

RESEARCH ARTICLE

Stimulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmannia virescens*

Eatai Roth^{1,*}, Katie Zhuang², Sarah A. Stamper³, Eric S. Fortune³ and Noah J. Cowan¹

¹Johns Hopkins University, Department of Mechanical Engineering, Baltimore, MD 21218, USA, ²Duke University, Biomedical Engineering, Durham, NC 27708, USA and ³Johns Hopkins University, Department of Psychology and Brain Sciences, Baltimore, MD 21218, USA

*Author for correspondence (eatai@jhu.edu)

Accepted 10 December 2010

SUMMARY

The weakly electric glass knifefish, *Eigenmannia virescens*, will swim forward and backward, using propulsion from an anal ribbon fin, in response to motion of a computer-controlled moving refuge. Fish were recorded performing a refuge-tracking behavior for sinusoidal (predictable) and sum-of-sines (pseudo-random) refuge trajectories. For all trials, we observed high coherence between refuge and fish trajectories, suggesting linearity of the tracking dynamics. But superposition failed: we observed categorical differences in tracking between the predictable single-sine stimuli and the unpredictable sum-of-sines stimuli. This nonlinearity suggests a stimulus-mediated adaptation. At all frequencies tested, fish demonstrated reduced tracking error when tracking single-sine trajectories and this was typically accompanied by a reduction in overall movement. Most notably, fish demonstrated reduced phase lag when tracking single-sine trajectories. These data support the hypothesis that fish generate an internal dynamical model of the stimulus motion, hence improving tracking of predictable trajectories (relative to unpredictable ones) despite similar or reduced motor cost. Similar predictive mechanisms based on the dynamics of stimulus movement have been proposed recently, but almost exclusively for nonlocomotor tasks by humans, such as oculomotor target tracking and posture control. These data suggest that such mechanisms might be common across taxa and behaviors.

Key words: *Eigenmannia virescens*, neural control, prediction, smooth pursuit, system modeling, weakly electric fish.

INTRODUCTION

The term ‘smooth pursuit’ typically refers to visual target tracking behaviors in the oculomotor system of foveate animals, in particular primates (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 (SPEM) and catch-up saccades, the distinction between which is typically made in kinematic terms: SPEMs are composed of continuous eye trajectories with low limits on velocity and acceleration, whereas saccades are ballistic, short-duration motions thought to correct for discontinuous positional errors, which might accumulate during smooth pursuit (Becker and Fuchs, 1969; de Brouwer et al., 2001; Rashbass, 1961).

This paper addresses a similar behavior: refuge tracking in the weakly electric glass knifefish, *Eigenmannia virescens* (Cowan and Fortune, 2007; Rose and Canfield, 1993). At the task level, fish swim forward and backward to remain within a computer-controlled moving refuge. In performing this smooth-pursuit task, fish rely primarily upon two sensing modalities, vision and active electrosensation, whereas mechanosensory cues have been found to play a negligible role in refuge tracking (Rose and Canfield, 1993).

For active electrosensation, an electric organ generates an oscillatory electric field and voltage-sensitive receptors in the skin measure fluctuations of the near field to generate an electrosensory image of the refuge. Electroreceptors are distributed over the

surface of the body with a higher density of receptors at the head, not unlike the increase in density of photoreceptors in the fovea of visual systems (Carr et al., 2004). Propulsion is generated by a ribbon-like anal fin, allowing the fish to swim both forward and backward with little body bending and without changing heading. Thus, by modulating commands to ribbon-fin motor units, the fish stabilizes its velocity relative to the refuge as encoded by visual or electrosensory images or some fusion of the two. This behavior is analogous to visual tracking of a moving scene.

Visual tracking behavior in primates is voluntary with neural mechanisms closely associated to those responsible for task attention (Khurana and Kowler, 1987). Adaptation and prediction are salient features of these primate behaviors. For example, despite substantial visuomotor delays, SPEMs can achieve zero phase lag with respect to target trajectories and persist even during target blanking (Orban de Xivry et al., 2008). In tracking horizontal piecewise-constant velocity trajectories (irregular triangle waves in position) human subjects change eye velocity in anticipation of target turnaround; Barnes and Collins proposed that the behavior incorporates a model for expected (or minimum) turnaround times (Collins and Barnes, 2009). Orban de Xivry et al. observed that while tracking a circular trajectory with target blanking, smooth pursuit and catch-up saccades occurred during blanked periods and explained the phenomena as a result of predicted target dynamics based on a ‘velocity memory’ (Orban de Xivry et al., 2008). Most relevant to the present study, Shibata et al. proposed a neuroanatomically consistent model in which target dynamics (as described by a dynamical system with

variable parameters) are learned in an online sense and used to predict future target velocities (Shibata et al., 2005).

Here we investigate the role of adaptive and predictive neural control strategies in the smooth pursuit locomotor task of refuge tracking in *Eigenmannia*. Previously, the input–output frequency response for this behavior was characterized using an assay of sinusoidal refuge trajectories (Cowan and Fortune, 2007). In fitting a linear dynamical model to the empirical frequency response, a phase roll-off exceeding 90deg at high frequencies indicated that the simplest (lowest order) model for this behavior was second-order (analogous to a spring–mass–damper). Subsequently, they used the empirically observed input–output response to predict the sensorimotor transform, and showed how this prediction depends strongly on the underlying locomotor dynamics (or ‘plant’ in control theory terminology).

This analysis is, however, predicated on the assumption that the behavior can be suitably approximated by a linear dynamical model for some salient regime of stimuli (arguments for why one might expect this can be found in the Discussion). In this work, we also characterize the behavior using frequency response analysis, although here we concurrently assess the validity of the linearity assumption: by testing the tracking response to both pure sinusoidal trajectories of differing amplitudes as well as sum-of-sines (pseudo-random) trajectories, we directly test the scaling and superposition properties that define a linear system. Our results clearly refute earlier assumptions of linearity. Fish behavior was not linear for any stimulus regime (i.e. single and sum-of-sines trials exhibited different frequency response functions). Specifically, single-sine tracking behavior exhibited broadband reduction in phase lag and high-frequency attenuation of gain when compared with the corresponding components of sum-of-sines trials. This supports the hypothesis of prediction based on a learned model of target dynamics, as proposed by Shibata et al. for visual target tracking (Shibata et al., 2005). The response to sinusoid trajectories, which are ‘predictable’, had greater predictive phase compensation than that to sum-of-sines trajectories, which appear ‘random’.

MATERIALS AND METHODS

Adult knifefish of the species *Eigenmannia virescens* (Valenciennes 1842) were obtained through commercial vendors and housed communally. 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 of 150–250 μ S. An individual fish would be placed in the experiment tank and given adequate time (2h to 1 day) to acclimatize 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 were in compliance with guidelines established by the National Research Council and the Society for Neuroscience.

Experimental apparatus

The refuge was machined from a 15 cm segment of 2 inch stock polyvinyl chloride (PVC) pipe; the bottom of the pipe was milled away to allow video recording of the fish 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 less than 0.5 cm from the bottom of the tank. A linear stepper motor with 0.94 μ m resolution (IntelLiDrives, Inc., Philadelphia, PA, USA) driven by a Stepmotor controller (Copley Controls, Canton, MA, USA) actuated the refuge, moving it forward and backward along specified velocity trajectories. Video recordings, 14-bit with 1280 \times 1024 resolution, were captured from below the refuge using a pco.1200s high-speed camera (Cooke Corp., Romulus, MI, USA) with a Micro-Nikkor 60 mm f/2.8D lens (Nikon Inc., Melville, NY, USA). For single-sine and sum-of-sines trials, video was captured at 50 frames s^{-1} ; for stimulus-switching adaptation trials, video was captured at 80 frames s^{-1} . The camera was controlled using the Camware software package (Cooke Corp.) from a standard PC. Custom MATLAB (The Mathworks Inc., Natick, MA, USA) 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, USA; Fig. 1).

Experiments

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 switched between sum of sines and single sine. One additional fish was presented with a set of sum-of-sines 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 of 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 amplitude of a given trajectory refers not to the distance but rather to the maximum




Fig. 1. (A) Experiment apparatus. The data acquisition board (DAQ) sends synchronized commands to the linear actuator (1; prescribing the trajectory) and the high-speed camera (4; triggering exposures). Riding smoothly along a set of guide rails and rigidly linked to the actuator, a rigid mast (2) suspends a PVC refuge near the bottom of the aquarium. Video is captured from below via an angled mirror (3) 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. Red and blue squares indicate the features tracked on the fish and refuge, respectively.

velocity associated with that stimulus. This is for three reasons. First, each animal might 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). Third, previous experiments (Cowan and Fortune, 2007), as well as preliminary experiments for the present study, suggest that the animals can exhibit saturation-like nonlinearities in tracking performance at high velocity amplitudes rather than positional amplitudes; as described in the 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 cm s⁻¹) and frequencies, and sums of these sinusoids. Refuge excursion frequencies (f) were drawn from the set of the first thirteen prime harmonics of 0.05 Hz, that is $f=k \times 0.05$ 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 composed of all frequency components with equal velocity amplitude (0.6, 0.8 or 1.2 cm 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 ($T=20$ 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 that persists) was inverted. As a result, averaging any two consecutive trials yielded an average input that 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 videos using 'custom code implement' 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.

Tests for linearity

Coherence analysis

Coherence, C_{vz} , the ratio of the squared cross-spectral density, R_{vz} , of two signals, $v(t)$ and $z(t)$, and the product of the respective power spectral densities, R_{vv} and R_{zz} :

$$C_{vz}(\omega) = \frac{|R_{vz}(\omega)|^2}{R_{vv}(\omega)R_{zz}(\omega)}, \quad (1)$$

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 that 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 operator, suppose the input–output pair $[u(t), y(t)]$ is related by:

$$y(t) = f(t) * u(t) + h(t) * m(t) \xrightarrow{\mathcal{F}} Y(\omega) = F(\omega) U(\omega) + H(\omega) M(\omega), \quad (2)$$

where the system $f(t)$ filters the input $u(t)$, and $h(t)$ filters a process noise $m \sim N(0, v^2)$. 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) $n_{v,z} \sim N(0, \sigma_{v,z}^2)$:

$$v(t) = u(t) + n_v(t) \xrightarrow{\mathcal{F}} V(\omega) = U(\omega) + N_v(\omega),$$

$$z(t) = y(t) + n_z(t) \xrightarrow{\mathcal{F}} Z(\omega) = Y(\omega) + N_z(\omega). \quad (3)$$

In the absence of noise ($\sigma_v, \sigma_z, v=0$) a linear dynamical system yields input–output pairs with unity coherence. Because noise variances appear only as additive terms in the denominator of the coherence function (Eqn 1), any noise introduced to the system or measurements diminish coherence, as shown here:

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

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.

Frequency response functions

A Bode plot is a graphical representation of the input–output frequency response function (FRF) of a linear dynamical system. In a Bode plot, the FRF is described using both the gain (scaling describing the level of amplification or attenuation) and relative phase imparted by a system. Input–output signal pairs that share the same Bode plot could be said to be resultant of (at least) qualitatively similar linear systems. Empirical Bode plots were generated for all trials. A fast Fourier transform (FFT) was applied to both input and output velocity signals. For single-sine trajectories we located the frequency at which the energy of the input signal peaks, ω_0 . We evaluated the output:input ratio of the FFT values at this point, $F(\omega_0)$, and calculated gain (magnitude) and phase from the resultant complex number, $|F(\omega_0)|$ and $\angle F(\omega_0)$, respectively. For sum-of-sines trials, we calculated the output:input ratio at the frequencies corresponding to the 13 greatest local maxima (excluding endpoints of the FFT) of the energy of the input signal. We verified in all cases that the 13 peaks indeed corresponded to the first 13 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 (Fig. 2). Each distribution represented the system response to a class of inputs (either single sine or sum of sines) at a set frequency. Single-sine trials yielded one point in the distribution corresponding

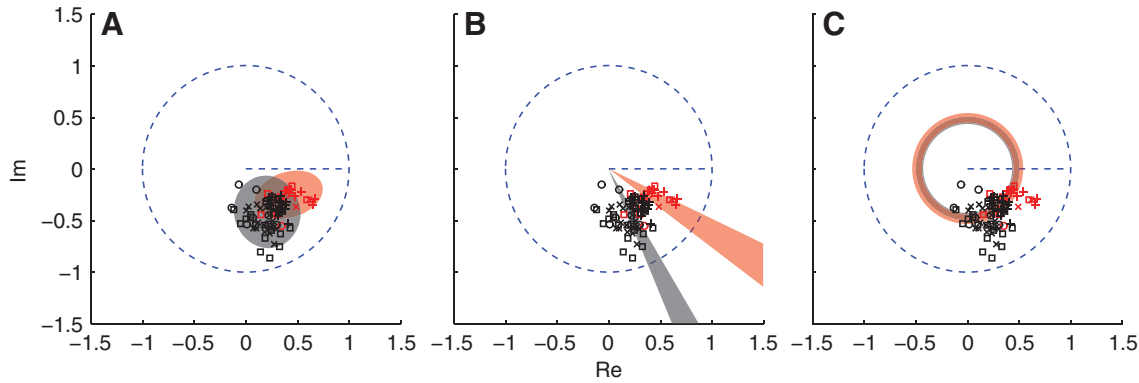


Fig. 2. For a given frequency, the system response can be characterized by a point on the complex plane, $\alpha e^{j\phi}$, where the magnitude of α is the gain of the system and ϕ is the phase shift, with positive phase measured counter clockwise from the positive real axis. The circle of unit magnitude (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) Phasor distribution. A Gaussian probability density function was fitted in the complex plane at each frequency; to illustrate this, the 95% covariance ellipse for single (red) and sum of sines (black) is shown for 2.05 Hz. (B) Phase 95% confidence. The phase confidence interval is the conic region over which the probability density function (PDF) integrates to 0.95. (C) Magnitude 95% confidence. Similarly, the magnitude confidence interval (95%) of the estimate of the mean is the annulus over which the PDF integrates to 0.95.

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 (Fig. 2A). 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 (Fig. 2B); the confidence interval for phase of the estimated mean was the minimal conic region over which the PDF integrates to 0.95 (Fig. 2C).

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 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 Eqn 5:

$$(A, B) = \arg \min_{A, B} \left\{ \sum_{i=1}^n [\alpha_i \sin(\omega_i t) + \beta_i \cos(\omega_i t)] - y(t) \right\}^2, \quad (5)$$

The frequencies $\{\omega_1, \dots, \omega_n\}$ are known and $A = [\alpha_1, \dots, \alpha_n]$ and $B = [\beta_1, \dots, \beta_n]$ are solved for in a least squares sense. Although 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 Eqn 6, we solved for the magnitude and phase of the refuge and fish as in Eqn 7:

$$\begin{bmatrix} A^T \\ B^T \end{bmatrix} = [\sin(\omega_i t) \quad \cos(\omega_i t)]^T [y(t)].$$

$$M \sin(\omega_i t + \phi) = \underbrace{M \cos(\phi)}_{\alpha} \sin(\omega_i t) + \underbrace{M \sin(\phi)}_b \cos(\omega_i t), \quad (6)$$

$$M_i = \sqrt{\alpha_i^2 + \beta_i^2},$$

$$\phi_i = \arctan_2(\beta_i, \alpha_i), \quad (7)$$




Fig. 3. A velocity-scaled repetition of a sum-of-sines trial with fish and refuge trajectories depicted in blue and red, respectively. The individual sinusoidal components have an amplitude of (A) 0.8 cm s^{-1} and (B) 1.2 cm s^{-1} .

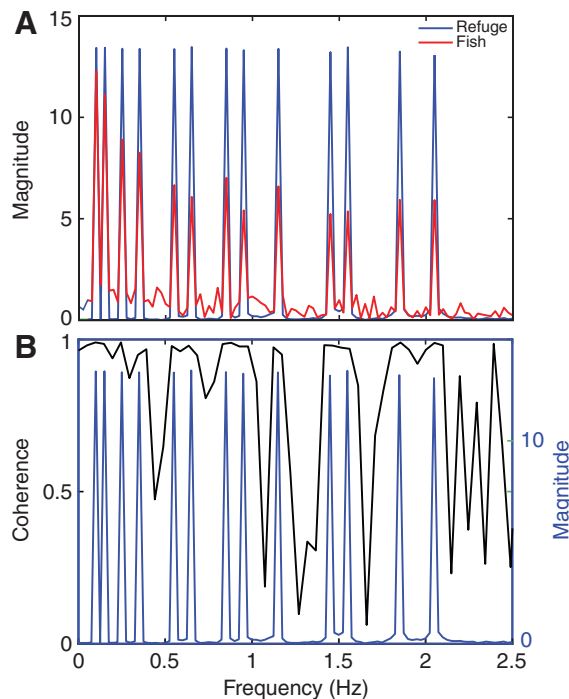


Fig. 4. (A) The magnitudes of the input (refuge) and output (fish) fast Fourier transforms (FFTs). (B) The coherence between refuge and fish trajectories (black) and the magnitude of the refuge trajectory FFT. The near unity coherence suggests fish and refuge trajectories are related by a linear dynamical system.

where \arctan_2 is the four quadrant version of the arctangent function. Gain and relative phase were then calculated as the ratio $M_{\text{fish}}/M_{\text{refuge}}$ 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 $f(t)$ is 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.

RESULTS

Responses to moving refuge 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 Fig. 3.

For each trial, we computed the magnitude of the Fourier components for input (refuge velocities) and output (fish velocities) as shown in Fig. 4A. In all instances, 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. Sum-of-sines trials consistently had coherences near unity at the stimulus frequencies (Fig. 4B). 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 measure of signal-to-noise ratio and not a measure of absolute gain (Fig. 4B).

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 the refuge-tracking behavior.

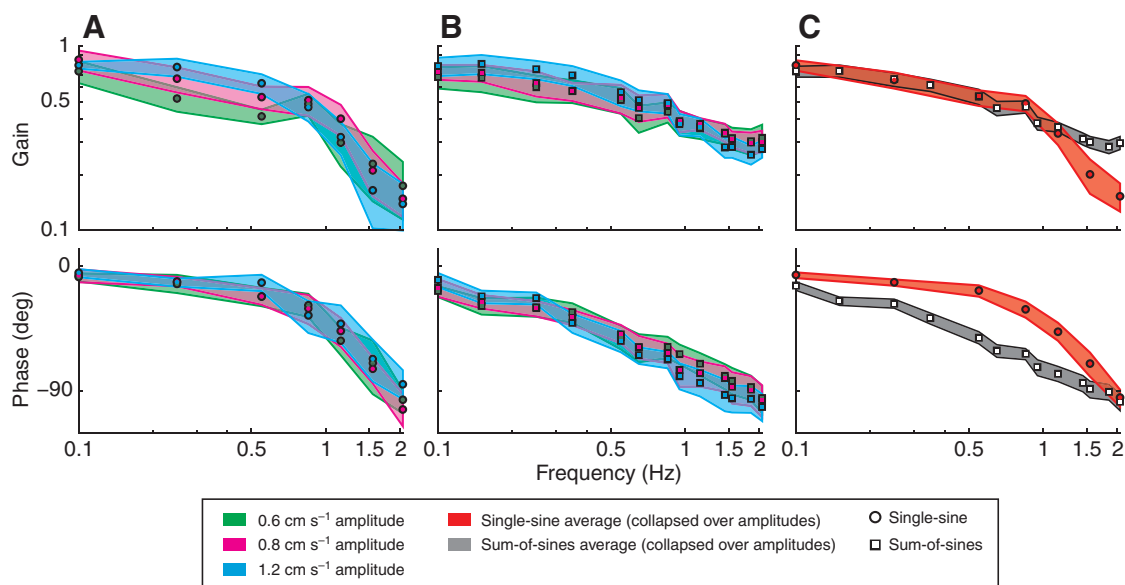


Fig. 5. The Bode plot describes the system response (gain and phase shift) over a range of frequencies. Bode plots for stimuli of different velocity amplitudes are compared for (A) single-sine trials and (B) sum-of-sines trials. (C) Average responses (collapsed across scaling), comparing sum-of-sines (black) and single-sine (red) trials, are depicted with confidence intervals calculated as described in Fig. 2. For a linear system, these estimates would agree.

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 and 1.2 cm s^{-1}) for each stimulus type (single sinusoid and sum of sines). The Bode plots for each velocity amplitude are shown in Fig. 5A for single sines and Fig. 5B 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 Fig. 5A,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 superposition property. This was done by comparing single-sine with 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 (Fig. 5) 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-sines stimuli (Fig. 5). Because the Bode plots are different across stimulus category, superposition therefore fails. A single linear model cannot account for the responses to both categories of input. 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 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 (Fig. 6). 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 Fig. 5C 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 shown in Fig. 5C 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, probably because of the nonlinearity revealed by the FRF data (Fig. 5). The consequence of the nonlinearity between stimulus categories is that fish perform better in response to predictable stimuli than to unpredictable stimuli.

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

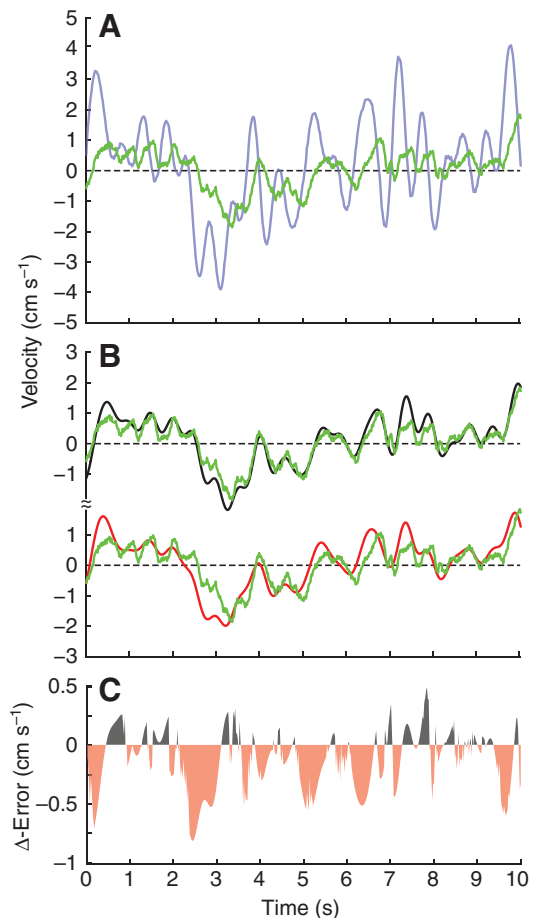


Fig. 6. An example demonstrating that frequency response functions generalize better within stimulus class than across classes. Assuming superposition holds, the frequency response functions (FRFs) in Fig. 5C generated from four fish, can be used to predict the response to an arbitrary input for a fifth fish. (A) 10 s of a sum-of-sines 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-sines 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. If the system were linear, an assay of single-sine experiments would be sufficient for predicting the response to the sum-of-sines stimulus.

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 and single sine, we present phase analyses at three timescales: one 30 s window, six 5 s windows, and a continuous phase estimate (see Fig. 7).

At the coarsest level, the asymptotic phase – calculated as the FFT estimate of phase from the second half (30) of each stimulus regime – reveal phase lag reductions of 10.8, 13.9 and 18.7 deg, which is less than the mean 33.9 deg 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-

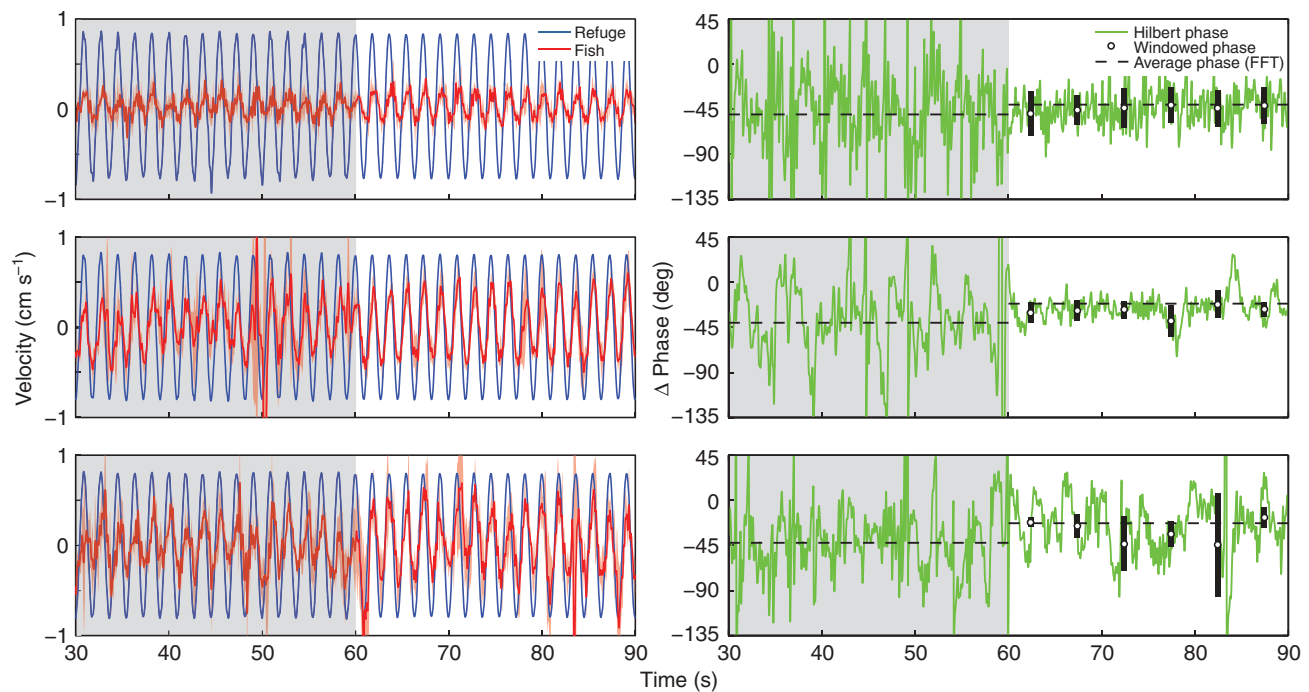


Fig. 7. In switched stimulus trials, fish ($N=3$ corresponding to the rows above) adapt to a change from sum-of-sines to single-sine stimuli. On the left are the mean of eight trials; sum-of-sines stimuli (interval shaded in grey) were designed so that the mean input would be purely sinusoidal. On the 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 s (as black error bars). The asymptotic phase over each interval was calculated using the Fourier transform method used to generate Fig. 5.

overlapping 5 s windows. For the first two fish, a trend seems to emerge, possibly suggesting an exponential decay to the asymptotic phase. However, because of 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 argument of the analytic signal to generate a continuous estimate of phase (shown in green) (Revzen and Guckenheimer, 2008). Although this approach yields a noisy estimate it highlights an important trait of the transition: the variance of the phase estimate is lower in the single-sine regime. This may partially be attributed to the method used for phase estimation but we suspect that this reduction in variance results at least in part from changes to the behavior that occur during adaptation to the switch in the stimulus.

Adaptation to single-sine stimuli reduces tracking error

Having observed categorically different FRFs elicited by single-sine and sum-of-sines 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 with the response to 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 (Fig. 8A). The intersection of the unit circle with the positive real axis, the point $1+i0$, 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 Fig. 8A, the mean phasor for each frequency is plotted for both single-sine and sum-of-sines stimuli. At every frequency compared, there is less error in the responses to the single-sine stimuli. Excluding the frequencies 0.1 and 0.85 Hz, these improvements in tracking were achieved despite a reduction of gain (Table 1). The gain (the distance between the empirical phasor and the origin $0+i0$) provides an indication of effort or energy expended during tracking. For the two frequencies 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 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 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.

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

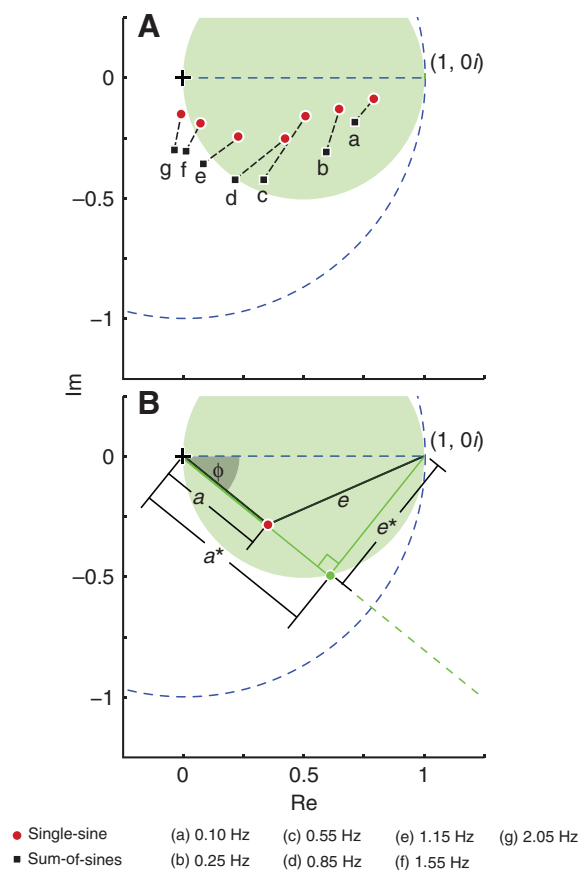


Fig. 8. (A) Frequency response phasors shift to decrease tracking error in the transition from sum-of-sines to single-sine stimuli. (B) For a frequency response characterized by gain a and phase lag ϕ , denoted by the red dot above, the magnitude of the sensory slip e is the distance from the frequency response point to the perfect tracking point $1+0i$. For a fixed phase ϕ , the minimal error $e^* = \sin(\phi)$ for $\phi \in (-\pi/2, \pi/2)$ is achieved by a gain of $a^* = \min(\cos(\phi), 0)$ (depicted as the green dot). For $\phi \notin (-\pi/2, \pi/2)$ the minimum achievable error is 1. The locus of minimum-error responses given fixed phase lag (or lead) is denoted by a circle with diameter equal to one and 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.

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*.

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

Both categories of stimuli – single sine 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-sines 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-sines stimuli than for single-sine stimuli. In this sense, single-sine stimuli are fundamentally more predictable than sum-of-sines stimuli, which we treat as pseudo-random.

Furthermore, the pseudo-random sum-of-sines stimuli are complex periodic waveforms with a long (20 s) 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-sines stimuli generalized (Fig. 6).

For mid-ranged frequencies, the gain of single-sine and sum-of-sines responses are approximately the same, but the single-sine phase lags are substantially reduced compared with the corresponding components of the sum-of-sines response (Fig. 5C). This corresponds to a substantial decrease in tracking error with little to no change in the motor effort. Moreover, complex-plane analysis (Fig. 8) reveals that at high frequencies, single-sine responses show 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 $1+i0$).

Thus, at all frequencies fish exhibit the same or less tracking error with approximately the same or less motor effort when presented with single-sine stimuli (Fig. 8). 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.

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

Table 1. Phase, gain and error differences between sum-of-sines and single-sine trials as depicted in Fig. 8A

	Frequency (Hz)	Error reduction (%)	Gain difference (%)	Phase difference (deg)
a	0.10	31.7	8.0	-7.7
b	0.25	25.3	-0.9	-15.4
c	0.55	33.8	-1.2	-33.9
d	0.85	28.8	3.9	-31.9
e	1.15	17.9	-7.9	-30.2
f	1.55	8.2	-33.4	-18.1
g	2.05	5.3	-48.7	-2.9

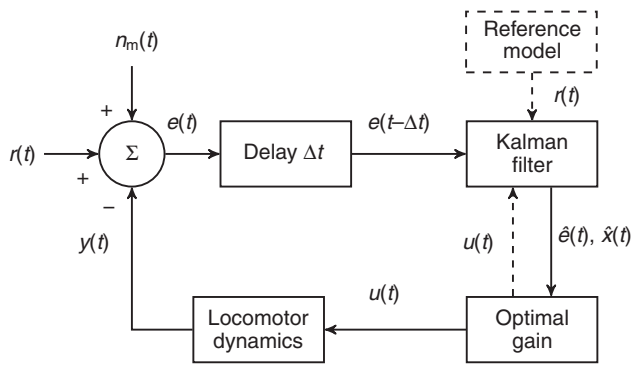


Fig. 9. A proposed model of tracking behavior that accounts for the categorical differences in phase and gain between predictable and pseudo-random stimuli. The reference signal $r(t)$ is an arbitrary refuge trajectory; the fish internally estimates this refuge trajectory as the output of an internal dynamical model. The animal's sensory systems measure the tracking error, $e(t)$, which represents the relative velocity of the refuge with respect to the fish. This measurement is corrupted by noise, $n_m(t)$, as shown. However, because of sensory and transmission delays amounting to Δt , 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 $\dot{\hat{x}}(t)$, 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 $y(t)$.

signal despite the sensorimotor delay. However, for trajectories that 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. 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 (Fig. 9) the 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 on the basis of 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 (Shadmehr and Mussa-Ivaldi, 1994; Wolpert et al., 1995)]. Although 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 a 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 assessed 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 that as the dynamics of ribbon-fin propulsion become better understood (Sefati et al., 2010; Shirgaonkar et al., 2008) our model (Fig. 9) can be used to generate quantitative predictions for refuge-tracking behavior.

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 with unpredictable stimuli. However, consider the tracking behavior transfer function on the complex plane (Fig. 2). 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 $1+0i$, indicates perfect tracking. The magnitude of the error signal (the perceived sensory slip) is measured as the distance 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 $1+0i$ than the sum-of-sines trials. At any given phase lag (or lead) ϕ , error is minimized by a gain of $\max(0, \cos(\phi))$ (Fig. 8B). 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 deg, 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.

A diversity of image stabilization behaviors

Optomotor responses in primates, including target tracking (in which a sensory image is stabilized on the fovea) and optokinetic nystagmus (OKN; tracking a moving broad-field stimulus) are perhaps the most commonly studied sensory-image-stabilization behaviors. OKN is recognized as a separate behavior from target tracking, although both are comprised of alternating epochs of slow

and fast eye motions. During the slow phase of OKN (the optokinetic response; OKR), eye trajectories are arguably indistinguishable from smooth pursuit. The fast phase of the OKN is composed of saccadic motions opposite to the direction of pursuit used to recenter the gaze. SPEMs are volitional whereas OKN is an involuntary response to broad-field motion. Both primates and humans show OKN and target tracking behaviors, but non-foveate and lateral-eyed animals exhibit OKN, although they lack SPEM (Büttner and Büttner-Ennever, 2006).

Although they are two distinct behaviors with independent neural pathways, target tracking and OKN are similar in many ways. Both share the common task-level goal of stabilizing a moving image (albeit narrow-field for tracking and broad-field for OKN). Kinematically, these behaviors are characterized by periods of smooth continuous motion interjected with abrupt corrective saccades. Most importantly, both target tracking and the OKR rely on sensory feedback. These attributes, however, are hardly unique to the oculomotor system.

Indeed, an image in the nervous system is simply a neurally coded representation of an exogenous stimulus – the output of a sensory transformation. This notion of image includes a wide range of sensory signals [e.g. in wall following, a cockroach antenna encodes a signal, hence an image, representing head-to-wall distance (Camhi and Johnson, 1999)] and sensory signals at higher brain centers, which may have already been transformed through neural processes (e.g. the spatial integral of retinal slip over the fovea). Image stabilization, therefore, refers to the class of feedback control policies in which the image – the sensory signal representing a moving exogenous stimulus – is stabilized to a sensory goal (typically target fixation or zero sensory slip) *via* motor output.

In this sense, image-stabilization tasks are ubiquitous across taxa: optomotor yaw regulation in *Drosophila*, in which a fly generates saccadic turning moments to stabilize a visual scene (Götz, 1968; Heisenberg and Wolf, 1988; Reiser and Dickinson, 2008); high-speed antennal wall following in the American cockroach, *Periplaneta americana* (Cowan et al., 2006; Lee et al., 2006; Lee et al., 2008); human posture control, in which leg muscles generate forces in response to proprioceptive, visual and vestibular cues to maintain balance (Carver et al., 2005; Jeka et al., 2004; Kiemel et al., 2006); visual prey capture in the tiger beetle (Gilbert, 1997); flower tracking in the hawkmoth, *Manduca sexta*, in which moths attempt to maintain a constant relative position with respect to a moving flower during feeding (Sprayberry and Daniel, 2007).

Image stabilization represents a model framework that can be used to describe a broad set of behaviors. Essentially, the image-stabilization description can be applied to those behaviors aimed at reducing sensory error or slip *via* closed-loop control. Often behaviors are compared on the basis of morphological similarity of the motor plant (e.g. ocular target tracking and OKN that manifest in the same mechanical system or tracking behaviors in *Drosophila* as in *Manduca*). Unified at the task level by the image-stabilization framework, we are more inclined to interpret the similarity of behaviors on the basis of their control strategies. Hence, we draw comparisons between refuge tracking in *Eigenmannia* and target tracking in the primate oculomotor system despite apparent biomechanical differences.

A role for linear models in describing image-stabilization behaviors

The frequency response analyses used in previous studies on image stabilization behaviors (Carver et al., 2005; Cowan and Fortune, 2007; Gilbert, 1997; Götz, 1968; Heisenberg and Wolf, 1988; Jeka et al.,

2004; Kiemel et al., 2006; Reiser and Dickinson, 2008; Sprayberry and Daniel, 2007) 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 to be described by such a seemingly restrictive set of models? Admittedly, nonlinearities manifest in many of the biological subsystems that 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 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 *E. 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 (Fig. 9) 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 that constrain future nonlinear models for the full behavior. Future work can address the mechanisms responsible for these adaptations.

Extending frequency analyses to other image-stabilization tasks

Similar assays to those described in this paper could be used to identify 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 obscure the task-level states of interest. For most animal behaviors of interest, the motor dynamics are cyclical (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 ~10) 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) whereas 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 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 sample of behaviors in which cyclical motor plants are controlled to achieve stationary sensory goals. But in the interest of identifying neural control algorithms, it is sometimes useful to isolate the task-level states from the ‘artifacts’ that can be introduced by such cyclic motor dynamics.

For the cases above, systems theory of cyclical 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 cyclical dynamics, in essence transforming an equilibrium cycle (or limit cycle) into an equilibrium point. Once the kinematic data are transformed to align these Floquet coordinates, data captured from different phases of a ‘stride’ can be compared using techniques such as those described above (Revzen, 2009; Revzen and Guckenheimer, 2008). 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 datum 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).

LIST OF ABBREVIATIONS

FFT	fast Fourier transform
FRF	frequency response function
OKN	optokinetic nystagmus
OKR	optokinetic response
PD	proportional-derivative
PDF	probability density function
SPEM	smooth pursuit eye movement

ACKNOWLEDGEMENTS

We thank David Solomon and Jean-Jacques Orban de Xivry for their insights regarding visual smooth pursuit and Rachel Jackson for assistance with collecting and analyzing preliminary data. This material is based upon work supported by the National Science Foundation (NSF) under grants 0543985 and 0817918, and the Office of Naval Research under grant N000140910531. E.R. was supported by an NSF Graduate Research Fellowship and an Achievement Rewards for College Scientists Scholarship. S.A.S. was supported by an NSF Graduate Research Fellowship. K.Z. received support through the Johns Hopkins Provost's Undergraduate Research Award and an NSF Research Experiences for Undergraduates supplement.

REFERENCES

- Alexander, R. (1995). Simple models of human movement. *Appl. Mech.* **48**, 461.
 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. *Vis. Res.* **9**, 1247-1258.
 Blickhan, R. (1989). The spring-mass model for running and hopping. *J. Biomech.* **22**, 1217-1227.
 Büttner, U. and Büttner-Ennever, J. (2006). Present concepts of oculomotor organization. *Prog. Brain Res.* **151**, 1-42.

- Camhi, J. M. and Johnson, E. N. (1999). High-frequency steering maneuvers mediated by tactile cues: antenna wall-following in the cockroach. *J. Exp. Biol.* **202**, 631-643.
 Carr, C., Maler, L. and Sas, E. (2004). Peripheral organization and central projections of the electrosensory nerves in gymnotiform fish. *J. Comp. Neurol.* **211**, 139-153.
 Carver, S., Kiemel, T., van der Kooij, H. and Jeka, J. J. (2005). Comparing internal models of the dynamics of the visual environment. *Biol. Cybern.* **92**, 147-163.
 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.
 Cowan, N. J. and Fortune, E. S. (2007). The critical role of locomotion mechanics in decoding sensory systems. *J. Neurosci.* **27**, 1123-1128.
 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.
 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.
 Fuchs, A. F. (1967). Saccadic and smooth pursuit eye movements in the monkey. *J. Physiol.* **191**, 609.
 Gilbert, C. (1997). Visual control of cursorial prey pursuit by tiger beetles (Cicindelidae). *J. Comp. Physiol. A* **181**, 217-230.
 Götz, K. (1968). Flight control in *Drosophila* by visual perception of motion. *Biol. Cybern.* **4**, 199-208.
 Harris, C. and Wolpert, D. (1998). Signal-dependent noise determines motor planning. *Nature* **394**, 780-784.
 Heisenberg, M. and Wolf, R. (1988). Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *J. Comp. Physiol. A* **163**, 373-388.
 Hitschfeld, A., 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.
 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.
 Khurana, B. and Kowler, E. (1987). Shared attentional control of smooth eye movement and perception. *Vis. Res.* **27**, 1603-1618.
 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, K. P. and Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* **10**, 319-326.
 Kuo, A. D. (2005). An optimal state estimation model of sensory integration in human postural balance. *J. Neural Eng.* **2**, S235-S249.
 Lee, J., Lamperski, A., Schmitt, J. and Cowan, N. J. (2006). Task-level control of the lateral leg spring model of cockroach locomotion. In *Fast Motions in Biomechanics and Robotics: Optimization and Feedback Control* (ed. M. Diehl and K. Mombaur), Vol. 340 of *Lecture Notes in Control and Information Sciences*, pp. 167-188. Heidelberg: Springer-Verlag.
 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, S., Morris, E. and Tychsen, L. (1987). Visual motion processing and sensory-motor integration for smooth pursuit eye movements. *Annu. Rev. Neurosci.* **10**, 97-129.
 Orban de Xivry, J., Missal, M. and Lefevre, P. (2008). A dynamic representation of target motion drives predictive smooth pursuit during target blanking. *J. Vis.* **8**, 6.
 Rashbass, C. (1961). The relationship between saccadic and smooth tracking eye movements. *J. Physiol.* **159**, 326.
 Reiser, M. B. and Dickinson, M. H. (2008). A modular display system for insect behavioral neuroscience. *J. Neurosci. Methods* **167**, 127-139.
 Revzen, S. (2009). *Neuromechanical Control Architectures of Arthropod Locomotion*. PhD thesis, University of California, Berkeley, CA, USA.
 Revzen, S. and Guckenheimer, J. M. (2008). Estimating the phase of synchronized oscillators. *Phys. Rev. E* **78**, 51907-51918.
 Rose, G. J. and Canfield, J. G. (1993). Longitudinal tracking responses of *Eigenmannia* and *Sternopygus*. *J. Comp. Physiol. A* **173**, 698-700.
 Sastry, S. (1999). *Nonlinear Systems*. New York: Springer.
 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.
 Shadmehr, R. and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208-3224.
 Shibata, T., Tabata, H., Schaal, S. and Kawato, M. (2005). A model of smooth pursuit in primates based on learning the target dynamics. *IEEE Trans. Neural Netw.* **18**, 213-224.
 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.
 Sprayberry, J. D. H. and Daniel, T. L. (2007). Flower tracking in hawkmoths: behavior and energetics. *J. Exp. Biol.* **210**, 37-45.
 Wolpert, D. M., Ghahramani, Z. and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science* **269**, 1880-1882.