

A Task-level Model for Optomotor Yaw Regulation in *Drosophila Melanogaster*: A Frequency-Domain System Identification Approach

Eatai Roth¹, Michael B. Reiser², Michael H. Dickinson³, and Noah J. Cowan¹

Abstract—Fruit flies adeptly coordinate flight maneuvers to seek, avoid, or otherwise interact with salient objects in their environment. In the laboratory, tethered flies modulate yaw torque to steer towards a dark vertical visual stimulus. This stripe-fixation behavior is robust and repeatable, making it a powerful paradigm for the study of optomotor control in flies. In this work, we study stripe fixation through a series of closed-loop perturbation experiments; flies are observed stabilizing moving stripes oscillating over a range of frequencies. A system identification analysis of input–output data furnishes a frequency response function (FRF), a nonparametric description of the behavior. We parameterize this FRF description to hypothesize a Proportional-Integral-Derivative (PID) control model for the fixation behavior. Lastly, we revisit previous work in which discrepancies in open- and closed-loop performance in stripe fixation were used to support the *reafference principle*. We demonstrate that our hypothesized PID model (with a modest biologically plausible nonlinearity) provides a more parsimonious explanation for these previously reported discrepancies.

Insect behavior has long inspired scientists and engineers to discover and recreate nature’s implementation of sensory processing, motor control, learning and navigation. Though insect behavior manifests from evolved mechanisms and an economy of neural architecture (merely 100,000 neurons comprise the fruit fly nervous system), the performance achieved by these animals far exceeds that of even the most impressive robots they have inspired (*e.g.* RHex [1], microflying machines [2], jumping robots [3]). Flies are remarkably adept agents within their environments, behaving robustly in response to unpredictable and changing surroundings, consolidating sensory cues from a diverse suite of modalities (*e.g.* visual, olfactory, haltere, etc.) to control a repertoire of flight behaviors. In this work, we model the feedback control of flight for stripe-fixation, an optomotor stabilization behavior in which flies modulate flight forces to orient a moving target (a laterally moving vertical stripe).

In order to constrain the dimensionality of the identification problem, we use a tethered preparation in which flies are rigidly attached to a fixture and presented with a fictive visual stimuli. This preparation has become a popular experimental paradigm for studying this class of behaviors. In free flight, naturalistic visual stimuli are extremely rich (*e.g.* luminance, contrast, spatial content, object motion, motion coherence, etc. can all be salient factors which contribute to the response) and the motor output is described kinematically with (at least) six degrees of freedom (three for rotation and

three for translation, not to mention head, body, and limb movements). In the tethered preparation, the visual stimulus and the kinematic output can each be constrained to a single degree of freedom, hence reducing the task-level locomotor behavior to a single-input–single-output system [4]–[6].

Optomotor yaw regulation has been studied extensively towards understanding many aspects of sensorimotor processing: comparing responses to small- and wide-field visual stimuli [7], [8], parsing cross-modality sensory integration [9]–[11], etc. However, empirically derived predictive models for optomotor behaviors are rare in the literature. Poggio and Reichardt [12] modeled the stripe-fixation behavior (in *Musca domestica*) in a stochastic differential equation framework, inferring the model through the propagation of motor noise given different stimulus conditions (*e.g.* salience or noise). Theobald *et al.* [13] modeled the optomotor response to wide-field point-cloud stimuli both rotating (roll, pitch, yaw) and translating (thrust, slip, lift) in each of the primary axes using binary noise sequences (impulses in optic flow velocities) and a time-domain correlation analysis to recover the impulse response function.

We approach the system identification of stripe-fixation through a series of perturbation experiments and frequency-domain analyses. In closed-loop experiments, fly motor output is measured and stabilizes stripe position through negative feedback making fixation (zero angular displacement and zero angular slip) a stable equilibrium point. All systems, in a neighborhood of (hyperbolic) equilibria, can be approximated by their local linearizations [14]. So despite the nonlinear mechanisms which contribute to the behavior (*e.g.* nonlinear flight dynamics, sensory tuning curves, motor saturations, etc), at the task-level, we assume the behavior can be captured by a linear frequency response function (transfer function). In this work, we fit a proportional-integral-derivative (PID) control model to our empirical frequency response data. Notions of PD and PID control have been referenced in past work [15], [16]; the PID structure furnishes biological hypotheses as to the contributions of different sensory information and we use this model to consolidate observations from previous studies [7], [8], [17]. Lastly, we revisit work by Heisenberg and Wolf in which a comparison between closed- and open-loop responses revealed surprising qualitative differences in performance between the two conditions. They attribute the observed discrepancies to the *reafference principle* [18]. We recreate the replay experiments and demonstrate that our hypothesized PID model may explain their observations, presenting a parsimonious alternative to the reafference principle. In terms of

¹Department of Mechanical Engineering, Johns Hopkins University, Baltimore, MD 21218 (eatai@jhu.edu)

²Howard Hughes Medical Institute Janelia Farm, Ashburn, VA 20147

³Department of Biology, University of Washington, Seattle, WA 98195

the measured frequency bandwidth and the generalizability of the resultant linear model to prior published observations, this is perhaps the most comprehensive task-level model of the stripe-fixation behavior to date.

I. MATERIALS AND METHODS

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

A. Experimental apparatus

The flight arena (Figure 1) was composed of 44 modular LED panels (each panel consisting of an 8 × 8 grid of LEDs) [19], arranged in a cylinder (4 panels high and 11 around the circumference) subtending 330 deg with a 30 deg gap at the rear. In this configuration, the circumferential pixel-to-pixel distance was 3.75 deg. Four levels of gray-scale dithering allowed for apparent motion with a minimum increment of 0.9375 deg. The visual scene consisted of two dark stripes, positioned antipodally on the cylinder, each subtending 30 deg. Positional error was measured as the angle from the fly’s sagittal plane to the front-most stripe considered the fixation target.

The fly was positioned in the center of the arena. Illuminated from above with infrared light, the fly cast a shadow on a sensor below; the wingbeat analyzer inferred wingbeat frequency (WBF) and left and right wingbeat amplitude (LWBA and RWBA) from the spatiotemporal pattern of the shadow. The wingbeat amplitude asymmetry (RWBA – LWBA = ΔWBA, measured in volts) was assumed to be proportional to the fly’s intended yaw torque [20].

A simple one degree-of-freedom model for the yaw dynamics is $J\ddot{\theta} + B\dot{\theta} = \tau$, where θ is the yaw angle, J is the moment of inertia, B is the linear rotational damping, and τ is the wing-generated torque. The damping time constant J/B , compared to the duration of a maneuver, determines whether the damping forces dominate [21]: if J/B is small then second-order model can be further reduced to first order. Indeed, recent modeling suggests damping dominates [22], thus we assume that the yaw dynamics further reduce to $\dot{\theta} \propto \tau$. This justifies, to some degree, the long-standing tradition of simply scaling ΔWBA and treating it as the angular velocity of the fly. We implemented this as

$$\dot{\theta} = K\Delta WBA, \quad (1)$$

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

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




Fig. 1. *Apparatus* (1) The LED arena displays an oscillating vertical bar, eliciting a tracking response from the fly. (2) Illuminated from above, the fly casts a shadow on a sensor below and the wingbeat analyzer extracts the salient kinematic descriptors from this wingstroke silhouette. (3) The differential amplifier subtracts the fly motor output (ΔWBA) from the reference trajectory (prescribed in software) to generate the error velocity (sensory slip). (4) The Flight Arena controller integrates velocity error and displays an updated stripe position.

B. Experiments

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

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

$$A(\omega) = (0.0153 + 0.0044(\omega))^{-1} \text{ deg} . \quad (2)$$

This relation prescribed an angular amplitude that decreases with increased frequency but an angular velocity that increases with frequency, addressing physical limitations of the animal behavior and constraints imposed by the LED display. At high frequencies, the velocity was constrained from above by the motor output of the fly (free-flight saccadic rotations are estimated at $1800 \text{ deg} \cdot \text{s}^{-1}$ [21] and the limitations on smooth tracking are significantly lower) and the angular amplitude was constrained from below by the resolution of the LED arena (peak-to-peak travel must be greater than a

single pixel). At 11.5 Hz, this relation yielded a peak-to-peak amplitude of 6 deg (approximately 1.5 pixels translation) and a maximum velocity of 216.8 deg · s⁻¹. At low frequencies, we bounded the amplitude of the exogenous reference to 60 deg, though the displayed error signal was allowed to exceed this bound as a result of feedback.

Replay paradigm: In 1988, Heisenberg and Wolf explored the role of feedback in processing exogenous (reference) and reafferent (self-generated) motion stimuli using an error replay experiment. This experiment paradigm compares the responses to two sequential perturbations, first in closed-loop and then in open-loop. In the first presentation, the fly's generated torque (or for our set-up, the ΔWBA) stabilizes the stripe position through negative feedback and the error signal, the displayed motion stimulus, is recorded. In the second presentation, this recorded error is replayed as the motion stimulus in the absence of any stabilizing feedback (*i.e.* the stimulus is unaffected by the fly motor output).

The replay experiments were recreated for oscillations of 0.1 Hz (as used in [15]) with the same visual stimulus and position-frequency relation (2) used in the system identification assay. Flies ($N = 10$, distinct from the sample used for identification) were presented three periods (30 s) of oscillation in closed-loop, followed by 3 s of reward fixation and then the open-loop replay. Flies completed six repetitions of closed- and open-loop pairings.

C. Analysis

The frequency analysis and subsequent linear transfer function model faithfully describe the stripe-fixation behavior for some neighborhood about the fixation equilibrium (stripe position and velocity are both zero). Acknowledging this caveat, we discard data points for which the stripe is not reasonably frontal, those for which the magnitude of the error signal exceeds 60 deg (Figure 1 *upper left*). Additionally, we ignore data for which the instantaneous wingbeat frequency falls below the mean WBF (as calculated per individual across all presentations), eliminating data for which the fly was fatigued or otherwise unmotivated.

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

$$(\alpha, \beta) = \arg \min_{\alpha, \beta} \sum_{t \in T} \left\{ \left(\sum_{i=1}^k \alpha_i \sin(\omega_i t) + \beta_i \cos(\omega_i t) \right) - y(t) \right\}^2. \quad (3)$$

The magnitude and phase of each component are then calculated as $M_i = \|\alpha_i \beta_i\|_2$ and $\phi_i = \arctan_2(\beta_i, \alpha_i)$;

denoting the behavioral response as a transfer function $F(\omega)$, gain is calculated as the output–input ratio of magnitudes, $|F| = M_{out}/M_{in}$, and phase relation as the difference of output–input phases, $\angle F = \phi_{out} - \phi_{in}$. For chirp signals, we perform a short-time LSSA, with non-overlapping windows of 10 s for $t \in (0 \ 40]$ and 4 s for $t \in (40 \ 120]$. The least-squares minimization in (3) is modified to accommodate time-varying frequency and signal amplitude where $\theta(t)$ and $\omega(t) = \frac{d\theta}{dt}$ are prescribed and $A(\omega)$ is determined by (2). The short-time LSSA assumes that for sufficiently small deviations in frequency (corresponding directly to window size for the chirp stimulus) the frequency response $F(\omega)$ is nearly constant:

$$(\alpha, \beta) = \arg \min_{\alpha, \beta} \sum_{t \in T} \left\{ \alpha A(\omega) \sin(\theta(t)) + \beta A(\omega) \cos(\theta(t)) - y(t) \right\}^2. \quad (4)$$

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

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

where

$$F_{\Delta}(\omega, \alpha) = \frac{F(\omega, \alpha)}{F^*(\omega)}. \quad (5)$$

The weight $W_1 \in \mathbb{R}$ is a free parameter chosen to favor data according to frequency; negative values of W_1 more heavily penalize low frequency errors and positive values penalize errors at high frequencies. $W_2 \in [0 \ 1]$ scales the relative costs contributed by gain and phase.

II. RESULTS

A. Empirical frequency responses

We describe the system as an input–output relation from the stripe position (measured as the azimuthal angle from the sagittal plane of the fly) to the fly motor output (measured as a voltage proportional to the difference in wingbeat amplitudes, ΔWBA). An example of the motor response (*green*) to a sum-of-sines trajectory (*black*) is shown in Figure 2A, both in the time domain (top) and the frequency domain (bottom). Though the fly motor output has significant stochastic components (noise and drift), comparing the magnitude of the frequency spectra of the input–output pair as calculated by a fast Fourier transform, allows us to disambiguate the response to the moving stimuli from motor noise. For this trial, the stimulus trajectory is a sum of sinusoids at frequencies of 1 and 3.5 Hz. In the frequency domain, the fly's motor output comprises two distinct peaks coinciding with the stimulus frequencies; we consider these peaks to be the response to the stimulus and disregarded

other spectral content as extraneous motion. Though in many trials, these response peaks are easily discernible, this is not always the case. We proceed under the assumption that for each trial, the response to the stimulus is the response measured at the stimulus frequencies. In the remaining analyses, we use LSSA (see Section I-C) to extract only the frequency components of interest (those present in the stimulus), understanding that, at times, the response data may be highly corrupted by extraneous motor output, or worse, that a responsive behavior is not present at all and we are sampling only stochastically driven motor output.

In the first suite of experiments, flies were recorded fixating a dark vertical stripe oscillating with sinusoidal, sum-of-sines, and logarithmic chirp trajectories. The empirical FRF (shown in *black* in Figure 2B) is calculated from the chirp stimulus using a moving window LSSA over the frequency range 0.1–11 Hz. The response to pure sinusoids (*filled circles*) and sums-of-sines (*open circles*) are superimposed on the FRF. The response to sinusoids coincides well with the chirp response; this is expected since for any small window of time (small with respect to the rate of change of frequency), the chirp stimulus resembles a pure sine and for any small band of frequencies (with caveats discussed in Section I) the system response should be similar. More importantly, the sums-of-sines responses provide frequency response estimates consistent with the empirical FRF, substantiating the superposition property and further motivating a linear model.

B. A linear transfer function model

The FRF reveals low-pass dynamics, as has been repeatedly observed and reported in the literature [7]. The high-frequency decay (estimated as 18.7 dB · decade⁻¹ from measurements at 3.65 and 9.33 Hz) suggests a transfer function with relative degree of one. Phase rolls off faster than expected for the assumed relative degree, which likely results from a delay in sensorimotor processing. Also notably, the gain plot exhibits a distinctive notch at 1 Hz. In order to capture these salient features, a linear model would require at minimum three poles, two zeros and a delay. The following transfer function from positional error $E(s)$ to motor output $Y(s)$ was fit to the empirical FRF as described in Section I-C:

$$\frac{Y(s)}{E(s)} = \exp\{-0.032s\} \frac{0.181s^2 + 1.23s + 8.68}{s^3 + 20.6s^2 + 277s + 1098}. \quad (6)$$

The transfer function has a real pole at -5.72 and complex poles at $-7.43 \pm 11.70i$. Dividing numerator and denominator polynomials by the real pole, we reformulate (6) into a proposed PID control model:

$$\underbrace{(s^2 + 14.85s + 192.1)}_{\text{motor dynamics}} Y(s) = \exp(-0.032s) \underbrace{\left(0.181s + 0.196 + \frac{7.55}{s + 5.72}\right)}_{\text{sensory weighting - PID (leaky integrator)}} E(s). \quad (7)$$

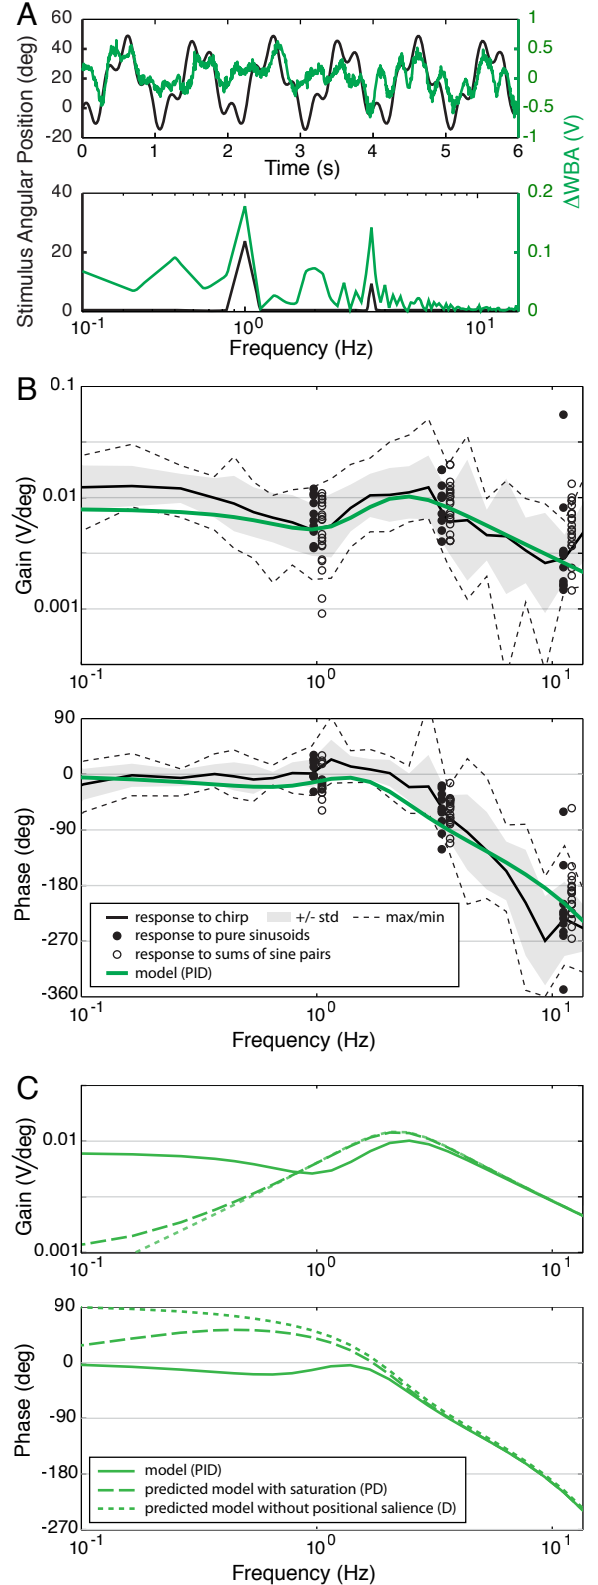


Fig. 2. (A) In the frequency domain (*bottom*), the motor response (*green*) to stimulus motion (*black*) is discernable from extraneous motor noise. (B) The FRF empirically calculated from chirp stimuli is compared to the responses to single and summed sine waves. Agreement between these responses motivates a linear transfer function model (*green*). (C) A saturating integrator model results in a PD model with phase and gain differences in accordance with Heisenberg and Wolf's replay experiments. A derivative (optic flow) model generalizes the model for wide-field stimuli.

While the FRF is unique for any given data set, many such parameterized models may be hypothesized, each furnishing a different mechanistic interpretation. In fact, there are fundamental limits in our ability to tease apart sensory dynamics from downstream processing due to the likelihood of pole-zero cancellations [23]. Notwithstanding this limitation, we parse the transfer function so that the left-hand and right-hand side equations describe the motor dynamics and sensory processing, respectively. Under this interpretation, the motor plant is described as a forced second-order (spring-mass-damper) system driven by a neural control signal. In the right-hand side equation, the control signal is calculated as a weighted sum of sensory measurements of sensory slip, positional error, and accumulated error. The fit visual-motor delay is 32.1 ms; this delay has been previously estimated at 40 ms for *Drosophila* [15], [24].

This model framework consolidates several previously published observations. Flies exhibit increased motor output for progressive (front-to-back) stripe motion compared to regressive (back-to-front) motion [8], [17]. This observation is consistent with the proposed PID control policy. In response to progressive motion (*e.g.* a stripe to the left moving leftward), the positional and derivative responses contribute constructively to the motor output; conversely, for regressive motion (*e.g.* a stripe to the left moving rightward), the positional and derivative responses contribute opposing torque commands yielding a reduced net torque. Additionally, Duistermars *et al.* observed that for wide-field (repeating stripe with period of 30 deg) visual scenes at 0.1 Hz, motor output was phase advanced to the stimulus motion by a quarter period while for small-field (single stripe) the input and output phases were practically locked. Assuming that this spatial periodicity suppresses the majority of positional and integral error (modulo 30 deg), we propose a derivative (optic flow) model for yaw regulation of this scene (Figure 2C in *short dash*); as observed by Duistermars *et al.*, at low frequencies this response to the wide-field scene phase leads the single-stripe (PID model) response by approximately 90 deg. We note that our model is inconsistent with their observations of response gain.

C. Revisiting the replay paradigm

For linear *stable* optomotor dynamics, the response to the replay experiment (as described in Section I-B) is expected to be qualitatively similar to the original closed-loop response excepting that in open-loop, there is no mechanism for the attenuation of measurement noise. Moreover, it can be shown that measurement noise attenuated in the closed-loop presentation reappears undiminished in the replay response. But Heisenberg and Wolf [15] observed qualitative differences between the closed-loop and replay responses: replay responses were phase leading by 50.9 deg (SEM = 48.6 deg, $N = 10$) and attenuated by approximately 50-65%. This discrepancy between the closed-loop and replay responses was presented as evidence corroborating the efference principle, that the disagreement between a feed-forward expectation (efference) and the sensory measurement

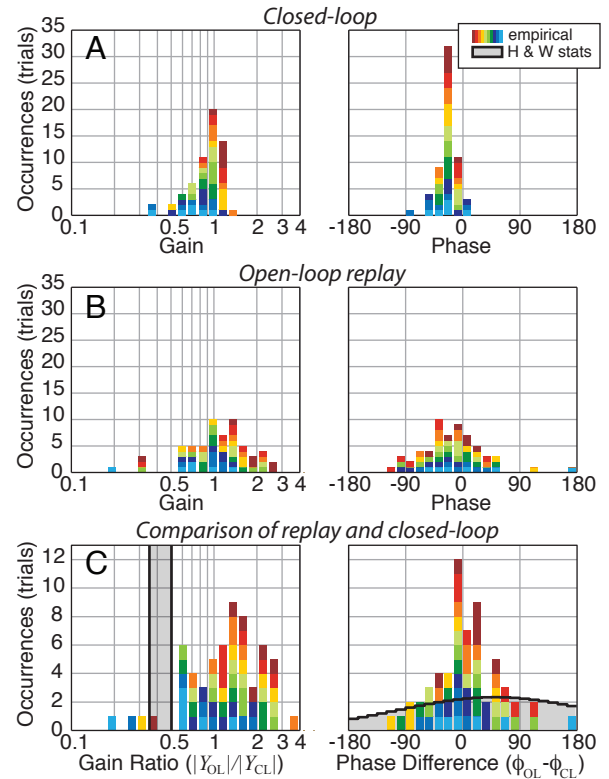


Fig. 3. Empirical histograms of gain and phase for (A) closed-loop and (B) open-loop replay trials illustrate the sensitivity relation between the two feedback conditions. Colors correspond to trials from different individuals with the warmer colors assigned to flies with higher mean gain. (C) We present the ratios of the two responses and compare the results to Heisenberg and Wolf's reported statistics [15].

of self-motion (reafference) initiates a categorical switch in behavior. Further, Heisenberg and Wolf concluded that such differences could not be the consequence of a linear system.

In recreating the replay experiments, we achieved the results predicted by a linearity assumption. Open-loop responses, both in terms of wingbeat asymmetry and frequency, were qualitatively similar to closed-loop observations but with greater variability (Figure 3). Though increased variance in open-loop makes it difficult to argue whether Heisenberg and Wolf's reported differences were significant, we proffer an explanation derived from our proposed PID model. If we allow the integrator to saturate (a nonlinearity inherent in any physical instantiations of an integrator), then the system responds quite differently about equilibrium when compared to the saturated regime. In closed-loop, every internal state is maintained at equilibrium and as a result, the integrator contributes to the output proportional to the accumulated error (which varies with time). In open-loop, small biases in noise or initial condition (*e.g.* measurement bias, asymmetry in baseline wingbeat amplitudes, fly misalignment, etc), which would be mitigated in closed-loop, can cause the integrator to drift to saturation. Once saturated, the integrator no longer conveys changes in the accumulated error, contributing a constant offset to the motor output, effectively silencing the integral path in sensory processing. Figure 3C compares

our proposed PID model about equilibrium (*solid line*) to a model prediction for a saturated state (*long dash*); over a broad range of low frequencies, the saturated model exhibits attenuation and phase-lead when compared to the equilibrium model, the same differences observed by Heisenberg and Wolf in comparing replay and closed-loop responses.

DISCUSSION

This work proposes a transfer function model derived from a broad-spectrum frequency analysis. It has been observed repeatedly and reaffirmed in this work that the fly optomotor response is low-pass; performance in reference-tracking tasks is strong for slow oscillations and diminishes rapidly for frequencies greater than 2 Hz. For this reason, the overwhelming majority of the literature focuses on the low-frequency response, where behavior is robust and easily discernable during experiments. However, from a control theoretic perspective, low frequency responses are least informative when the underlying system is low-pass.

Excepting differences in gain, for low frequencies (relative to the cut-off frequency), the majority of linear low-pass systems exhibit qualitatively similar responses. To differentiate between competing models—to determine attributes such as model order, cut-off frequency, damping, and delay—requires observation beyond the cut-off frequency, where performance wanes.

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

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

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

$$\frac{\partial F}{\partial G} = \frac{1}{(1 - G(\omega))^2}. \quad (8)$$

We see in (8), that the open-loop optomotor response is most sensitive to performance near $G(\omega) = 1$ (perfect tracking performance). The heightened sensitivity in this regime has two important implications: small measurement errors will yield highly variable open-loop model predictions, and, conversely, a diversity of open-loop models could have generated the observed closed-loop performance. For this common feedback topology, sensitivity improves to unity as performance degrades. This trade-off must be a critical consideration for system identification, to find the regime of stimuli which elicit observable behavior and informative model constraints, performance and sensitivity near unity.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 0845749. ER was supported by an NSF Graduate Research Fellowship and an Achievement Rewards for College Scientists Scholarship.

REFERENCES

- [1] U. Saranlı, M. Buehler, and D. E. Koditschek, “RHex: A simple and highly mobile hexapod robot,” *Int. J. Robot. Res.*, vol. 20, no. 7, pp. 616–631, 2001.
- [2] R. Wood, “The first flight of an insect-sized robotic fly,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, San Diego, CA, USA, Oct 29–Nov 2 2007, p. 2555.
- [3] S. Bergbreiter and K. Pister, “Design of an autonomous jumping microrobot,” in *Proc. IEEE Int. Conf. Robot. Autom.*, april 2007, pp. 447–453.
- [4] N. J. Cowan, J. Lee, and R. J. Full, “Task-level control of rapid wall following in the American cockroach,” *J. Exp. Biol.*, vol. 209, no. 9, pp. 1617–1629, 2006.
- [5] N. J. Cowan and E. S. Fortune, “The critical role of locomotion mechanics in decoding sensory systems,” *J. Neurosci.*, vol. 27, no. 5, pp. 1123–1128, 2007.
- [6] E. Roth, K. Zhuang, S. A. Stamper, E. S. Fortune, and N. J. Cowan, “Stimulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmannia virescens*,” *J. Exp. Biol.*, vol. 214, no. 7, pp. 1170–1180, Apr 2011.
- [7] B. Duistermars, M. Reiser, Y. Zhu, and M. Frye, “Dynamic properties of large-field and small-field optomotor flight responses in *Drosophila*,” *J. Comp. Physiol. A*, vol. 193, pp. 787–799, 2007.
- [8] G. Geiger and D. Nessel, “Visual processing of moving single objects and wide-field patterns in flies: Behavioural analysis after laser-surgical removal of interneurons,” *Biol. Cybern.*, vol. 44, pp. 141–149, 1982.
- [9] M. A. Frye and M. H. Dickinson, “Motor output reflects the linear superposition of visual and olfactory inputs in *Drosophila*,” *J. Exp. Biol.*, vol. 207, no. 1, pp. 123–131, 2004.
- [10] A. Sherman and M. H. Dickinson, “Summation of visual and mechanosensory feedback in *Drosophila* flight control,” *J. Exp. Biol.*, vol. 207, no. 1, pp. 133–142, 2004.
- [11] S. A. Budick, M. B. Reiser, and M. H. Dickinson, “The role of visual and mechanosensory cues in structuring forward flight in *Drosophila melanogaster*,” *J. Exp. Biol.*, vol. 210, no. 23, pp. 4092–4103, 2007.
- [12] W. Reichardt and T. Poggio, “Visual control of orientation behaviour in the fly. part I. a quantitative analysis,” *Q Rev Biophys.*, vol. 9, no. 3, pp. 311–75, 428–38, Aug. 1976.
- [13] J. C. Theobald, D. L. Ringach, and M. A. Frye, “Dynamics of optomotor responses in *Drosophila* to perturbations in optic flow,” *J. Exp. Biol.*, vol. 213, no. 8, pp. 1366–1375, 2010.
- [14] S. Sastry, *Nonlinear Systems*. Springer, 1999.
- [15] M. Heisenberg and R. Wolf, “Reafferent control of optomotor yaw torque in *Drosophila melanogaster*,” *J. Comp. Physiol. A*, vol. 163, no. 3, pp. 373–388, May 1988.
- [16] L. Ristroph, A. J. Bergou, G. Ristroph, K. Coumes, G. J. Berman, J. Guckenheimer, Z. J. Wang, and I. Cohen, “Discovering the flight autostabilizer of fruit flies by inducing aerial stumbles,” *Proc. Nat. Acad. Sci.*, 2010.
- [17] G. Maimon, A. D. Straw, and M. H. Dickinson, “A simple vision-based algorithm for decision making in flying *Drosophila*,” *Curr. Biol.*, vol. 18, no. 6, pp. 464–470, 2008.
- [18] E. Holst and H. Mittelstaedt, “Das reafferenzprinzip,” *Naturwissenschaften*, vol. 37, pp. 464–476, 1950.
- [19] M. B. Reiser and M. H. Dickinson, “A modular display system for insect behavioral neuroscience,” *J. Neurosci. Methods*, vol. 167, no. 2, pp. 127–139, Jan. 2008.
- [20] L. F. Tammero, M. A. Frye, and M. H. Dickinson, “Spatial organization of visuomotor reflexes in *Drosophila*,” *J. Exp. Biol.*, vol. 207, no. 1, pp. 113–122, 2004.
- [21] S. N. Fry, R. Sayaman, and M. H. Dickinson, “The aerodynamics of free-flight maneuvers in *Drosophila*,” *Science*, vol. 300, no. 5618, pp. 495–498, Apr. 2003.
- [22] T. L. Hedrick, B. Cheng, and X. Deng, “Wingbeat time and the scaling of passive rotational damping in flapping flight,” *Science*, vol. 324, no. 5924, pp. 252–255, 2009.
- [23] S. G. Carver, T. Kiemel, N. J. Cowan, and J. J. Jeka, “Optimal motor control may mask sensory dynamics,” *Biol. Cybern.*, vol. 101, no. 1, pp. 35–42, Jul. 2009.
- [24] R. C. Hardie and P. Raghu, “Visual transduction in *Drosophila*,” *Nature*, vol. 413, pp. 186–193, Sep. 2001.