

Learning of Novel Dynamics in *Eigenmannia virescens*
Refuge Tracking Task

by

Dominic G. Yared

A thesis submitted to The Johns Hopkins University in conformity with the
requirements for the degree of Master of Science in Engineering.

Baltimore, Maryland

May, 2020

© Dominic G. Yared 2020

All rights reserved



Abstract

The ability to learn novel dynamics is a characteristic we exhibit our whole lives. From learning to ride a bicycle, to walking with crutches, we are constantly experiencing new physical settings that require a certain method of control. The goal of my research was to examine and evaluate this capability in *Eigenmannia virescens*, the weakly electric glass knifefish. We sought to understand the plasticity of a fish's controller; namely, to quantify the fish's change in internal control scheme as it adapted to new dynamics. In particular, we were testing motor learning in *E. virescens*.

We evaluated the fish's controller by having it track a moving target, which was a small open-sided cage placed in the fish tank. We moved the target using two different methods. In the first method, "open-loop" tracking, the target moved according to a predetermined trajectory. This method presented the fish with dynamics it had already experienced throughout its whole life, allowing us to establish a baseline evaluation of tracking performance. In the second method, we used a camera and fish tracking software to make the fish's own

ABSTRACT

position affect the target's movement; incrementally changing the parameters associated with this "closed-loop" method, we were able to simulate a new form of dynamics for the fish. We used principles of control theory and frequency domain analysis to measure the resulting change in the fish's controller as it learned to compensate for these novel dynamics. We found the fish reduced their damping coefficient to increase the stability of the overall system.

To place this work in a broader context, we believe that the fundamental principles at work in *Eigenmannia virescens*' controllers can generalize to many animals, including humans. By achieving a greater understanding of how a fish's brain can change in response to new dynamics, we can potentially open new avenues in understanding neurological disorders in humans.

Primary Reader and Advisor: Noah J. Cowan

Secondary Reader: Eric S. Fortune

Acknowledgments

I would like to begin by thanking Dr. Noah Cowan for the abundant insight and inspiration he has provided me throughout this whole process. I am incredibly fortunate to have had the opportunity to work with such a thoughtful advisor, and I look forward to hearing about all the fantastic work that will come from his lab and the people he has inspired.

I would also like to express my deepest gratitude to Mr. Yu Yang, for the endless help, support, and knowledge he has given me during my Masters. My research and this thesis would not be where it is without his help. Without a doubt, he has been one of the best teachers I have had at Johns Hopkins.

Finally, I would like to thank the entire LIMBS Laboratory for the welcoming and enriching environment they have created. Working with such intelligent and caring people has been a true blessing.

Dedication

I would like to dedicate this thesis to my parents, Reem and Wael. I am forever grateful for their constant love and support, and for the opportunities they have opened up for me.

Contents

Abstract	ii
Acknowledgments	iv
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Contributions	3
2 Materials and Methods	5
2.1 Experimental Setup	5
2.2 Experiment Design: Open vs. Closed Loop	8
2.3 Input Signal	14
2.4 Experimental Procedure	16
2.5 Parameter Selection	20

CONTENTS

2.6 Data Processing Methodology	27
3 Results and Analysis	29
3.1 Bode Plot Analysis	29
3.2 Bode Plot Trend Analysis – Damping Coefficient	42
3.3 Root Locus Analysis	46
3.4 Washout	52
4 Conclusion and Future Work	56
Bibliography	59
Vita	63

List of Tables

2.1 k and b values used in learning experiments	26
---	----

List of Figures

2.1	Fish swimming in refuge with real time (RT) feedback. $r(t)$ is the location of the refuge, $y(t)$ is the location of the fish.	7
2.2	<i>Eigenmannia virescens</i> Refuge tracking system – $G(s)$	10
2.3	Block diagram of closed loop fish tracking system	12
2.4	Open- and closed-loop experimental setups	13
2.5	Plot of reference signal, 60 seconds in total, with the first 10 s "ramp-up", last 10 s "ramp-down" and middle 40 s as per equation (2.2).	15
2.6	Closed-loop experimental procedure	19
2.7	Dominant pole as a function of k and b parameters	24
2.8	Dominant pole contour plot with final experimental values	26
3.1	Bode plot of first Key experiment	31
3.2	Bode plot of second Key experiment	37
3.3	Bode plot of Bada experiment	39
3.4	Bode plot of Hope experiment	40
3.5	Increase in gain at 1.15 Hz as a function of k value for Key, Bada and Hope	41
3.6	Decrease in gain at 0.1 Hz as a function of k value for Key, Bada and Hope	42
3.7	Gain progression at 1.1 Hz as a function of k value for Owain, Frank, and Mac	43
3.8	Mass-Spring-Damper System	43
3.9	Frequency response of varying damping coefficient in 2nd order system	45
3.10	Dominant pole from first Key experiment	49
3.11	Dominant poles for Hope and Bada.	51
3.12	Washout results for both Key experiments	53
3.13	Bada washout result	55
3.14	Hope washout result	55

Chapter 1

Introduction

The weakly electric fish *Eigenmannia virescens* has been widely researched from a control systems perspective. The fish's ability to move equally well forward and backwards with its ribbon fin [1,2], as well as its propensity to follow a moving target, makes it an ideal candidate for such analysis. Up to this point, the majority of the research [3-8] performed on *E. virescens* has examined its ability to track an object in an "open-loop" context – this to say, following a target moving along a predetermined trajectory, independently of the fish [3,4]. Work has begun, however, in evaluating the fish performing in a "closed-loop" setting [9,10] – an experimental method where the position of the fish affects the position of the target it is attempting to track. Due to the shift in dynamics this presents, it provides a fertile ground for understanding the plasticity of the fish's brain. Namely, we can quantitatively examine how the fish adapts

CHAPTER 1. INTRODUCTION

to novel dynamics. The research I am presenting in this thesis builds on the work performed by Biswas et al [10] and Yoshida et al [9] to further explore this concept.

In order to establish a baseline to which we could compare learned behavior, we began all fish experiments with open-loop trials. As mentioned before, all of these trials featured a target moving on a predetermined trajectory, completely independently from the fish. Each trial lasted 60 s, and was performed five times to provide an average representation of a certain fish’s baseline tracking performance. In our closed loop experiments, we used a high-speed camera and fish recognition software to pass the fish’s position through a high-pass filter, before sending it as a command to the target. We used a high-pass filter to further perturb the dynamics in a controlled manner. By changing the parameters of the filter throughout the course of the experiment, we were able to change the stability of the fish-target system, and evaluate the fish’s response. This phase of the experiment was where we hypothesized the fish would tune their controllers, and learn a new sensorimotor transform for tracking. We performed five trials for each of five different sets of filter parameters, with each trial lasting 60 s.

We finished each experiment by testing the fish once again on open-loop trials. This allowed us to see if any after-effects remained in the fish from the learning trials. Presence of an aftereffect in the open-loop experiments

CHAPTER 1. INTRODUCTION

provided a strong indication that the fish had indeed tuned their controllers during the closed-loop phase. We hypothesized that any aftereffect would be “washed-out”, or unlearned, over a timescale of about 15 minutes, consistent with cerebellar learning and washout [11].

In total, we tested this procedure on $N = 6$ fish. We found that the fish did indeed change their controllers to adapt to the novel closed-loop dynamics. In response to the instability we introduced through our high pass filter, we found that the fish reduced their damping ratio adding some phase lead in order to increase the stability of the fish-target system. We modeled the fish as second order linear systems with delay in order to reach this conclusion. In particular, we found these results strongly displayed in three of our six fish. The data from the remaining three fish were more variable, but were not inconsistent with learning. While this sample size may not be large enough to draw definitive conclusions about cerebellar-like learning in *E. virescens*, we believe it is noteworthy enough to warrant further research.

1.1 Contributions

In this research effort, my main contribution was designing, implementing, and testing a new form of closed loop experiment with six fish. This included developing the theoretical framework for choosing a new set of high-pass filter

CHAPTER 1. INTRODUCTION

parameters, modifying existing LabVIEW architecture to implement a continuous experiment, and testing it with six different fish. In all of these steps, I was greatly aided by my colleague Yu Yang. He guided me through the development of system stability analysis, helped me modify our experimental program, and helped me collect fish data. I cannot overemphasize how crucial he was to this project. Lastly, I would like to mention the vital contribution Balázs Vágvölgyi had in explaining and helping me modify the LabVIEW system. His expertise in computer science was indispensable to this research.

Chapter 2

Materials and Methods

2.1 Experimental Setup

Our experimental setup consisted of a fish tank, a plastic refuge (shuttle) connected to a linear stepper motor, a camera, two computers, and a mirror, similar to that reported in several previous studies [7, 9, 10, 12]. Our experimental setup functioned as follows: a command was sent from one computer to the stepper motor, via a Stepmotor controller (Copley Controls, Canton, MA, USA) [4], to move at a certain speed in a given direction. The stepper motor was positioned directly above the fish tank using scaffold previously constructed from extruded aluminum parts (80/20 Inc, Columbia City, IN, USA), with the refuge rigidly attached to a sliding mechanism, and descending into the tank. The fish followed the refuge in response to its linear motion. A mirror

CHAPTER 2. MATERIALS AND METHODS

placed below the tank at a 45 degree angle reflected the image of the bottom of the fish and the refuge into a camera, which then fed the image into a separate computer. In this manner, we obtained live video feed of the fish moving back and forth in the tank as it maintained its position within the refuge.

Our refuge was machined from a gray PVC 2"x2" rectangular tube, cut at a length of 8.7" [7]. The bottom surface of the refuge was removed so as to allow the image of the fish to pass through the bottom of the tank and reach the camera. In addition, the refuge had a series of parallel rectangular holes cut in each side, providing salient visual and electrosensory cues to the fish. The fish we used were 10 – 15 cm long weakly electric knifefish, *Eigenmannia virescens*, and were obtained from commercial vendors and housed according to published guidelines [13]. During experiments, they were put in freshwater tanks kept at 78°F, with a conductivity of 15–150 $\mu\text{S}/\text{cm}$. The fish were transferred to the experimental tank one day prior to experimentation for acclimation. We used a pco.1200s high speed camera (Cooke Corp., Romulus, MI, USA), with a AF Micro-Nikkor 60mm f/2.8D lens (Nikon Inc., Melville, NY, USA) [4], and recorded at a frame rate of 25 fps. A diagram showing our fish and refuge can be seen in Figure 2.1.

Critical to our system was a feedback controller implemented on a PC-based LabVIEW system (National Instruments, Austin, TX, USA) [10]. The system incorporated a USB-6221 Multifunction DAQ (National Instruments, Austin,

CHAPTER 2. MATERIALS AND METHODS

TX, USA) [5] and an image capture card. Each video image was processed in real time using a custom NI LabVIEW program, and the location of the fish was then processed via experimentally specified feedback dynamics (see section for details). The basic concept of this system was first implemented in the laboratory by Kyle Yoshida during a 2017 Research Experiences for Undergraduates summer internship in Prof. Cowan's laboratory [9]. This Master's essay builds on the preliminary data from this study.

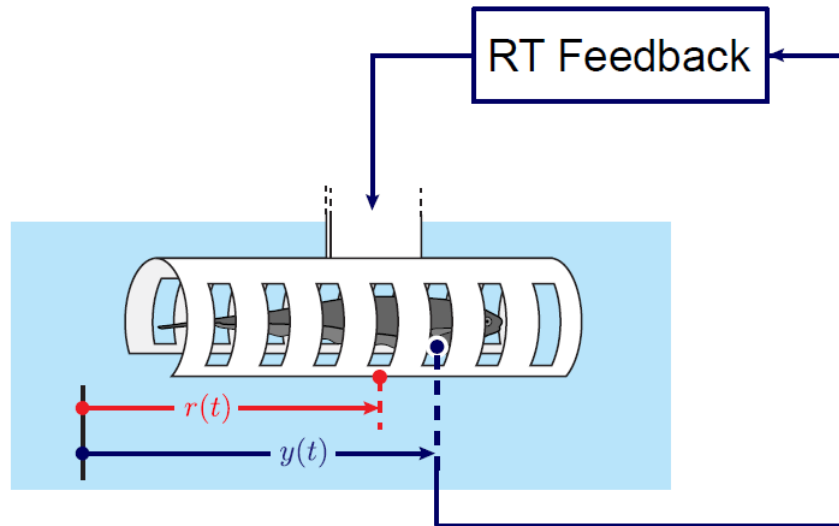


Figure 2.1: Fish swimming in refuge with real time (RT) feedback. $r(t)$ is the location of the refuge, $y(t)$ is the location of the fish.

2.2 Experiment Design: Open vs. Closed Loop

In our experiments, we employed both “open-loop” and “closed-loop” experiments to evaluate learning in the fish. We alternated between these two modes to provide a change in the dynamics that the fish was experiencing, and as a result force the fish to change the way it controls its movements to remain stable. To elaborate on these modes, an open loop experiment is one in which the refuge is controlled by a computer to follow a predetermined path, and its position over time is completely independent from the motion of the fish. This case represents the “baseline” tracking scenario. Native to the Amazon river basin, these fish routinely track objects both in the lab (staying “hidden” in refuges in their housing tanks) and in their natural habitat, hiding among root systems and reed grasses [14]. In general, their natural tracking is, like in this first type of experiment, “open loop”: the position of the objects the fish are tracking are not dependent on the motion of the fish. As a result, when the fish “see” (and sense electrically) an object at a certain position, they determine the “error” in their position. This means they compare where they are with where they aim to be, and then move accordingly. In this manner, the fish are relying on their own feedback, via their eyes and electrosense, to correct positional error [10]. In this open loop baseline case, this is exactly the case that the fish are used to

CHAPTER 2. MATERIALS AND METHODS

addressing. The fish sees the refuge, determines the error in its position, and moves forward or backwards to stay inside the refuge. As an aside, we believe the fish desired to stay inside the refuge because it provided a dark environment that made the fish “feel” safer. While this was not a perfect motivator, and the fish occasionally did leave the refuge, the refuge-seeking behavior was robust enough to allow collecting many trials of usable data.

The block diagram describing the fish’s open loop transfer function, first introduced by Cowan and Fortune [3], can be seen in Figure 2.2. In this diagram, the input, in red, is the position of the refuge, $r(t)$. This is treated as the desired position for the fish. The $C(s)$ block represents the fish’s controller [3, 4, 12], a simplified model of the central nervous system (CNS). This is the decision making component of the block diagram, and what we believe the fish changes when it is learning a new control scheme. The output signal from the controller is the modulation of muscular commands, which get sent to the next block, the plant, $P(s)$ [2, 8]. The plant represents the dynamics present in the mechanical system, including the muscular movement as a result of the controller signal. We are also including the resulting fluid dynamics in the plant; in this way $P(s)$ captures the dynamics of the movement of the entire fish in response to the modulation of muscle commands that drive the fish’s ribbon fin [2, 8]. In order for the plant to change throughout our experiment, this would necessitate the fish changing the overall patterns of muscle activation, or for the fluid

CHAPTER 2. MATERIALS AND METHODS

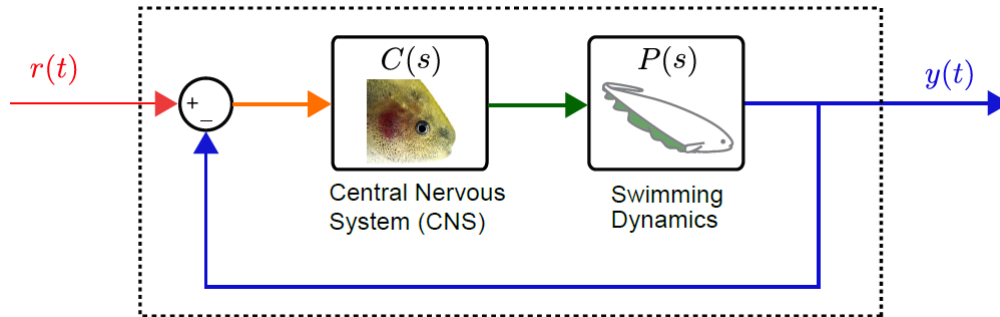


Figure 2.2: *Eigenmannia virescens* Refuge tracking system – $G(s)$

dynamics to significantly change. These possibilities are neglected in our analysis and experiments. Finally, the signal output from the plant is the position, $y(t)$ of the fish itself. As mentioned before, this signal gets fed back through the fish’s error correction mechanisms, its eyes and electrosense (as in the prior studies cited above, mechanosensory feedback is ignored), and gets subtracted from the refuge position to generate an error signal [15]. This error signal is the input to the fish’s controller.

The second experimental mode was “closed-loop tracking”, and resulted in the fish undergoing a learning process. In this mode, we used a live video stream and fish tracking software to feed the fish’s position back to the refuge through experimentally specified dynamics. The refuge’s position was therefore a function of both the fish’s location and a predetermined trajectory. The reason we employed this method was to introduce a new form of dynamics for the fish to experience. When the fish swam forward within the refuge, depending on the parameters of our setup, the refuge would move a certain way in response

CHAPTER 2. MATERIALS AND METHODS

to motion of the fish. This presented an entirely different form of motion for the fish, because the fish was tracking an object which in turn was tracking the fish itself. From the perspective of the fish, this novel feedback could make it more challenging for the fish to remain within the refuge. Indeed, for certain parameter values, the feedback could result in the refuge moving away from the fish faster than the fish itself was swimming. As a result, this could appear to the fish that the faster it swam towards the refuge, the farther away the refuge would move. At the risk of anthropomorphising, this would be very counterintuitive for the fish to experience, and contributed to the novelty of the dynamics.

Owing back to the original way that Kyle Yoshida performed his experiments [9], we chose to feed back the fish's movement through a high-pass filter in our closed-loop experiments. The main reason behind this was to further perturb the dynamics the fish was experiencing. Due to the high-pass filter, when the fish slowly changed directions to move towards the target, and therefore moved at a low frequency, the filter prevented the signal from being fed back to the refuge. This resulted in the experiment more closely resembling the open loop case, which was what the fish was used to. We hypothesized that this would therefore be easier for the fish, and would result in more accurate tracking. However, if the fish rapidly changed directions, thereby moving at a high frequency, the fish's position would pass through the filter (amplified by

CHAPTER 2. MATERIALS AND METHODS

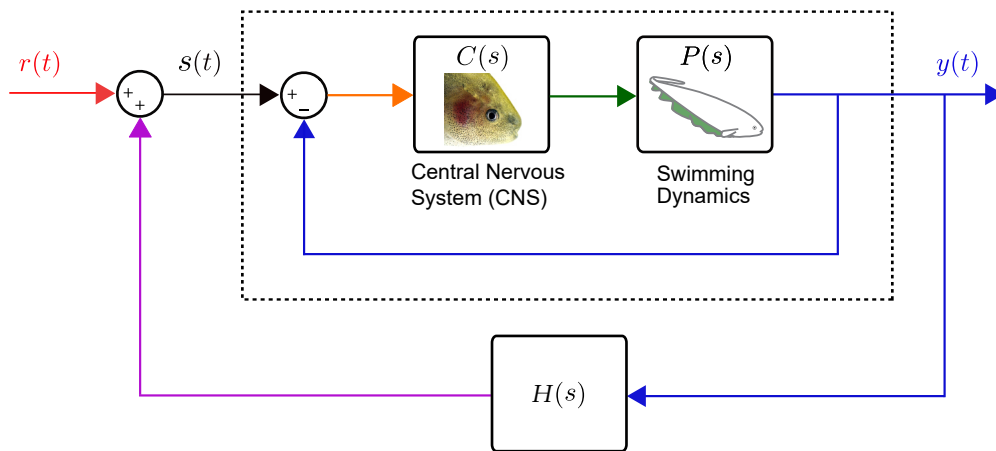


Figure 2.3: Block diagram of closed loop fish tracking system

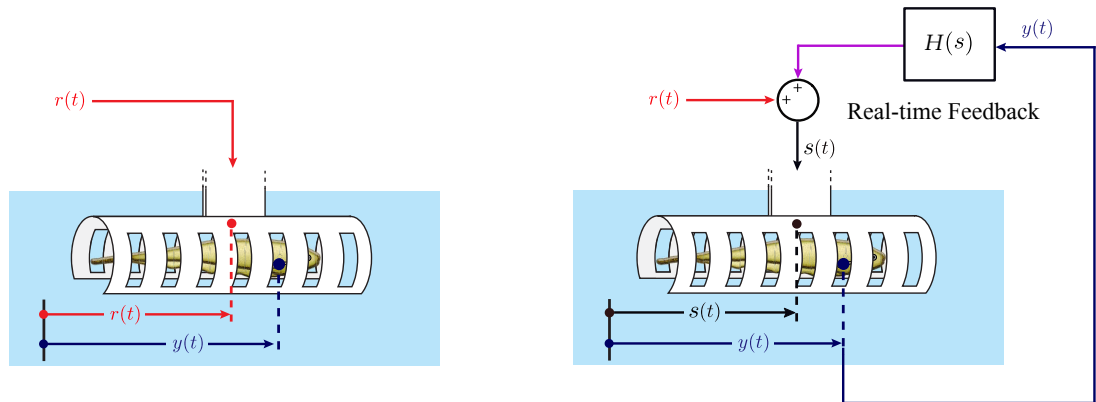
a certain gain), and be fed back to the refuge. The refuge would move rapidly forward or backward as a direct result of the fish's motion, potentially undermining the fish's ability to stay within the refuge. Naturally, this was much more difficult for the fish to track, especially if its control system remained unchanged. Thus, we hypothesized that the fish would learn to adapt its control system to more accurately track the refuge in these changed conditions. A block diagram illustrating the closed loop transfer function can be seen in Figure 2.3.

The block inside the dashed black line represents the transfer function of the fish itself, and is the same diagram as was explained in the open-loop case. The $H(s)$ block represents the high pass filter:

$$H(s) = \frac{ks}{s + b} \quad (2.1)$$

CHAPTER 2. MATERIALS AND METHODS

In our filter, k determines the gain of the filter, and b determines the cutoff frequency. The implications of these parameters will be discussed later. The output of the fish's position, measured by the camera, gets fed back through the filter and added to our reference trajectory. This reference trajectory is a pre-computed sinusoidal signal, and is the same input as the refuge position in the open loop case. In the closed loop case, however, the reference trajectory gets added to the filtered fish position, to produce the actual refuge position, $s(t)$. This is therefore what the fish actually sees, and is the input to the fish's transfer function. Another visual representation of our two experimental setups can be seen in Figure 2.4.



(a) Open-loop experiment. Only $r(t)$, the reference signal, determines the trajectory of refuge.

(b) Closed-loop experiment. Both reference signal $r(t)$ and filtered fish position determine the trajectory of refuge $s(t)$.

Figure 2.4: Open- and closed-loop experimental setups

2.3 Input Signal

Our reference trajectory was a sinusoid composed of the sum of 13 different sinusoids at different frequencies. The goal of this reference was to simulate a pseudo-random signal for the fish [4]. We wanted this reference to be as unpredictable as possible, so that the fish would be actively tracking the refuge at all times, as opposed to predicting the movement [4]. The dynamics involved in following a repetitive, predictable trajectory are different than the ones involved in active tracking, and would feature the fish's controller in a different manner. With a sum-of-sines signal however, we were able to determine the fish's frequency response at a range of frequencies, without having to worry about the predictability involved with running a constant sinusoid test at each frequency.

In the construction of our signal, each component of the overall sinusoid was at a different, prime-valued frequency. This prevented there from being repetition or interference among components of the signal. The amplitude of each sinusoidal component was also multiplied by $1/\text{frequency}$ to ensure a constant velocity for all frequency components [4]. The phase at each frequency ϕ_i was randomized. A formula describing our signal is as follows:

$$r(t) = \sum_i \frac{1}{2\pi \cdot 0.05k_i} \cos(2\pi \cdot 0.05k_i t + \phi_i) \quad k_i = 2, 3, 5, 7, \dots, 41 \quad (2.2)$$

CHAPTER 2. MATERIALS AND METHODS

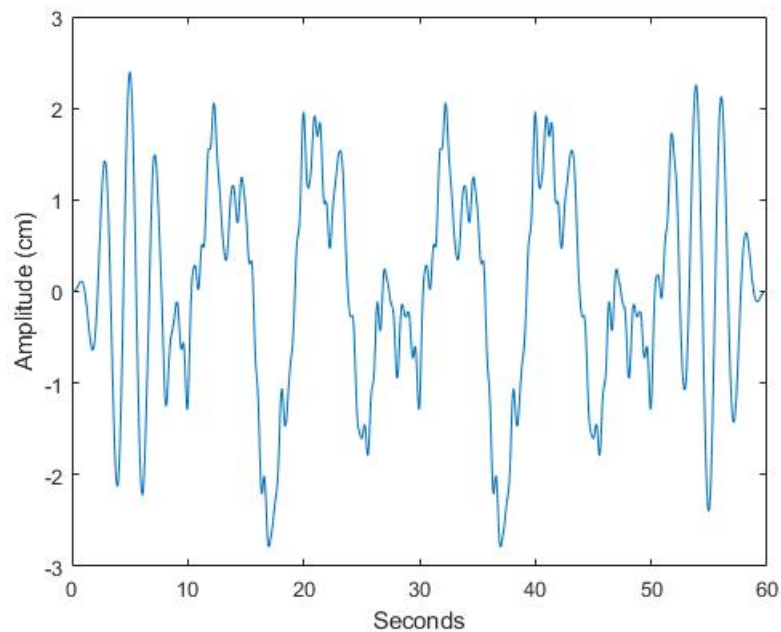


Figure 2.5: Plot of reference signal, 60 seconds in total, with the first 10 s "ramp-up", last 10 s "ramp-down" and middle 40 s as per equation (2.2).

A sample of the reference trajectory can be seen in the figure below. The signal has a duration of 60 seconds, and a maximum amplitude of 5.19 cm. The first and last ten seconds of the signal feature a "ramp up" and "ramp down" component, which is a gradual transition between the sum of sines signal and no movement. Between 10s and 50s, the pure sum of sines signal is expressed. The sum of sines signal can be seen in Figure 2.5.

2.4 Experimental Procedure

In order to maximize the learning effect, we implemented three stages in our experiments. Our first stage consisted of running the fish through five trials of open-loop tracking. This phase was for the purpose of ascertaining a baseline tracking ability. This essentially allowed us to quantify the “normal” operating condition of the fish, and was the comparison point for all subsequent experiments, as well as relate it to the extensive prior literature on refuge tracking dynamics in these fish. Each open-loop trial was 60s long, and was followed by a 20s rest period. The period allowed the fish to rest and regain energy, primarily for the sake of maintaining tracking consistency throughout the course of the hour-long experiment. Otherwise, the fish might fatigue and lose interest in following the refuge. This would consequently affect the tracking behavior, adding another variable to our experiment, making it difficult to isolate the real cause of tracking change. As a result, we tried to mitigate fatigue via a rest period.

In the second phase, we switched over to closed-loop. The closed-loop experiments consisted of a total of 25 trials, each trial 60 s in duration followed by a 20 s rest period. Again, we sought to keep a consistent experimentation method as with the open loop, in order to minimize the number of variables and highlight the sole effect of changing the dynamics. The 25 trials included 5 sets

CHAPTER 2. MATERIALS AND METHODS

of 5 trials, with each set of 5 trials associated with a distinct set of parameters on our high pass filter. This method of testing allowed us to increment through a set of parameters in our filter, and measure the associated tracking response in the fish. We designed our LabVIEW script to include all 25 trials into one single experiment, with rest periods and parameter changes written into the script. A more detailed explanation of the reasoning behind these parameters will be elaborated upon in the next section.

A key goal during the second phase was to maintain feedback at all times. As explained before, our closed-loop experiments relied on feeding the fish's movement through a high pass filter, and sending it as a command to the refuge. We chose to include this feedback during the rest periods as well, when there was no reference sent to the shuttle. We did this to prevent the fish from experiencing the baseline tracking dynamics at any point during our learning experiments. When a rest period is occurring, the fish is still moving back and forth at small amplitudes, and is constantly adjusting its body; if the shuttle were completely stationary, it would constitute a form of open loop tracking, as the fish would be tracking a target that moves (or does not move in this case) completely independently from the fish itself. As a result, the fish would be "interrupting" its learning process, reverting back to its baseline tracking controller. In so doing, the fish might actually learn to rapidly switch between open-loop and closed-loop control modes, making it hard for us to ascertain the

CHAPTER 2. MATERIALS AND METHODS

time course of learning (and unlearning or washout) of the new control law needed for closed loop tracking. To prevent this from occurring, we chose to keep sending the fish's filtered position to the shuttle even during the rest period where the reference signal $r(t) = 0$. As a result, the fish was still tracking a nominally stationary target that was only moving according to the closed loop dynamics, but in a low amplitude manner that would not fatigue the fish. This is analogous to postural balance: even when you are attempting to stand upright in a stationary manner, small deviations must constantly be corrected through feedback in accordance with the unstable pendular dynamics.

In the third and final stage, we switched back to open-loop. The purpose of this phase was to measure any possible after effect that the learning experiments had on the fish. We hypothesized that if the fish had indeed learned a new method of control, this would carry over into its "normal" method of tracking for a certain amount of time. The fish would need some time to "unlearn" their new controller, and would have to readapt to baseline tracking. Therefore, by comparing the open-loop trials of the fish immediately after the learning experiments with the open-loop data from before the first phase (baseline open-loop), we could determine how much the fish's controller had actually changed. This final open-loop phase included running several sets of open-loop trials at certain time increments after the end of the learning experiments. The open loop phase began with running 3 open loop trials (1 minute each, followed

CHAPTER 2. MATERIALS AND METHODS

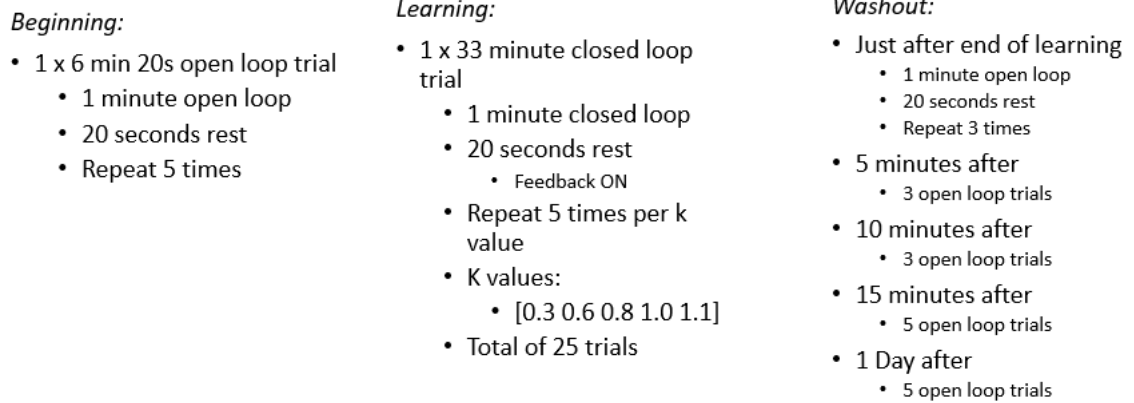


Figure 2.6: Closed-loop experimental procedure

by 20 seconds of rest) immediately after the end of the learning experiments. We expected these trials to most clearly show the after-effects of learning on the fish. We then performed 3 more trials at 5 minutes, 10 minutes, and 15 minutes after the end of learning. These time intervals were chosen under the hypothesis that full cerebellar adaptation in humans tends to occur on the time scale of roughly 15 to 20 minutes [11]. More explicitly, we expected that during these time intervals, the fish would gradually unlearn the new dynamics, and after 15 minutes of open loop tracking would have entirely adapted back to their baseline tracking abilities. Finally, we performed one last set of 5 open loop trials one day after the learning experiments, to test whether the fish had completely washed out any of the learned tendencies. We expected this data set to be virtually identical to the baseline open-loop data (phase 1). A visual of the three phases of our experiment can be seen in Figure 2.6.

2.5 Parameter Selection

In designing the feedback learning experiments, our aim was to increasingly perturb the tracking dynamics throughout the course of the experiment. According to our hypothesis, the fish would gradually tune their controllers to adapt to the new dynamics. We achieved this increase in dynamics perturbation through manipulation of the high-pass filter parameters. See [2.3](#); to remind the reader, the portion of the block diagram inside the dashed black line represents the fish itself, and will be referred to as $G(s)$ throughout this essay. The $H(s)$ block represents the high-pass filter, which was added to the system via the experimental setup. The transfer function of the entire system, which we will call $F(s)$, can be expressed as the output, the fish position, divided by the input, the reference trajectory:

$$F(s) = \frac{Y(s)}{R(s)} \quad (2.3)$$

We can also express the transfer function of the entire system as a function of the high pass filter parameters, k and b , in the following manner:

$$F(s) = \frac{G(s)(s + b)}{s + b - k \cdot s \cdot G(s)} \quad (2.4)$$

This transfer function is dependent on both k , the filter's gain, and b , the filter's cutoff frequency. The gain essentially determines how much the fish's

CHAPTER 2. MATERIALS AND METHODS

position is amplified by when it is fed back through the filter. For positive gain values, the gain moves the refuge in the same direction that the fish is moving. We hypothesized that a higher positive gain value would be more difficult for the fish to track, because it would move the shuttle farther in response to a certain movement by the fish. The cutoff frequency, b , determines what range of frequencies are attenuated by the filter. We hypothesized that lower b values would present a more difficult system for the fish to track. This is because a lower b value allows lower and lower frequencies to get passed through the filter, and therefore affect the refuge trajectory. For the fish, this means that a higher percentage of its movement will get fed back to the refuge, and the system will be more distinct from the open loop case.

In preliminary experiments by Kyle Yoshida, which were repeated on a new animal for this essay, we initially chose to keep the gain, k , constant, and only vary the cutoff frequency, b , throughout our learning trials. As previously stated, we believed the broader range of frequencies would present a sufficient challenge for the fish, and force them to tune their controllers. However, our preliminary experiments showed little change in $G(s)$, the fish transfer function (which includes both plant $P(s)$ and controller $C(s)$, but does not include the artificial feedback $H(s)$). This meant that the fish would react in roughly the same manner to the refuge regardless of the change in dynamics, which was a vexing result.

CHAPTER 2. MATERIALS AND METHODS

To analyze why the fish did not adapt to significant changes in closed loop dynamics, we chose to evaluate the dynamics through the lens of system stability. The stability of a transfer function can be evaluated by using its dominant pole. The poles of a linear system determine if the system decays or grows exponentially. Right-half-plane (RHP) poles result in exponential growth, and therefore system instability, whereas left-half-plane (LHP) poles result in exponential decay. The dominant pole is the pole with the most positive real component, and determines the overall stability of the entire system. For a system to be stable, its dominant pole must have a negative real component. To perform this analysis on our closed loop system, we need to represent the system in Figure 2.3 as a proper rational function. In order to do this, we modeled, $G(s)$, the fish's transfer function, as a 2nd order transfer function, of the form:

$$G(s) = \frac{A}{s^2 + Bs + C} \quad (2.5)$$

This is a commonly used model for the fish refuge tracking system [3,8]. We did this by gathering data from a certain fish across several trials, and fitting its empirical transfer function [16,17] to this model. As an example, using experimental data from one of our best fish, we obtained the following model for $G(s)$:

CHAPTER 2. MATERIALS AND METHODS

$$G(s) = \frac{17.35}{s^2 + 5.75s + 17.86} \quad (2.6)$$

Plugging this expression for $G(s)$ into $F(s)$ we obtain the following expression:

$$F(s) = \frac{17.35s + 17.35b}{s^3 + (5.75 + b)s^2 + (5.75b - 17.35k + 17.86)s + 17.86b} \quad (2.7)$$

The roots of the denominator polynomial of $F(s)$ will determine the stability of the system. We then proceeded to plot the value of the dominant pole of $F(s)$ as a function of a given k and b value, to determine the direct impact that these parameters had on system stability. We sampled k in the range from -3 to 3, and b in the range from -1 to 2.5. See Figure [2.7a](#). Consider a plane parallel to the $k - b$ plane, at 0 (z -axis). Any point on the surface below this plane represents a stable system, whereas any point above the plane represents an unstable system.

A contour map of this surface is shown in Figure [2.7b](#). Here, each blue contour line represents a particular dominant pole value; all contour lines with positive values are unstable, and contour lines with negative values are stable. It can be seen in the center of the plot that a dashed red line was added to represent the parameter-space location of our prior experiments. As previously mentioned, in our learning experiments we initially chose to keep the gain, k ,

CHAPTER 2. MATERIALS AND METHODS

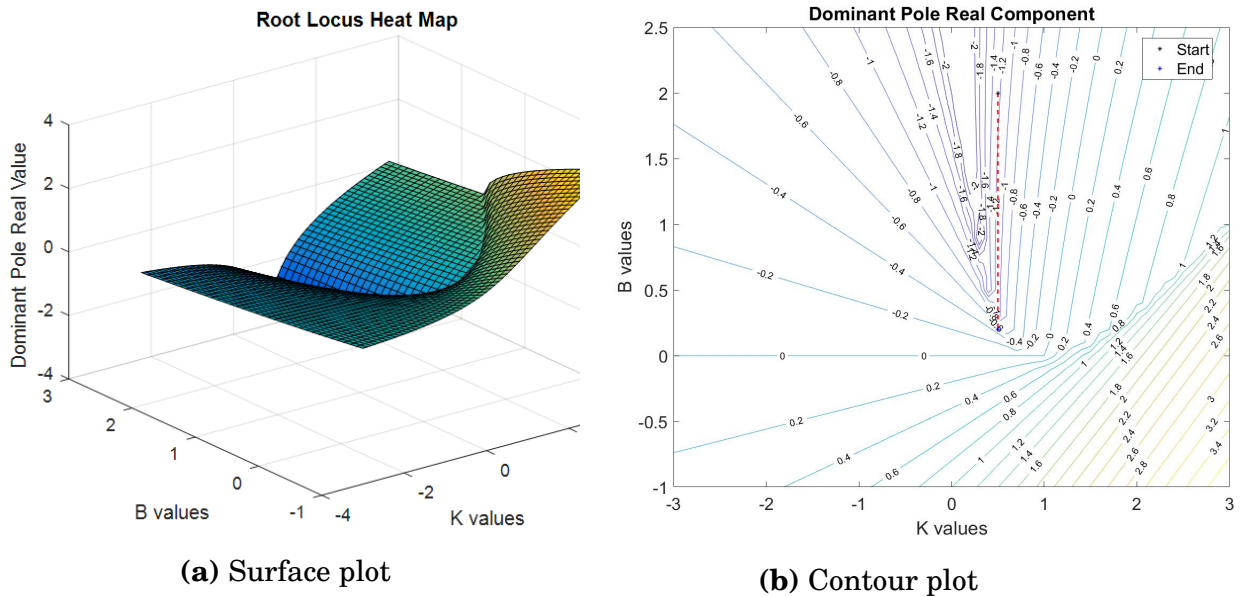


Figure 2.7: Dominant pole as a function of k and b parameters

at a constant value of 0.5. We gradually decreased b , the cutoff point of our filter, from 2.0 down to 0.2. The dashed red line correspondingly begins at $k = 0.5$ and $b = 2$, and drops down vertically to the point $k = 0.5$ $b = 0.2$. Upon closer examination of the red line's location on the contour plot, it is clear that this set of parameters does little to change the stability of the system. By and large, keeping a constant gain of 0.5 and altering the cutoff value b between 2 and 0.2 results in a path that travels roughly parallel to the contour lines. This means that the dominant pole of our system most likely remained around -1.2 throughout the duration of our experiments. While the pole value does appear to rise to -0.7 when b reaches its final value of 0.2, this proved to be too small a change to the stability of the system. This was reflected in our initial experimental results, as the data processing showed little change in the fish's

CHAPTER 2. MATERIALS AND METHODS

controller.

We strove to make the experiments more challenging by choosing variation in parameters of $H(s)$ that would increase the instability of the model. We hypothesized that parameters leading to greater instability in the nominal model would force the fish to change their controllers to maintain stability. Using the dominant pole graph as a reference, we kept b constant, and changed the gain, k , instead. Along the $b = 1$ line, and for k values greater than 0.5, it can be seen that changing the gain while keeping the cutoff frequency constant results in a path that travels transverse to the contour lines. As a result, the system stability dramatically changes as a function of k . Plotting this new parameter trajectory on the contour plot yields Figure 2.8.

The new set of k and b parameters are represented by the red arrow. The arrow begins at $k = 0.5$ and $b = 1$, and ends at $k = 1.3$, $b = 1$. At these final parameter values, the contour plot predicts that an un-adapted system would border on instability, given that the contour value is 0. Through experimentation, we confirmed the veracity of this estimate. In some trials, the system did indeed become increasingly unstable as the k value increased, and at a k value of 1.3, the motion of the fish and refuge grew, resulting in the refuge hitting one of its end bounds, thereby prematurely ending the trial. Due to this, we chose to scale back our experimental range of parameters, and settled on the final values seen in Table 2.1. Future experiments can evaluate whether the higher

CHAPTER 2. MATERIALS AND METHODS

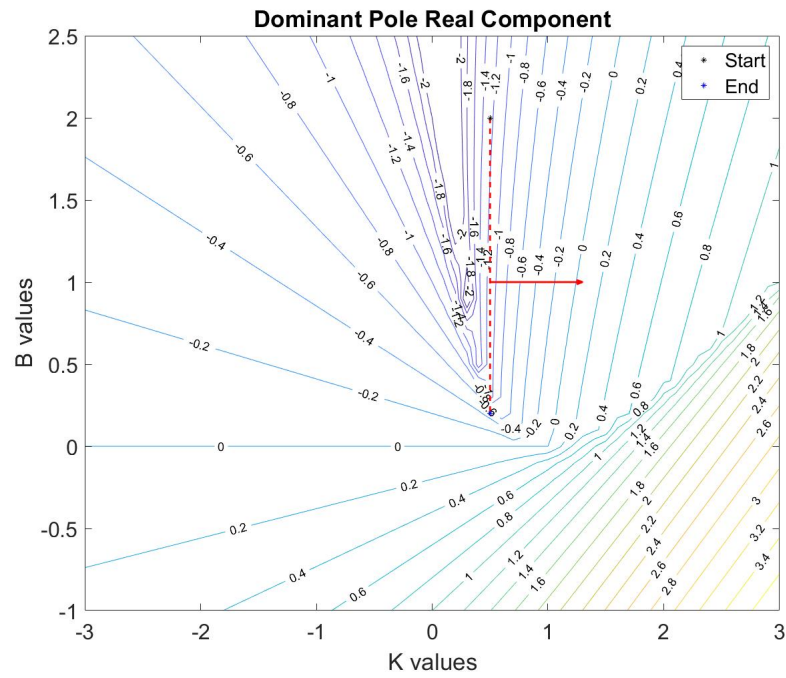


Figure 2.8: Dominant pole contour plot with final experimental values

gain values can lead to a stable adapted controller.

k value	b value
0.3	1
0.6	1
0.8	1
1.0	1
1.1	1

Table 2.1: k and b values used in learning experiments

2.6 Data Processing Methodology

In general, our data processing largely consisted of transferring our time-domain data into the frequency domain, and then generating Bode plots of our data to evaluate the frequency response. This was all performed in MATLAB. This began with combing through the fish position data, which was stored in csv files. As previously mentioned, each 60 second trial consisted of 10 seconds of “ramp-up” at the beginning, 40 seconds of sum-of-sines movement, followed by 10 seconds of “ramp-down” at the end. As a result, we only used the middle 40 seconds of every trial for further analysis.

Given that the 40 second data segment actually consisted of two repeated 20 second trajectories, we were further able to break each trial into two segments. The purpose of this segmentation was to conserve more data in case of tracking loss. A tracking loss occurred when the fish tracking software could not recognize the fish, and would result in false data in the fish position file. These losses usually lasted for less than 5 seconds, yet would force us to eliminate an entire 40 second trial. By halving each trial however, we were able to eliminate 20 seconds instead in case of a tracking loss.

After removing all erroneous trials, we converted the time domain data for the fish and the reference trajectory into the frequency domain. This was performed using MATLAB’s `fft` function, which performs a discrete fourier trans-

CHAPTER 2. MATERIALS AND METHODS

form (DFT). For open-loop experiments, we computed the transfer function $G(s)$ of the fish tracking system directly by taking the ratio of the output and the input frequency domain data. For closed-loop experiments, we first divided the frequency domain fish output signal by the frequency domain reference signal to obtain $F(s)$, the transfer function of the entire closed-loop system. We then calculated $G(s)$ according to the following equation:

$$G(s) = \frac{F(s)}{1 + H(s)F(s)} \quad (2.8)$$

Given that multiple trials were performed, we averaged the results for both open- and closed-loop experiments. A deeper analysis of our experimental results will be elaborated upon in the next section.

Chapter 3

Results and Analysis

In this chapter, I will analyze the experimental results obtained from the six fish we tested. Of the six fish, three showed very promising learning results, while the results from the remaining three were less clear. I analyzed the data through two different methods; Bode plot analysis as well as root-locus analysis. The majority of my analysis will focus on the promising fish data.

3.1 Bode Plot Analysis

With a Bode plot, we are able to assess the frequency response of a system. More specifically, a Bode plot shows both the gain of the system as well as phase of the system across a range of frequencies. The gain, or magnitude, is defined as the magnitude of the output signal divided by the magnitude of the

CHAPTER 3. RESULTS AND ANALYSIS

input signal. This describes how much the system amplifies the input signal. Given that this is a ratio of magnitudes, it is a unitless value.

The phase describes how closely synchronized the period of the output is with the period of the input. A phase lag means that the output signal is behind in its cycle compared to the input signal, whereas a phase lead means that the output signal is ahead in its cycle relative to the input. In the case of our experiment, we determined the frequency response of $G(s)$, the transfer function of a given fish. As explained before, the fish can be represented as a system that takes in a visual input of the refuge position, and outputs its own position. For reference, the block diagram describing $G(s)$ can be seen in Figure 2.2. The gain of the Bode plot for our fish describes the amplitude of the fish's movements as it tracks the refuge. In practical terms, this describes how much the fish "overshoots" or "undershoots" the refuge's motion.

In a perfect tracking system, the gain would be equal to 1 at all frequencies with zero phase lag, i.e. the magnitude of the fish's movement would be exactly equal to the magnitude of the refuge's movement, and the fish's movement would be perfectly synchronized with the refuge. A zero-phase-lag system necessitates no delays, which of course is impossible. There will always be nonzero delays due to response time and other factors contributing to the inertia of the system. However, we can draw the general conclusion that the closer the phase shift is to zero, the better the performance of the fish. With these

CHAPTER 3. RESULTS AND ANALYSIS

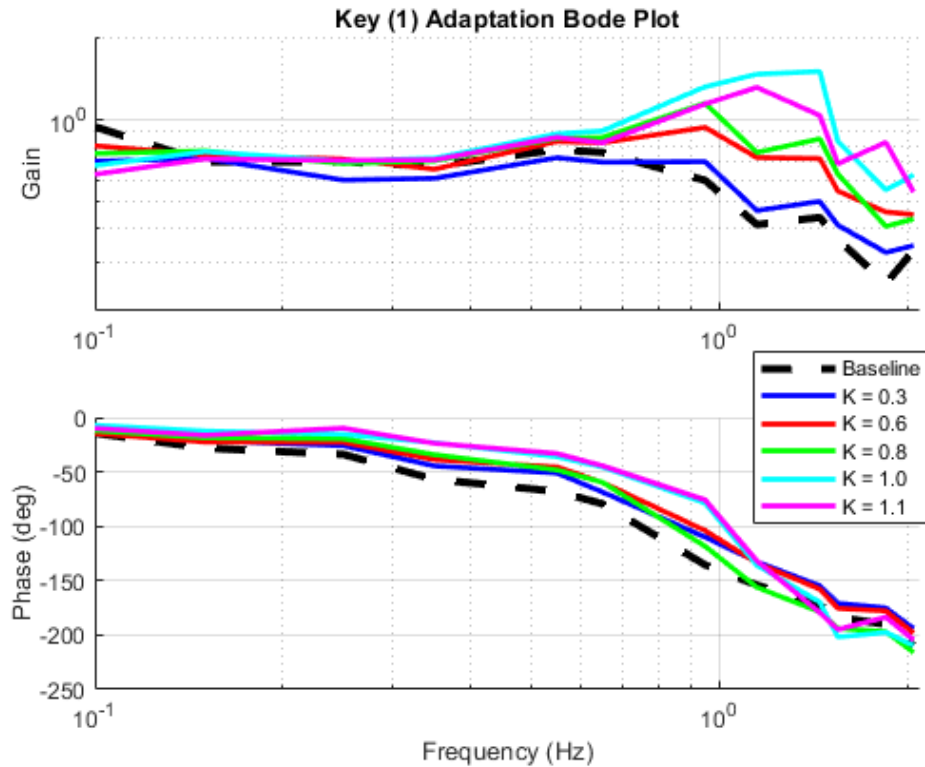


Figure 3.1: Bode plot of first Key experiment

two metrics in mind, gain and phase shift, we will examine the performance of our best tracking fish.

The Bode plot of our best tracking fish, named Key, can be seen in Figure [3.1](#). This figure shows the data from the learning portion of the experiment, which features closed-loop dynamics. In this plot, the baseline data is shown in the dashed black line. This baseline represents the fish's data obtained from open loop-tracking experiments, which occurred prior to the learning experiments.

In Key, as well as all of the fish we tested, the baseline data is characterized

CHAPTER 3. RESULTS AND ANALYSIS

by a maximum gain at low frequencies. This data point reveals that the fish is best able to track the refuge during the point in the refuge's trajectory when it is most slowly changing directions. This is logical, as it requires the least amount of control input from the fish to follow an object that is moving both slowly and consistently. In this instance, Key was able to attain a gain of nearly 1 when tracking the portion of the trajectory oscillating at 0.1 Hz. Between 0.2 to 0.4 Hz, however, the baseline gain appears to drop down to roughly 0.6. Interestingly, the gain then rises up to a local maximum of about 0.7, at a frequency of about 0.6 Hz. This local increase in gain at a higher frequency is a counter-intuitive phenomenon, yet one that we saw in many fish. Beyond this local maximum, the gain begins to drop off more sharply. This trend is logical, as beyond a certain frequency, it becomes increasingly difficult for the fish to follow the refuge as it changes directions more rapidly.

The baseline phase shift is more easily characterized, as it decreases monotonically as the frequency increases. Again, this is a logical trend, as the refuge becomes more difficult to follow under faster directional changes. When the refuge moves at higher and higher frequencies, by definition the period of the refuge's movements shortens. As a result, assuming the response time of the fish is relatively constant throughout the trial, the response time of the fish then accounts for a larger portion of the refuge's period. This results in a larger phase lag, as represented by the phase line dropping off across increasing fre-

CHAPTER 3. RESULTS AND ANALYSIS

quencies.

In terms of gain, the clearest trend we observed from the learning experiments was an increase in the gain value at high frequencies. More specifically, this spike in gain is most pronounced between frequencies of 1.0 and 1.1. In this range, the increase in gain of $G(s)$ appears to be correlated to the increase in the k parameter value for that trial. As shown in the gain plot, as the k parameter increases in value, from 0.3 to 1.1, the spike around a frequency of 1 Hz appears to get more pronounced. This trend holds true for nearly every k value - it should be noted that the gain associated with the $k = 1.0$ trial spikes higher than the gain associated with the 1.1 trial, yet these values are very similar. For trials of $k = 0.3$ and $k = 0.6$, the gain attains values of 0.7 and 0.9 at 1.0 Hz. In this latter trial at $k = 0.6$, a value of 0.9 indicates that the fish is tracking more accurately at this frequency, given that the baseline gain at 1.0 Hz is around 0.5. However, as the k values continue increasing, we see the gains increase above 1, as the fish appears to be “overshooting”. For k values of 1.0 and 1.1, the gain appears to achieve a maximum value of 1.5 and 1.3 respectively.

The physical interpretation of this behavior is that the fish is reacting more forcefully at higher frequencies to the input signal. Potential reasons why this occurs will be discussed in the following sections.

The converse to the increase in gain value at high frequencies is a slight

CHAPTER 3. RESULTS AND ANALYSIS

decrease in gain values at the lowest frequencies. It can be seen on Figure [3.1](#) that at a frequency of 0.1 Hz, the trend just described seems to reverse itself. The gain of the fish seems to decrease as the k value increases. Trials with $k = 1.1$ and $k = 1.0$ feature the lowest gain, valued at around 0.5, whereas the baseline gain attains a value of around 0.9. Physically, this corresponds to the fish “undershooting” the refuge’s movement when it is moving at the lowest frequency.

The crossover point between the decreased gain at low frequencies and the increased gain at high frequencies appears to land between 0.2 and 0.3 Hz. At this point, the gain from the higher k valued experiments increases and overtakes the gain from the lower k experiments. This can be viewed as the transition point between the two phenomena.

In the phase plot, an analogous correlation between k value and frequency response can be found. Across the majority of the frequency range, as the k value increases, the phase lag appears to decrease. This trend is most stark in the middle of the frequency range, at around 0.5 Hz. At this point, the progression of increasing k values directly follows the decrease in phase lag. While this is a tight distribution, it can be clearly seen that the largest phase lag occurs in the baseline case, and this steadily decreases as the k values increase. However, as with before, there is a crossover frequency, at which point the trend seems to reverse itself. At around 1.15 Hz, the higher k value trials appear

CHAPTER 3. RESULTS AND ANALYSIS

to have a larger phase lag than the lower k value trials. Accounting for the entire frequency range, the phase data appears to fit into a more pronounced reverse “S” shape as the k value increases. In a more quantitative sense, as the k value increases, the increase in phase lag around 1 Hz occurs more quickly. This means that for a higher k value, the portion of the frequency range below the crossover zone will have a smaller lag, whereas the portion above the crossover zone will have a larger phase lag. This forces a steeper dropoff during the crossover zone as the k value increases. In the case of Key, the crossover zone appears to occur around 1.15 Hz.

This experiment was executed again with Key. After our first data collection with Key, we decided to change the refuge we were using for fish software recognition purposes. All subsequent fish experiments were performed with the new shuttle. As a result, we decided to test Key a second time with the new shuttle to force the only variable between fish trials to be the fish itself. In this manner, we can directly compare trials between fish knowing no aspect of our experimental setup changed. However, it should be noted that we have no reason to believe that the different shuttle would have changed anything about the results. The shuttles had the same dimensions, and likely provided the fish a comparable sense of security.

The results from the second Key experiment, using the new shuttle (a gray, rectangular PVC tube, as described in Chapter 2.1) can be seen in the Bode

CHAPTER 3. RESULTS AND ANALYSIS

plot in Figure 3.2. In general, Key's data set seen below with the new shuttle displays the same general trends. As before, the gain plot shows a dip in magnitude at the lowest frequencies associated with an increase in k value. The baseline gain has a value of nearly 1 at a 0.1 Hz. However, at this same frequency, the gain for the $k = 1.1$ trial attains a value of 0.5, the lowest seen across any trial. Conversely, a spike is seen in the gain values at high frequencies as k increases. At just over 1 Hz, the gain of the $k = 1.1$ trial achieves a value of 1.24, the highest of any trial. At 1 Hz, the gain of the baseline is the lowest of any trial. The intermediary trials are scattered roughly proportionately throughout this range.

It should be noted that the trends seen in the gain plot are not as clearly defined in this trial with Key as they were in the previous trial. Many of the lines cross each other, making the correlation between k value and gain less definitive. This may be due to the overuse of this particular fish. Under repeated testing, the fish can lose interest in the task, and begin to track more poorly. It is possible this had an impact on our experimental data.

In the phase plot, the previously seen relationship between k value and phase is even more accentuated. Beginning with the baseline data, each successive k value trial strays further from the baseline beyond 0.5 Hz. Above this frequency and below the crossover point, the higher k value trials seem to have less phase lag. This is remarkably clear around 0.7 Hz, as the amount

CHAPTER 3. RESULTS AND ANALYSIS

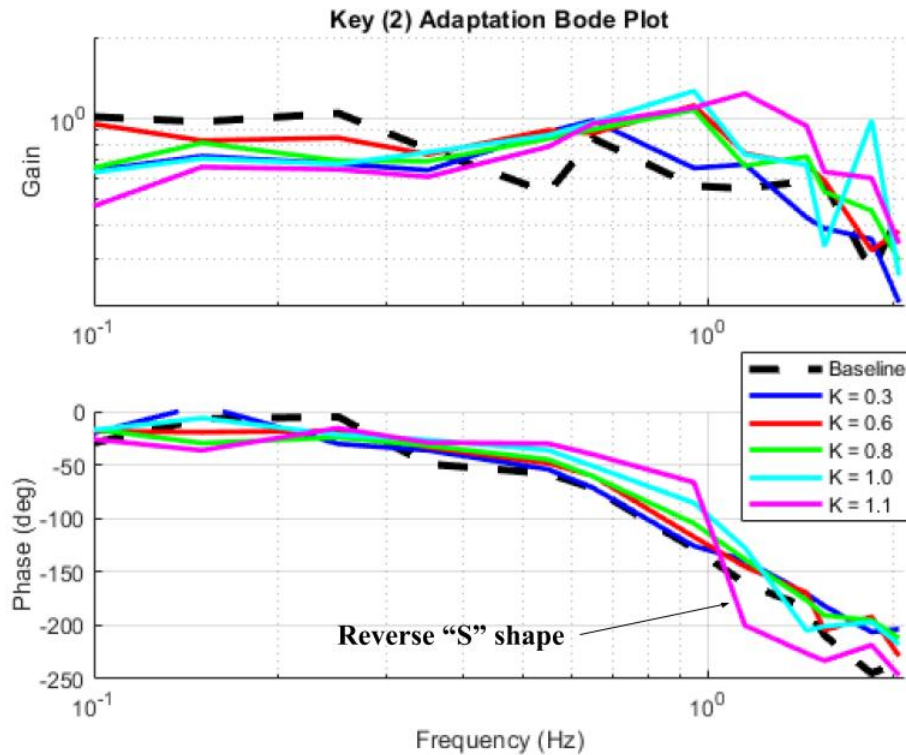


Figure 3.2: Bode plot of second Key experiment

of phase lag directly follows the k values. At 0.7 Hz, the baseline and $k = 0.3$ trial are nearly equal with the highest phase lag, followed by $k = 0.6$, $k = 0.8$, $k = 1.0$, and $k = 1.1$. The crossover frequency – the point at which higher k value trials shift to having a higher phase lag – appears to occur around 1.3 Hz. At this point, higher k value trials exhibit a far steeper dropoff in phase. This is particularly evident in the $k = 1.1$ trial, as the purple curve shows a pronounced reverse “S” shape – far more so than the lower k value trials. The next most pronounced dropoff occurs in the $k = 1.0$ trial. On the other end of

CHAPTER 3. RESULTS AND ANALYSIS

the spectrum, the most gradual change in phase lag occurs in the baseline case. Interestingly, while the gain plot of this experiment performed on Key showed a less clear trend than in the previous experiment, the phase plot appears to show a stronger correlation than before.

Another promising data set was obtained from the fish Bada. As seen in Figure 3.3, Bada showed many of the similar learning characteristics we saw with Key. Beginning with the gain plot, a dramatic change can be seen throughout the evolution of the experiment. At low frequencies, the baseline case peaks at a gain of around 1. The highest k value trials, however, achieve a maximum gain of 0.5. This equates to the low frequency gain halving as the k value increases to its maximum value. In the high frequency regions, however, the gains associated with k values of 1.0 and 1.1 achieved values of 1.23 and 1.07 respectively. Compared to the baseline gain of 0.39 at this frequency, the higher k values resulted in a nearly three-fold increase in the gain.

As before, the phase plot shows a similar steepening of the phase lag at high frequencies. Though not as exaggerated as Key's $k = 1.1$ trial, the $k = 1.1$ trial in this data set displays the strongest reverse "S" shape perturbation. The crossover point appears to occur just above 1 Hz.

Our third clear data set, resulting from testing a fish named Hope, can be seen in Figure 3.4. Hope's data shows all of the hallmark characteristics we saw with Key and Bada. The gain shows a decrease in magnitude at lower fre-

CHAPTER 3. RESULTS AND ANALYSIS

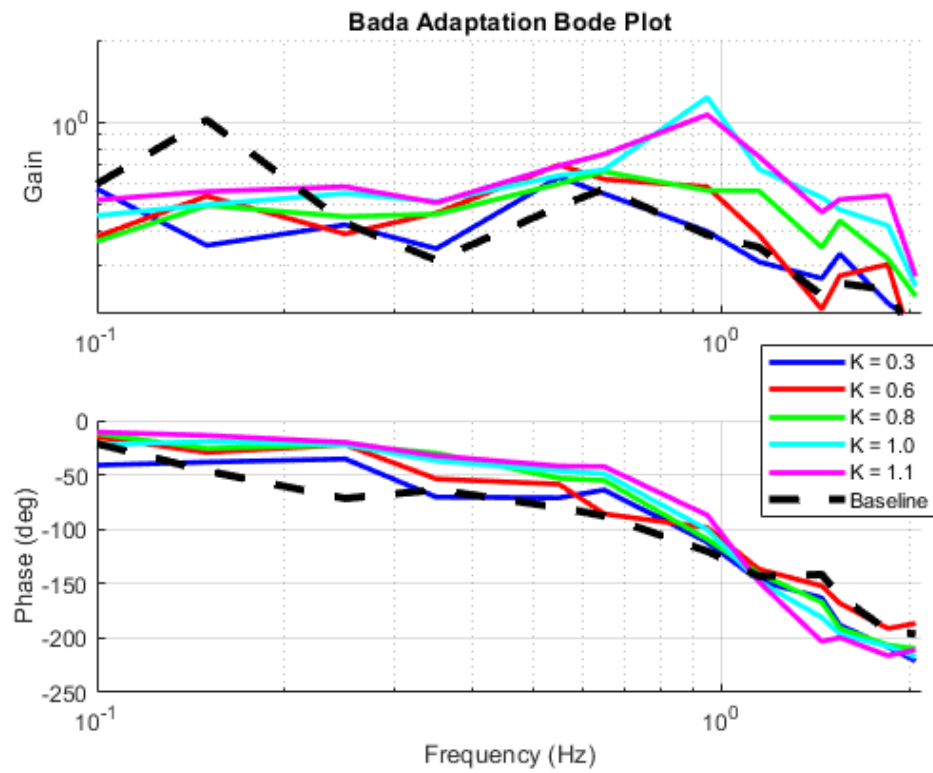


Figure 3.3: Bode plot of Bada experiment

CHAPTER 3. RESULTS AND ANALYSIS

quencies, and spike in magnitude around 1 Hz, as the k value increases. This is most clearly seen at 0.1 Hz and 1.26 Hz (the lowest and highest frequencies respectively), as the sequence of lines on the plot directly matches the sequence of k values.

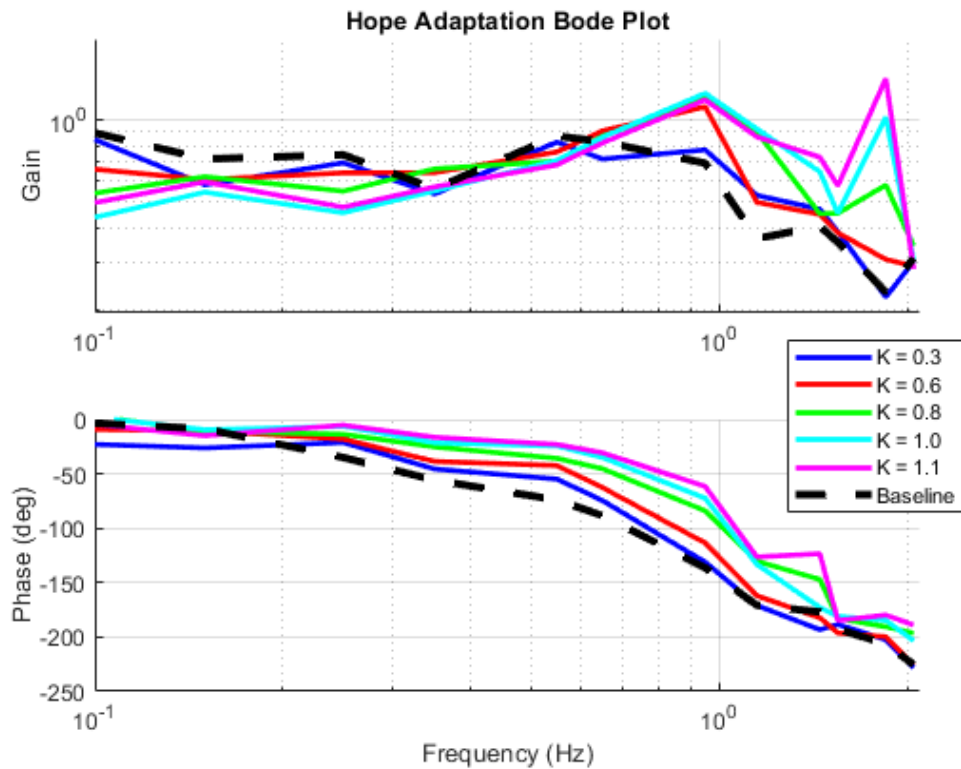


Figure 3.4: Bode plot of Hope experiment

Once again, the phase plot shows the distinct progression we have seen previously. With this specific fish, higher k value trials resulted in less phase lag for nearly the entirety of the frequency spectrum, as can be seen above. This trend is particularly clear in this data set, as the distribution of the phase lines on the plot matches the order of their associated k values. The steepening

CHAPTER 3. RESULTS AND ANALYSIS

of the phase lines can be seen increasing as the k values themselves increase. This phase dropoff, as before, is maximized in the highest k value case, as $k = 1.1$ displays a sharp increase in the magnitude of phase lag.

We can plot the change in gain at a certain frequency as a function of k value for each fish. Figure 3.5 shows the general increase in gain at 1.1 Hz for Key, Bada, and Hope.

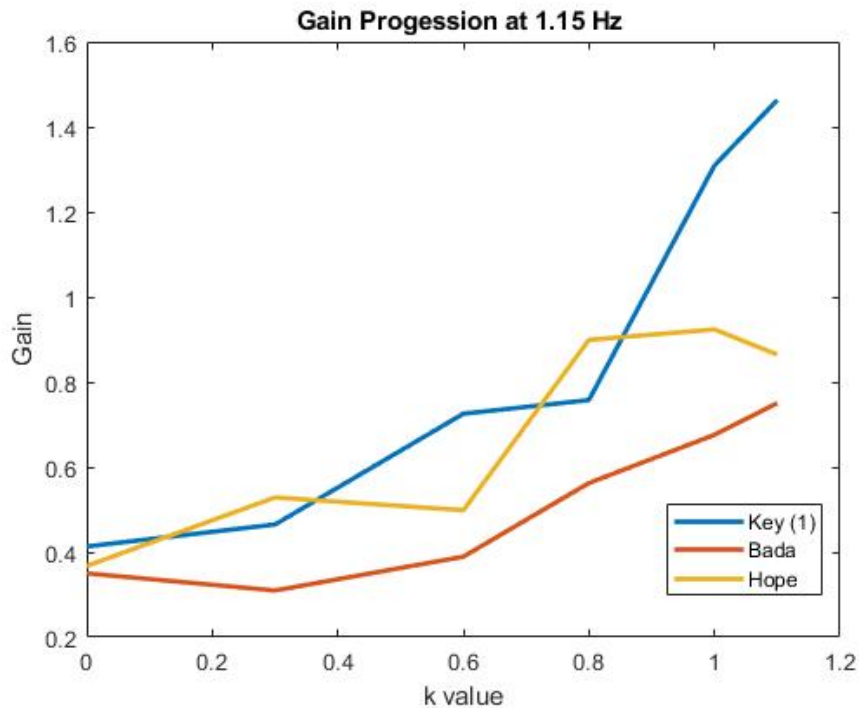


Figure 3.5: Increase in gain at 1.15 Hz as a function of k value for Key, Bada and Hope

Likewise, we can visualize the converse of this trend at 0.1 Hz, as seen in Figure 3.6.

The results from our remaining three fish, Owain, Frank, and Mac, were less clear. For comparison, we can plot the progression in gain at 1.1 Hz for

CHAPTER 3. RESULTS AND ANALYSIS

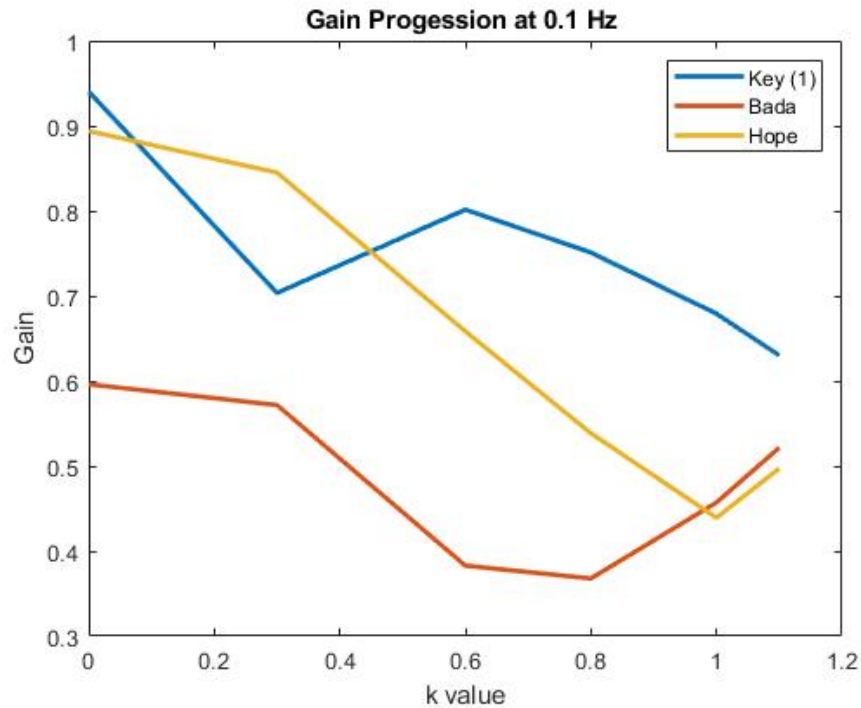


Figure 3.6: Decrease in gain at 0.1 Hz as a function of k value for Key, Bada and Hope

these three fish, as we did in Figure 3.5. This can be seen in Figure 3.7.

3.2 Bode Plot Trend Analysis – Damping Coefficient

The changes in both the gain and phase of the bode plots seen during the learning experiments are linked, and can be quantified. In essence, they both reflect a change in the parameters of the 2nd order system we are approximating the fish as. As mentioned before, we modeled $G(s)$, the transfer function

CHAPTER 3. RESULTS AND ANALYSIS

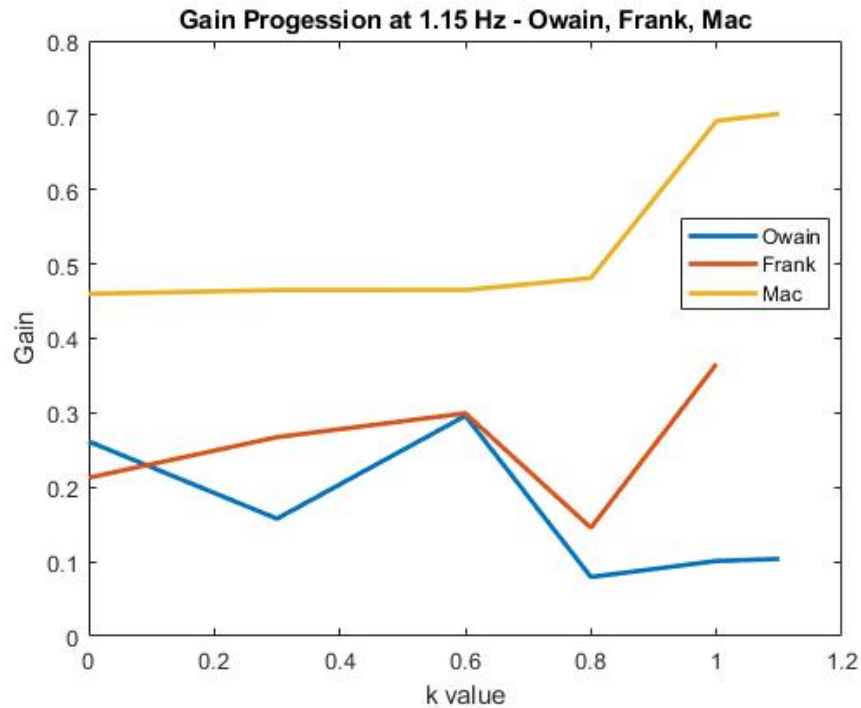


Figure 3.7: Gain progression at 1.1 Hz as a function of k value for Owain, Frank, and Mac

describing the fish, as a 2nd order system shown in equation (2.5).

This general second order transfer function can be applied to systems that involve forces proportional to the first and second derivatives of the position. A common example of this is a spring-mass-damper system [18], as illustrated in Figure 3.8.

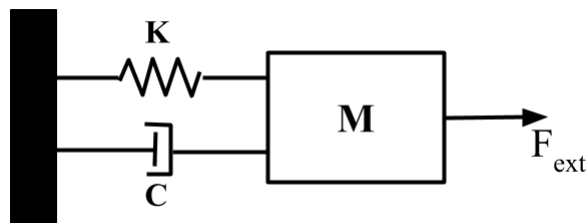


Figure 3.8: Mass-Spring-Damper System

CHAPTER 3. RESULTS AND ANALYSIS

This system features a mass of weight M , connected to a rigid surface via a spring with spring constant K , and a viscous damper, with damping coefficient of C . An external force F_{ext} is applied to the system. The sum of the forces is:

$$F_{\text{ext}}(t) = m\ddot{x}(t) + c\dot{x}(t) + kx(t) \quad (3.1)$$

Using a Laplace transform, and assuming a zero initial condition, we can express this system in the frequency domain:

$$X(s) = \frac{1}{ms^2 + cs + k}F(s) \quad (3.2)$$

In this sample system, the input is the forcing function $F_{\text{ext}}(s)$, which is a potentially varying time-dependent function. The output is $X(s)$, the position of the mass. Rearranging the frequency domain expression, have:

$$TF(s) = \frac{X(s)}{F(s)} = \frac{1}{ms^2 + cs + k} \quad (3.3)$$

We have now obtained the transfer function of the system. As with our fish, this is a second order system. This allows us to see how the parameters of the spring-mass-damper system affect the frequency response of the system, and draw an analogy to the fish dynamics.

Assuming the forcing function, our input, can be applied at a range of frequencies, we can generate a Bode plot of our system as we vary system param-

CHAPTER 3. RESULTS AND ANALYSIS

eters. A plot showing the frequency response varying with c , the coefficient multiplying the viscous damper, can be seen in Figure 3.9.

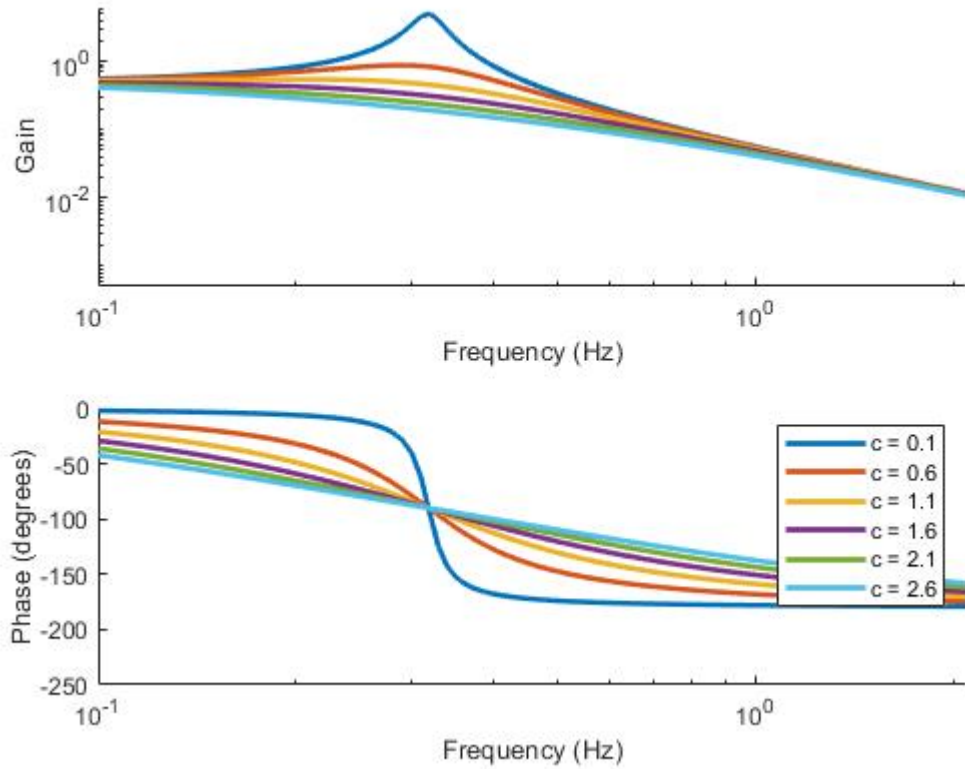


Figure 3.9: Frequency response of varying damping coefficient in 2nd order system

As we see in this plot, by varying the damping coefficient, we can generate a similar alteration in both the gain and phase plot as we observed during the fish learning experiments. More specifically, we can see that reducing the damping coefficient leads to a spike in the gain at mid to high frequencies. This is precisely the trend that occurred in the fish gain when the k value of a trial was increased. Furthermore, in the phase plot, we see that a reduction in c

CHAPTER 3. RESULTS AND ANALYSIS

results in the same reverse “S” shape curve that occurred with higher k values in our fish experiments. This leads us to the conclusion that the fish was indeed adapting during the learning experiments by changing its damping coefficient. Namely, the higher the k value we used on our high pass filter, the more the fish reduced its effective damping coefficient. Of course, this is a physical analogy describing what we believe to be a neurological change. This analogy has to be interpreted in the context of how the fish changes its controller: a fish has no viscous damper to alter, but instead changes the way it controls its movement. The reason for why this reduction in damping coefficient is beneficial for the fish, as it adapts to new dynamics, will be explored in the following section.

3.3 Root Locus Analysis

In the same manner that we chose the parameters of the high pass filter, we can use the poles, or roots, of our system to determine the system stability. In that analysis, we examined how, when we varied the parameters of our feedback, the system stability would change assuming the fish did not change its dynamics. This allowed us to choose a means by which to change the feedback to make it difficult for the fish to maintain stability. We showed strong evidence that the fish do indeed change their feedback.

Given that the fish control system does change, it presents an opportunity to

CHAPTER 3. RESULTS AND ANALYSIS

examine the closed-loop stability, but this time taking into account the change in performance of the animal. Specifically, we evaluate the stability of the evolving system transfer function to determine the effect of the reduced damping coefficient we identified above. As before, we can determine our system stability by finding the roots of the transfer function's denominator. In this case, we are evaluating the entire system, including both the fish transfer function as well as the added closed loop dynamics. This system, which we will call $F(s)$, is represented in Figure [2.3](#). The denominator of the whole system's transfer function can be seen in the equation below:

$$s^4 + \left(\frac{2}{\tau} + B + b\right) s^3 + \left(b \cdot B + C + A \cdot k + \frac{2(B + b)}{\tau}\right) s^2 + \left(b \cdot C + \frac{2(C + b \cdot B - A \cdot k)}{\tau}\right) s + \frac{2 \cdot b \cdot C}{\tau} \quad (3.4)$$

In this expression, s is the complex frequency, and b and k are the parameters of the high pass filter. As explained, these parameters were set for every given trial. The τ term derives from the time delay of the system, which we modeled with the following first order Pade approximation:

$$e^{-s\tau} \approx \frac{2 - \tau s}{2 + \tau s} \quad (3.5)$$

A , B and C represent the parameters for the second order model of $G(s)$, the

CHAPTER 3. RESULTS AND ANALYSIS

transfer function of the fish, as shown in equation (2.5). All of these parameters, with the exception of k and b , were estimated in MATLAB.

Using Bode plots of our frequency domain data, we fitted second order polynomials (with delay) to each fish transfer function $G(s)$ associated with a pair of k and b values (each line on the Bode plot). We then plugged these A , B , and C parameters into the general denominator expression seen above, and determined the dominant poles of that expression. This produced a dominant pole for each transfer function corresponding to a k and b value pair, using our experimental results for $G(s)$.

We next determined the dominant poles of the theoretical transfer function that the fish would have had given no change in $G(s)$, its own transfer function. To do this, we used the A , B , and C parameters we calculated from the open-loop baseline case, and plugged them into our denominator expression. By keeping the A , B and C parameters constant, and only changing the k and b values, we can see what kind of stability would have arisen if the fish had not deviated from its baseline tracking method (open-loop $G(s)$) in any way. This is what we call the predicted dominant poles. Each pole is a predicted value of how the system would have behaved given a certain k and b value, and assuming no change at all in the fish's controller. We plotted these poles in the complex plane. When plotting the dominant pole, it can be expressed in up to two points - complex conjugate pairs. These conjugate pairs have the same

CHAPTER 3. RESULTS AND ANALYSIS

real component, and opposite signs of their imaginary component. However, as explained before, it is the real component of the pole which determines its exponential behavior. The more positive the real component of a dominant pole is, the less quickly it will decay, and therefore the more unstable the system is. A system becomes unstable when its dominant pole has a positive real component. The plots of the actual dominant poles and the predicted dominant poles of Key can be seen in Figure 3.10. As stated above, Figure 3.10b represents the

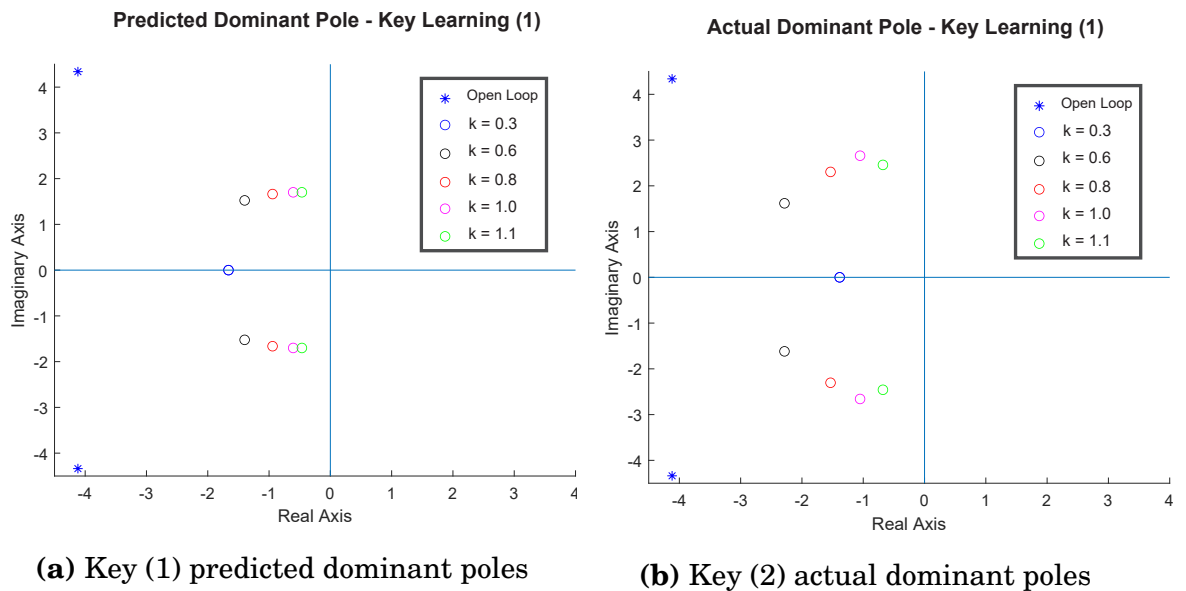


Figure 3.10: Dominant pole from first Key experiment

estimated dominant poles that Key displayed for each learning trial. Figure 3.10a shows what the corresponding predicted poles are.

As can be seen, none of the poles had positive real components. However, the most important conclusion we can draw from these plots is the shift towards stability that occurred in the learning experiments. For each pair of

CHAPTER 3. RESULTS AND ANALYSIS

dominant poles associated with a certain k value in the predicted poles plot, the corresponding poles in the actual poles plot shows a shift away from the y axis – meaning the dominant pole became more negative. Therefore, we are seeing that the new parameters the fish develops during the closed loop learning experiments serve to make the overall system more stable. This is especially stark in the points corresponding to the $k = 1.2$ trial. According to the prediction, if the fish used its baseline controller, the system would be bordering on instability, as the real component of the pole is nearly positive. In the actual dominant pole plot, however, we can see that this point is well clear of the y -axis, and is a stable pole.

This gives us some insight into why the fish changed in the manner that it did in response to the new dynamics it was experiencing: it was creating a more stable system. A more stable system results in the fish reaching its target more easily. This is due to the exponential decay of a system – the more stable it is, the more quickly it will decay to a constant state. In terms of the fish system, this corresponds to the fish and refuge converging to a stationary state more quickly. In other words, the fish renders a system more “trackable” by increasing its stability. The dominant pole plots for Bada and Hope can be seen in Figure [3.11](#).

The plots for Hope and Bada show the same characteristic shift towards increased stability. While these had some outlier points, in particular the poles

CHAPTER 3. RESULTS AND ANALYSIS

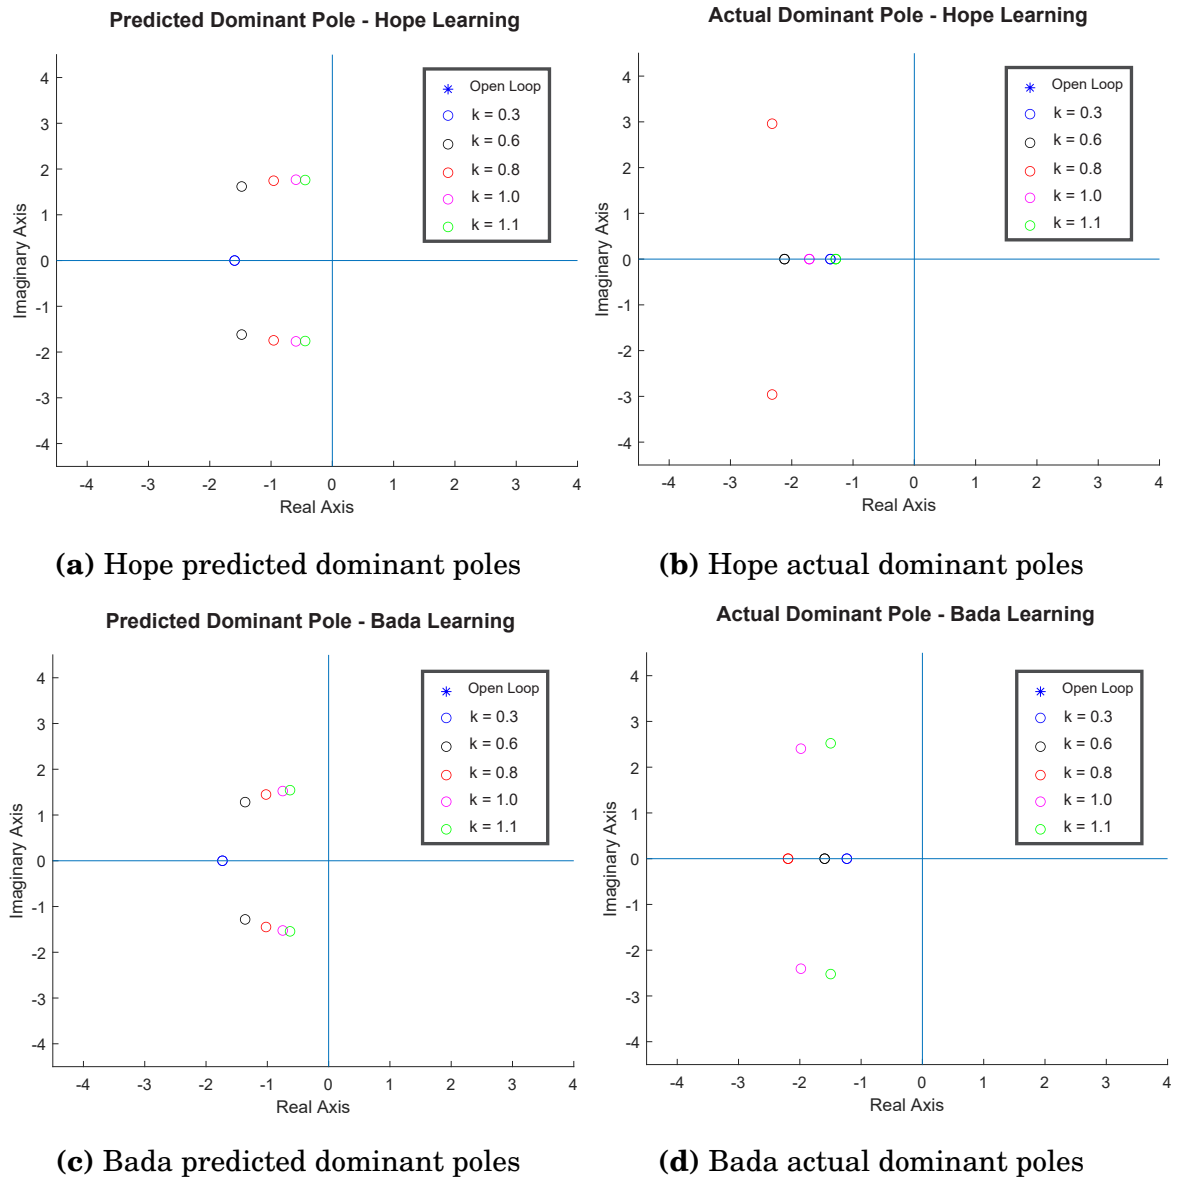


Figure 3.11: Dominant poles for Hope and Bada.

associated with $k = 1.3$, the general trend remains.

3.4 Washout

We were able to examine the after-effects of the learning trials on the fish by performing “wash-out” experiments. As previously explained, these were open loop experiments that revealed to us if the baseline tracking behavior had been altered by the learning experiments. The time dependency of these washout experiments was crucial – we expected the after effects of the learning trials to be most strongly pronounced immediately after they occurred, and to gradually fade over time. As a result, we performed these open loop washout tests over the course of the 15 minutes immediately following the learning experiments, spaced out at 5 minute intervals. We finished by performing one last set of open loop trials one day after the learning experiments, at which point we expected any learned behavior to have been completely washed out. The results from our three fish can be seen in Figures [3.12](#), [3.13](#), and [3.14](#).

The results confirmed our hypothesis that the learning experiments would have an after effect on the fish’s controller. This is most clearly seen in the open loop data taken just after the end of the learning experiments. This data is represented in red on all plots. The baseline data is represented by the dashed black line on all plots. This was collected before the learning experiments, and describes the natural tracking behavior of the fish. As is particularly evident on the bottom Key plot, the “just after” data carries the hallmark increase in

CHAPTER 3. RESULTS AND ANALYSIS

gain value around 1 Hz, and the decrease in gain value at lower frequencies.

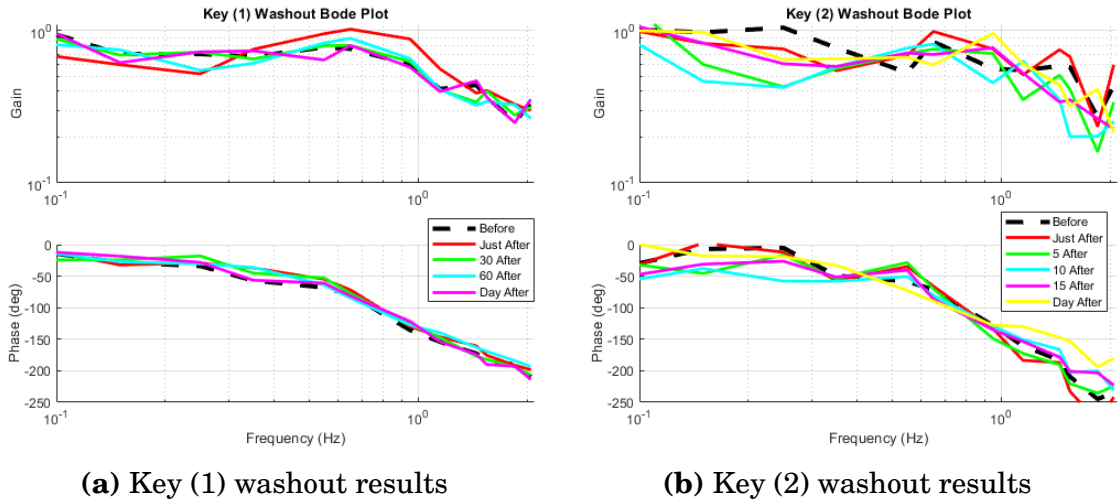


Figure 3.12: Washout results for both Key experiments

This implies that the damping coefficient that Key had learned during the closed loop experiments was still reduced compared to its baseline. This effect is also present in the Key (2) plot, yet less pronounced. However, beyond the “just after” data set, washout effects are very difficult to detect in the gain plots. Between 5 and 15 minutes, the data lines frequently cross, and do not necessarily show a noticeable change from the baseline. It should be noted, all the fish data collected from 5 to 15 minutes after the learning does show some level of reduction of the gain at low frequencies, which was seen in the learning experiment. Furthermore, the data for Bada shows a maintained increase in gain at high frequencies for both 5 and 10 minute trials. Both of these trials partially retain the characteristics learned from the closed-loop experiments. However, for the most part, there is a lack of gain increase at high frequencies

CHAPTER 3. RESULTS AND ANALYSIS

among trials later in the washout experiment.

Interestingly, in all of our washout experiments, the washout effect was far less present in the phase shift of the fish compared to the gain. For both experiments with Key, there is very little difference in the phase shift over time. All trials appear to be tightly distributed around the baseline data.

Finally, with all three fish, the open-loop results taken a day after the learning experiments show a complete washout of the learned behavior. The gain and phase of the day-after results match very closely with the data taken before the learning experiments, as expected.

The general lack of retention of learned behavior beyond a few minutes may reflect a miscalculation in our estimation of the washout time scale. Given the lifetime of baseline (open-loop) tracking that the fish experience, it is entirely possible that the fish “un-learn” the closed loop dynamics far more quickly than we expected. Fifteen minutes may have been too long a time period to expect a gradual change; this change may have occurred in its entirety across a few minutes. Despite this, the clear after effects of learning seen in the fish tested just after the closed loop experiments show that the fish had indeed changed their tracking in a tangible way. Regardless of the time scale with which this after-effect was washed out, the fish had learned a new way of control.

CHAPTER 3. RESULTS AND ANALYSIS

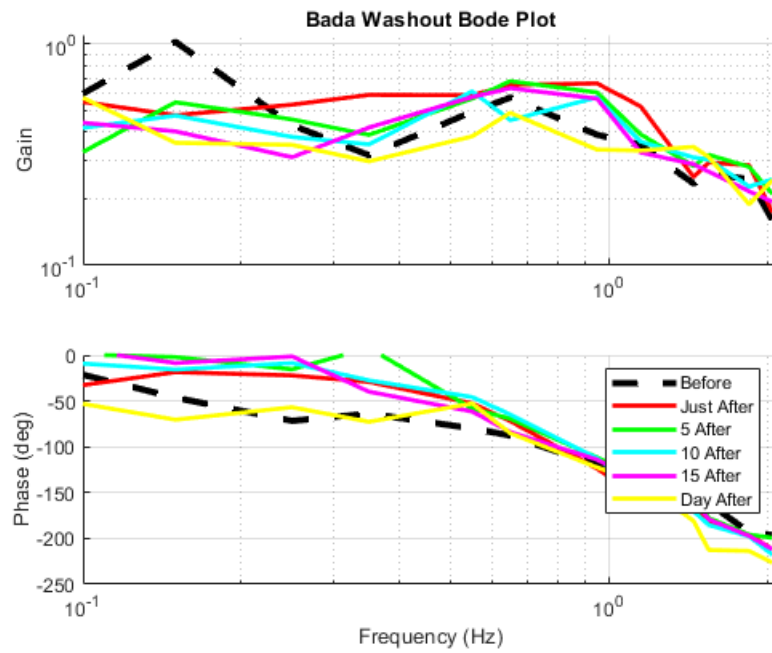


Figure 3.13: Bada washout result

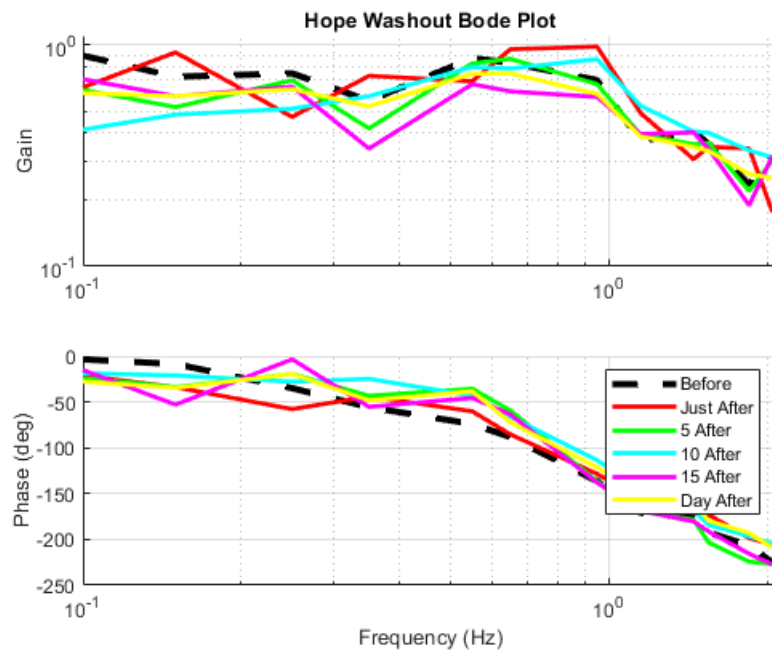


Figure 3.14: Hope washout result

Chapter 4

Conclusion and Future Work

Through our use of open- and closed-loop tests, we found that three of our six fish noticeably reduced the damping coefficients of their transfer functions, when modeled as second order linear systems with delay. They exhibited this behaviour while performing in a closed loop system with a high pass filter, where the filter parameters were manipulated to decrease the stability of the system. The damping coefficient reduction response served to increase the overall stability of the fish-refuge system, which can be linked to better tracking performance. The remaining three fish showed less clear results. The lack of a consistent trend in their Bode plots made it difficult to draw a quantitative conclusion regarding their ability to learn novel dynamics.

In our three clear results, a subtle washout trend was seen, as the after effects of the learning experiments were most evident in the open loop trials

CHAPTER 4. CONCLUSION AND FUTURE WORK

taken just after the end of the closed loop experiments. This served to more concretely substantiate the idea that the fish did adapt to the novel dynamics by changing their controller. However, our hypothesis of this washout process occurring steadily over the course of about 15 minutes proved to be an overestimate. The fish appeared to be free of after effects from the learning experiments in a shorter time frame. We suspect that the lifetime of open-loop control for the fish may have led to a shorter washout timescale.

Going forward, this research can be taken in many directions. One potential path for future study is to replay the shuttle trajectory exhibited during the closed-loop experiments, in an open-loop experiment. If the shuttle position were recorded for each closed-loop experiment, and then played back to the fish in an open-loop setting, we would be able to compare the fish's controller in open- and closed-loop modes in response to the exact same input [10]. While we already compare the fish's tracking performance in open- and closed-loop cases in my current research, this method would allow us to do so with an identical input signal to the fish.

Another direction for future research is to change the way in which we vary the high-pass filter parameters. In this paper, in order to change the stability of the entire fish-refuge system, we only varied the k value of the high-pass filter. This altered the gain of the filter. In the future, we can explore the effects of changing the b value – the cutoff point – as well. Furthermore, it is worth

CHAPTER 4. CONCLUSION AND FUTURE WORK

examining if the fish only tunes its controller in response to a change in overall system stability, or if the values of the individual k and b parameters are significant. More specifically, we can use the dominant pole level-set approach to identify multiple sets of parameters that result in the same system stability. By choosing parameters that exist on the same contour line, as shown in Figure 2.8, we can change the values in our high pass filters without affecting the theoretical stability of the system. By testing these sets of values, we can determine if overall system stability is indeed the sole driving factor causing a fish to tune its controller.

Bibliography

- [1] S. Sefati, I. Neveln, M. A. MacIver, E. S. Fortune, and N. J. Cowan, “Counter-propagating waves enhance maneuverability and stability: a bio-inspired strategy for robotic ribbon-fin propulsion,” in *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, 2012, pp. 1620–1625.
- [2] S. Sefati, I. D. Neveln, E. Roth, T. R. Mitchell, J. B. Snyder, M. A. MacIver, E. S. Fortune, and N. J. Cowan, “Mutually opposing forces during locomotion can eliminate the tradeoff between maneuverability and stability,” *Proceedings of the National Academy of Sciences*, vol. 110, no. 47, pp. 18 798–18 803, 2013.
- [3] N. J. Cowan and E. S. Fortune, “The critical role of locomotion mechanics in decoding sensory systems,” *Journal of Neuroscience*, vol. 27, no. 5, pp. 1123–1128, 2007.
- [4] E. Roth, K. Zhuang, S. A. Stamper, E. S. Fortune, and N. J. Cowan, “Stim-

BIBLIOGRAPHY

- ulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmannia virescens*,” *Journal of Experimental Biology*, vol. 214, no. 7, pp. 1170–1180, 2011.
- [5] S. A. Stamper, E. Roth, N. J. Cowan, and E. S. Fortune, “Active sensing via movement shapes spatiotemporal patterns of sensory feedback,” *Journal of Experimental Biology*, vol. 215, no. 9, pp. 1567–1574, 2012.
- [6] B. Nixon, I. Uyanik, Y. Yang, and N. Cowan, “Sensory salience affects sensorimotor delay in the tracking response of the glass knifefish,” in *Integrative and Comparative Biology*, vol. 59. Oxford Univ Press Inc Journals Dept, 2001 Evans RD, Cary, NC 27513 USA, 2019, pp. E378–E378.
- [7] I. Uyanik, S. A. Stamper, N. J. Cowan, and E. S. Fortune, “Sensory cues modulate smooth pursuit and active sensing,” *Frontiers in behavioral neuroscience*, vol. 13, p. 59, 2019.
- [8] I. Uyanik, S. Sefati, S. A. Stamper, K.-A. Cho, M. M. Ankarali, E. S. Fortune, and N. J. Cowan, “Variability in locomotor dynamics reveals the critical role of feedback in task control,” *Elife*, vol. 9, p. e51219, 2020.
- [9] K. T. Yoshida, I. Uyanik, E. Sutton, and N. J. Cowan, “A new experimental system to test how the brain learns novel locomotion dynamics,” in *Integrative and Comparative Biology*, vol. 58. Oxford Univ Press Inc Journals Dept, 2001 Evans RD, Cary, NC 27513 USA, 2018, pp. E455–E455.

BIBLIOGRAPHY

- [10] D. Biswas, L. A. Arend, S. A. Stamper, B. P. Vágvölgyi, E. S. Fortune, and N. J. Cowan, “Closed-loop control of active sensing movements regulates sensory slip,” *Current Biology*, vol. 28, no. 24, pp. 4029–4036, 2018.
- [11] A. J. Bastian, “Understanding sensorimotor adaptation and learning for rehabilitation,” *Current opinion in neurology*, vol. 21, no. 6, p. 628, 2008.
- [12] E. E. Sutton, A. Demir, S. A. Stamper, E. S. Fortune, and N. J. Cowan, “Dynamic modulation of visual and electrosensory gains for locomotor control,” *Journal of The Royal Society Interface*, vol. 13, no. 118, p. 20160057, 2016.
- [13] É. M. Hitschfeld, S. A. Stamper, K. Vonderschen, E. S. Fortune, and M. J. Chacron, “Effects of restraint and immobilization on electrosensory behaviors of weakly electric fish,” *ILAR journal*, vol. 50, no. 4, pp. 361–372, 2009.
- [14] S. A. Stamper, “Making sense: Weakly electric fish modulate sensory feedback via social behavior and movement,” Ph.D. dissertation, Johns Hopkins University, 2012.
- [15] N. J. Cowan, M. M. Ankarali, J. P. Dyhr, M. S. Madhav, E. Roth, S. Sefati, S. Sponberg, S. A. Stamper, E. S. Fortune, and T. L. Daniel, “Feedback control as a framework for understanding tradeoffs in biology,” *American Zoologist*, vol. 54, no. 2, pp. 223–237, 2014.

BIBLIOGRAPHY

- [16] L. Ljung, Ed., *System Identification (2nd Ed.): Theory for the User*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1999.
- [17] A. K. Tangirala, *Principles of System Identification: Theory and Practice*. Crc Press, 2014.
- [18] K. J. Åström and R. M. Murray, *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.

Vita

Dominic Yared graduated with his B.S. degree in Mechanical Engineering from the Johns Hopkins University in May 2019. He will graduate with his M.S.E degree in Robotics from the Johns Hopkins University in May 2020, as part of the concurrent B.S./M.S program. His research interests include controls and dynamics, particularly as they apply to biological systems.