming. This locomotor strategy puts forwards and backwards swimming on a continuum, reducing the control strategy largely to a single parameter (the ratio of fin recruited in each direction) and increases maneuverability and stability about the hovering equilibrium (when the counter-propagating waves cancel to zero net thrust) [Sefati et al.(2012)].

### 3.1.2 Behaviors

The neural mechanisms of electroreception are best understood in the context of a social behavior, the jamming avoidance response (JAR). In the JAR, fish modulate their EOD frequencies to avoid low frequency amplitude modulations (AM) caused by the interaction of their EOD with those of a conspecific [Watanabe and

### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

Takeda(1963), Bullock et al.(1972)]. For a grouping of two animals, the AM fluctuates at a frequency equal to the difference of the two individual EOD frequencies, the beat frequency. Low frequency beats are thought to be detrimental to motion detection. This behavior represents an unstable equilibrium; when two fish have similar EOD frequencies, each will adjust their frequency to increase the gap, hence raising the AM beat frequency. Through electrophysiological experiments using fictive stimuli to emulate the EOD of a conspecific, Heiligenberg and Bastian have mapped the complete neural pathway for the JAR behavior, identified the neurons responsible for amplitude and phase encoding in the torus semicircularis, and proposed and validated an algorithm describing how a fish might estimate differential frequency (positive or negative) with respect to the nearby conspecific [Heiligenberg and Bastian(1980), Bastian and Heiligenberg(1980), Heiligenberg(1991)].

The refuge-tracking behavior examined in this research is one of several such behaviors observed in knifefish. In an earlier observed behavior, the electromotor response, fish swim side to side to adjust their lateral distance to a rod moved perpendicularly to the side of the body [Heiligenberg(1973), Bastian(1987)]. Similar kinematic behaviors have been studied in freely swimming fish performing prey capture [Maciver et al.(2001)]. The longitudinal refuge-tracking behavior was first described by Rose and Canfield [Rose and Canfield(1993a)] and later modeled by Cowan and Fortune [Cowan and Fortune(2007)]. The model, derived from observations of the tracking behavior in response to sinusoids of various frequencies, described the

task-level dynamics as a second order model. Cowan and Fortune demonstrate how this task-level model imposes mutual constraints between the mechanical plant and the neuromechanical controller; if the mechanics of swimming are assumed to be inertial (second order) the task-level model prescribes a high-pass neural policy, while an assumption of over-damped swimming dynamics requires a low-pass control policy in order to satisfy the task-level dynamics. In this work, we extend these preliminary experiments, observing tracking responses to a broader range of frequencies and stimulus categories.

### 3.1.3 Our approach

Initially, this work was intended as a validation of the linearity assumption for refuge-tracking (Cowan and Fortune constructed their model based on the response to single sinusoids only). However, in comparing the responses between different categories of stimuli (single and sum-of-sines) we observed a categorical nonlinearity, improved performance in tracking sinusoids [Roth et al.(2011)]. In the following sections, we investigate the nature of this nonlinearity and postulate the underlying mechanisms which might generate the discrepancy between the single-sine and sum-of-sine responses. The methodology developed for these studies has been extended to research aimed at disentangling the contributions of the visual and electrosensory modalities and the behavioral consequences when one of these senses is inhibited or degraded [Stamper et al.(2012)].

## 3.2 Materials and Methods

### 3.2.1 Animal husbandry and preparation

Adult knifefish of the species *Eigenmannia virescens* were obtained through commercial vendors and housed in community tanks. Animal husbandry followed published guidelines for the care and use of Gymnotiform fishes [Hitschfeld et al.(2009)]. For both community and experiment tanks, water was maintained at a temperature of approximately 27 °C and a conductivity in the range of 150–250  $\mu$ S. An individual fish would be placed in the experiment tank and given adequate time (2 h–1 d) to acclimate to the environment and enter the refuge. All experimental procedures with animals were approved by the animal care and use committee at the Johns Hopkins University, and are in compliance with guidelines established by the National Research Council and the Society for Neuroscience.

### 3.2.2 Experimental apparatus

The refuge was machined from a 15 cm segment of 2 inch stock PVC pipe; the bottom of the pipe was milled away to allow the fish to be video recorded from below and a series of windows 0.625 cm in width and equally spaced at 2.5 cm intervals were machined into the side of the pipe to provide visual and electrosensory cues. The refuge was positioned to be less than 0.5 cm from the bottom of the tank. A linear

### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

stepper motor with  $0.94 \mu\text{m}$  resolution (IntelLiDrives, Inc, Philadelphia, PA) driven by a Steynet motor controller (Copley Controls, Canton, MA) actuated the refuge, moving it forward and backward along specified velocity trajectories. A pco.1200s high-speed camera (Cooke Corp, Romulus, MI) with a Micro-Nikkor 60 mm f/2.8D lens (Nikon Inc., Melville, NY) captured 14-bit video with  $1280 \times 1024$  resolution from below. For single-sine and sum-of-sine trials, video was captured at  $50 \text{ frames} \cdot \text{s}^{-1}$ ; for stimulus-switching adaptation trials, video was captured at  $80 \text{ frames} \cdot \text{s}^{-1}$ . The camera was controlled using the Camware software package (Cooke Corp, Romulus, MI) from a standard PC. Custom Matlab (The Mathworks Inc., Natick, MA) scripts were used to generate and log trials as well as to synchronize actuator trajectories and camera shutter triggering via a USB-6221 Multifunction DAQ (National Instruments, Austin, TX) (Figure 3.2).




Figure 3.2: (A) *Experiment apparatus*. The data acquisition board sends synchronized commands to (1) the linear actuator (prescribing the trajectory) and (4) the high-speed camera (triggering exposures). Riding smoothly along a set of guide rails and rigidly linked to the actuator, (2) a rigid mast suspends a PVC refuge near the bottom of the aquarium. Video is captured from below via (3) an angled mirror and images are subsequently ported back to the PC via CamLink. (B) *Coordinate system*. Distinct patches are tracked using an SSD algorithm (custom Matlab code). Positions and velocities of these patches are measured from a fixed reference.

### 3.2.3 Experiment Design

Naïve individual fish (n=4) were presented with a variety of refuge trajectories composed of sinusoids, including single-sine and sum-of-sines stimuli. An additional set of naïve fish (n=3) were presented trajectories that switch between sum-of-sines to single-sines. One additional fish was presented with a set of sum-of-sine trajectories the responses to which were used for cross validation of the frequency response function models described below. To reduce the occurrence of startle responses, before each individual trial, animals were presented with 10 s of band limited noise refuge motion, and further, each stimulus amplitude was gradually ramped up at the beginning of the trial and down at the end of the trial (10 s ramp duration) to prevent abrupt onset and offset refuge movements. Together, these eliminated startle responses to the stimuli.

The stimuli are described in relation to velocity rather than position. Thus, throughout this paper, the word *amplitude* of a given trajectory refers not to the *distance* but rather to the maximum *velocity* associated with that stimulus. This is for three reasons. First, each animal may maintain an arbitrary absolute position within the refuge, creating an artificial DC offset in position but not velocity. Second, the sensory receptors are high pass, so that they encode velocity of movement rather than position [Cowan and Fortune(2007)]. Finally, previous experiments [Cowan and Fortune(2007)] as well as preliminary experiments for the present study suggest that

### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

the animals can exhibit saturation-like nonlinearities in tracking performance at high velocity amplitudes rather than positional amplitudes; as described in Results, the velocity amplitudes selected for our experiments avoid these saturation nonlinearities which simply define the performance boundaries of the animal and are not the focus of this work.

The sinusoidal stimuli were presented at a variety of velocity amplitudes  $\{0.6, 0.8$  and  $1.2 \text{ cm} \cdot \text{s}^{-1}\}$  and frequencies, and sums of these sinusoids. Refuge excursion frequencies were drawn from the set of the first thirteen prime multiples of 0.05 Hz, that is  $f = k \times 0.05 \text{ Hz}$  with  $k \in \{2, 3, 5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41\}$ . For single-sine trials, every other frequency was selected,  $f \in \{0.1, 0.25, 0.55, 0.85, 1.15, 1.55, 2.05 \text{ Hz}\}$ . Sum-of-sines trials were comprised of all frequency components with equal velocity amplitude (0.6, 0.8 or  $1.2 \text{ cm} \cdot \text{s}^{-1}$ ) and randomized phase. Consequently, when waveforms summed constructively, significantly higher velocities would be achieved. These periodic signals appear pseudo-random within a single period ( $T = 20 \text{ s}$ ) of the stimulus.

We also explored the time scales of adaptation between single-sine and sum-of-sines trajectories. Fish ( $n=3$ ) were presented with eight trials of longer stimuli (120 s duration) that switched between sum-of-sines and single-sine trajectories. In the first minute, fish were subjected to a sum-of-sines stimulus; at 60 s all but the 0.55 Hz frequency component were discontinued. The transition between stimulus types was instantaneous, but sum-of-sine frequency components were phase shifted to ensure

continuous velocity at the switch. In addition, between trials, the gain of the sum-of-sine frequency components (excluding the single component which persists) was inverted. As a result, averaging any two consecutive trials yielded an average input which was purely sinusoidal. Analysis was performed on these time-averaged trial pairs. This proved helpful in estimating phase transitions, because the averaged response to such pairs of stimuli was dominated by the frequency component of interest.

Positions of both the fish and the refuge were extracted from video using custom code implemented in Matlab. Volitional or exploratory behaviors within the refuge were included in the data set. Though infrequent, trials with excess volitional movement (*e.g.* the fish left the refuge or reversed rostrocaudal orientation within the refuge) were omitted from further analysis.

### 3.2.4 Analysis

#### 3.2.4.1 Frequency response functions

Empirical Bode plots were generated for all trials. A fast Fourier transform (FFT) was applied to both input and output velocity signals to compute the discrete-time Fourier transform (DTFT). For single-sine trajectories, we located the frequency at which the energy of the input signal peaks,  $\omega_o$ . We evaluated the output : input ratio of the DTFT values at this point,  $F(\omega_o)$ , and calculated gain (magnitude) and phase from the resultant complex number,  $|F(\omega_o)|$  and  $\angle F(\omega_o)$ , respectively. For sum-of-sines trials, we calculated the output : input ratio at the frequencies corresponding to



the thirteen greatest local maxima (excluding endpoints of the DTFT) of the energy of the input signal. We verified in all cases that the thirteen peaks indeed corresponded to the first thirteen prime multiples of the base frequency.

Confidence intervals in Bode plots were calculated from the distributions of output:input ratios (phasors) on the complex plane (Figure 3.3). Each distribution represented the system response to a class of inputs (either single-sine or sum-of-sine) at a set frequency. Single-sine trials yielded one point in the distribution corresponding to the stimulus frequency; sum-of-sine trials yielded a point for every constituent frequency. Fitting a Gaussian probability density function (PDF) to each cluster, we calculated the standard error and the associated PDF of the estimated mean (Figure 3.3(a)). The 95% confidence interval of the magnitude of the estimated mean was calculated as the minimum-area annulus over which the PDF integrates to 0.95 (Figure 3.3(b)); the confidence interval for phase of the estimated mean was the minimal conic region over which the PDF integrates to 0.95 (Figure 3.3(c)).

#### 3.2.4.2 Continuous phase estimation

For stimulus-switching trials, frequency response analysis was performed on the trial-averaged input-output pair. The mean phase for the sum-of-sines and single-sine intervals was calculated as described above. Assuming that the beginning of the single-sine interval represents a period of transition, the mean phase for this regime is calculated over the final 30 s to give a better approximation of the asymptotic phase

CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*




Figure 3.3: For a given frequency, the system response can be characterized by a point on the complex plane,  $\alpha e^{i\phi}$ , where the magnitude of  $\alpha$  is the gain of the system and  $\phi$  the phase shift with positive phase measured counter clockwise from the positive real axis. The circle of unit magnitude (representing unity gain) and the positive real axis (representing zero phase shift) are denoted in blue. Each trial yields one estimate for the system response at each frequency. (A) We fit a Gaussian probability density function in the complex plane at each frequency; to illustrate, the 95% covariance ellipse for single (*red*) and sum-of-sines (*black*) is shown for 2.05 Hz. (B) The phase confidence interval is the conic region over which the PDF integrates to 0.95. (C) Similarly, the magnitude confidence interval (95%) of the estimate of the mean is the annulus over which the PDF integrates to 0.95.

value.

A coarse estimation of the phase was calculated as it changed with time. For a 5.0 s moving window in time, the best fit sum-of-sines trajectory was fit using (3.1):

$$\begin{aligned}
 (\alpha, \beta) &= \arg \min_{\alpha, \beta} \sum_t \left\{ \left( \sum_{i=1}^k \alpha_i \sin(\omega_i t) + \beta_i \cos(\omega_i t) \right) - y(t) \right\}^2 \\
 \begin{bmatrix} A^T \\ B^T \end{bmatrix} &= [\sin(\omega_i t) \quad \cos(\omega_i t)]^\dagger [y(t)].
 \end{aligned} \tag{3.1}$$

The frequencies  $\{\omega_1, \dots, \omega_n\}$  are known and  $A = [\alpha_1, \dots, \alpha_n]$  and  $B = [\beta_1, \dots, \beta_n]$

CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

are solved for in a least squares sense. Though the trial-averaged response is ideally sinusoidal, we use the best fit sum-of-sines trajectory to account for any residual frequency components not entirely eliminated through averaging. Using the trigonometric identity in (3.2), we solved for the magnitude and phase of the refuge and fish as in (3.3):

$$M_i \sin(\omega_i t + \phi_i) = \underbrace{M_i \cos(\phi_i)}_{\alpha} \sin(\omega_i t) + \underbrace{M_i \sin(\phi_i)}_{\beta} \cos(\omega_i t) , \quad (3.2)$$

$$M_i = \sqrt{\alpha_i^2 + \beta_i^2} , \quad (3.3)$$

$$\phi_i = \arctan_2(\beta_i, \alpha_i) ,$$

where  $\arctan_2$  is the four quadrant version of the arctangent function. Gain and relative phase were then calculated as the ratio  $\frac{M_{\text{fish}}}{M_{\text{refuge}}}$  and the difference  $\phi_{\text{fish}} - \phi_{\text{refuge}}$ , respectively.

The finer estimate of the instantaneous phase was computed as the argument of the analytic signal,  $f(t) + i\mathcal{H}(f(t))$ , where  $f(t)$  represents either the input or output time signal and  $\mathcal{H}(\cdot)$  denotes the Hilbert transform. This method, however, is highly sensitive to noise in the time-domain signal.

### 3.3 Results

#### 3.3.1 Responses to motion stimuli are coherent

As previously reported [Cowan and Fortune(2007)], fish robustly followed the experimentally controlled movements of the refuge by swimming backwards and forwards. The swimming of the fish was strongly correlated with movements of the refuge, and as a result the movement of the fish exhibited strong coherence to the stimulus trajectory. This result held for each category of stimulus that was tested, including predictable sine wave stimuli, sum-of-sines stimuli, and more complex stimuli. An example response to a sum-of-sines stimulus is shown in Figure 3.4.

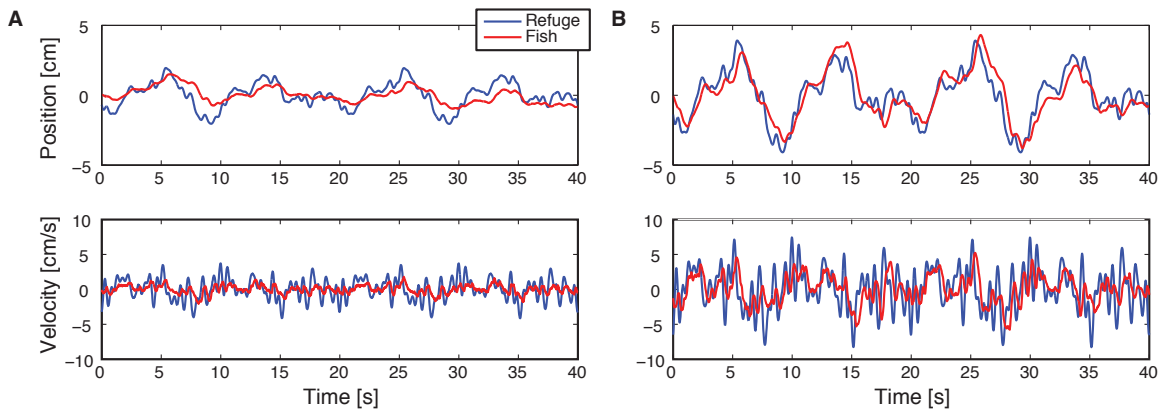


Figure 3.4: A velocity-scaled repetition of a sum-of-sines trial with fish and refuge trajectories depicted in blue and red respectively: (A) the individual sinusoidal components have an amplitude of  $0.8 \text{ cm} \cdot \text{s}^{-1}$  and in (B)  $1.2 \text{ cm} \cdot \text{s}^{-1}$ .

For each trial, we computed the magnitude of the Fourier components for input (refuge velocities) and output (fish velocities) as shown in Figure 3.5(A). In all in-

### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

stances, peaks in output power correspond to peaks in input power. These strong relationships confirm that the fish is tracking the stimulus, and that the fish's movements are not the result of other potential behaviors such as exploratory movements. And this behavior is remarkably *clean* of ancillary frequency content. In a subsequent set of experiments (not in the scope of this thesis), fish were observed performing the refuge-tracking task with diminished sensory information [Stamper et al.(2012)]; visual cues were limited by illumination (two conditions, light and dark) and electrosensation was attenuated through changes in water conductivity (high conductivities collapse the extent of the EOD). In dark conditions, fish perform a noticeable auxiliary behavior in which a broad-frequency oscillatory motion is superimposed on the smooth pursuit response (Figure 3.6A). On a trial-by-trial basis, we exclude the data at the stimulus frequency and average the remaining data to reconstruct the spectrum of the emergent oscillations. In fact, this volitional movement is present in both light and dark conditions (and for all conductivity levels), though the magnitude of oscillations is heavily attenuated in the presence of visual sensory information (Figure 3.6B).

Sum-of-sines trials consistently had coherences near unity at the stimulus frequencies (Figure 3.5(B)). Note that for frequencies not present in the stimulus (i.e. between peaks) the coherence value is not informative (the input-output relationship is dominated by noise). It is also important to note that coherence remains near unity even at high frequencies where tracking performance diminishes, because coherence is a

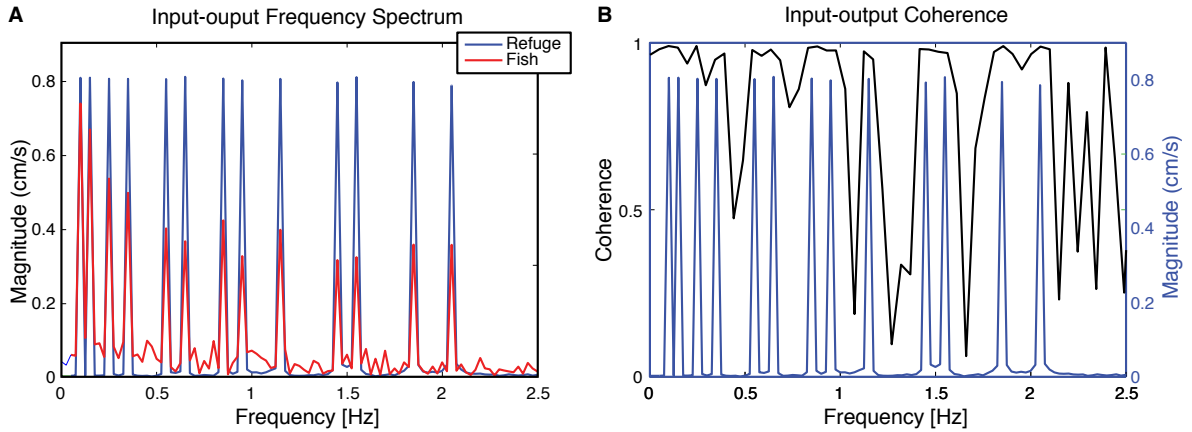


Figure 3.5: (A) The magnitudes of the input (refuge) and output (fish) DTFTs. (B) The coherence between refuge and fish trajectories (black) and the magnitude of the refuge trajectory DTFT. The near unity coherence suggests fish and refuge trajectories are related by a linear dynamical system.

measure of signal-to-noise ratio and not a measure of absolute gain (Figure 3.5(B)).

Strong coherence for each stimulus–response pair suggests that the tracking behavior may be described by linear dynamics. We examined whether one linear dynamical system can indeed adequately describe all input–output pairs across stimulus categories. If so, a small subset of input–output pairs could furnish a predictive linear model for refuge tracking behavior.

### 3.3.2 Linear models do not generalize across stimulus classes

Linearity of a system is defined by two properties: scaling and superposition. To test scaling, we presented three velocity amplitudes,  $\{0.6, 0.8, 1.2 \text{ cm} \cdot \text{s}^{-1}\}$ , for each stimulus type,  $\{\text{single-sinusoid, sum-of-sines}\}$ . The Bode plots for each velocity amplitude

CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

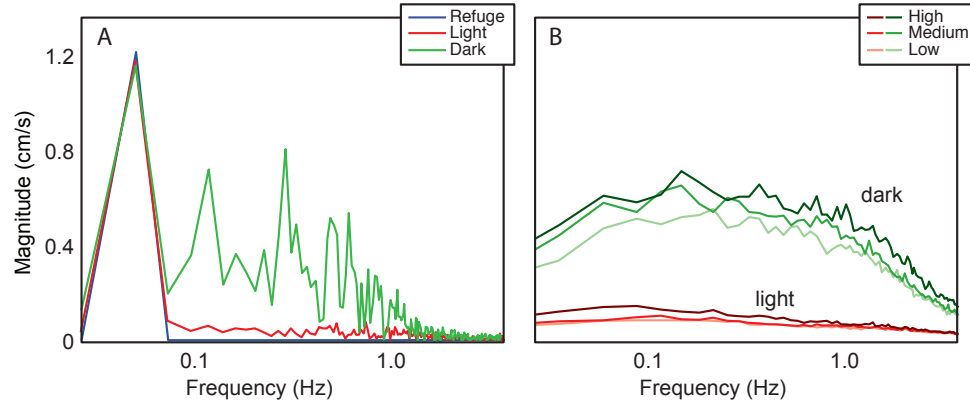


Figure 3.6: (A) A comparison of the output (fish) DTFTs for a single-sine trial performed in light and dark. (B) The trial-averaged DFT of volitional motion reveals the effect of lighting and conductivity conditions on the ancillary emergent oscillations (figure adapted from [Stamper et al.(2012)]).

are shown in Figure 3.7(A) for single-sines and Figure 3.7(B) for sum-of-sines.

In general, the scaling property cannot hold for an arbitrarily large regime of stimuli. Thus, based on previous work [Cowan and Fortune(2007)] we examined a biologically relevant range of velocity amplitudes. Over this range of velocity amplitudes, the phase response curves within each of the two stimulus classes were remarkably invariant, as shown in Figure 3.7(A,B). Amplitudes were also generally consistent with the scaling property, although some differences can be seen for single-sine stimuli in the range of 0.25 – 0.55 Hz. Despite the noted discrepancies in gain, within a fixed stimulus type, changes in trajectory amplitude do not suggest categorical changes in the response. Taken together, the amplitude and phase responses strongly suggest that tracking behavior scales linearly with input over the range of velocity amplitudes tested.

Having demonstrated the scaling property in these data, we next examined the

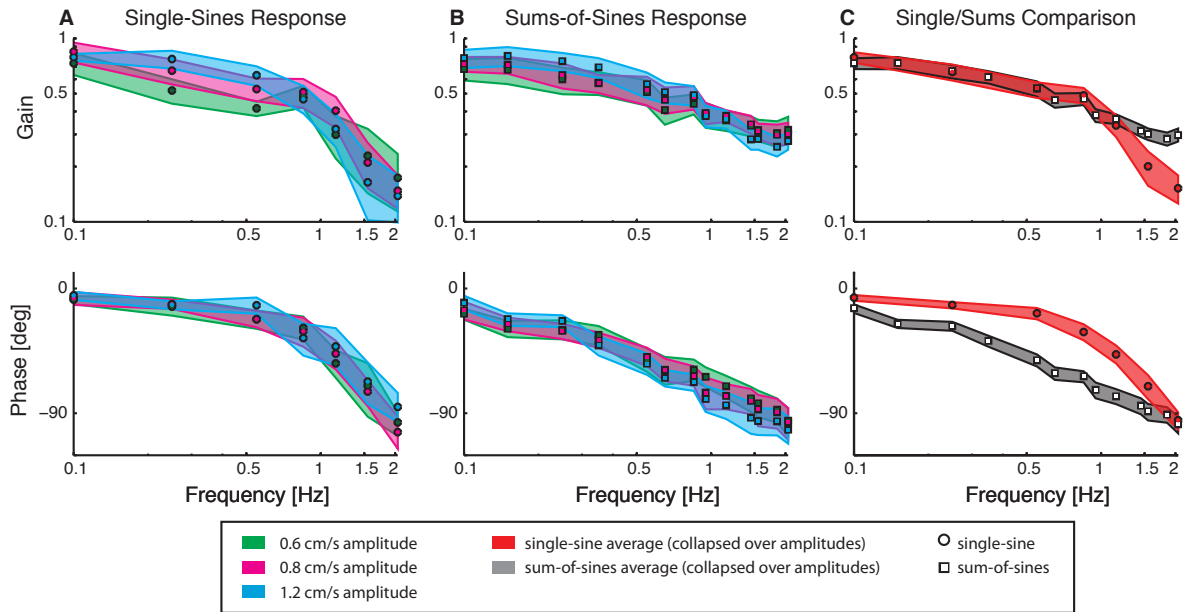


Figure 3.7: The Bode plot describes the system response (gain and phase shift) over a range of frequencies. Bode plots for stimuli of different velocity amplitudes are compared for (A) single-sine trials and (B) sum-of-sines trials. (C) Average responses (collapsed across scaling) comparing sum-of-sines (black) and single-sine (red) trials are depicted with confidence intervals calculated as described in Figure 3.3. For a linear system, these estimates would agree.

superposition property. This was done by comparing single-sine to sum-of-sines data. If superposition holds, the responses to single-sine inputs should predict sum-of-sines responses. In other words, if superposition holds for these data, the Bode plots from the two stimulus categories should be identical. Interestingly, the Bode plots (Figure 3.7) for the two stimulus categories exhibit unmistakable differences: responses to single-sine stimuli exhibited lower phase lag at mid-range frequencies and greater attenuation at high frequencies than responses to sum-of-sine stimuli (Figure 3.7). Because the Bode plots are different across stimulus categories, superposition therefore fails. A single linear model cannot account for the responses to both categories of in-



### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

put. However, when analysis is limited to either single-sine or sum-of-sines trials, the high coherence and low variance of frequency response estimates suggest that a linear system might be useful in describing this behavior within each stimulus category.

What is the consequence of the mismatch in FRFs/Bode plots in terms of their predictive power? For a linear system, the linear model furnished by one FRF can be used to predict the temporal response of the system to the same or a different stimulus category (Figure 3.8). We used this technique as a mechanism to understand the differences between the linear models for each stimulus category. To do this, we used the single-sine FRF to make predictions of the responses of the fish to sum-of-sines stimuli. Next we compared these predictions to the actual responses of the fish. Specifically, the average single-sine and sum-of-sines FRFs shown in Figure 3.7(C) were used to predict the response of a different fish (not included in the FRF data) to individual sum-of-sines stimuli. For each of the 15 trials, the sum-of-sines FRF model from Figure 3.7(C) predicted the response with less root-mean-squared error than the single-sine FRF model; the mean improvement was 36.7%, the minimum improvement was 12.0%, and the maximum improvement was 64.3%. As expected, the FRF from single sine data does not generalize to spectrally different stimuli, likely due to the nonlinearity revealed by the FRF data (Figure 3.7). The consequence of the nonlinearity between stimulus categories is that fish perform better to predictable stimuli than to unpredictable stimuli.

### 3.3.3 Fish adapt to changes in stimulus

We next investigated the time course of the transition between the two responses, focusing specifically on the response to stimuli at 0.55 Hz where phase shows maximal change. However, current methods for estimating the time-varying phase of a signal often yield noisy or unreliable results for short time intervals, and in fact instantaneous estimation of phase for dynamical systems remains an area of active research (*e.g.* [Revzen and Guckenheimer(2008)]). Towards wholly describing the transition between sum-of-sines to single sines, we present phase analyses at three time-scales: one 30 s window, six 5 sec windows, and an instantaneous/continuous phase estimate (see Figure 3.9). At the coarsest level, the asymptotic phase—calculated as the DTFT estimate of phase from the second half (30 s) of each stimulus regime—reveal phase lag reductions of 10.8, 13.9, and 18.7°, which is less than the mean 33.9° reduction observed in the first population of fish. A more refined view of the adaptation is captured by dividing the 30 s following transition into six consecutive non-overlapping 5 s windows. For the first two fish, a trend seems to emerge, possibly suggesting an exponential decay to the asymptotic phase. However, due to the variance of the phase estimate (standard deviation shown as *black* error bars), any estimate of a time constant for such decay would be tenuous. For the third fish, volitional movement and/or other sources of motion noise yield phase estimation at this time scale that is unreliable. At the most refined time scale, we use the ar-




Figure 3.8: An example demonstrates that frequency response functions generalize better within stimulus class than across classes. Assuming superposition holds, the FRFs in Figure 3.7(C), generated from four fish, can be used to predict the response to an arbitrary input for a fifth fish. (A) Ten seconds of a sum-of-sine stimulus (*blue*) and the fish's response (*green*). (B) A comparison of predictions made by different FRF models. The sum-of-sines prediction (*black*) closely matches the fish's performance (*green*). The single-sine prediction (*red*) is worse than for the sum-of-sines FRF. (C) The difference between the single-sine and sum-of-sine prediction errors. Negative values (in *red*) indicate time intervals for which the single-sine FRF model has greater error than the sum-of-sines model. Predominantly, the sum-of-sines model better predicts the fish's actual response. Were the system linear, an assay of single-sine experiments would be sufficient for predicting the response to the sum-of-sines stimulus.

CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

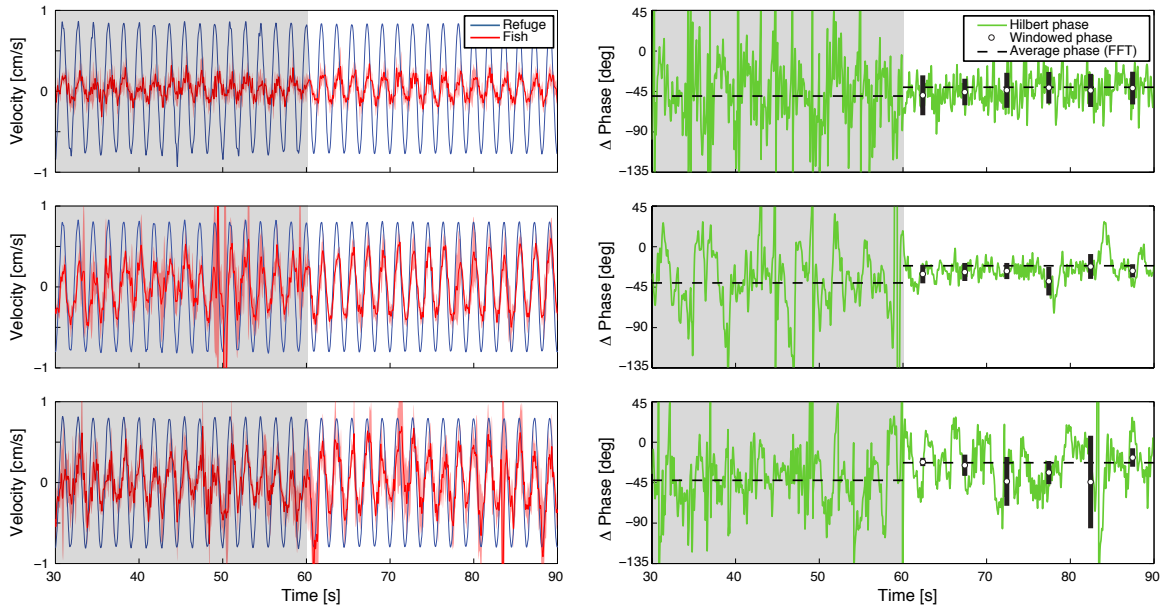


Figure 3.9: In switched stimulus trials, fish ( $n=3$  corresponding to the rows above) adapt to a change from sum-of-sines to single-sine stimuli. On the *left*, the mean of eight trials; sum-of-sines stimuli (interval shaded in *grey*) were designed so that the mean input would be purely sinusoidal. On the right, the instantaneous phase estimate is calculated from the averaged response using the argument of the analytic signal (the Hilbert phase shown in *green*). A coarser phase estimate is shown for non-overlapping windows of 5 sec (as *black* error bars). The asymptotic phase over each interval is calculated using the Fourier transform method used to generate Figure 3.7.




Figure 3.10: An abrupt transition between lighting conditions elicits an immediate and categorical change in behavior. Broad spectrum oscillations emerge superimposed on the smooth pursuit trajectory (figure adapted from [Stamper et al.(2012)]).

gument of the analytic signal to generate a continuous estimate of phase (shown in *green*) [Revzen and Guckenheimer(2008)].

Though this approach yields a noisy estimate it highlights two important traits of the transition: the time constant of decay is on the order of several seconds to tens of seconds and the variance of the phase estimate is lower in the single-sine regime. Though we cannot make an accurate estimate of the transition in dynamics, we interpret the long gradual change as the consequence of a slow learning or adaptation process. To emphasize this point, we can compare the transition of tracking behavior at an abrupt change in lighting condition (Figure 3.10 from [Stamper et al.(2012)]); in response to this change, the fish exhibits a complete behavioral switch within the first 1 or 2 s. Back to our sum-to-single-sine switching experiments, the reduced variance may partially be attributed to the method used for phase estimation, but we suspect that this reduction in variance results at least in part due to changes to the behavior that occur during adaption to the switch in the stimulus.

### **3.3.4 Adaptation to single-sine stimuli reduces tracking error**

Having observed categorically different FRFs elicited by single-sine and sum-of-sine stimuli, we hypothesized that this nonlinearity was indicative of a stimulus-mediated adaptation. In this section we explore the benefits of such an adaptation: whether tracking performance for single-sine stimuli improves compared to the response to

CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

Table 3.1: Phase, gain and error differences between sum-of-sine and single-sine trials as depicted in (Figure 3.11(A)).

| Frequency (Hz)             | 0.10 | 0.25  | 0.55  | 0.85  | 1.15  | 1.55  | 2.05  |
|----------------------------|------|-------|-------|-------|-------|-------|-------|
| Error reduction (%)        | 31.7 | 25.3  | 33.8  | 28.8  | 17.9  | 8.2   | 5.3   |
| Gain difference (%)        | 8.0  | -0.9  | -1.2  | 3.9   | -7.9  | -33.4 | -48.7 |
| Phase difference (degrees) | -7.7 | -15.4 | -33.9 | -31.9 | -30.2 | -18.1 | -2.9  |

sum-of-sines stimuli and the energetic trade-offs of improved performance. In order to address these questions, we consider yet another representation of the frequency response, as complex phasors.

When considering the tracking behavior in terms of phasors on the complex plane, gain is measured as the distance from the origin and phase measured as the angle (counter clockwise) from the positive real axis. Hence, unity gain is represented by a unit circle and zero phase corresponds to the positive real axis (Figure 3.11(A)). The intersection of the unit circle with the positive real axis, the point  $1 + i \cdot 0$ , indicates perfect tracking. The magnitude of the error signal (the sensory slip) is measured as the distance between the empirical phasor (represented as a point on the complex plane) and the perfect tracking point.

In Figure 3.11(A), the mean phasor for each frequency is plotted for both single-sine and sum-of-sines stimuli. At every frequency compared, the responses to the single-sine stimuli exhibited less error. Excluding the frequencies 0.1 and 0.8 Hz, these improvements in tracking were achieved despite a *reduction* of gain (Table 3.1). The gain (the distance between the empirical phasor and the origin  $0 + i \cdot 0$ ) provides an indication of effort or energy expended during tracking. For the two frequencies

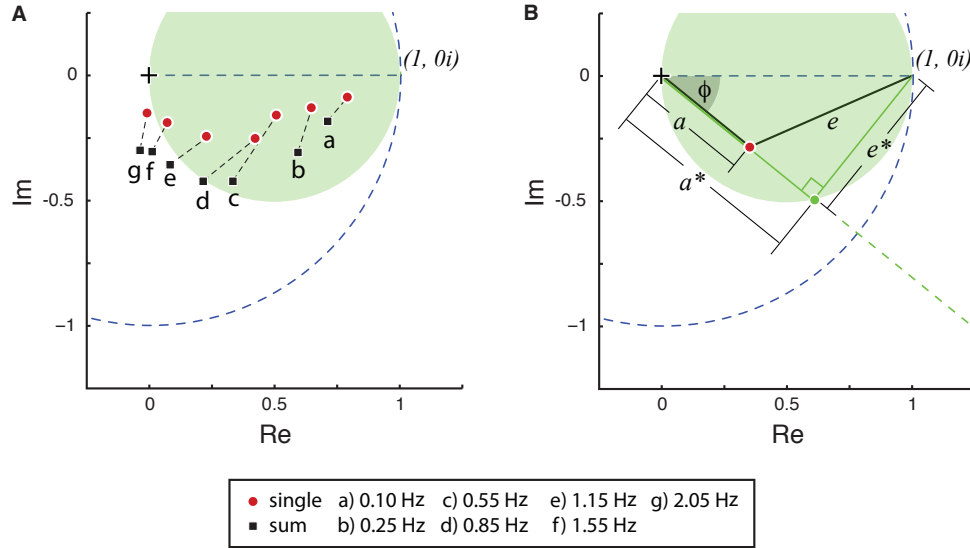


Figure 3.11: (A) Frequency response phasors shift to decrease tracking error in the transition from sum-of-sine to single-sine stimuli. (B) For a frequency response characterized by gain  $a$  and phase lag  $\phi$ , denoted by the *red* dot above, the magnitude of the sensory slip  $e$  is the distance from the frequency response point to the perfect tracking point  $1 + i \cdot 0$ . For a fixed phase  $\phi$ , the minimal error  $e^* = \sin(\phi)$  for  $\phi \in (-\frac{\pi}{2}, \frac{\pi}{2})$  ( $e^* = 1$  otherwise) is achieved by a gain of  $a^* = \min(\cos(\phi), 0)$  (depicted as the *green* dot). The locus of minimum-error responses given fixed phase lag/lead is denoted by a circle unity diameter centered at the point  $(0.5, 0i)$ ; within this circle (shaded *green*), there is a trade-off between increased error and savings in expended energy.

where gain increased, it increased only 8.0% and 3.9%, and these increases were not statistically significant (two-sample, one-sided  $t$ -test,  $P = 0.1206$  and  $P = 0.3343$ , respectively). This was consistent with performance in the frequency range from 0.1 to 1.15 Hz (see Table 3.1), where gain remained relatively constant (within  $\pm 10\%$ ) while error was reduced dramatically (18-32%). In contrast, at the highest frequencies tested, 1.55 Hz and 2.05 Hz, fish dramatically reduced their effort maintaining a small but statistically significant improvement in tracking error ( $P = 0.0218$  and  $P = 0.0151$ , respectively). These results show that there was a frequency-dependent shift

in the trade-off between effort and tracking error.

## 3.4 Discussion

*Eigenmannia virescens* exhibit a switch in tracking performance depending on the category of the refuge trajectory—a simple sinusoid versus a more complex sum-of-sines. This nonlinear switch results in reduced tracking error to simpler sinusoidal stimuli despite an often dramatic reduction in motor effort. This concomitant decrease in tracking error and motor effort suggests adaptive and predictive neural mechanisms for locomotor control in *Eigenmannia*.

### 3.4.1 Responses to single-sine and sum-of-sine stimuli

Both categories of stimuli—single-sines and sum-of-sines—are fundamentally deterministic. So, why then are fish able to track single-sine stimuli so much better and with less motor effort at each frequency? Intuitively, single-sine stimuli are more predictable than sum-of-sine stimuli. More formally, as the number of parameters of a signal increases, noisy measurements—which are inescapable—lead to greater variance in parameter estimates. Thus, given the same amount of measurement data, computational algorithms that extrapolate sensory measurements of stimuli will perform worse for sum-of-sine stimuli than for single-sine stimuli. In this sense, single-sine stimuli are fundamentally more predictable than sum-of-sine stimuli, which we treat



### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

as pseudo-random.

Further, the pseudo-random sum-of-sine stimuli are complex periodic waveforms with a long (20 second) period. To avoid the potential for long-term learning of these stimuli, the relative phases of each component sinusoid were randomized from trial to trial, thus creating distinct temporal trajectories that nevertheless had identical spectral content. Importantly, the response to these distinct sum-of-sine stimuli generalized (Figure 3.8).

For mid-ranged frequencies, the gain of single-sine and sum-of-sine responses are about the same, but the single-sine phase lags are substantially reduced compared to the corresponding components of the sum-of-sine response (Figure 3.7C). This corresponds to a substantial decrease in tracking error with little-to-no change in the motor effort. Moreover, complex-plane analysis (Figure 3.11) reveals that at high frequencies, single-sine responses exhibit a dramatic reduction in motor effort (the high-frequency responses are much closer to the origin of the complex plane) and a simultaneous decrease in tracking error (the responses are closer to the point  $1 + i \cdot 0$ ).

Thus, at all frequencies fish exhibit the same or less tracking error with about the same or less motor effort when presented with single-sine stimuli (Figure 3.11). The decreases in tracking error are generally associated with reduced phase lag for single sines, and the decrease in motor effort (which occurs at high frequencies, where there is substantial phase lag for both single and sum-of-sines) is generally associated with lower gain.

### 3.4.2 An internal model predicting refuge movement explains phase discrepancies

Phase profiles are consistent between trials when the stimulus regime is fixed but shift categorically between the two different stimulus types. Specifically, for single-sine stimuli, fish exhibit reduced phase lag, but surprisingly this decrease in phase occurs with little-to-no change in gain for frequencies up to 1 Hz. Thus, we suspect a predictive mechanism—in which stimulus dynamics are included in the state estimate—to be responsible for this disparity between single-sine and sum-of-sines phase responses.

Neural delays introduce inherent phase lags between the sensory stimulus (input) and locomotor action (output). But if a stimulus were sufficiently predictable, the nervous system could, in principle, compensate for these delay-induced phase lags by extrapolating the stimulus trajectory forward in time. This would enable the neural control system to act upon an estimate of the current-time stimulus signal despite the sensorimotor delay. However, for trajectories which evolve randomly, this prediction is inaccurate, requiring the system to rely heavily on the delayed sensory measurements to calculate the appropriate motor response. Hence the internal delays manifest as phase lag. The Kalman filter, a state-estimation algorithm common to many engineering applications, provides a flexible framework for discussing prediction in the context of sensory and motor uncertainty (*e.g.* the postural control problem

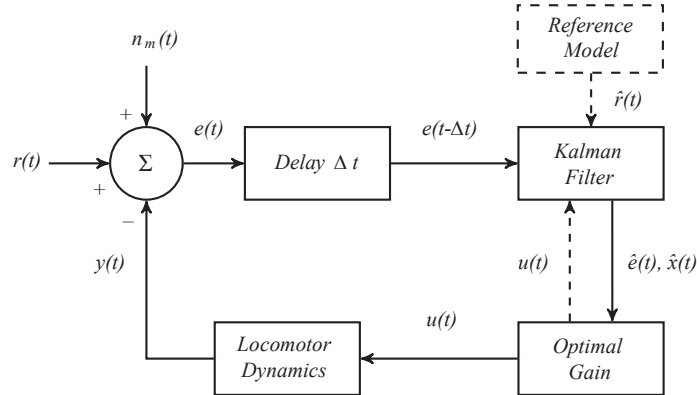


Figure 3.12: A proposed model of tracking behavior accounts for the categorical differences in phase and gain between predictable and pseudo-random stimuli. The reference signal  $r(t)$  is an arbitrary refuge trajectory; the fish internally estimates this refuge trajectory as the output of a noise-driven dynamical model. A noisy measurement of the sensory slip,  $e(t)$ , represents the relative velocity of the refuge with respect to the fish. However, due to sensory and transmission delays amounting to  $\Delta t$ , control actions at time  $t$  must be determined using outdated sensory slip information. The Kalman Filter uses this outdated sensory measurement in conjunction with the internal reference model to estimate the *current* state of the system,  $\hat{x}(t)$  and  $\hat{e}(t)$ , which includes the position and velocity for the fish and the relative position and velocity of the refuge. This time-corrected estimate is used to determine the control signal  $u(t)$  sent to muscles along the ribbon fin which, in turn, result in changes to the fish position and velocity  $y(t)$ .

in [Kuo(2005)]). The Kalman filter generates the optimal state estimate by reconciling two streams of information: a belief about what the state of the system should be (as predicted by an internalized model of the system dynamics) and sensory measurements. Each of these streams of information, the model-based prediction and the measurement, bear their own sources of uncertainty: process noise determines the extent to which the evolution of the system states is affected by randomness—in effect, the unpredictability of the system—and measurement noise degrades the reliability of observed quantities. In our proposed model of refuge tracking (Figure 3.12), the

### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

internalized model includes a stochastic dynamical model of refuge motion in addition to the locomotor dynamics of the animal. The Kalman filter reweights the contributions of these two streams based on the relative sizes of the measurement and process noise variances. Thus, the internalized model makes a less effective prediction about the state evolution for these pseudo-random stimuli, requiring the nervous system to rely on the delayed sensory measurement as described above.

Prediction in motor control often refers to an adaptive model of internal system states (*e.g.* estimating the position or orientation of your hand during a reaching task without visual feedback [Wolpert et al.(1995), Shadmehr and Mussa-Ivaldi(1994)]). Though this kind of prediction would be pertinent to refuge tracking, this is *not* the kind of predictive mechanism we suspect here. Rather, we contend that the nervous system predicts—using an internalized dynamical model—the movement of the exogenous signal. Similar stimulus prediction has been described in terms of probabilistic representations of target locations in a pointing task [Körding and Wolpert(2006)] and in terms of the anticipation of the time of direction reversal in an visual target tracking behavior [Collins and Barnes(2009)].

Similar to our proposed model, Carver *et al.* investigated whether the dynamics of a moving visual scene are estimated for human posture control [Carver et al.(2005)]. They compared three different dynamical models for the external scene and assess how well prediction schemes incorporating these models might reproduce empirical data. Through a broad parameter search they found that, even for an optimized set of

parameters, the models they considered in which external dynamics were estimated did not satisfactorily capture qualitative features of the empirical data. The data presented in this work suggest that similar prediction-of-dynamics models should be revisited in the context of the refuge tracking behavior in *Eigenmannia*. Using input–output FRF models to reverse engineer the sensorimotor transform requires a sufficiently representative model of the locomotor dynamics [Cowan and Fortune(2007)], so as the dynamics of ribbon-fin propulsion become better understood [Shirgaonkar et al.(2008), Sefati et al.(2010)] our model (Figure 3.12) can be used to generate quantitative predictions for refuge tracking behavior.

### 3.4.3 Gain discrepancies indicate improved tracking for predictable stimuli

At the higher frequencies, we observed a significant reduction in gain for single-sine presentations. Typically, this attenuation would be interpreted as a worsening of tracking performance which would deceptively suggest that at high frequencies fish do poorly at tracking predictable stimuli compared to unpredictable stimuli. However, consider the tracking behavior transfer function on the complex plane (Figure 3.3). In this representation, unity gain is represented as the dashed unit circle; zero phase is designated by the dashed line along the positive real axis. The intersection of the unit circle with the positive real axis, the point  $1 + i \cdot 0$ , indicates perfect tracking. The magnitude of the error signal (the perceived sensory slip) is measured as the distance

### CHAPTER 3. REFUGE-TRACKING IN *EIGENMANNIA VIRESCENS*

between the empirical transfer function (represented as a point on the complex plane) and the perfect tracking point.

For 2.05 Hz, we see that the distribution of single-sine trials is on average closer to  $1 + i \cdot 0$  than the sum-of-sines trials. At any given phase lag (or lead)  $\phi$ , error is minimized by a gain of  $\max(0, \cos(\phi))$  (Figure 3.11(B)). In these analyses, gain represents a normalized velocity (the ratio of fish and refuge velocities) and therefore might serve as an indicator of expended energy. Subscribing to this interpretation, gains lower than the minimal-error gain compromise error for energetic savings; higher gains are suboptimal in both error and energetic cost. For phase lags greater than  $90^\circ$ , error is minimized at zero gain. Despite the immediate interpretation that reduced gain indicates reduced tracking performance, in the high-frequency regime reduced gain improves tracking performance with respect to sensory slip. Hence, for predictable stimuli (single sines), the controller adapts to reduce both error and energetic cost.

# Chapter 4

## Optomotor Yaw Regulation in

### *Drosophila melanogaster*

Time flies like an arrow.  
Fruit flies like a banana.

---

Groucho Marx

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*



Figure 4.1: A *tethered fruit fly*. Merely one hundred thousand neurons packaged into a 3 mm body (the specimen shown is compared to a dime), the fruit fly is a remarkably adept agent within its environment. In the laboratory, fruit fly optomotor response persists even when the animal is tethered to a tungsten wire.

## 4.1 Why the Fruit Fly?

Fruit flies serve as model systems in many branches of biological research: genetics, biomechanics, neuroscience, developmental biology and behavioral biology to name a few. There is a significant body of literature describing the vast repertoire of behaviors and the neurophysiological mechanisms involved. In large part, the ubiquity of fruit flies as model systems (particularly for the study of genetics) is due to the ubiquity of fruit flies; fruit flies can be found almost anywhere, have short life-cycles, and reproduce rapidly. But for the study of behavior, the draw of *Drosophila* is more than just availability and convenience.

Of the myriad behaviors, optomotor stabilization behaviors—those in which in-



## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

sects modulate flight in response to visual stimuli, including stripe/object fixation and large-field stabilization—have been studied extensively in flies. From a neurophysiological perspective, fruit flies represent a desirable balance of system complexity (*Drosophila melanogaster* comprises merely 100,000 neurons) and behavioral richness and robustness even under aggressive experimental manipulation (such as the tethered preparation in Figure 4.1). This is not to suggest that fruit flies are simple, only that the neural circuitry which governs these behaviors must be in some ways parsimonious. Despite the frugality of neural hardware, many behaviors and underlying mechanisms generalize to other not-so-simple animals; in fact, the prevailing neural model for the elementary motion detection unit (the Reichardt-Hassenstein model [Reichardt(1961)]) was inspired by experiments on beetles and later validated for flies, but analogues of this correlation-based detector have been extended to motion processing across many taxa (Chacron *et al.* have discovered neural mechanisms analogous to the Reichardt-Hassenstein detector in weakly electric fish [Chacron et al.(2009)]).

### 4.1.1 Mechanisms

The sensorimotor control of flight represents the convergence of sensory cues (visual, vestibular, haltere, olfactory) to modulate wing stroke kinematics to generate turn in roll, pitch, and yaw or translate in thrust or lift. In our experiments, tethered flies (Figure 4.1) modulate yaw torque (via changes in muscle activation) to stabilize

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

a visual scene moving about the azimuth, a closed-loop single-input–single-output control topology. For the purpose of this brief review, we limit discussion to the mechanisms of visual motion processing and yaw torque modulation, to reflect the constraints imposed by our experiment paradigm.

In Section 3.1.1, we described the processing of the electrosensory image: electroreceptors project onto the electrosensory lateral line lobe (ELL) —where parallel topographic maps segregate receptive fields by spatial content—which subsequently projects onto the torus semicircularis where direction selectivity (and velocity) are encoded. In many ways, the fly visual system is analogous to the electrosensory system of fish. Photoreceptors in the retina project onto the lamina which in turn projects onto the medulla which in turn projects onto the lobula complex (composed of the lobula and lobula plate). The lamina, medulla and lobula are columnar structures which preserve a retinatopic map [Resh and Cardé(2009),Krotki(2011),Weber(2011)]. These columnar structures are themselves stratified and the functional sensitivities of different layers have been identified. In the medulla, strata respond preferentially either to wide-field (striped drum) or localized motion (single bar) [Bausenwein and Fischbach(1992)]. Downstream of the medulla, scene motion is encoded in the lobula complex, specifically in the sixty lobula plate tangential cells (LPTCs) [Krotki(2011),Weber(2011)]. LPTCs are motion-selective interneurons, categorized by their preferred direction as either horizontal system (HS) cells (responsive to regressive or progressive motion horizontally along the azimuth) and VS cells (responsive

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

to vertical motion, either up or down from the equator) [Buchner et al.(1984)]. More recently, vertical system (VS) cells have been found to respond to complex patterns of optic flow, vector fields congruent with rotational egomotion [Krapp et al.(1996), Franz and Krapp(2000), Borst and Haag(2007), Weber et al.(2008)]. The physiological evidence for the representation of optic flow in the lobula plate has steered models of motion control to rely solely on motion (velocity) cues; these bio-inspired optomotor control paradigms have been demonstrated on robot platforms [Zbikowski(2005), Humbert et al.(2005), Conroy et al.(2009)].

In the stripe-fixation behavior, flies modulate yaw torque in response to moving visual stimuli. In our experiments, we measure a kinematic parameter (the difference between left and right wingbeat amplitudes,  $\Delta\text{WBA}$ ) to infer yaw torque and subsequently turning velocity. However, there is no definitive agreement as to the wing kinematics responsible for yaw torque generation and the dynamics of the turning response. Yaw torque is regulated by a number of kinematic factors, primarily left-right asymmetries in wingbeat amplitude and angle of attack [Zanker(1990)]. Dickinson *et al.* advocate that the relative timing of the wing reversal from downstroke to upstroke is the dominant factor in modulating torque—this “ventral flip” is strongly correlated to WBA [Dickinson et al.(1993)]. Bergou *et al.* contend that asymmetric adjustment to wing pitch more significantly contributes to yaw torque [Bergou et al.(2010)]. Historically, wingbeat amplitude asymmetry is the standard kinematic indicator of yaw torque [Götz et al.(1979)]. This kinematic parameter provides an observable measure

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

from which we can infer torque, but to close the loop from motor output (torque) to visual scene (angle), we must integrate the forced dynamics of yaw rotation.

A simple one degree-of-freedom model for the yaw dynamics is  $J\ddot{\theta} + B\dot{\theta} = \tau$ , where  $\theta$  is the yaw angle,  $J$  is the moment of inertia,  $B$  is the linear rotational damping, and  $\tau$  is the wing-generated torque assumed to be proportional to  $\Delta\text{WBA}$ . The damping time constant  $J/B$ , compared to the duration of a maneuver, determines whether the damping forces dominate [Fry et al.(2003)]: if  $J/B$  is small then the second-order model can be further reduced to first order. During a turning maneuver, the rotation of the body with respect to the inertial frame imparts an asymmetric velocity component to each wing. This asymmetry generates a flapping counter torque (FCT) opposed to the direction of rotation [Hedrick et al.(2009)]. Indeed, recent modeling suggests damping dominates (the FCT-induced half life of rotational velocity is approximately two wing strokes) [Hedrick et al.(2009), Cheng et al.(2010)], thus we assume that the yaw dynamics further reduce to  $\dot{\theta} \propto \tau$ . This justifies, to some degree, the long-standing tradition (and our choice of feedback control policy) to simply scale  $\Delta\text{WBA}$  and treating it as the angular velocity of the fly.

### 4.1.2 Behaviors

Optomotor yaw regulation has been studied extensively towards understanding many aspects of sensorimotor processing: comparing responses to small- and wide-field visual stimuli [Duistermars et al.(2007), Geiger and Nässel(1982)], parsing cross-modality

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

sensory integration [Frye and Dickinson(2004), Sherman and Dickinson(2004), Budick et al.(2007)], etc. However, empirically derived predictive models for optomotor behaviors are rare in the literature. Theobald *et al.* [Theobald et al.(2010)] modeled the optomotor response to wide-field point-cloud stimuli both rotating (roll, pitch, yaw) and translating (thrust, slip, lift) in each of the primary axes using binary noise sequences (impulses in optic flow velocities) and a time-domain correlation analysis to recover the impulse response function.

Poggio and Reichardt proposed a framework for modeling the optomotor response in flies (*Musca domestica*) [Reichardt and Poggio(1976)]. The framework assumed a model structure composed of two parts: a second-order model of turning dynamics (governing the transformation of torque to angular position, velocity and acceleration) and a control policy prescribing torque as a function of the time-varying visual scene. The control policy was further parsed into a position-dependent function of the scene, a velocity-dependent function of scene motion, and noise. In control theoretic parlance, Poggio and Reichardt proposed a second-order plant stabilized by a PD controller. The analysis, however, focused on the stochastic nature of the differential equation, inferring the model through the propagation of motor noise given different stimulus conditions (*e.g.* salience or noise) absent any reference trajectory. The framework was comprehensive in the sense that it permitted visual scenes with different spatio-temporal properties. However, in focusing on the stochastic over the deterministic contributions of the behavior, the analysis becomes complex, unintu-

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

itive and uninviting. Most importantly, assuming an *a priori* model structure, it is not apparent how Poggio and Reichardt’s model could be used to generate new hypotheses and inform new biological experiments.

### 4.1.3 Our approach

We address the system identification of stripe-fixation through a series of perturbation experiments and frequency-domain analyses. In the frequency domain, the deterministic responses to sinusoidal perturbations are easily parsed from the spectrum of the motor output. No model structure is assumed *a priori*, but we fit a proportional-integral-derivative (PID) control model to our empirical frequency response data. Notions of PD and PID control have been referenced in past work [Heisenberg and Wolf(1988),Ristroph et al.(2010)]. Though recent optic flow models disregard the contribution of positional cues (let alone the integral of positional cues); the PD and PID structures furnish biological hypotheses as to the contributions of different sensory information and we use our derived model to consolidate observations from previous studies [Duistermars et al.(2007), Geiger and Nässel(1982), Maimon et al.(2008)]. In terms of the measured frequency bandwidth and the generalizability of the resultant linear model to prior published observations, this is perhaps the most comprehensive task-level model of the stripe-fixation behavior to date. Lastly, we revisit work by Heisenberg and Wolf in which a comparison between closed- and open-loop responses revealed surprising qualitative differences in performance between the two conditions.

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

They attribute the observed discrepancies to the *reafference principle* (summarized in Section 4.4.3) [von Holst and Mittelstaedt(1950)]. We recreate the replay experiments and proffer a parsimonious alternative to the reafference principle.

### 4.2 Materials and Methods

In order to constrain the dimensionality of the identification problem, we use a tethered preparation in which flies are rigidly attached to a fixture and presented with a fictive visual stimuli. This preparation has become a popular experimental paradigm for studying this class of behaviors. In free flight, naturalistic visual stimuli are extremely rich (*e.g.* luminance, contrast, spatial content, object motion, motion coherence, etc. can all be salient factors which contribute to the response) and the motor output is described kinematically with (at least) six degrees of freedom (three for rotation and three for translation, not to mention head, body, and limb movements). In the tethered preparation, the visual stimulus and the kinematic output can each be constrained to a single degree of freedom, hence reducing the task-level locomotor behavior to a single-input–single-output system [Cowan et al.(2006), Cowan and Fortune(2007), Roth et al.(2011)].

### 4.2.1 Animal husbandry and preparation

Large adult fruit flies, female *Drosophila melanogaster* at 2-3 days post-eclosion, were selected for all experiments. Flies were anesthetized at temperatures of 3-4°C and tethered to tungsten wire at the anterior end of the thorax using UV-cured cement. Additionally, the head was glued to the thorax to eliminate movement during experiments, fixing the head to the stationary frame of the arena.

### Experimental apparatus

The flight arena (Figure 4.2) was composed of 44 modular LED panels (each panel consisting of an  $8 \times 8$  grid of LEDs) [Reiser and Dickinson(2008)], arranged in a cylinder (4 panels high and 11 around the circumference) subtending 330 deg with a 30 deg gap at the rear. In this configuration, the circumferential pixel-to-pixel distance was 3.75 deg. Four levels of gray-scale dithering allowed for apparent motion with a minimum increment of 0.9375 deg. The visual scene consisted of two dark stripes, positioned antipodally on the cylinder, each subtending 30 deg. Positional error was measured as the angle from the fly's sagittal plane to the front-most stripe considered the fixation target. The fly was positioned in the center of the arena. Illuminated from above with infrared light, the fly cast a shadow on a sensor below; the wingbeat analyzer inferred wingbeat frequency (WBF) and left and right wingbeat amplitude (LWBA and RWBA) from the spatiotemporal pattern of the shadow. The



CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

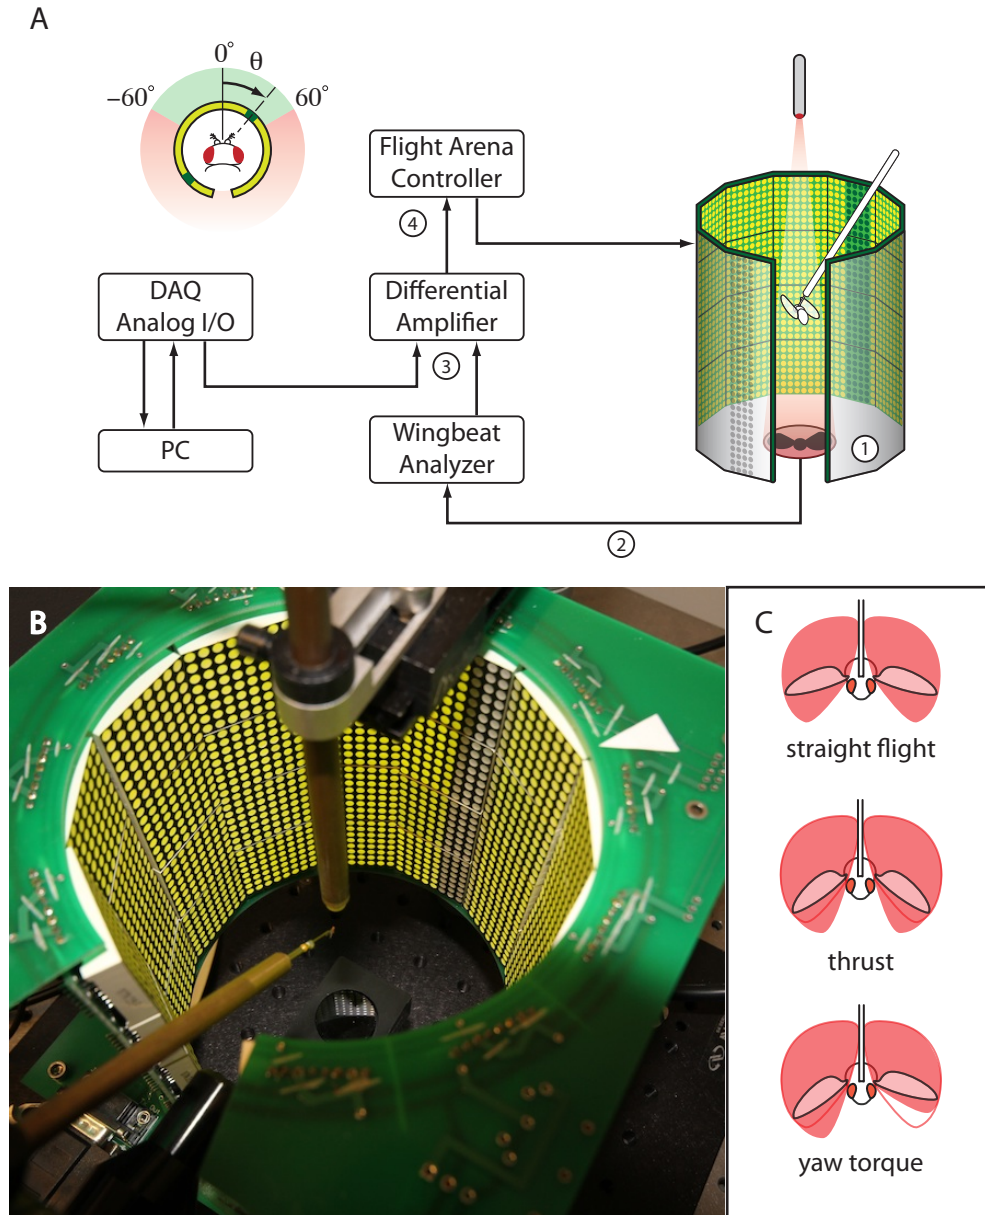


Figure 4.2: (A) (1) The LED arena displays an oscillating vertical bar, eliciting a tracking response from the fly. (2) Illuminated from above, the fly casts a shadow on a sensor below and the wingbeat analyzer extracts the salient kinematic descriptors from this wingstroke silhouette. (3) The differential amplifier subtracts the fly motor output ( $\Delta$ WBA) from the reference trajectory (prescribed in software) to generate the error velocity (sensory slip). (4) The Flight Arena controller integrates velocity error and displays an updated stripe position. (B) A photograph of the arena shows the fly suspended above the WBA sensor which measures changes in wing stroke kinematics illustrated in (C).

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

wingbeat amplitude asymmetry ( $RWBA - LWBA = \Delta WBA$ , measured in volts) was assumed to be proportional to the fly’s intended yaw torque [Tammero et al.(2004)] which in turn is proportional to the yaw angular velocity as per Section 4.1.1. We implemented this as

$$\dot{\theta} = K \cdot \Delta WBA, \quad (4.1)$$

where  $K = 239.1 \text{ deg} \cdot \text{s}^{-1} \cdot \text{V}^{-1}$  was tuned by hand to achieve a robust, closed-loop stripe fixation behavior.

The velocity error signal was calculated as the difference between the reference trajectory velocity and the wingbeat asymmetry; the flight arena controller integrated the velocity error signal (calculating the positional error) and updated the corresponding stripe image.

### 4.2.2 Experiment design

#### System identification assay

Flies ( $N = 10$ ) were observed stabilizing an array of moving-stripe stimuli comprising pure sinusoidal, sum-of-sines and chirp trajectories. Between presentations, flies were given a simple “reward” fixation task intended to maintain motivation and provide consistent initial conditions. This set of trajectories was repeated three times with trial order randomized within an iteration; for the large majority of trajectories, flies maintained the tracking behavior through all three presentations.

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

Sinusoidal trajectories were presented at frequencies of 1, 3.5, and 11.5 Hz for durations of 6 s; sum-of-sines trajectories encompassed every pair-wise sum of sinusoids from this same set of frequencies and the same duration. The frequencies were selected to be mutually prime, so that for sum-of-sines trials, the harmonics of the response to any frequency component would not coincide with any other stimulus (fundamental) frequency. The logarithmic chirp stimulus spanned frequencies from 0.05-11.5 Hz with frequency increasing continuously over a 120 s duration. The angular amplitude  $A(\omega)$  (and consequently the maximum angular velocities <sup>1</sup>) for all trajectories was selected as a function of frequency (in  $\text{rad} \cdot \text{s}^{-1}$ ):

$$A(\omega) = (0.0153 \text{ deg}^{-1} + 0.0044 \text{ s} \cdot \text{deg}^{-1}(\omega))^{-1} \text{ deg} . \quad (4.2)$$

This relation prescribed an angular amplitude that decreases with increased frequency but an angular velocity that increases with frequency, addressing physical limitations of the animal behavior and constraints imposed by the LED display. At high frequencies, the velocity was constrained from above by the motor output of the fly (free-flight saccadic rotations are estimated at  $1800 \text{ deg} \cdot \text{s}^{-1}$  [Fry et al.(2003)] and the limitations on smooth tracking are significantly lower) and the angular amplitude was constrained from below by the resolution of the LED arena (peak-to-peak travel must be greater than a single pixel). At 11.5 Hz, this relation yielded a peak-to-peak amplitude of 6 deg (approximately 1.5 pixels translation) and a maximum velocity of

---

<sup>1</sup>Angular amplitude and angular velocity refer to the spatial coordinates on the cylindrical arena, described in units of degrees or  $\text{deg} \cdot \text{s}^{-1}$  respectively. The frequency of oscillation of the moving image will be denoted (unconventionally) by  $\omega$  and measured in units of Hz or  $\text{rad} \cdot \text{s}^{-1}$ .

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

216.8 deg  $\cdot$  s<sup>-1</sup>. At low frequencies, we bounded the amplitude of the exogenous reference to 60 deg, though the displayed error signal was allowed to exceed this bound as a result of feedback.

### **Replay paradigm**

In 1988, Heisenberg and Wolf explored the role of feedback in processing exogenous (reference) and reafferent (self-generated) motion stimuli using an error replay experiment. This experiment paradigm compares the responses to two sequential perturbations, first in closed-loop and then in open-loop. In the first presentation, the fly's generated torque (or for our set-up, the  $\Delta$ WBA) stabilizes the stripe position through negative feedback and the error signal, the displayed motion stimulus, is recorded. In the second presentation, this recorded error is replayed as the motion stimulus in the absence of any stabilizing feedback (*i.e.* the stimulus is unaffected by the fly motor output). The replay experiments were recreated for oscillations of 0.1 Hz (as used in [Heisenberg and Wolf(1988)]) with the same visual stimulus and position-frequency relation (4.2) used in the system identification assay. Flies ( $N = 10$ , distinct from the sample used for identification) were presented three periods (30 s) of oscillation in closed-loop, followed by 3 s of reward fixation and then the open-loop replay. Flies completed six repetitions of closed- and open-loop pairings.

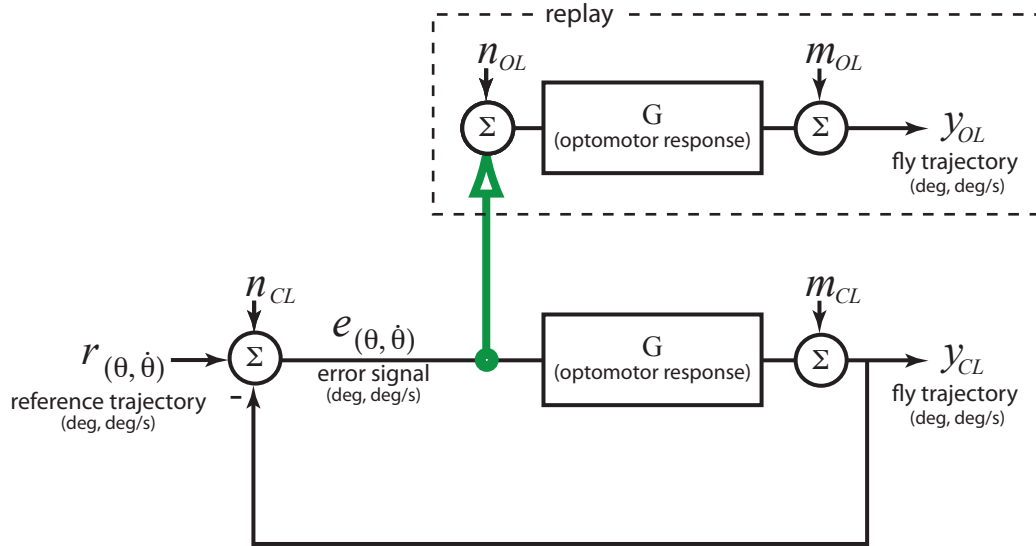


Figure 4.3: *The replay experiment.* In this sequence of experiments, the fly is first presented with a stripe fixation task in closed-loop. The error signal (the displayed visual scene for a tethered preparation) is recorded in the closed-loop presentation and subsequently replayed to the fly without feedback stabilization. For a stable linear system, the motor outputs in closed-loop and replay,  $y_{CL}$  and  $y_{OL}$  respectively, should be qualitatively similar excepting  $y_{OL}$  will be noisier.

### 4.2.3 Analysis

The frequency analysis and subsequent linear transfer function model faithfully describe the stripe-fixation behavior *for some neighborhood* about the fixation equilibrium (stripe position and velocity are both zero). Acknowledging this caveat, we discard data points for which the stripe is not reasonably frontal, those for which the magnitude of the error signal exceeds 60 deg (Figure 4.2 *upper left*). Additionally, we ignore data for which the instantaneous wingbeat frequency falls below the mean WBF (as calculated per individual across all presentations), eliminating data for which the fly was fatigued or otherwise unmotivated.

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

The discrete-time Fourier transform (DTFT) of the complete data reveals the coincidence of peaks in the input and output power spectra (*i.e.* the fly motor output has significant power at the stimulus frequencies). The DTFT, however, is not amenable to incomplete or irregularly sampled data sets. In lieu of the DTFT, we apply a least squares spectral analysis (LSSA) to both the input and output signals. For pure sinusoids and sum-of-sines trajectories, we fit the coefficients to cosine–sine pairs only at the frequencies of interest (just as in (3.1)) but only for the admissible data samples (at times  $t \in T$ ):

$$(\alpha, \beta) = \arg \min_{\alpha, \beta} \sum_{t \in T} \left\{ \left( \sum_{i=1}^k \alpha_i \sin(\omega_i t) + \beta_i \cos(\omega_i t) \right) - y(t) \right\}^2. \quad (4.3)$$

The magnitude and phase are calculated as in (3.3). For chirp signals, we perform a short-time LSSA, with non-overlapping windows of 10 s for  $t \in (0 \ 40]$  and 4 s for  $t \in (40 \ 120]$ . The least-squares minimization in (4.3) is modified to accommodate time-varying frequency and signal amplitude where  $\theta(t)$  and  $\omega(t) = \frac{d\theta}{dt}$  are prescribed and  $A(\omega)$  is determined by (4.2). The short-time LSSA assumes that for sufficiently small deviations in frequency (corresponding directly to window size for the chirp stimulus) the frequency response  $F(\omega)$  is nearly constant. This assumption fails if the system response changes drastically within a small frequency band (*e.g.* near the

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

natural frequency of an underdamped system):

$$(\alpha, \beta) = \arg \min_{\alpha, \beta} \sum_{t \in T} \left\{ \alpha A(\omega) \sin(\theta(t)) + \beta A(\omega) \cos(\theta(t)) - y(t) \right\}^2. \quad (4.4)$$

A parameterized transfer function model  $F(\omega, \alpha)$ , a function of frequency  $\omega$  and parameters  $\alpha$ , is fit in the least-squares sense to the empirical frequency response  $F^*(\omega)$  (at the discrete frequencies in the set  $\Omega$ ) represented in the log space (as discussed in Section 2.2.2):

$$\alpha = \arg \min_{\alpha} \sum_{\omega \in \Omega} \omega^{W_1} \left\{ W_2 (\log |F_{\Delta}(\omega, \alpha)|)^2 + (1 - W_2) (\angle F_{\Delta}(\omega, \alpha))^2 \right\},$$

where

$$F_{\Delta}(\omega, \alpha) = \frac{F(\omega, \alpha)}{F^*(\omega)}. \quad (4.5)$$

The weight  $W_1 \in \mathbb{R}$  is a free parameter chosen to favor data according to frequency; negative values of  $W_1$  more heavily penalize low frequency errors and positive values penalize errors at high frequencies. A non-negative weighting function  $W_1(\omega)$  could be substituted for the term  $\omega^{W_1}$  to more selectively weight frequency bands (*e.g.* to favor capturing a particular feature in the data or to down weight noisy or unreliable measurements).  $W_2 \in [0, 1]$  scales the relative costs contributed by gain and phase.

## 4.3 Results

### 4.3.1 Empirical frequency responses

We describe the system as an input–output relation from the stripe position (measured as the azimuthal angle from the sagittal plane of the fly) to the fly motor output (measured as a voltage proportional to the difference in wingbeat amplitudes,  $\Delta\text{WBA}$ ). An example of the motor response (*green*) to a sum-of-sines trajectory (*blue*) is shown in Figure 4.5A, both in the time domain (top) and the frequency domain (bottom). Though the fly motor output has significant stochastic components (noise and drift), comparing the magnitude of the frequency spectra of the input–output pair as calculated by a fast Fourier transform, allows us to disambiguate the response to the moving stimuli from motor noise. For this trial, the stimulus trajectory is a sum of sinusoids at frequencies of 1 and 3.5 Hz. In the frequency domain, the fly’s motor output comprises two distinct peaks coinciding with the stimulus frequencies; we consider these peaks to be the response to the stimulus and disregard other spectral content as extraneous motion. Though in many trials, these response peaks are easily discernable, this is not always the case. We proceed under the assumption that for each trial, the response to the stimulus is the response measured at the stimulus frequencies. In the remaining analyses, we use LSSA (see Section 4.2.3) to extract only the frequency components of interest (those present in the stimulus), un-



CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*




Figure 4.4: In the frequency domain (*bottom*), the motor response (*green*) to stimulus motion (*blue*) is discernable from extraneous motor noise.

derstanding that, at times, the response data may be highly corrupted by extraneous motor output, or worse, that a responsive behavior is not present at all and we are sampling only stochastically driven motor output.

In the first suite of experiments, flies were recorded fixating a dark vertical stripe oscillating with sinusoidal, sum-of-sines, and logarithmic chirp trajectories. The empirical FRF (shown in *black* in Figure 4.5B) is calculated from the chirp stimulus using a moving window LSSA over the frequency range 0.1–11 Hz. The response to pure sinusoids (*blue*) and sums-of-sines (*green*) are superimposed on the FRF. The response to sinusoids coincides well with the chirp response; this is expected since for any small window of time (small with respect to the rate of change of frequency), the chirp stimulus resembles a pure sine *and* for any small band of frequencies (with caveats discussed in Section 4.2.3) the system response should be similar. More

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

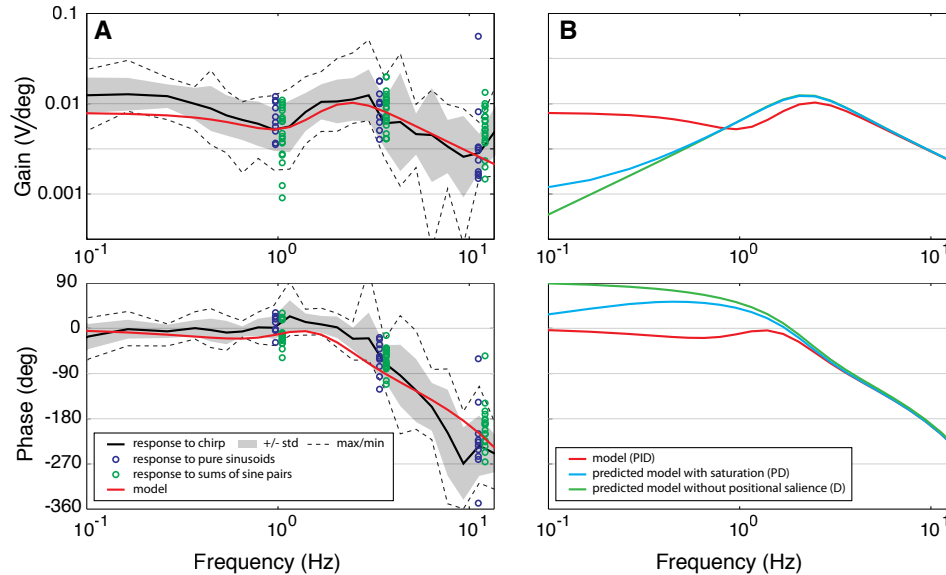


Figure 4.5: (A) The FRF empirically calculated from chirp stimuli is compared to the responses to single and summed sine waves. Agreement between these responses motivates a linear transfer function model (*red*). (B) A saturating integrator model results in a PD model with phase and gain differences in accordance with Heisenberg and Wolf’s replay experiments [Heisenberg and Wolf(1988)]. A derivative model relying only on optic flow generalizes the model for wide-field stimuli.

importantly, the sums-of-sines responses provide frequency response estimates consistent with the empirical FRF, substantiating the superposition property and further justifying linear modeling.

### 4.3.2 A linear transfer function model

The FRF reveals low-pass dynamics, as has been repeatedly observed and reported in the literature [Duistermars et al.(2007)]. The high-frequency decay (estimated as  $18.7 \text{ dB} \cdot \text{decade}^{-1}$  from measurements at 3.65 and 9.33 Hz) suggests a transfer function with relative degree of one. Phase rolls off faster than expected for the assumed

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

relative degree, which likely results from a delay in sensorimotor processing. Also notably, the gain plot exhibits a distinctive notch at 1 Hz. In order to capture these salient features, a linear model would require at minimum three poles, two zeros and a delay. The following transfer function from positional error  $E(s)$  to motor output  $Y(s)$  was fit to the empirical FRF as described in Section 4.2.3:

$$\frac{Y(s)}{E(s)} = \exp\{-0.032s\} \frac{0.181s^2 + 1.23s + 8.68}{s^3 + 20.6s^2 + 277s + 1098}. \quad (4.6)$$

The transfer function has a real pole at -5.72 and complex poles at  $-7.43 \pm 11.70i$ . Dividing numerator and denominator polynomials by the real pole, we reformulate (4.6) into a proposed PID control model:

$$\begin{aligned} & \underbrace{(s^2 + 14.85s + 192.1)}_{\text{motor dynamics}} Y(s) & (4.7) \\ & = \exp(-0.032s) \underbrace{\left( 0.181s + 0.196 + \frac{7.55}{s + 5.72} \right)}_{\text{sensory weighting - PID (leaky integrator)}} E(s). \end{aligned}$$

While the FRF is unique for any given data set, many such parameterized models may be hypothesized, each furnishing a different mechanistic interpretation. In fact, there are fundamental limits in our ability to tease apart sensory dynamics from downstream processing due to the likelihood of pole-zero cancellations [Carver et al.(2009)]. Notwithstanding this limitation, we parse the transfer function so that the left-hand and right-hand side equations describe the motor dynamics and sensory

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

processing, respectively. Under this interpretation, the motor plant is described as a forced second-order system (analogous to a spring-mass-damper) driven by a neural control signal (Figure 4.6). In the right-hand side equation, the control signal is calculated as a weighted sum of sensory measurements of sensory slip, positional error, and accumulated error. The best fit visual-motor delay is 32.1 ms; this delay has been previously estimated at 40 ms for *Drosophila* [Hardie and Raghu(2001), Heisenberg and Wolf(1988)].

### 4.4 Discussion

The empirically derived model deviates from the prevailing optic flow models in the contribution of positional error in the control policy. While the positional (and integral) terms of the PID controller may be instantiated as integrals of a measured velocity (or optic flow), we believe that position is a distinct measured quantity (referred to as the flicker hypothesis in [Poggio and Reichardt(1981)]). Consider two stripe-fixation tasks administered with the same velocity trajectory but different initial positions; we've observed that flies compensate for the positional bias and both stripes are fixated. Optic flow models have been formulated to reflect physiological representations of optic flow in the lobula plate. However, this physiological evidence *for* optic flow does not preclude the contribution of positional cues to optomotor control (and *should not* preclude position from our models). Our PID model consolidates several previously published observations.

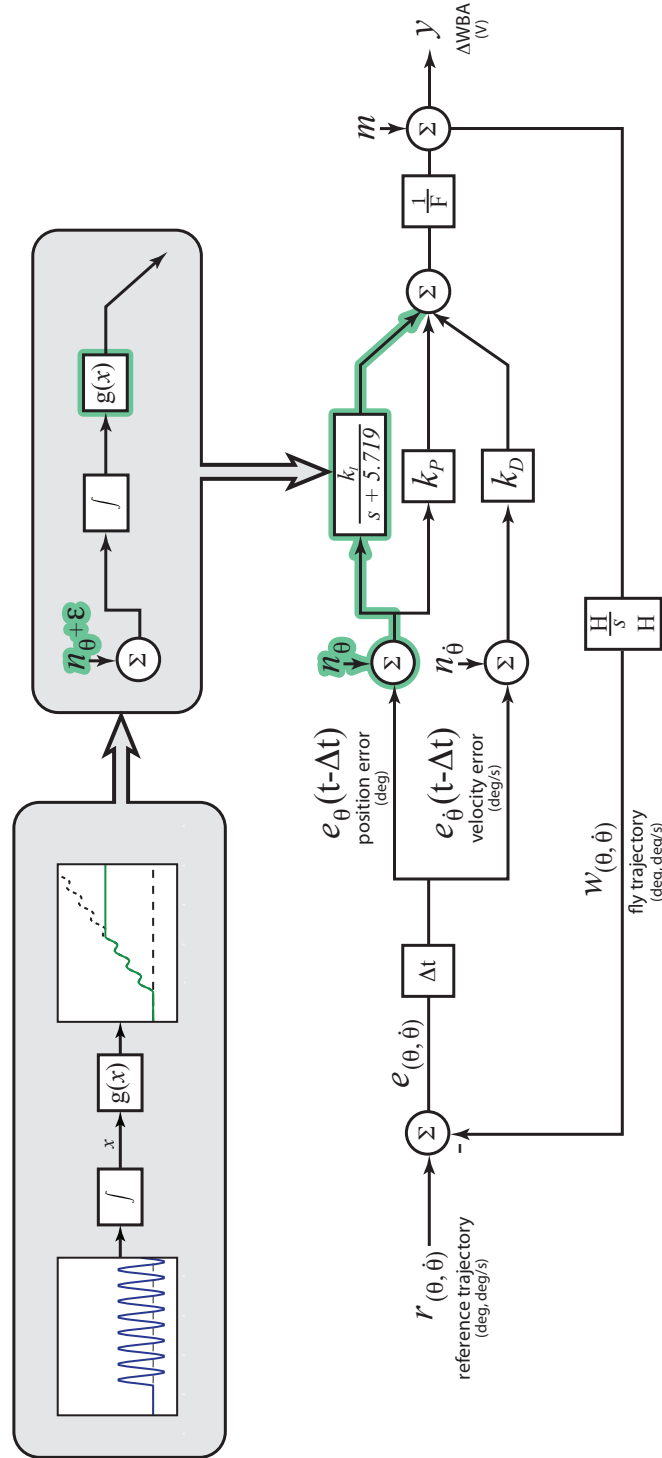


Figure 4.6: A block diagram illustrates the PID hypothesis proposed in (4.7). In the grey insets, we substitute the leaky integrator with a pure integrator and saturating nonlinearity to address differences between closed- and open-loop responses (discussed in Section 4.4.3).

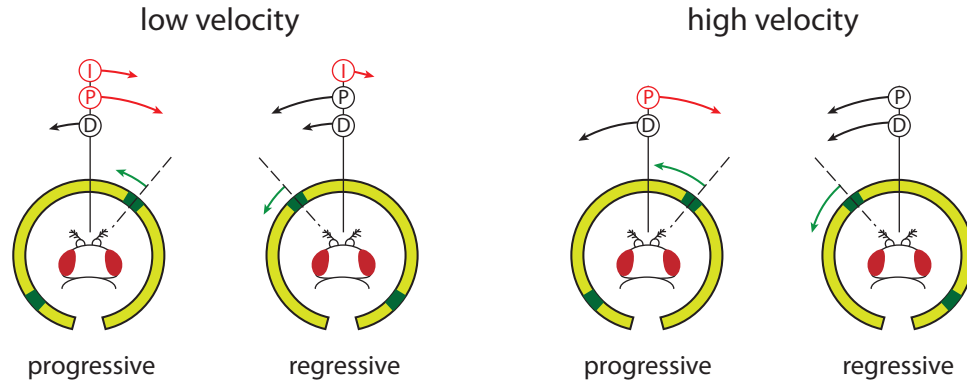


Figure 4.7: The PID control hypothesis explains the disparity in torque output in response to progressive and regressive stripe motion. For progressive motion, the proportional and derivative channels contribute opposing terms; in response to regressive motion, these terms add constructively. At low velocities, the integral term mitigates the difference between the progressive and regressive torque responses.

#### 4.4.1 Progressive-regressive response concurs with PID control

Flies exhibit increased motor output for regressive (front-to-back) stripe motion compared to progressive (back-to-front) motion (this difference motivated the progressive-regressive control hypothesis described in [Poggio and Reichardt(1981)]). The disparity in torque production in response to progressive and regressive motion stimuli becomes more pronounced with increased velocity. This difference has been observed in behavioral experiments [Geiger and Nüssel(1982), Maimon et al.(2008)] and in physiological measurements from the lobula complex [Buchner et al.(1984)].

This phenomenon is consistent with PID control policies. For the a leftward moving stimulus (as in Figure 4.7), we can decompose the total torque response as the sum of the individual contributions of the proportional, integral and derivative chan-

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

nels of the PID controller. The proportional channel contributes a torque command proportional to the position of the moving stripe and independent of stripe velocity; in the illustration below, the torque would be rightward during the progressive portion and leftward for regressive motion. The derivative response would be constant (and commensurate to the velocity) throughout, contributing a leftward torque in for both the progressive and regressive portions. Conversely, the integral channel only contributes a rightward torque; the positional error accumulates during the progressive motion portion, reaches a maximum when the stripe crosses the zero degree axis, and then subsequently wanes with regressive motion.

The difference in responses is largely explained by the relative torque contributions of the derivative (velocity) and proportional (position) channels of the PID controller. In response to regressive motion (*e.g.* a stripe to the left moving leftward), the proportional and derivative responses contribute constructively to the motor output (Figure 4.7); conversely, for progressive motion (*e.g.* a stripe to the left moving rightward), the proportional and derivative responses contribute opposing torque commands yielding a reduced net torque. As velocity increases, the derivative channel contributes a more significant torque command, increasing the magnitude of the response to regressive motion and diminishing the response to progressive motion, widening the gap between the two responses. At low velocities, the integral term contributes a significant rightward torque which increases the response to progressive motion and reduces the response to regressive motion, narrowing the bias between

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

the two responses. As velocity increases, positional error is integrated over shorter durations, attenuating (or eliminating) the contribution of the integral channel (Figure 4.7 *right*). The velocity dependence of both the derivative and integral terms explains the experimentally observed trends in the responses to these two categories of stimuli.

### 4.4.2 Positional cues explain differences between wide-field and small-field responses

Physiologically, wide-field and small-field scenes elicit responses in different strata of the medulla [Bausenwein and Fischbach(1992)]. Behaviorally, Duistermars *et al.* observed that for wide-field (repeating stripe with period of 30 deg) visual scenes at 0.1 Hz, motor output was phase advanced to the stimulus motion by a quarter period while for small-field (single stripe) the input and output phases were practically locked. Assuming that this spatial periodicity suppresses the majority of positional and integral error (modulo 30 deg), we propose a derivative-only (optic flow) model for yaw regulation of this scene (Figure 4.5C in *green*); as observed by Duistermars *et al.*, at low frequencies this response to the wide-field scene phase leads the single-stripe (PID model) response by approximately 90 deg. We note that our model is inconsistent with their observations of response gain. However, the spatial extent of the motion stimulus could be a factor in modulating the gain of the response. This factor is not captured in the proposed model (generated using small-field stimuli



CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

exclusively) and may account for the difference between the model prediction and the experimental observations.

### 4.4.3 Replay experiments yield predicted results

For linear *stable* optomotor dynamics, the response to the replay experiment (as described in Section 4.2.2) is expected to be qualitatively similar to the original closed-loop response excepting that in open-loop, there is no mechanism for the attenuation of measurement noise:

$$\begin{aligned}
 Y_{CL} &= \frac{G}{1+G}(R + N_{CL}) + \frac{1}{1+G}M_{CL} \\
 E &= \frac{1}{1+G}(R + N_{CL}) - \frac{1}{1+G}M_{CL} \\
 Y_{OL} &= \frac{G}{1+G}(R + N_{CL}) - \frac{G}{1+G}M_{CL} + G N_{OL} + M_{OL}
 \end{aligned} \tag{4.8}$$

At frequencies for which tracking performance is good,  $\frac{G}{1+G} \approx 1$  implies  $|G| \gg 1$ . As a result, the open-loop replay is extremely sensitive to the process noise,  $N_{OL}$ , and measurement noise attenuated in the closed-loop presentation,  $M_{CL}$ , reappears undiminished in the replay response. But Heisenberg and Wolf [Heisenberg and Wolf(1988)] observed qualitative differences between the closed-loop and replay responses, not merely noise: replay responses were phase leading by 50.9 deg (SEM = 48.6 deg,  $N = 10$ ) and attenuated by approximately 50-65%. This discrepancy between the closed-loop and replay responses was presented as evidence corroborating the reafference principle.

## CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

The reafference principle was proposed by von Holst and Mittelstaedt [von Holst and Mittelstaedt(1950)] to address a specific problem: if reflexive behaviors yield stabilization to some equilibrium, how do animals make voluntary motions away from equilibrium? The prevailing theory was that, during voluntary motion, stabilizing reflexes were inhibited. But this is not the case, since animals exhibit stabilizing reflexes concurrent with voluntary motion. To resolve this problem, von Holst and Mittelstaedt proposed that efferent signals cause an imbalance in state which is stabilized by the reflex (that is voluntary efferents introduce a new set point which is stabilized to by the same mechanism as reflexive behavior). An efference copy is compared to the afferent stream; the residual (error) distinguishes exafference (perception of external motion) from reafference (perception of egomotion). Nowadays, efference typically refers to a feedforward computation of the state estimate (like the use of the Kalman filter for prediction discussed in Section 3.4.2). However, the reafference framework proposed by von Holst and Mittelstaedt can be posited as a feedback system (without feedforward) [Powers(1989)], so it is often difficult to confirm efference copy through external observation, without direct measurement of the efferent signal.

Heisenberg and Wolf proposed that, in replay, the mismatch between perceived motion and efferent commands for volitional movement triggered a switch in behavior mediated by a higher center. This switch results in gated intervals of reflex suppression (a hybrid of the reafference principle and the reflex theory it supplanted). Further, they concluded that such differences could not be the consequence of a linear

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

system and could not be identified by a frequency domain analysis [Heisenberg and Wolf(1988)]:

...the modulation of yaw torque in response to the sinusoidal oscillation of visual patterns is influenced by the frequency and polarity of body saccades and by the presence or absence of reafferent visual stimuli. A schematic input–output analysis using a continuum of oscillation frequencies and amplitudes would thus not lead to satisfactory characterization of the optomotor controller.

Contrary to this hypothesis, in recreating the replay experiments, we achieved the results predicted by a linearity assumption. Specifically, open-loop responses were qualitatively similar to closed-loop observations but with greater noise, hence variability (Figure 4.8). Though increased variance in open-loop observations makes it difficult to argue whether Heisenberg and Wolf’s observed differences were significant, we proffer an explanation derived from our proposed PID model. If we allow the integrator to saturate (a nonlinearity inherent in any physical instantiations of an integrator), then the system would respond quite differently about equilibrium when compared to the saturated regime.

The predictions in (4.8) are predicated on the assumption that the system is both linear and stable. Though our hypothesized model uses a leaky integrator (low-pass filter), a pure integrator would violate this assumption: integrators are not stable (only marginally stable in the asymptotic sense and unstable in the bounded-input–bounded-output sense). Moreover, any physical realization of an integrator must have a saturation limit, since real signals must be bounded. In closed-loop, every internal state is maintained at equilibrium and as a result, the integrator contributes to the

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

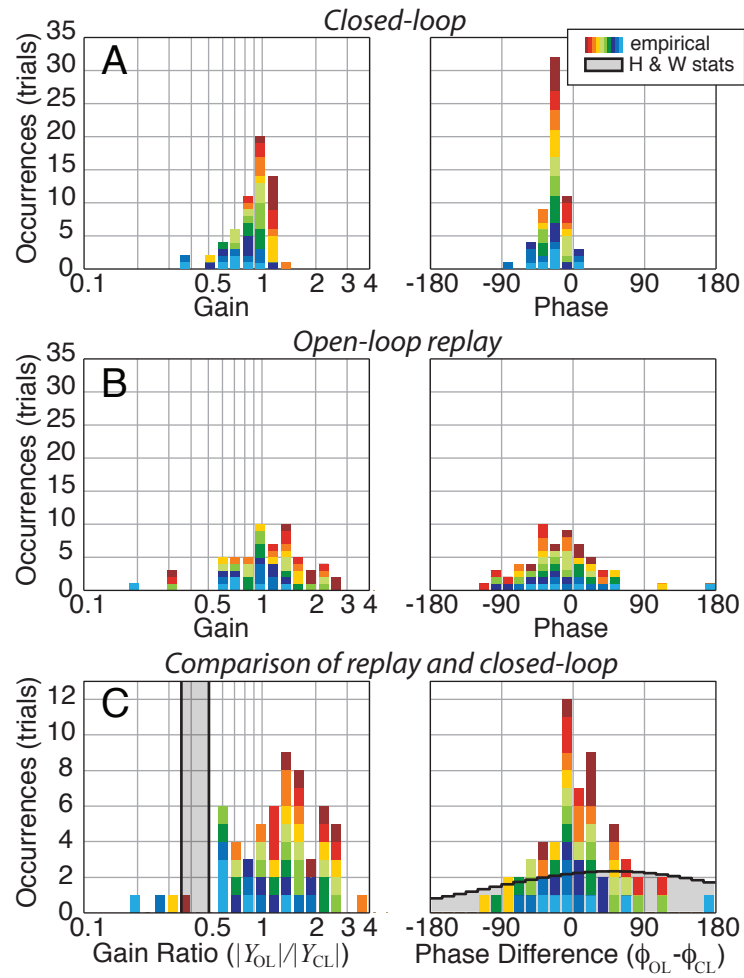


Figure 4.8: Empirical histograms of gain and phase for (A) closed-loop and (B) open-loop replay trials illustrate the sensitivity relation between the two feedback conditions. Colors correspond to trials from different individuals with the warmer colors assigned to flies with higher mean gain. (C) We present the ratios of the two responses and compare the results to Heisenberg and Wolf’s reported statistics [Heisenberg and Wolf(1988)].

CHAPTER 4. OPTOMOTOR YAW REGULATION IN *DROSOPHILA MELANOGASTER*

output proportional to the accumulated error (which varies with time). In open-loop, small biases in noise or initial condition (*e.g.* measurement bias, asymmetry in baseline wingbeat amplitudes, fly misalignment, etc), which would be mitigated in closed-loop, can cause the integrator to drift to saturation. Once saturated, the integrator no longer conveys changes in the accumulated error, contributing a constant offset to the motor output, effectively silencing the integral path in sensory processing.

Figure 4.5B compares our proposed PID model (with leaky integrator) (*red*) to a model prediction for a saturating integrator in the saturated state (*cyan*); over a broad range of low frequencies, the saturated model exhibits attenuation and phase-lead, the same differences observed by Heisenberg and Wolf in comparing replay and closed-loop responses. This is a weak illustration using our hypothetical model, but many control topologies could yield this effect. For example, suppose that the PID control was implemented as parallel PI controllers on sensed position and velocity signals. The position-dependent term would then be a weighted sum of the measured position signal and the integrated velocity signal (Heisenberg and Wolf assert such an integrator in [Wolf and Heisenberg(1990)]). In open-loop, the integrated velocity signal would saturate, effectively reweighting the gain on position. In summary, instability and nonlinearity *can* yield qualitatively different results in closed- and open-loop, reiterating the importance of closed-loop observation of behavior.

# Chapter 5

## Conclusion

In this dissertation, we have proposed a general framework for system identification of image-stabilization tasks in animals. We prescribe the assay of closed-loop experiments and analyses required to elicit these behaviors, differentiate between stimulus-mediated behaviors and ancillary responses (and between behavior and non-behavior) and generate predictive and generative models. The proposed methods were demonstrated on two model image-stabilization behaviors, stripe fixation in *Drosophila* and refuge-tracking in *Eigenmannia*. In both applications, we discovered aspects of the behavior which were contrary to accepted notions.

For these closed-loop behaviors, we expect the dynamics to be (approximately) linear local to an equilibrium, the consequence of Hartman-Grobman linearization theory [Sastry(1999)]. This is not the case for the refuge-tracking behavior in glass

## CHAPTER 5. CONCLUSION

knife-fish. Comparing responses to single-sine and sum-of-sine trajectories, we observe that superposition failed and demonstrate that the qualitative differences persisted for stimuli at different scales of velocity. The behavior is fundamentally nonlinear for any neighborhood about the equilibrium. The nature of the differences in performance suggest a predictive mechanism and optimal control policy. No prior evidence for this prediction mechanism has been reported in the literature.

Unlike for fish, in fruit flies, the stripe-fixation responses to sinusoids, chirps and sums-of-sines trajectories corroborate the linearity assumption. Fitting a transfer function model to the frequency response function, we see that positional cues (and integral of position cues) from the visual scene play a significant role in determining the fly's motor output. This is not the first mention of the role of position-dependent contributions in optomotor control, but the prevailing models rely primarily (if not solely) on velocity information. The proposed PID model also generalizes to explain several previously observed phenomena.

Task-level descriptive models can serve to constrain mechanistic models. In a bottom-up approach, mechanistic models are constructed from known building blocks. Top-down models not only serve to define permissible mechanisms, but identify necessary computations. That is to say, task-level models can both falsify a proposed mechanism *and* provide insight into what missing building blocks must do/be.

## 5.1 Future work: A Systems Approach to Modularity

The reductionist approach aims to describe a system mechanistically, as the joint contribution of a set of more elementary components. In many ways, reductionism reflects the goals of behavioral biology and neuroscience. But in Section 1.3.1, we illustrate the drawbacks of a reductionist methodology. Towards the reductionist aim, how might we identify behavioral building blocks in a top-down approach?




Figure 5.1: A sensorimotor transform from input  $R$  to output  $Y$  is split into arbitrary blocks  $A$  and  $B$  by a measurement (or injected perturbation) at  $S$ .

Multi-input–multi-output experimentation is a natural extension of this research and can be used to parse the sensorimotor transform into constituent components. Consider the contrived sensorimotor block diagram in Figure 5.1 which maps input  $R$  to output  $Y$ . In the diagram, the sensorimotor transform is divided into constituent models  $A$  and  $B$  by some measurement (kinematic or electrophysiological as discussed in Section 1.3.2) or perturbation (sensory or mechanical)  $S$ , suppose a perturbation. We could fit models for the sensorimotor transform  $Y/R$  as well as the response to the perturbation  $Y/S$  and derive the constituent models, as follows:



CHAPTER 5. CONCLUSION

$$\left. \begin{array}{l} Y/R = \frac{AB}{1+AB}, \text{ given } S = 0 \\ Y/S = \frac{B}{1+AB}, \text{ given } R = 0 \end{array} \right\} \rightarrow \begin{array}{l} A = \frac{Y/R}{Y/S} \\ B = \frac{Y/S}{1-Y/R} \end{array} \quad (5.1)$$

The derivation follows similarly for a measurement  $S$ . Extrapolate this example by adding more measurements and perturbations (either simultaneously or in a sequence of experiments), and the sensorimotor transform begins to reveal its ingredients. Further, by measuring or perturbing within the loop, it is possible to observe or excite dynamics which are not discernable elsewhere in the system (in control theory, this results from a pole-zero cancellation between constituent models [Carver et al.(2009)]). Additionally, biomechanical models derived from physical principles, measured empirically, or computed through numerical simulation could be used to naturally and intuitively parse the sensorimotor transform.

Applying this top-down methodology, the goal is to progress to increasingly low-level descriptive models. But what are the natural boundaries between models? How might we parse the sensorimotor transform into functional modules (individually packaged descriptive models) in a meaningful way? In the reductionist approach, atomic mechanistic models (often observed from some physical biological mechanism) are networked to synthesize phenomena. In contrast, the top-down functional modules are simply those models which recur across a spectrum of similar behaviors; they are not necessarily the building blocks of behavior or even representative of a particular physical structure. Common to the two approaches is the notion that individual

## CHAPTER 5. CONCLUSION

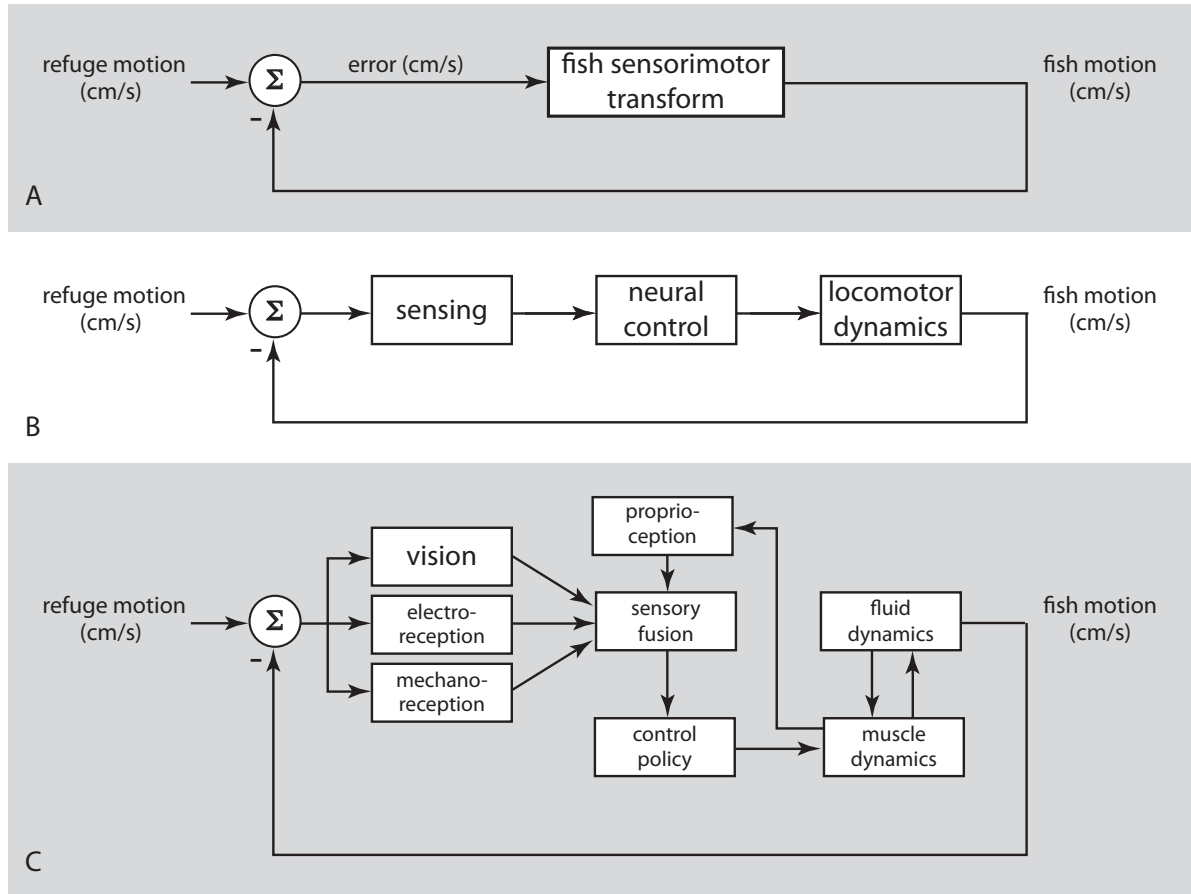


Figure 5.2: Successive decompositions of the refuge-tracking behavior.

modules contribute to many behaviors. Using commonality as the defining notion for a functional module, identifying a parsimonious set of functional modules and the topological relations between them requires intra-behavioral comparisons. In any animal, many categorically different behaviors recruit the same sensory and locomotor modalities. A parallel analysis of these behaviors could reveal salient functional modules.

Assuming we have parsed the sensorimotor transform into a set of smaller con-

## CHAPTER 5. CONCLUSION

stituent models, how do we infer the set of functional modules and the network topology of the system model? In a block diagram describing a system (consider the cartoon in Figure 5.2(c)), constituent components may be connected in a limited number of ways: in series (sensory fusion to control policy), in parallel (vision, electroreception, and mechanoreception), or feedback (muscle dynamics to proprioception which eventually flows back to muscle dynamics). Each of these relationships can be reduced to a single unique transformation (*e.g.*  $A$  and  $B$  in series yield the single transform  $AB$ ,  $A$  and  $B$  in parallel yield  $A + B$ ,  $A$  with  $B$  in the feedback path yields  $\frac{A}{1+AB}$ ). Conversely, a single transformation can be decomposed into any of the above networks, but obviously not uniquely; each choice of local network topology furnishes a subspace of feasible constituent systems. The set of functional modules and the corresponding network topology might be approached as an optimization problem across data sets and behaviors. The choice of functional modules and topology could be selected from the space of feasible solutions in order to minimize the set of functional modules (parsimony) or maximize the recurrence of particular modules or maximize the use of suspected topologies, etc. The choice of cost function is not yet apparent, but the optimization problem can be constrained by any known constituent models (*e.g.* models derived from biomechanical observations and first principles). This inter-behavioral analysis may provide a framework for comparing different behaviors in the same animal, thus reconciling many past phenomenological results into a more general model for behavior.

# Bibliography

- [Alexander(1995)] **Alexander, R.** (1995). Simple models of human movement. *Appl. Mech.* **48**, 461.
- [Babineau et al.(2007)] **Babineau, D., Lewis, J. E. and Longtin, A.** (2007). Spatial acuity and prey detection in weakly electric fish. *PLoS Comp. Biol.* **3**, e38.
- [Bastian(1987)] **Bastian, J.** (1987). Electrolocation in the presence of jamming signals: behavior. *J. Comp. Physiol. A* **161**, 811–824.
- [Bastian et al.(2002)] **Bastian, J., Chacron, M. and Maler, L.** (2002). Receptive field organization determines pyramidal cell stimulus-encoding capability and spatial stimulus selectivity. *J. Neurosci.* **22**, 4577–4590.
- [Bastian and Heiligenberg(1980)] **Bastian, J. and Heiligenberg, W.** (1980). Neural correlates of the jamming avoidance response of eel. *J. Comp. Physiol. A* **136**, 135–152.

## BIBLIOGRAPHY

- [Bausenwein and Fischbach(1992)] **Bausenwein, B. and Fischbach, K.** (1992). Activity labeling patterns in the medulla of *drosophila melanogaster* caused by motion stimuli. *Cell and tissue research* **270**, 25–35.
- [Becker and Fuchs(1969)] **Becker, W. and Fuchs, A. F.** (1969). Further properties of the human saccadic system: eye movements and correction saccades with and without visual fixation points. *Vision Res.* **9**, 1247–1258.
- [Bergbreiter and Pister(2007)] **Bergbreiter, S. and Pister, K.** (2007). Design of an autonomous jumping microrobot. In *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 447 –453.
- [Bergou et al.(2010)] **Bergou, A., Ristroph, L., Guckenheimer, J., Cohen, I. and Wang, Z.** (2010). Fruit flies modulate passive wing pitching to generate in-flight turns. *Phys. Rev. Letters* **104**, 148101.
- [Berman and Maler(1999)] **Berman, N. and Maler, L.** (1999). Neural architecture of the electrosensory lateral line lobe: adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering. *J. Exp. Biol.* **202**, 1243–1253.
- [Blickhan(1989)] **Blickhan, R.** (1989). The spring-mass model for running and hopping. *J. Biomech.* **22**, 1217–27.
- [Borst and Haag(2007)] **Borst, A. and Haag, J.** (2007). Optic flow processing in

## BIBLIOGRAPHY

the cockpit of the fly. In *Invertebrate Neurobiology*, volume 49, pp. 101–122. Cold Spring Harbor Monograph Archive.

[Braitenberg(1986)] **Braitenberg, V.** (1986). *Vehicles: Experiments in synthetic psychology*. The MIT press.

[Brigandt and Love(2008)] **Brigandt, I. and Love, A.** (2008). Reductionism in biology. *Stanford encyclopedia of philosophy* .

[Buchner et al.(1984)] **Buchner, E., Buchner, S. and Bülthoff, I.** (1984). Deoxyglucose mapping of nervous activity induced in *drosophila* brain by visual movement. *J. Comp. Physiol. A* **155**, 471–483.

[Budick et al.(2007)] **Budick, S. A., Reiser, M. B. and Dickinson, M. H.** (2007). The role of visual and mechanosensory cues in structuring forward flight in *drosophila melanogaster*. *J. Exp. Biol.* **210**, 4092–4103.

[Buehler et al.(2005)] **Buehler, M., Playter, R. and Raibert, M.** (2005). Robots step outside. In *Proc. Int. Symp. Adaptive Motion of Animals and Machines*, pp. 1–4.

[Bullock et al.(1972)] **Bullock, T. H., Hamstra, R. H. and Scheich, H.** (1972). The jamming avoidance response of high-frequency electric fish, I & II. *J. Comp. Physiol.* **77**, 1–48.

[Carver et al.(2005)] **Carver, S., Kiemel, T., van der Kooij, H. and Jeka, J. J.**

## BIBLIOGRAPHY

- (2005). Comparing internal models of the dynamics of the visual environment. *Biol. Cybern.* **92**, 147–163.
- [Carver et al.(2009)] **Carver, S. G., Kiemel, T., Cowan, N. J. and Jeka, J. J.** (2009). Optimal motor control may mask sensory dynamics. *Biol. Cybern.* **101**, 35–42.
- [Carver et al.(2008)] **Carver, S. G., Roth, E., Cowan, N. J. and Fortune, E. S.** (2008). Synaptic plasticity can produce and enhance direction selectivity. *PLoS Comp. Biol.* **4**.
- [Chacron and Fortune(2010)] **Chacron, M. and Fortune, E.** (2010). Subthreshold membrane conductances enhance directional selectivity in vertebrate sensory neurons. *J. Neurophysiol.* **104**, 449–462.
- [Chacron et al.(2009)] **Chacron, M., Toporikova, N. and Fortune, E.** (2009). Differences in the time course of short-term depression across receptive fields are correlated with directional selectivity in electrosensory neurons. *J. Neurophysiol.* **102**, 3270–3279.
- [Chacron et al.(2003)] **Chacron, M. J., Doiron, B., Maler, L., Longtin, A. and Bastian, J.** (2003). Non-classical receptive field mediates switch in a sensory neuron’s frequency tuning. *Nature* **423**, 77–81.
- [Cheng et al.(2010)] **Cheng, B., Fry, S., Huang, Q. and Deng, X.** (2010). Aero-

## BIBLIOGRAPHY

- dynamic damping during rapid flight maneuvers in the fruit fly *drosophila*. *J. Exp. Biol.* **213**, 602–612.
- [Clark et al.(2001)] **Clark, J., Cham, J., Bailey, S., Froehlich, E., Nahata, P., Cutkosky, M. et al.** (2001). Biomimetic design and fabrication of a hexapedal running robot. In *Proc. IEEE Int. Conf. Robot. Autom.*, volume 4, pp. 3643–3649. IEEE.
- [Collins and Barnes(2009)] **Collins, C. J. S. and Barnes, G. R.** (2009). Predicting the unpredictable: weighted averaging of past stimulus timing facilitates ocular pursuit of randomly timed stimuli. *J. Neurosci.* **29**, 13302–13314.
- [Conroy et al.(2009)] **Conroy, J., Gremillion, G., Ranganathan, B. and Humbert, J.** (2009). Implementation of wide-field integration of optic flow for autonomous quadrotor navigation. *Autonomous Robots* **27**, 189–198.
- [Cowan and Fortune(2007)] **Cowan, N. J. and Fortune, E. S.** (2007). The critical role of locomotion mechanics in decoding sensory systems. *J. Neurosci.* **27**, 1123–1128.
- [Cowan et al.(2006)] **Cowan, N. J., Lee, J. and Full, R. J.** (2006). Task-level control of rapid wall following in the American cockroach. *J. Exp. Biol.* **209**, 1617–1629.
- [Curet et al.(2011)] **Curet, O., Patankar, N., Lauder, G. and MacIver, M.**



## BIBLIOGRAPHY

- (2011). Mechanical properties of a bio-inspired robotic knifefish with an undulatory propulsor. *Bioinspiration & Biomimetics* **6**, 026004.
- [de Brouwer et al.(2001)] **de Brouwer, S., Missal, M. and Lefevre, P.** (2001). Role of retinal slip in the prediction of target motion during smooth and saccadic pursuit. *J. Neurophysiol.* **86**, 550–558.
- [Dickinson et al.(1993)] **Dickinson, M. H., Lehmann, F. O. and Gotz, K. G.** (1993). The active control of wing rotation by *drosophila*. *J. Exp. Biol.* **182**, 173–189.
- [Duistermars et al.(2007)] **Duistermars, B., Reiser, M., Zhu, Y. and Frye, M.** (2007). Dynamic properties of large-field and small-field optomotor flight responses in *Drosophila*. *J. Comp. Physiol. A* **193**, 787–799.
- [Epstein et al.(2006)] **Epstein, M., Colgate, J. and MacIver, M.** (2006). Generating thrust with a biologically-inspired robotic ribbon fin. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pp. 2412–2417. IEEE.
- [Fearing et al.(2000)] **Fearing, R., Chiang, K., Dickinson, M., Pick, D., Sitti, M. and Yan, J.** (2000). Wing transmission for a micromechanical flying insect. In *Proc. IEEE Int. Conf. Robot. Autom.*, volume 2, pp. 1509–1516. IEEE.
- [Fortune and Cowan(2010)] **Fortune, E. and Cowan, N.** (2010). Robot behavior

## BIBLIOGRAPHY

- [Franz and Krapp(2000)] **Franz, M. and Krapp, H.** (2000). Wide-field, motion-sensitive neurons and matched filters for optic flow fields. *Biol. Cybern.* **83**, 185–197.
- [Fry et al.(2003)] **Fry, S. N., Sayaman, R. and Dickinson, M. H.** (2003). The aerodynamics of free-flight maneuvers in *Drosophila*. *Science* **300**, 495–498.
- [Frye and Dickinson(2004)] **Frye, M. A. and Dickinson, M. H.** (2004). Motor output reflects the linear superposition of visual and olfactory inputs in *drosophila*. *J. Exp. Biol.* **207**, 123–131.
- [Fuchs(1967)] **Fuchs, A. F.** (1967). Saccadic and smooth pursuit eye movements in the monkey. *J. Physiol.* **191**, 609.
- [Full and Koditschek(1999)] **Full, R. J. and Koditschek, D. E.** (1999). Templates and anchors: neuromechanical hypotheses of legged locomotion on land. *J. Exp. Biol.* **202**, 3325–3332.
- [Geiger and Nüssel(1982)] **Geiger, G. and Nüssel, D.** (1982). Visual processing of moving single objects and wide-field patterns in flies: Behavioural analysis after laser-surgical removal of interneurons. *Biol. Cybern.* **44**, 141–149.
- [Gilbert(1997)] **Gilbert, C.** (1997). Visual control of cursorial prey pursuit by tiger beetles (Cicindelidae). *J. Comp. Physiol. A* **181**, 217–230.

## BIBLIOGRAPHY

- [Götz(1968)] **Götz, K.** (1968). Flight control in *Drosophila* by visual perception of motion. *Biol. Cybern.* **4**, 199–208.
- [Götz et al.(1979)] **Götz, K., Hengstenberg, B. and Biesinger, R.** (1979). Optomotor control of wing beat and body posture in drosophila. *Biol. Cybern.* **35**, 101–112.
- [Hafting et al.(2005)] **Hafting, T., Fyhn, M., Molden, S., Moser, M. and Moser, E.** (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806.
- [Hardie and Raghu(2001)] **Hardie, R. C. and Raghu, P.** (2001). Visual transduction in *drosophila*. *Nature* **413**, 186–193.
- [Harris and Wolpert(1998)] **Harris, C. and Wolpert, D.** (1998). Signal-dependent noise determines motor planning. *Nature* **394**.
- [Harrison et al.(2010)] **Harrison, R., Kier, R., Leonardo, A., Fotowat, H., Chan, R. and Gabbiani, F.** (2010). A wireless neural/emg telemetry system for freely moving insects. In *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, pp. 2940–2943. IEEE.
- [Hedrick et al.(2009)] **Hedrick, T. L., Cheng, B. and Deng, X.** (2009). Wingbeat time and the scaling of passive rotational damping in flapping flight. *Science* **324**, 252–255.

## BIBLIOGRAPHY

- [Heiligenberg(1973)] **Heiligenberg, W.** (1973). Electrolocation of objects in the electric fish *Eigenmannia* (rhamphichthyidae, gymnotoidei). *J. Comp. Physiol. A* **87**, 137–164.
- [Heiligenberg(1991)] **Heiligenberg, W.** (1991). *Neural nets in electric fish*. Cambridge, MA: MIT Press.
- [Heiligenberg and Bastian(1980)] **Heiligenberg, W. and Bastian, J.** (1980). The control of *eigenmannia*'s pacemaker by distributed evaluation of electroreceptive afferences. *J. Comp. Physiol. A* **136**, 113–133.
- [Heisenberg and Wolf(1988)] **Heisenberg, M. and Wolf, R.** (1988). Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *J. Comp. Physiol. A* **163**, 373–388.
- [Hitschfeld et al.(2009)] **Hitschfeld, E., Stamper, S., Vonderschen, K., Fortune, E. and Chacron, M.** (2009). Effects of restraint and immobilization on electrosensory behaviors of weakly electric fish. *ILAR J* **50**, 361–372.
- [Holmes et al.(2006)] **Holmes, P., Full, R. J., Koditschek, D. and Guckenheimer, J.** (2006). The dynamics of legged locomotion: Models, analyses, and challenges. *SIAM Rev.* **48**, 207–304.
- [Humbert et al.(2005)] **Humbert, J., Murray, R. and Dickinson, M.** (2005). Pitch-altitude control and terrain following based on bio-inspired visuomotor

## BIBLIOGRAPHY

- convergence. In *2005 AIAA Guidance, Navigation, and Control Conference and Exhibit*, pp. 1–12.
- [Ijspeert et al.(2007)] **Ijspeert, A., Crespi, A., Ryczko, D. and Cabelguen, J.** (2007). From swimming to walking with a salamander robot driven by a spinal cord model. *Science* **315**, 1416–1420.
- [Jeka et al.(2004)] **Jeka, J., Kiemel, T., Creath, R., Horak, F. and Peterka, R.** (2004). Controlling human upright posture: velocity information is more accurate than position or acceleration. *J. Neurophysiol.* **92**, 2368–2379.
- [Kiemel et al.(2006)] **Kiemel, T., Oie, K. S. and Jeka, J. J.** (2006). Slow dynamics of postural sway are in the feedback loop. *J. Neurophysiol.* **95**, 1410–1418.
- [Körding and Wolpert(2006)] **Körding, K. P. and Wolpert, D. M.** (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn Sci* **10**, 319–326.
- [Krapp et al.(1996)] **Krapp, H., Hengstenberg, R. et al.** (1996). Estimation of self-motion by optic flow processing in single visual interneurons. *Nature* **384**, 463–466.
- [Krotki(2011)] **Krotki, M. A. J.** (2011). *Lobula Plate Tangential Cells in Drosophila melanogaster; Response Properties, Synaptic Organization and Input Channels*. Ph.D. thesis, Ludwig Maximilians Universität.

## BIBLIOGRAPHY

- [Kuo(2005)] **Kuo, A. D.** (2005). An optimal state estimation model of sensory integration in human postural balance. *J Neural Eng* **2**, S235–S249.
- [Lee et al.(2008)] **Lee, J., Sponberg, S. N., Loh, O. Y., Lamperski, A. G., Full, R. J. and Cowan, N. J.** (2008). Templates and anchors for antenna-based wall following in cockroaches and robots. *IEEE Trans. Robot.* **24**, 130–143.
- [Lisberger et al.(1987)] **Lisberger, S., Morris, E. and Tychsen, L.** (1987). Visual motion processing and sensory-motor integration for smooth pursuit eye movements. *Annual Review of Neuroscience* **10**, 97–129.
- [Ljung(1999)] **Ljung, L.** (1999). *System identification*. Wiley Online Library.
- [Maciver et al.(2001)] **Maciver, M. A., Sharabash, N. M. and Nelson, M. E.** (2001). Prey-capture behavior in gymnotid electric fish: Motion analysis and effects of water conductivity. *J. Exp. Biol.* **204**, 543–557.
- [Maimon et al.(2010)] **Maimon, G., Straw, A. and Dickinson, M.** (2010). Active flight increases the gain of visual motion processing in *drosophila*. *Nature neuroscience* **13**, 393–399.
- [Maimon et al.(2008)] **Maimon, G., Straw, A. D. and Dickinson, M. H.** (2008). A simple vision-based algorithm for decision making in flying *drosophila*. *Curr. Biol.* **18**, 464 – 470.

## BIBLIOGRAPHY

- [Mason and Burdick(2000)] **Mason, R. and Burdick, J.** (2000). Experiments in carangiform robotic fish locomotion. In *Proc. IEEE Int. Conf. Robot. Autom.*, volume 1, pp. 428 –435 vol.1.
- [Nelson et al.(2002)] **Nelson, M. E., MacIver, M. A. and Coombs, S.** (2002). Modeling electrosensory and mechanosensory images during the predatory behavior of weakly electric fish. *Brain Behav. and Evol.* .
- [Oie et al.(2002)] **Oie, K. S., Kiemel, T. and Jeka, J. J.** (2002). Multisensory fusion: simultaneous re-weighting of vision and touch for the control of human posture. *Cog Brain Res.* **14**, 164–176.
- [O’Keefe and Dostrovsky(1971)] **O’Keefe, J. and Dostrovsky, J.** (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain research* .
- [Pintelon and Schoukens(2001)] **Pintelon, R. and Schoukens, J.** (2001). *System identification: a frequency domain approach*. Wiley-IEEE Press.
- [Poggio and Reichardt(1976)] **Poggio, T. and Reichardt, W.** (1976). Visual control of orientation behaviour in the fly: Part ii. towards the underlying neural interactions. *Quarterly reviews of biophysics* **9**, 377–438.
- [Poggio and Reichardt(1981)] **Poggio, T. and Reichardt, W.** (1981). Visual fix-

## BIBLIOGRAPHY

- ation and tracking by flies: mathematical properties of simple control systems. *Biol. Cybern.* **40**, 101–112.
- [Powers(1989)] **Powers, W.** (1989). Living control systems: Selected papers of William T. Powers .
- [Rashbass(1961)] **Rashbass, C.** (1961). The relationship between saccadic and smooth tracking eye movements. *J. Physiol.* **159**, 326.
- [Rasnow(1996)] **Rasnow, B.** (1996). The effects of simple objects on the electric field of *Apterionotus*. *J. Comp. Physiol.* **178**, 397–411.
- [Rasnow et al.(1989)] **Rasnow, B., Assad, C., Nelson, M. E. and Bower, J. M.** (1989). Simulation and measurement of the electric fields generated by weakly electric fish. *Advances in neural information processing systems 1* pp. 436–443.
- [Reichardt(1961)] **Reichardt, W.** (1961). Autocorrelation and the central nervous system. In *Sensory Communication* (ed. A. Rosenblith), pp. 303–318. MIT Press Cambridge.
- [Reichardt and Poggio(1976)] **Reichardt, W. and Poggio, T.** (1976). Visual control of orientation behaviour in the fly. part I. a quantitative analysis. *Q Rev Biophys* **9**, 311–75, 428–38.
- [Reiser and Dickinson(2008)] **Reiser, M. and Dickinson, M.** (2008). A modular



## BIBLIOGRAPHY

- display system for insect behavioral neuroscience. *J. Neurosci. Methods* **167**, 127–139.
- [Resh and Cardé(2009)] **Resh, V. and Cardé, R.** (2009). *Encyclopedia of Insects*. Academic Press.
- [Revzen(2009)] **Revzen, S.** (2009). *Neuromechanical Control Architectures of Arthropod Locomotion*. Ph.D. thesis, University of California, Berkeley.
- [Revzen and Guckenheimer(2008)] **Revzen, S. and Guckenheimer, J. M.** (2008). Estimating the phase of synchronized oscillators. *Phys. Rev. E* **78**, 051907–051918.
- [Ristroph et al.(2010)] **Ristroph, L., Bergou, A. J., Ristroph, G., Coumes, K., Berman, G. J., Guckenheimer, J., Wang, Z. J. and Cohen, I.** (2010). Discovering the flight autostabilizer of fruit flies by inducing aerial stumbles. *Proc. Nat. Acad. Sci.* .
- [Rose and Canfield(1993a)] **Rose, G. J. and Canfield, J. G.** (1993a). Longitudinal tracking responses of *Eigenmannia* and *Sternopygus*. *J. Comp. Physiol. A* **173**, 698–700.
- [Rose and Canfield(1993b)] **Rose, G. J. and Canfield, J. G.** (1993b). Longitudinal tracking responses of the weakly electric fish, *Sternopygus*. *J. Comp. Physiol. A* **171**, 791–798.

## BIBLIOGRAPHY

- [Roth et al.(2008)] **Roth, E., Carver, S., Fortune, E. S. and Cowan, N. J.** (2008). Mechanisms for encoding velocity in the electrosensory system of weakly electric fish. In *Proc. Int. Symp. Adaptive Motion of Animals and Machines*. Cleveland, OH.
- [Roth et al.(2011)] **Roth, E., Zhuang, K., Stamper, S. A., Fortune, E. S. and Cowan, N. J.** (2011). Stimulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmannia virescens*. *J. Exp. Biol.* **214**, 1170–1180.
- [Saranli et al.(2001)] **Saranli, U., Buehler, M. and Koditschek, D. E.** (2001). RHex: A simple and highly mobile hexapod robot. *Int. J. Robot. Res.* **20**, 616–631.
- [Sastry(1999)] **Sastry, S.** (1999). *Nonlinear Systems*. Springer.
- [Sefati et al.(2010)] **Sefati, S., Fortune, E. S. and Cowan, N. J.** (2010). Counter-propagating waves in the ribbon fin of *Eigenmannia virescens* enhance maneuverability. In *The Society for Integrative and Comparative Biology*. Seattle, WA.
- [Sefati et al.(2012)] **Sefati, S., Neveln, I., MacIver, M. A., Fortune, E. S. and Cowan, N. J.** (2012). Counter-propagating waves enhance maneuverability and stability: a bio-inspired strategy for robotic ribbon-fin propulsion. In *Proc. IEEE Int. Conf. on Biomedical Robotics and Biomechatronics*. Rome, Italy.

## BIBLIOGRAPHY

- [Shadmehr and Mussa-Ivaldi(1994)] **Shadmehr, R. and Mussa-Ivaldi, F. A.** (1994). Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224.
- [Sherman and Dickinson(2004)] **Sherman, A. and Dickinson, M. H.** (2004). Summation of visual and mechanosensory feedback in drosophila flight control. *J. Exp. Biol.* **207**, 133–142.
- [Shirgaonkar et al.(2008)] **Shirgaonkar, A. A., Curet, O. M., Patankar, N. A. and Maciver, M. A.** (2008). The hydrodynamics of ribbon-fin propulsion during impulsive motion. *J. Exp. Biol.* **211**, 3490–3503.
- [Shumway(1989)] **Shumway, C. A.** (1989). Multiple electrosensory maps in the medulla of weakly electric gymnotiform fish. i. physiological differences. *J. Neurosci.* **9**, 4388–4399.
- [Sim and Kim(2012)] **Sim, M. and Kim, D.** (2012). Electrolocation of multiple objects based on temporal sweep motions. *Adaptive Behavior* .
- [Söderström and Stoica(1988)] **Söderström, T. and Stoica, P.** (1988). *System identification*. Prentice-Hall, Inc.
- [Solstad et al.(2008)] **Solstad, T., Boccara, C., Kropff, E., Moser, M. and Moser, E.** (2008). Representation of geometric borders in the entorhinal cortex. *Science* **322**, 1865–1868.

## BIBLIOGRAPHY

- [Sprayberry and Daniel(2007)] **Sprayberry, J. D. H. and Daniel, T. L.** (2007). Flower tracking in hawkmoths: behavior and energetics. *J. Exp. Biol.* **210**, 37–45.
- [Stamper et al.(2012)] **Stamper, S., Roth, E., Cowan, N. and Fortune, E.** (2012). Active sensing via movement shapes spatiotemporal patterns of sensory feedback. *J. Exp. Biol.* **215**, 1567–1574.
- [Tammero and Dickinson(2002)] **Tammero, L. and Dickinson, M.** (2002). The influence of visual landscape on the free flight behavior of the fruit fly *drosophila melanogaster*. *J. Exp. Biol.* **205**, 327–343.
- [Tammero et al.(2004)] **Tammero, L. F., Frye, M. A. and Dickinson, M. H.** (2004). Spatial organization of visuomotor reflexes in *Drosophila*. *J. Exp. Biol.* **207**, 113–122.
- [Theobald et al.(2010)] **Theobald, J. C., Ringach, D. L. and Frye, M. A.** (2010). Dynamics of optomotor responses in drosophila to perturbations in optic flow. *J. Exp. Biol.* **213**, 1366–1375.
- [Tsang et al.(2011)] **Tsang, W., Stone, A., Otten, D., Aldworth, Z., Daniel, T., Hildebrand, J., Levine, R. and Voldman, J.** (2011). Insect-machine interface: A carbon nanotube-enhanced flexible neural probe. *J. Neurosci. Methods* .

## BIBLIOGRAPHY

- [von der Emde(2006)] **von der Emde, G.** (2006). Non-visual environmental imaging and object detection through active electrolocation in weakly electric fish. *J. Comp. Physiol. A* **192**, 601–612.
- [von Holst and Mittelstaedt(1950)] **von Holst, E. and Mittelstaedt, H.** (1950). Das reafferenzprinzip. *Naturwissenschaften* **37**, 464–476.
- [Watanabe and Takeda(1963)] **Watanabe, A. and Takeda, K.** (1963). The change of discharge frequency by a.c. stimulus in a weak electric fish. *J. Exp. Biol.* **40**, 57–66.
- [Weber(2011)] **Weber, F.** (2011). *System Identification of Optic-Flow Processing Neurons in the Fly: Single Cell and Small Circuit Models*. Ph.D. thesis, Ludwig Maximilians Universität.
- [Weber et al.(2008)] **Weber, F., Eichner, H., Cuntz, H. and Borst, A.** (2008). Eigenanalysis of a neural network for optic flow processing. *New Journal of Physics* **10**, 015013.
- [Wolf and Heisenberg(1990)] **Wolf, R. and Heisenberg, M.** (1990). Visual control of straight flight in *drosophila melanogaster*. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* **167**, 269–283.
- [Wolpert et al.(1995)] **Wolpert, D. M., Ghahramani, Z. and Jordan, M. I.**

## BIBLIOGRAPHY

- (1995). An internal model for sensorimotor integration. *Science* **269**, 1880–1882.
- [Wood(2007)] **Wood, R.** (2007). The first flight of an insect-sized robotic fly. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pp. 2555–2555. San Diego, CA, USA.
- [Ye et al.(1995)] **Ye, S., Dowd, J. and Comer, C.** (1995). A motion tracking system for simultaneous recording of rapid locomotion and neural activity from an insect. *J. Neurosci. Methods* **60**, 199–210.
- [Zanker(1990)] **Zanker, J.** (1990). The wing beat of drosophila melanogaster. i. kinematics. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* pp. 1–18.
- [Zbikowski(2005)] **Zbikowski, R.** (2005). Fly like a fly [micro-air vehicle]. *Spectrum, IEEE* **42**, 46–51.

# Vita

Eatai Roth received a BFA in painting and art history from Washington University in 2001, a BS in mechanical engineering from the University of Pittsburgh in 2005, and an MS in mechanical engineering from Johns Hopkins University in 2008 where he is pursuing his PhD in the LIMBS (Locomotion in Mechanical and Biological Systems) Lab under the supervision of Prof. Noah Cowan. He was awarded the NSF Graduate Research Fellowship in 2006 and the ARCS Foundation Scholarship in 2009 and 2010. His research investigates the neuromechanical basis of behavior, how animals transform sensory signals into locomotor action. Working at the interface of control theory and neuroscience, Eatai develops dynamical models to describe the mechanisms responsible for robust animal behaviors. His non-research interests include art and architecture, American football, and wordplay. He is also the co-founder of the Mechanical Engineering Graduate Association and the Student Steering Committee for the Laboratory for Computational Sensing and Robotics.