Published in final edited form as:

J Neural Eng. 2018 April 01; 15(2): 026011–. doi:10.1088/1741-2552/aaa506.

Design strategies for dynamic closed-loop optogenetic neurocontrol *in vivo*

M F Bolus¹, A A Willats¹, C J Whitmire¹, C J Rozell², and G B Stanley¹

¹Wallace H. Coulter Department of Biomedical Engineering, Georgia Institute of Technology and Emory University, Atlanta, GA 30332, USA

²School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, USA

Abstract

Objective—Controlling neural activity enables the possibility of manipulating sensory perception, cognitive processes, and body movement, in addition to providing a powerful framework for functionally disentangling the neural circuits that underlie these complex phenomena. Over the last decade, optogenetic stimulation has become an increasingly important and powerful tool for understanding neural circuit function, owing to the ability to target specific cell types and bidirectionally modulate neural activity. To date, most stimulation has been provided in open-loop or in an on/off closed-loop fashion, where previously-determined stimulation is triggered by an event. Here, we describe and demonstrate a design approach for precise optogenetic control of neuronal firing rate modulation using feedback to guide stimulation continuously.

Approach—Using the rodent somatosensory thalamus as an experimental testbed for realizing desired time-varying patterns of firing rate modulation, we utilized a moving average exponential filter to estimate firing rate online from single-unit spiking measured extracellularly. This estimate of instantaneous rate served as feedback for a proportional integral (PI) controller, which was designed during the experiment based on a linear-nonlinear Poisson (LNP) model of the neuronal response to light.

Main results—The LNP model fit during the experiment enabled robust closed-loop control, resulting in good tracking of sinusoidal and non-sinusoidal targets, and rejection of unmeasured disturbances. Closed-loop control also enabled manipulation of trial-to-trial variability.

Significance—Because neuroscientists are faced with the challenge of dissecting the functions of circuit components, the ability to maintain control of a region of interest in spite of changes in ongoing neural activity will be important for disambiguating function within networks. Closed-loop stimulation strategies are ideal for control that is robust to such changes, and the employment of continuous feedback to adjust stimulation in real-time can improve the quality of data collected using optogenetic manipulation.

Keywords

closed-loop; control; optogenetics; thalamus; firing rate; *in vivo*

1. Introduction

The ability to control neuronal activity has deep clinical and scientific significance, ranging from treating movement disorders and epilepsy to understanding fundamental operating principles of single cells and intricately interconnected networks. Indeed, it was the use of a feedback controller—Cole’s voltage clamp—that enabled the pioneering studies of Hodgkin and Huxley and gave rise to our understanding of the ionic currents underlying the action potential [1]. Today, the neuroscience community is faced more acutely than ever with the task of dissecting the functions of neurons in the context of connected networks. To disentangle the roles of different cell types or structures under such conditions, systems neuroscience requires a set of tools for controlling neural activity at a meso-scale, between the extremes of stimulating single neurons and non-selectively manipulating large, diverse populations. Importantly, these control methodologies should also be robust to changes in ongoing activity in the brain which could otherwise be sources of unexplained experimental variability.

Over the last decade, optogenetic stimulation has emerged as a tool for understanding neural circuit function. Unlike electrical stimulation, optogenetic manipulations have the ability to target the expression of opsins genetically (*e.g.*, [2,3]) and/or anatomically (*e.g.*, single-cell [4] or retrogradely-labeled [5]). Given the flexibility of this technique and a maturing genetic toolbox, there is growing interest in the intersection between optogenetics and engineering control theory as a method for dissecting circuit function [6–8]. Notably, optogenetics lends itself particularly well to closed-loop control, pairing electrophysiological recordings with optical stimulation. In contrast with simultaneous electrical stimulation and recording which is plagued by stimulation artifacts, there is comparatively less concern for such artifacts corrupting measurements when using optical stimulation.

Control theory has already been brought to bear on the problem of manipulating neural activity for the purposes of halting seizures ([9–13]), reducing oscillatory activity in models of Parkinson’s ([14]), and artificially replicating LFP patterns naturally evoked in response to touch ([15]). To date, most so-called neurocontrol—whether by means of optical or electrical stimulation—has been conducted in an open-loop or on/off closed-loop fashion. In the latter case, stimulation is triggered by activity of interest (*e.g.*, [11, 12, 16–18]). In such applications, the stimulation that is delivered has been previously determined, through experimentation or using previously identified mathematical models. However, neuronal responses elicited by stimulation can vary across individuals, cells, and even over time. Rather than using feedback merely to trigger pre-determined stimulation, another strategy is to make continuous use of feedback to update stimulation in real-time. Recently, we demonstrated the first such use of closed-loop optogenetic control *in vivo* [19], where light intensity was adjusted to maintain a constant firing rate over time. While the first of its kind, the study did not offer a principled design methodology, nor did it extend the method beyond static reference tracking. In the present study we have developed and demonstrated a strategy for designing a proportional-integral (PI) controller for eliciting desired patterns of temporal rate modulation.

Here, we demonstrate that a simple control scheme can be effective for eliciting desired patterns of firing rate *in vivo*, and that the use of feedback confers a reduction in trial-to-trial variability as compared to open-loop stimulation. We have developed an approach for tuning the closed-loop control system for eliciting sinusoidally modulated patterns of firing rate, which includes an observer for estimating firing rate online from measured spikes as well as the controller itself. Finally, we demonstrate that this procedure can generalize to more complex, non-sinusoidal signals of interest.

2. Methods

We have developed a principled design strategy for closed-loop control of dynamic trajectories in neural firing patterns through the use of a proportional-integral (PI) controller and an exponential filter serving as an observer to estimate the latent firing rate from measured spikes. We utilized an *in vivo* rodent model where we recorded from and optically stimulated neurons in the ventral posteromedial (VPM) region of the thalamus that have been transfected with a depolarizing light-sensitive ion channel. The experimental preparation is outlined in Figure 1A, illustrating the single-unit thalamic recording, optical drive of opsin-expressing thalamic neurons, and the relationship of the thalamus to the whisker-driven afferent input from the periphery and the ascending and descending connections with the primary somatosensory cortex. The block diagram in Figure 1B illustrates the control framework, where the “neural system” (in this case, single unit in the thalamus) emits measured spiking patterns (n), which are in turn utilized to estimate firing rate ($\hat{\lambda}$). The difference in estimated firing rate and reference (or, desired) firing rate (λ_r) is defined as the online estimate of tracking error ($\hat{e}$). A conventional PI controller operates on this error signal, yielding a light input (u) which drives optogenetic excitation of the thalamic neurons.

2.1. Experimental Preparation

All procedures were approved by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology and were in agreement with guidelines established by the NIH. Experiments were carried out using female albino (Sprague-Dawley) rats. Expression of channelrhodopsin was targeted to excitatory neurons (rAAV5/CamKIIa-hChR2(H134R)-mcherry-WPRE-pA; UNC Vector Core, Chapel Hill, NC) in the ventro-posteromedial nucleus (VPM) of the thalamus by way of stereotactic injections ($3 \times 3 \times 5.2$ mm rostro-caudal $\times$ medio-lateral $\times$ depth) [20]. A 1 μ L volume of virus was injected at a rate of 0.1 μ L/min. The animals were given buprenorphine for pain management (0.03 mg/kg). Animals were then monitored daily following injection surgery. Wound clips were removed at 10–13 days post-surgery. Animals were allowed to recover and opsins allowed to express > 3 weeks. While VPM of the thalamus was the anatomical target, any optogenetically-driven neuron in thalamus that exhibited well-isolated spiking activity was considered a candidate for this study.

On the day of the experiment, rats were anesthetized using a cocktail of fentanyl (5 μ g/kg), midazolam (2 mg/kg), and dexmedetomidine (150 μ g/kg) delivered intravenously through the tail vein [20, 21]. Animal body temperature was maintained at approximately 37°C using a feedback-controlled heating pad. A 3 mm $\times$ 3 mm cranial window centered over the left

hemisphere at 3 mm rostral and caudal of bregma was created and the *dura mater* carefully removed. Single units were isolated in thalamus using an optrode: 80 μm , 2 M Ω tungsten electrode (FHC), coupled to a 200 μm optic fiber (Thorlabs). Blue light was conducted from an LED (470 nm, ThorLabs) to the thalamus via the optic fiber. Prior to the experiment, an optical power meter was used to measure the light intensity emitted from the tip of the fiber when peak command voltage (5 V) was sent to the LED driver. We found a linear relationship between command voltage and light intensity. Therefore, we estimated the light inputs reported in this study by scaling the command voltage accordingly.

All signals were recorded using a Tucker Davis Technologies (TDT) RZ2 Bioprocessor. Extracellular voltage was recorded at 24.414 kHz and bandpass filtered from 500 to 5000 Hz. Single unit spikes were manually thresholded, their dimensionality reduced using principal component analysis, and clustered in this reduced space using K-means. Measured spikes were then smoothed into an online estimate of instantaneous firing rate, and a proportional-integral (PI) controller was used to modulate the amplitude of optical stimulation of the neuron (Figure 1A&B). In post-hoc analysis, we found that optical artifacts were minimal, especially when considering the 500 to 5000 Hz band used for spike thresholding/sorting. We also found that spikes used for feedback during control epochs did not significantly differ from the spikes that occurred spontaneously in inter-trial-intervals.

At the conclusion of each experiment, animals were sacrificed using an overdose of sodium pentobarbital.

2.2. Reference Trajectories

2.2.1. Sinusoidal—To extend beyond the methodology laid out by [19] for maintaining a constant firing rate, we designed the control system to elicit sinusoidally-modulated rates of the form

$$\lambda_i = \sigma \sin(2\pi f_{\text{mod}} \Delta i) + \mu,$$

where σ , f_{mod} , μ , i , and Δ are the amplitude, modulation frequency, DC firing rate, time index, and sample period, respectively. Such a parametric time-varying reference is the ideal starting point for this design problem. Motivated by firing rates observed in the awake animal (see next section), all reference trajectories had mean firing rates of 20 spikes/s. To focus our investigation further, the sinusoidal reference trajectories used here were maximally modulated about the mean (*i.e.*, $\sigma = \mu$). The control system was tuned for 1, 5, or 10 Hz modulated patterns.

2.2.2. Non-sinusoidal Trajectory—In order to test the results of the design procedure on a non-sinusoidal reference trajectory of interest, we used an example of rate modulation observed in single-unit data recorded in the VPm of an awake rat (data from [22]).

Rhythmic spiking activity possibly related to the animal moving its whiskers was identified in a subset of trials. To inspect for rhythmic spiking, spike trains were smoothed using a Gaussian window with standard deviation (SD) of 20 ms and autocorrelograms calculated

for each trial. Putative “whisking” trials ($n = 3$) were identified by peaks in correlation located at lead/lag of 100 ms, corresponding to 10 Hz which is within the natural frequency range of whisking [23–25]. Spike trains from these trials were aligned such that their cross-correlations had a peak at zero-lag, and the resulting PSTH was smoothed with a Gaussian window of 20-ms SD, resulting in the reference trajectory shown in Figure 2C&D. The mean (*i.e.*, DC) firing rate of this signal was approximately 20 spikes/s, and 95% of the total power in this signal was between DC and 10 Hz. Note that approximately 50% of the power in this signal occurred at DC.

2.3. Observer Design

The closed-loop control system had two designed components: an observer and controller. An observer is a mathematical construct put in place to estimate state variables of interest for control when the state is not directly measured. In this context, the observer was tasked with estimating the latent firing rate from an evolving spike train measured extracellularly. Given a raw measure of instantaneous firing rate, n_i (either -1 or 0 if there is/is not a spike, where Δ is the sampling period), the observer yields an online estimate of the latent rate, $\hat{\lambda}_i$.

2.3.1. Fixed-bandwidth Smoothing—Fixed-bandwidth smoothing was used to estimate the firing rate online. An exponential window was chosen because it is causal and efficiently computed online as a first-order recursive filter. The filter was characterized by a single parameter, its decay time constant (τ):

$$\hat{\lambda}_i = \alpha \hat{\lambda}_{i-1} + (1 - \alpha) \frac{n_i}{\Delta}$$

where,

$$\alpha = \exp\left(-\frac{\Delta}{\tau}\right),$$

and i and Δ are the sample index and sample period, respectively.

2.3.2. Parametric Sweeps for Optimal Sinusoidal Rate Estimation—In designing the bandwidth of the exponential filter, the goal was to choose a filter which provided an appropriate amount of smoothing such that the underlying firing rate could be accurately recovered from measured spikes. Consistent with previous work (*e.g.*, [26–29]), the filter bandwidth was designed in simulation by minimizing the mean integrated squared error (MISE) between a ground truth firing rate pattern and the estimate:

$$\tau^* = \arg \min_{\tau} \left| \frac{1}{N_I} \sum_{i=1}^{N_I} (\lambda_i - \hat{\lambda}_i)^2 \right|,$$

where $\langle \cdot \rangle$ denotes trial-averaging and $\hat{\lambda}$ is the rate estimated by filtering the spike trains generated according to Poisson statistics from the ground truth rate, λ . Here, λ was taken to be the reference firing rate pattern of interest, λ_r . Depending on the properties of this pattern, the MISE-optimal estimator may smooth out the temporal modulation due to a scarcity of spikes within a period of oscillation. Koyama and Shinomoto [27] discussed such “divergent” solutions in the context of designing the optimal bin width for constructing a peristimulus time histogram (PSTH).

In the case of sinusoidal rates used here, the optimal filter time constant depended upon the expected number of spikes per period of the sinusoid, $(\mu N_{\text{trials}}) / f_{\text{mod}}$, and the degree of the modulation as defined by the ratio of amplitude to DC offset of the sinusoid, σ / μ , where N_{trials} is the number of trials used for estimation. To develop a parametric expression for the filter design, we determined the optimal time constant for different conditions of a simulated rate-modulated Poisson process.

In simulation, four parameters were swept: the modulation frequency of the driving sinusoid (f_{mod}), the mean firing rate (DC offset of the sinusoid, μ), the amplitude of the sinusoid around the mean (σ), and the number of trials used for estimation (N_{trials}). Note that at the conclusion of the design, we sought time constants that were optimized for single-trial estimation: *i.e.*, where $N_{\text{trials}} = 1$. For the rates used in this study, we focused on the specific cases where $\mu = 20$ spikes/s and $\sigma / \mu = 1$.

The MISE-optimal filters (τ^*) were calculated for a range of frequencies, mean rates, modulation intensities, and number of trials. In keeping with Koyama and Shinomoto [27], a power law relationship was fit to these data for τ^* to yield a tuning curve of estimated time constants, $\hat{\tau}$:

$$\hat{\tau} f_{\text{mod}} = b z^{-a},$$

where [27] defined z as

$$z \equiv \left(\frac{\mu}{f_{\text{mod}}} N_{\text{trials}} \right) \left(\frac{\sigma}{\mu} \right)^2.$$

2.4. Linear-nonlinear Poisson Model

Previously, an *ad hoc* approach to controller design demonstrated proof-of-principle feedback optogenetic control. Here, however, we sought a principled approach that generalized to dynamic reference trajectories. To finely tune the controller during experiments, we used a linear-nonlinear Poisson (LNP) model to approximate the response of the neural system to optical stimulation. The controller gains (K_p , K_i) were tuned around the LNP model in simulation. In this model structure, a linear system (denoted L) was cascaded with a static nonlinearity (N) to produce a latent firing rate which drives a Poisson spike generator (P) and emits spikes.

The optical stimulus was filtered through a feedforward linear system described by the kernel, $\mathbf{k}$:

$$\mathbf{x} = \mathbf{k}^T \mathbf{U},$$

where $\mathbf{U}$ is the stimulus design matrix in which the i^{th} column is a vector of stimulus history over a time window up to time i . The output of this linear system, $\mathbf{x}$, was then mapped through a static nonlinearity (*e.g.*, [30]):

$$\lambda_i(\theta, x_i) = \alpha \log [1 + \exp (gx_i + m)],$$

where $\theta = [m, g, \alpha]^T$ are parameters describing the static nonlinearity. m serves as a bias term reflecting neural firing that does not co-vary with stimulation (*i.e.*, spontaneous firing rate). g and α together set the effective DC gain of the model neuron's response to light stimulation as it approaches the asymptotically linear region of the curve, whereas their relative values set the knee of the nonlinearity.

During the experiment, an LNP model was fit to spiking data recorded in response to repeated presentations of a 5-second instantiation of optical uniform white noise. The range of this noise was titrated for each cell to avoid apparent depolarization block, but on average ranged between 0 and 8.5 mW/mm². The white noise stimulus was mean-subtracted and a light-to-spiking kernel, $\hat{\mathbf{k}}$, was estimated by reverse correlation ("rotated" or "whitened" spike-triggered averaging, *e.g.*, [31–33]):

$$\hat{\mathbf{k}} = (\mathbf{U}^T \mathbf{U})^{-1} \mathbf{U}^T \mathbf{n},$$

where $\mathbf{n}$ is the vector spike signal (1 or 0 if there is/is not a spike measured). To avoid ambiguity between the static gain of the kernel and the scaling factor applied to the output of the kernel in the nonlinearity (g), the kernel was normalized by its static gain: $\hat{\mathbf{k}}_1$. The stimulus was filtered with this normalized kernel, yielding $\mathbf{x}$, and the remaining parameters ($\hat{\theta}$) were fit by maximum likelihood:

$$\hat{\theta} = \arg \max_{\theta} \sum_j \sum_i n_{ij} \log [\lambda_{ij}(\theta, x_{ij}) \Delta] - \lambda_{ij}(\theta, x_{ij}) \Delta,$$

where i and j denote the time and trial index, respectively. Note that [34] fit LNP models by maximum likelihood, with the kernel parameters only *initialized* using the values predicted by whitened spike-triggered averaging (STA). In our application, we found that the kernels resulting from this approach were primarily scaled versions of the STA estimate, $\hat{k}$. Because parameter g accounts for this scaling, we found that it was unnecessary to re-estimate the kernel parameters in practice. This allowed for the model to be fit quickly and dependably during the experiment, where experimental viability is time-limited.

2.5. Controller Design

We implemented a proportional-integral (PI) controller, which was defined by two parameters, K_p and K_i . Given a reference trajectory of interest, a choice for the observer time constant (Section 2.3), and a model, these controller gains were tuned in simulation during the experiment.

2.5.1. Control Law—The PI controller was implemented in its parallel form:

$$\hat{e}_i = \lambda_i^r - \hat{\lambda}_i$$

$$u_i = K_p \hat{e}_i + K_i \sum_{k=0}^i \hat{e}_k \Delta,$$

where K_p and K_i are the proportional and integral gains, respectively.

Note that due to the fact the light delivered can neither be negative nor exceed the limits of the LED output, the variable u was bounded between $[0, 1]$ and scaled to the dynamic range of the LED driver ($5 \text{ V} = 1 \text{ A}$). For all simulations, the control signal was hard rectified and bounded as was done in hardware to ensure actuation was subject to the same limitations.

2.5.2. Objective Function—A common objective for controller design is to minimize the integrated squared tracking error (*e.g.*, in the context of PID control [35]). We found that tuning the controller to minimize the integrated squared error between the reference and online estimate of rate (*i.e.*, between λ_r and $\hat{\lambda}$) yielded undesired tracking behavior, especially at higher frequency references where the firing rate estimate is less accurate (see insets of Figure 3C, Section 3.2). Because observer estimates are on average slightly lagged and attenuated in amplitude compared to ground truth, if stimulation were optimized around this estimate, the resulting neural firing would lead the target and be of larger amplitude. To avoid this behavior, we instead tuned the controller by minimizing the squared error between the reference (λ_r) and the raw measure of instantaneous rate (n'), which neither imposes lag nor amplitude attenuation on the estimate.

While minimizing this raw tracking error ameliorates the aforementioned problems, the highly punctate nature of the n' estimate of rate and the large error incurred each time the neuron spikes ($\lambda_r - n' = -1$) mean that simply minimizing the integrated square of the error, $e = \lambda_r - n'$, yielded trivially low solutions for controller gains where the neuron never spikes. This motivated controller design in the frequency domain, where the effects of the spiking error can be down-weighted relative to important tracking criteria. Because we are willing to tolerate error that occurs at frequencies higher than those of interest for a control task, we avoided trivial low-gain solutions by weighting the squared error in the frequency domain according to the spectral content of the reference trajectory (w), thereby penalizing error at frequencies according to their significance for the intended control.

We calculated the frequency-domain amplitude of the tracking error, e , as well as that of the reference trajectory:

$$E[f] = |\mathcal{F}\{e\}|$$

$$R[f] = |\mathcal{F}\{\lambda_r\}|,$$

where $\mathcal{F}\{\cdot\}$ denotes the Fourier transform.

The spectral content of the reference was used to create frequency-dependent weights:

$$w[f] = \frac{R^2[f]}{\sum_{f=0}^{N_f} R^2[f]},$$

where N_f corresponds to the Nyquist frequency.

These weights were used to calculate a modified squared error metric where the frequency spectrum of the error was penalized as a function of importance for the control task:

$$J_{\text{fwt}}(K_p, K_i) = \left\langle \sum_{f=0}^{N_f} w[f] E^2[f] \right\rangle,$$

where $\langle \cdot \rangle$ denotes the across-trial average.

Finally, controller gains were chosen to minimize J_{fwt} , using a numerical solver (*ga*, Mathworks, Inc.):

$$[K_p^*, K_i^*] = \arg \min_{K_p, K_i} J_{\text{fwt}}(K_p, K_i).$$

2.6. Open-loop Stimulus Design

To assess the benefits of utilizing feedback with optogenetic stimulation, we designed open-loop stimuli for each control task (*i.e.*, sinusoids, and a more natural non-sinusoidal trajectory). In Figure 2A&B, we compare pulsatile open-loop stimulation (most commonly used stimulation mode) with closed-loop, continuous modulation of light amplitude. In this scenario, 5-ms pulses were generated by varying the frequency of a carrier wave according to the desired firing rate. To determine the appropriate pulse amplitude at the time of the experiment, a static logistic mapping was fit to the response to pulsatile inputs of varying amplitude, where light inputs were presented at the DC firing rate of the target trajectory (*i.e.*, 20 pulses/s). The amplitude of stimulation was chosen such that approximately 1 spike/pulse was generated on average.

In all other cases of open-loop control, a static mapping from light intensity to firing rate was estimated by fitting a logistic curve to the steady-state firing rate in response to 1.5-second step inputs of light at various amplitudes. Open-loop control signals were designed for sinusoidal and non-sinusoidal firing rate trajectories by inverting this curve (i.e., using this curve as a lookup table).

2.7. Disturbance

A load disturbance in the form of whisker stimulation was used to challenge both open-loop and closed-loop control strategies. Using a computer-controlled galvanometer motor [20,36,37] positioned approximately 10 mm from the animal's face, sensory white noise was applied to the thalamic neuron's principal whisker [20] at two seconds into the control epoch.

2.8. Offline Firing Rate Estimation

All reported firing rates were estimated offline ($\bar{\lambda}$) using an appropriate Gaussian window to smooth either single-trial (Figure 2A&B) or trial-averaged spike trains binned at 1-ms resolution. In the case of the slowly-modulated rates in Figure 2A&B, the same procedure used for designing the observer time constant was employed to estimate an MISE-optimal SD of a Gaussian window. Because accurate single-trial estimation is difficult at higher frequencies, for all other firing rate estimation reported here, a Gaussian window was used to smooth a trial-averaged peri-stimulus time histogram (PSTH). The SD of this filter was chosen for each reference trajectory (i.e., 5 Hz or non-sinusoidal) in the same way as before except now for multiple trials: the MISE-optimal width for recovering the reference rate from a PSTH of simulated Poisson spikes, averaging the same number of trials collected experimentally. Bands around these trial-averaged firing rate estimates are 95% confidence intervals for the smoothed PSTH, bootstrapped by sampling the trials with replacement.

2.9. Fano Factor

Trial-to-trial variability in spike count was quantified using the Fano factor [38] calculated in a 250-ms sliding window:

$$FF = \frac{\text{var}[N_{250}]}{\langle N_{250} \rangle},$$

where N_{250} is the spike count per 250-ms window of time and $\langle N_{250} \rangle$ indicates the across-trial average.

3. Results

Here, we develop a design strategy for closed loop control that could be applied to a range of different of neural circuits. We have applied this design strategy to the problem of closing the loop around the spiking activity of a single neuron in the somatosensory thalamus of the rat *in vivo*, as illustrated in Figure 1A. Specifically, the input to the “neural system” was light delivered by way of a fiber optic cable inserted deep into the brain, targetting the somatosensory thalamus. Blue light (470 nm) drove the depolarization of excitatory thalamic

neurons expressing channelrhodopsin (ChR2). Spiking activity of a single thalamic neuron was measured using a tungsten extracellular recording electrode bundled to the fiber optic cable (often referred to as an optrode).

The block-diagram of the control system is shown in Figure 1B. We applied light input u to the “neural system” or plant, whose activity was measured through neural spiking activity n . For the current application, we have taken time-varying trajectories of firing rate to be the control objective, and thus the feedback signal consisted of an observer’s estimation of instantaneous firing rate. A proportional-integral (PI) controller acted on an error signal defined as the difference between the reference/desired time-varying firing rate λ_r and the observer’s estimate of the instantaneous firing rate, $\hat{\lambda}$. Control was effected in the face of unobserved disturbances, d , which could take the form of uncontrolled inputs and/or changes in the dynamics of the system.

In this context, closed-loop control required the real-time sorting of spiking activity from recorded extracellular signals and estimation of instantaneous firing rate by an observer. This estimated rate was fed back, and a PI controller operates on the resulting error signal, modulating the light input to the optic fiber and optogenetically manipulating thalamic membrane polarization. Furthermore, the design procedure required identification of the dynamics of the neural system during the course of the experiment; combined with testing the controller design, this places serious demands on the duration of an experiment, which is typically limited to approximately 2–3 hours. The timeline and task demands are outlined in Figure 1C.

3.1. Examples of Closed- vs. Open-loop Stimulation

Before unpacking the technical details of the system design in subsequent sections, we first present examples representing the basic abilities of the closed-loop framework vs. open-loop stimulation strategies to track slowly-modulated reference firing rates (here, 1 Hz), reject exogenous disturbances, and track complex, biologically relevant trajectories in firing rate.

In the example for controlling a 1 Hz sinusoidal firing rate provided in Figure 2A, pulsatile input was used for open-loop stimulation (bottom). Note that for this instance of open-loop control, pulsatile inputs were used since this is the most common way to stimulate ChR2. Single-trial and trial-averaged estimates of firing rate are shown for closed-loop (top) and open-loop (bottom) cases. For the open-loop case, measured ring activity was not utilized in shaping the light input, but instead the amplitude of the light inputs were designed based on previous measurements of number of spikes elicited as function of pulse amplitude (see Section 2.6). In each scenario, the trial-averaged input (*i.e.*, light intensity) is plotted below the corresponding firing rate. The control epochs begin at time zero. Both closed- and open-loop stimulation strategies achieved the target firing behavior *on average* in the undisturbed scenario (Figure 2A, thick red and black versus green). However, single-trial estimates of firing rate (thin red lines, smoothed with a 120-ms SD Gaussian window) reveal that open-loop stimulation resulted in more variable rate trajectories than closed-loop for this example.

A major benefit of a closed-loop system is its capacity to react to changes in ongoing activity and reject disturbances, as illustrated in Figure 2B. For the same example thalamic

neuron in Figure 2A, we identified the whisker on the contralateral side of the animal's face to which the neuron responded most robustly, often referred to as the principal whisker". A whisker disturbance (see Section 2.7) begins at 2 seconds into the control epoch. During the application of the disturbance, closed-loop stimulation was able to adjust to maintain reasonable control of the 1 Hz trajectory (Figure 2B, top). Conversely, in the case of open-loop stimulation (Figure 2B, bottom), the firing rate was unsurprisingly increased well above the reference. It is also of interest to note that, while the same open-loop pulsatile stimulus was used in Figure 2A&B, the effectiveness of stimulation was weaker in the initial 2 seconds of the control epoch before the onset of the disturbance. Given that time elapsed between these two recordings, this phenomenon speaks to apparent non-stationarity in the system for which closed-loop stimulation is able to compensate, even over relatively short timeframes.

We also challenged the control framework with a non-sinusoidal, reference trajectory derived from spiking in an awake animal (Section 2.2.2), as shown for a different thalamic neuron in Figure 2C. As described in the Section 2.2.2, we separately recorded single unit firing activity in the VPM thalamus of the awake rat during active whisking, and used this to generate a more "naturalistic" non-sinusoidal firing rate trajectory to track in these separate experiments. Note that open-loop control performed qualitatively similarly to closed-loop control on average in this example, although closed-loop does provide modest improvements in tracking. However, in agreement with the sinusoidal trajectories above, closed-loop stimulation results in lower trial-to-trial variability, as shown by Fano factor (FF), where the variability of closed-loop controlled firing generally falls below open-loop.

These examples of control shown in Figure 2 depended upon the design and implementation of both the observer and controller elements of a closed-loop system, which we further detail below, before returning to further analyses of the control performance.

3.2. Observer Design

The observer was tasked with estimating the latent firing rate, given spiking activity measured online. As described in detail previously (Section 2.3), for simplicity of real-time implementation the observer was implemented in the form of an exponential filter with a single parameter, the decay time constant τ . Figure 3A shows the basic concept for the design strategy, where a Poisson spike generator was driven with a sinusoidal rate function, λ , and the observer estimated the latent rate from the spiking activity by filtering the spike train. The goal of the design procedure was to choose the filter time constant that optimally recovered the underlying sinusoidal rate. Shown are the results for an illustrative example where a simulated spike train was filtered with one of three values of the decay time constant τ (actual rate in green and estimated rates in purple shown on the right; bottom: too slow, top: too fast, middle: MISE-optimal).

As described in detail in Section 2.3, for sinusoidal modulation the optimal observer in this framework depended upon the baseline (DC) offset of the sinusoid (μ), the modulation frequency (f_{mod}), the amplitude (σ), and the number of trials used for estimation (N_{trials}), which together influenced the optimal filter bandwidth for recovering sinusoidal rate modulation from recorded spikes. Plotted on the top panel of Figure 3B are the MISE-

optimal values of the observer decay time constant, τ^* , as a function of the number of spikes/period of the sinusoid, for varying degrees of modulation (σ/μ). As expected, the MISE-optimal filter time constant decreased with increasing number of spikes per period and modulation intensity. We then used the quantity z defined by Koyama and Shinomoto [27] to reduce the dimensionality (see Section 2.3). The disparate curves in the top panel of Figure 3B then collapsed to a single curve when plotted as a function of z (Figure 3B, bottom). We fit a power law relationship for the optimal time constant using least-squares, the result of which is shown in red in Figure 3B (bottom): $a = 0.423$, $b = 0.389$.

In the case of the sinusoidal firing rates utilized in the remainder of this study, where modulation amplitude was maximal (*i.e.*, $\sigma/\mu = 1$), the MISE-optimal filter masked temporal modulation in favor of capturing the DC firing rate if there were fewer than 12 spikes per period. More generally, this corresponds to a regime where $z < 12$ in which there was a deficit in spiking data on a single-trial basis (denoted by the vertical gray line in Figure 3B, bottom). As such, these data points were excluded from the regression. For the rates used in the remainder of this study where $\mu = 20$ spikes/s, this corresponded to frequencies greater than 2 Hz (Figure 3C, light purple). By extrapolating the expression for observer time constant that was developed and fit in the regime where $z > 12$, we arrived at choices for time constants that did not smooth out temporal modulation at higher frequencies despite this fundamental limit while still providing a reasonable degree of smoothing, as seen in Figure 3C (insets). From the insets which provide examples of single-trial estimation of a 1 Hz, 5 Hz, and 10 Hz sinusoidal rates from spikes, it is apparent that when the modulation frequency reached 10 Hz for this DC offset, spikes were too infrequent per period of oscillation to recover the sinusoid accurately on a single-trial basis.

The net result of this procedure was a relationship between the observer filter time constant and the parameters of the sinusoidal rate which was valid in a regime where there was sufficient data to resolve the temporal modulation. When the rate estimation was data-impooverished, this relationship allowed us to estimate what the filter time constant would be if not for this fundamental limit. Therefore, while it is impossible to faithfully recover higher frequency modulation on a single-trial basis, this approach provides a principled choice of observer time constant that does not smooth out the rate modulation.

3.3. Controller Design

Given the design for observer time constant, we developed a PI controller design strategy for tracking sinusoidal trajectories of different frequencies (1, 5, 10 Hz). For principled design of the controller, we undertook a brief system identification step during the course of the experiment to fit a model for the neural system, consisting of a cascade of a linear filter, static nonlinear function and Poisson spike generator (LNP model), as will be described in more detail below. Controller gains were then tuned in simulation using the LNP model fit to data in place of the neural system (Figure 4A), followed by implementation in the experiment.

To tune the controller, we chose gains which minimized the squared error between the reference rate and the measured spiking signal. This noisy error signal was weighted in the frequency-domain according to the importance of a frequency for the intended control (see

Section 2.5.2). In the context of sinusoidal trajectories used here (where $\sigma = \mu$), the frequency-weighted squared error objective function (Section 2.5.2) evenly penalized error at DC and the modulation frequency.

To illustrate the tuning procedure, we used an experimentally fit LNP model for simulation. The frequency-weighted error metric is shown as a function of the proportional controller gain (K_p) and the integral controller gain (K_i) for a 5 Hz sinusoidal trajectory (Figure 4B). The result of the minimization was that there were optimal controller gains (circle symbol) which yielded improvement over two suboptimal examples (x & triangle symbols). This is shown more explicitly with simulated examples (Figure 4C, bottom), where desired (green) and achieved (black) firing rates are plotted for the three sets of control parameters. In the middle row of Figure 4C, the resulting error spectra for these three tunings are compared to the error expected for a Poisson spike generator (PSG) that has been driven at the reference rate (light grey). The spectrum for the reference-driven PSG shows the error that results purely as a function of random spiking, rather than off-target firing rate modulation. Because of the use of feedback, the spikes resulting from simulated closed-loop control are not truly random, leading to less error at low frequencies than would be expected in the Poisson case (black vs. grey error spectra). The peaks in error at $2f_{\text{mod}}$ (here, 10 Hz) in Figure 4C occur because the controlled neuron can spontaneously fire even when the reference rate is zero, resulting in errors at the troughs of each period. This in combination with error that occurs at peaks of the sinusoidal reference results in some power at double the modulation frequency. In comparison to the other two examples, the optimal tuning clearly reduces the error at both DC and f_{mod} . For the sinusoidal trajectory used here, the tuning procedure involved minimizing the combined error at DC and modulation frequency, illustrated in Figure 4C (top) as the square-root of the frequency-weighted squared error metric (J_{fwt}).

3.4. Model Accuracy and Closed-loop Performance

In the development of a strategy for closed-loop controller design, the above results relied upon knowledge of the dynamics relating the light input for optogenetic stimulation and the neuronal firing. In an experimental context, we must therefore estimate a model that captures this relationship between the light input, u , and the measurement, n , (Figure 5A) for use in controller design. We utilized a band-limited “white-noise” optical input for driving the thalamic neuron experimentally and fit a simple linear-nonlinear-Poisson (LNP) cascade to capture these dynamics (see Section 2.4). An example LNP fit is shown in Figure 5B. Note that the kernel is plotted time-reversed for visualization.

To understand how accurate these LNP models were, we tested their ability to predict the input-output relationship in the operating regime observed during a control epoch. Upon investigating how the LNP models responded to control inputs used during closed-loop stimulation experiments, we found that the models produced varying levels of success in predicting the observed firing activity of the thalamic neurons. Shown in Figure 5C is one such representative example where light inputs used experimentally for tracking a 5 Hz sinusoidal target trajectory were presented to the neuron’s corresponding LNP model. The response of the LNP model neuron (red) did not accurately predict the experimental rate (black). Given a wealth of literature suggesting difficulty in obtaining predictive models of

neuronal firing *in vivo* beyond the sensory periphery (e.g., [39,40]), this finding is perhaps unsurprising.

However, because these models were used for tuning the controller gains in simulation, it is more important to assess their predictive capability in the context of closed-loop operation. The use of feedback for control conferred some robustness to model inaccuracy, and although the LNP model was in general not a good predictor of neuronal firing activity in open-loop, the model produced outputs that agreed with experimental results when simulated in the context of closed-loop control. Figure 5D provides an example where the same model neuron in Figure 5C was simulated in the closed-loop system parameterized by the observer time constant (τ) and controller gains used experimentally. In this case, the firing rate of the model neuron (blue) more closely matched the experimentally-realized firing (black).

3.5. Robustness of Control to Model Inaccuracy

Given the observation that the LNP models used for controller tuning were generally poor predictors of experimental data in the open-loop sense, it is important to know to what extent the controller performance was robust to modelling error in this context. Therefore, we set out to determine more systematically through simulation how robust control performance was to inaccuracy in two identified LNP model parameters: gain and bias. Taking an LNP fit to experimental data, the PI controller was tuned around it, as would be done during an experiment. We then tested these controller parameters on perturbed versions of the original LNP model, where the static gain and the bias term of the linear component of the LN model were changed systematically. To quantify the distance between the performance of the original and perturbed systems, the percentage change in tracking performance was calculated by comparing the error (J_{fwt}) for the perturbed models to that of control around the original LNP. Further, to assess the effect of model inaccuracy on “steady-state” tracking performance, J_{fwt} was calculated from 1 second onward for 5-second control epochs. To inspect for added benefit of closed-loop stimulation as opposed to open-loop, the tracking error was assessed for the perturbed models in both closed- and open-loop contexts. For the latter, light inputs used to control the original LNP were presented to the perturbed versions of the model in open-loop.

Decrements in control performance due to model inaccuracy are illustrated in Figure 6 for tracking a 5 Hz sinusoidal reference, using an example LNP fit to experimental data. The schematic in Figure 6A shows the two parameters that were varied: the static gain (g) and the bias term (m) of the LNP model (see Section 2.4). Altering g increases or decreases the apparent ‘sensitivity’ of the neuron to optical drive. The bias term reflects the baseline firing rate in the absence of optical drive. Contrary to Section 2.4, for this analysis, LNP models were fit with α equal to 1 to control for differences in the knee of the nonlinearity. Note also that in the analysis shown here, the shape of the kernel was retained. The parameters g and m were perturbed over a range extending from five times smaller to larger than the original values for which the controller was designed and the decrements in controller performance quantified (Figure 6B).

While there is interaction between the effects of the two parameters on control performance, changing the bias led to binary effects, as the performance quickly transitioned from optimal to very poor, moving from left to right in Figure 6B. This observation is made plain when holding the static gain at the nominal value and fractionally changing the bias (Figure 6C). Note that the identified bias term was negative. As the bias term approached zero (*i.e.*, x_m decreases), this increased the baseline firing rate of the model neuron. Conversely, as the bias became more negative than the original value (x_m increases), this mapped the output of the linear portion of the LN model to the highly nonlinear portion of the static nonlinearity.

In Figure 6C, the change in tracking performance for a 5 Hz sinusoid is shown for both closed-loop stimulation (blue) as well as open-loop replay of the stimulation pattern used to manipulate the original (*i.e.*, unperturbed) model (red). Except at extremes, closed-loop control was very robust to inaccurate estimation of the bias term. At these extremes, either the neuron's baseline firing rate exceeded the reference firing rate or the output of the linear component of the model was mapped to a regime of the nonlinearity where the LED or other light source was incapable of delivering inputs intense enough to raise the firing rate. In both of these scenarios, there would be little if anything a controller could do to salvage performance. Figure 6D provides simulated examples for when the neuron being controlled (blue) was two times less biased (top) or two times more negatively biased (bottom) than the neuron for which the controller was tuned (black). In the first case, the baseline firing rate of the neuron in blue was too high to be effectively controlled. Conversely, in the case where the bias of the neuron was more negative (bottom), it took more light input (not shown) and, therefore, more time for the system to reach "steady state" behavior. However, once the controlled system reached steady state, the tracking performance (blue) was identical to the that of original model neuron the controller was designed around (black).

Holding the bias term at the original value and instead perturbing the static gain, Figure 6E reveals that while the impact of inaccurate estimation of static gain on controller performance was more graded, closed-loop control of time-varying trajectories was less robust to inaccurate estimation of gain than to the bias term. As before, it was the case that closed-loop control (blue) offers a buffer against the effects of such model inaccuracy as compared to open-loop stimulation (red). Figure 6F provides examples for when the controller was tuned around a model that was two times more sensitive (top) or less sensitive (bottom) to light than the neuron being controlled in blue. In the case where the neuron being controlled was less sensitive than the model around which the controller was designed (top), the controller achieved the correct DC firing rate, but was more weakly modulated than the target oscillatory activity. Conversely, where the neuron being controlled was more sensitive to optical drive, there were periodic overshoots of the reference. Unlike the case of the bias term, there are relatively simple actions that could be taken in the future to ameliorate errors due to inaccurate gain estimation, including online re-estimation.

3.6. Control Summary: Tracking a Sinusoidal Trajectory

Figure 2 provided a single example of tracking at 1 Hz. Here, we expand on this by presenting results for eliciting a 5 Hz sinusoidal pattern in thalamic neurons recorded in separate animals. Figure 7A provides an experimental example for tracking a 5 Hz sinusoid

at “steady-state” 3 seconds into a control epoch. In this example, controller performance reached the level predicted in simulation for the LNP model used for design (“Cell 3” in Figure 7B, left vs. middle). Again, the tracking error was quantified as the frequency-weighted squared error (J_{fwt}) between the raw measurement of rate (n/s) and the reference (λ_r), either for the LNP model used to design the controller (Figure 7B, left) or for the experiment (middle, right). In the case of the sinusoidal trajectories used here, J_{fwt} reflects even penalty placed on error at DC and the modulation frequency. As for the perturbation analysis (Section 3.5), this metric was calculated from 1 second onward for 5-second control epochs. 95% confidence intervals for the metric were calculated for simulated Poisson firing at the reference rate and plotted as a light grey horizontal band (Figure 7B).

In Figure 7B, each colored line corresponds to a different cell recorded in a separate animal, and the black bars represent the population average in simulated closed-loop (left), experimental closed-loop (middle), and experimental open-loop (right). Here, open-loop stimulation was a continuously modulated light signal, designed as previously detailed (Section 2.6). Simulated control around the LNP model used for tuning provides a lower bound for the control error (CL LNP, left). While the level of performance achieved in simulation was not achieved experimentally (Figure 7B, middle), reasonable control was nonetheless achieved for all neurons in closed-loop. In contrast, open-loop control was not as robust, with the designed stimulation sometimes resulted in much worse performance than the closed-loop scenario (Figure 7B, right).

The control loop was also tuned for and tested on 1 and 10 Hz sinusoidal trajectories. The tracking error of closed- vs. open-loop control is provided for all frequencies tested (*i.e.*, 1, 5, 10 Hz) in Figure 7C. When all sinusoidal tasks are considered, closed-loop control provides significantly lower tracking error ($p < 0.05$, Wilcoxon signed rank test).

3.7. Control Summary: Tracking a Non-Sinusoidal Trajectory

For the non-sinusoidal rate trajectory estimated from spiking in the awake animal, 95% of the power in the reference signal was below 10 Hz. We found that of control systems designed for 1, 5, or 10 Hz sinusoidal trajectories with the same DC firing rate (20 spikes/s) as this non-sinusoidal signal, controller gains and observer time constants designed for the 10 Hz trajectory outperformed lower frequencies (not shown).

Figure 8A provides an example where the system was tuned for a 10 Hz sinusoidal trajectory and tasked with tracking the non-sinusoidal pattern of rate modulation (“Cell 5” in Figure 8B), as well as its corresponding error spectrum (Figure 8A, bottom). For comparison, the error spectrum of simulated Poisson firing at the reference rate is plotted in light grey. After a period of approximately 750 ms, the neuron’s firing rate faithfully followed the reference. A controller/observer tuned for a 10 Hz sinusoidal trajectory did approximately as well as within-experiment simulations predicted (Figure 8B). Again, controller performance was quantified using the frequency-weighted squared error. Note that for this non-sinusoidal reference trajectory, 50% of its power lies at DC and 45% of the remaining power ranges from DC to 10 Hz; therefore, half the penalty is placed on achieving the correct average firing rate, while the remaining half of the penalty is primarily exerted at and below 10 Hz. In all cases, open-loop stimulation proved less effective than closed-loop.

Furthermore, as was first shown qualitatively for a single cell in Figure 2C, the response to open-loop stimulation was also more variable than in the case of closed-loop stimulation. We observed this phenomenon across cells ($n=5$), where the spike-count variability, as measured by time-averaged Fano factor calculated in a 250-ms sliding window, was greater in open-loop than in closed-loop (Figure 8C).

4. Discussion

The advent of new tools for measuring and manipulating activity within complex neural circuits opens up a wealth of possibilities for more interactive electrophysiological experiments, where feedback is used to inform stimulation continuously. Previously, we laid the groundwork for using optogenetic control to manipulate neural activity with continuous use of feedback [19]. Specifically, the previous study was a proof-of-concept concerned with the demonstration of holding firing rates constant over long control epochs. Further, while holding firing rate steady is certainly of scientific interest for probing the functions of neural circuitry near steady-state, generalization of the closed-loop control approach to time-varying patterns of firing would open up a range of new applications and lines of investigation. Therefore, in the present study we have developed and demonstrated a design strategy for using closed-loop stimulation to track time-varying patterns of firing rate.

Ranging from intracellular current injection to electrical and optogenetic stimulation, there has been previous work using open- and closed-loop control strategies to manipulate neural firing patterns (*e.g.* [7, 8, 15, 41–44]). Because spiking is often thought of as the information currency of the nervous system, such studies are usually concerned with achieving a target train of spike times. However, as a matter of practical concern, some jitter in the elicited spike times is tolerated, whether explicitly as in [43] or by means of a quadratic penalty between a spike time and the next desired spike in a train [44]. For this reason, the problem of controlling spike times can be reconceptualized as a problem of controlling a time-varying pattern of firing rate, where the timescale of rate variation depends upon the tolerated jitter in spike timing. Therefore, while we have chosen to tackle the problem of manipulating instantaneous firing rate, the design methodology laid out here could in principle be applied to the problem of controlling spike times, where the allowable jitter sets the timescale over which desired spike times are smoothed into a rate function that is fed to the control loop as a reference (*e.g.*, [41]).

However, we have presently tasked a control system with tracking not precise spike timing, but relatively slow patterns of rate modulation (1–10 Hz). Therefore, in this study we are not exercising control over spiking at timescales faster than approximately 100 ms. Notable rate modulation in this slower range of timescales occurs across neural systems, including in the context of active sensation [45,46], hippocampal theta-phase precession [47], and movement [48]. The ability to insert such patterns of rate modulation would enable causal investigation into how instantaneous rate affects the function of such systems. Moreover, the basic framework developed here could be extended to finer timescales, opening up additional avenues of investigation that revolve around questions of precision timing.

The choice to use firing rate as the controlled output of the system necessitated an observer designed for estimating this quantity from noise-corrupted measurements in the form of spikes. In the interest of simplicity needed for widespread adoption of this technique, we designed the observer in a model-free manner, without taking into account any dependence of the rate on inputs to the system. The obvious benefit of this approach is that it allows an experimenter to design this part of the control loop prior to the experiment, and it is simple to implement. However, given the point process nature of the measurements, this approach is not ideal for recovering time-varying rates, depending on the timescale of rate modulation. This is especially true of low firing rate regimes: *e.g.*, at modulation frequencies greater than approximately 2 Hz, with an average rate of 20 spikes/s (Figure 3), where estimates become increasingly poor.

Given the difficulty of accurate rate estimation at higher frequencies of rate modulation, we took the practical approach of designing the controller with the observer in the loop. That is to say, we did not make the customary assumption that we could design the controller for the case where the latent state/controlled output (here, λ) was known. Instead, a Poisson spike generator and observer were included in the simulated control loop when tuning the PI controller. By minimizing the error between the measured spiking and the reference at frequencies of interest for control, we endeavored to safeguard performance of the final closed-loop system against deleterious effects of imperfect online rate estimation.

Another design choice made in this study was to modulate the amplitude of optical input to the system directly. While channelrhodopsin is most often stimulated using pulsatile inputs, our choice to use feedback to modulate the amplitude of light was made on the basis that it requires the fewest design decisions for the experimenter. In contrast to continuous modulation of amplitude, manipulating pulsatile inputs would require a mapping between the control signal and pulse amplitude, width, and frequency. That said, there is nothing preventing the details of the methodology laid out here to be applied to the modulation of a single pulse parameter such as amplitude or width. We expect that pulsatile and continuously-modulated modes of optical stimulation will have differential effects on higher-order aspects of activity such as local population synchrony [19]. While a robust way to stimulate channelrhodopsin-expressing cells, pulsatile inputs may not suit all control applications, such as manipulating subthreshold neuron polarization. Another alternative to continuous modulation or modulation of pulsatile inputs would be to wait to update the stimulation intensity each time the neuron spikes, as in [42]. In either case, the current methodology may still be applied if the simulations used for controller tuning were altered to reflect the chosen implementation of the control signal.

In this study, we have demonstrated the successful use of a single-degree-of-freedom controller for tracking a time-varying trajectory of interest. In such a system, the designed controller is necessarily a compromise between feeding forward an optimally transformed version of the reference and making best use of feedback to attenuate the influence of disturbances. In this case, the goal of the design procedure was tracking a time-varying reference trajectory. The success of such a simple approach is noteworthy in that it indicates scaling of the reference trajectory is a reasonable first-order approach to stimulus design for the desired rate modulation, while the feedback signal ensured tracking at or near DC by

providing the appropriate bias to this input. The downside of the use of a single-degree-of-freedom controller and tuning for trajectory tracking is that there are no explicit constraints put on the response to load disturbances or other changes in the system. Satisfying requirements on both reference tracking and disturbance rejection independently begs for a two-degree-of-freedom control strategy [49].

Since the predominant manner in which stimulation is applied to neural systems is through either open-loop control or on/off closed-loop control [6, 11, 12, 16, 17, 50], we have compared the effectiveness of our current approach to an open-loop strategy. The principal reason one might utilize feedback continuously would be to adjust for disturbances, whether these come in the form of unmeasured inputs to the system or changes in the underlying dynamics. Although not an explicit goal of controller tuning in this study, the use of feedback does, of course, grant some robustness to such disturbances. As shown previously for the static reference case [19] and again here in the case of a time-varying reference (Figure 2), the use of feedback enables effective control even in the face of uncontrolled input to the system in the form of sensory drive. However, a noteworthy limitation of the current single excitatory opsin approach is the inability to actively inhibit neural activity. As a consequence, disturbances which raise the firing rate above the desired rate at timescales of interest for control will not be effectively rejected. This scenario will necessitate a two-input approach, whether it be expression of both inhibitory and excitatory opsins or expression of opsins targeted to inhibitory neurons. Notably, this extension will require additional modeling for capturing effects of inhibitory inputs (whether direct or indirectly mediated through inhibitory neurons) as well as design of multi-input controllers.

Aside from the ability to reject disturbances, another difference between the results of closed- vs. open-loop stimulation highlighted in Figure 2 and again in Figure 8 is a reduction in across-trial variability provided by reactively updating stimulation in real-time. Given that stimulation of neural systems is often plagued by response variability, this effect of closed-loop stimulation will certainly be a positive for many applications. However, it is also the case that there is naturally-occurring variability in spiking within a given cell and across a population [38,51,52]. There may be scenarios, such as artificial stimulation aimed at naturalistic perception, in which variability will be an important goal for control. Since adding response-variability will be an easier task than reducing it, the current approach represents a good initial step toward a more nuanced control of neural activity.

We have shown that it is feasible to undergo a brief period of system identification followed by controller tuning in simulation within the timespan of an experiment. The fact that the LNP models used for this process were often poor predictors of the input-output relationship in the open-loop sense and yet closely predicted the performance of the closed-loop system highlights the need to assess the goodness of a model in the context of its intended use. If, for example, the use of a mathematical model is purely as a tool for design rather than for making mechanistic inferences, then more complicated models that may be more accurate but also more difficult to estimate may prove unnecessary.

Since the controller was designed around an identified model for the system, it is important to ask how sensitive, or conversely how robust, control system performance is to model

inaccuracy. If the resulting design is sensitive to small deviations in estimated parameters, either the modeling/system identification approach or the control scheme should be modified accordingly. Here, we have found that even with often inaccurate LNP models, we achieved reasonable controller performance (Figures 7&8). Although closed-loop stimulation does relax the requirements for model accuracy, we have also shown that something as simple as inaccurate estimation of static gain of the LNP model can have deleterious effects on control performance in the case of a time-varying reference signal. Importantly, changes in response properties such as gain and baseline firing rate often occur over time in neural systems in the contexts of sensory adaptation [53–56], level of arousal [57], *etc.* It is noteworthy that changes in response gain observed in thalamic neurons in the context of sensory input (*e.g.*, [53, 56]) are often on par or greater than those observed here in the context of optogenetic stimulation. Therefore, the question of controller robustness will be important moving forward.

5. Conclusions

In the present study, we have found that when appropriately tuned, the simple exponential filter and PI controller framework used by Newman *et al* [19] to maintain a constant firing rate can also be effective for eliciting dynamic firing rate trajectories. The strength of this approach lies in its relative simplicity, which should render it readily adoptable for the neuroscience community. With the availability of an open-source platform for real-time control of electrophysiology (RTXI, [58]), the design methodology laid out here should enable widespread application of closed-loop control to optogenetics experiments.

Of particular note is the fact that the simple models used for controller tuning could be poor predictors of the light-to-spiking transformation and still prove useful. This phenomenon highlights that it is not necessary in all contexts to fit the best possible mathematical model. Instead, modelling decisions should be made in light of the intended end use, where considerations such as expedience or mathematical tractability may be of great importance. Indeed, it is possible that in the context of this study, an even simpler model structure could have sufficed. Moreover, the finding that closed-loop optogenetic stimulation can reduce spike-count variability at the timescales investigated highlights the technique's utility for more reliable electrophysiological recordings and opens the door to control strategies that make variability explicit goals of the design process. Taken together, the methodology developed in this study will lay the groundwork for more refined use of optogenetic stimulation and may enable a new class of experiments aimed at elucidating the functional roles of neural populations in networks.

Acknowledgments

This work was supported by the NIH National Institute of Neurological Disorders and Stroke Grants R01NS085447 and R01NS048285. MFB is supported by NSF Graduate Research Fellowship under Grant No. DGE-1650044. AAW and CJW were supported by NIH Training Grant T90DA032466. CJW is supported by the NIH NRSA Pre-doctoral Fellowship (NIH NINDS Grant F31NS089412). CJR is supported by the Georgia Tech Neural Engineering Center, NSF Grant No. CCF-1409422, and James S. McDonnell Foundation Grant No. 220020399. GBS is supported under NIH NINDS Grants R01NS085447 and R01NS048285.

References

1. Hodgkin, Alan L., Huxley, Andrew F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*. 1952; 117(4):500. [PubMed: 12991237]
2. Cardin, Jessica A., Carlén, Marie, Meletis, Konstantinos, Knoblich, Ulf, Zhang, Feng, Deisseroth, Karl, Tsai, Li-Huei, Moore, Christopher I. Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*. 2009; 459(7247):663. [PubMed: 19396156]
3. Madisen, Linda, Mao, Tianyi, Koch, Henner, Zhuo, Jia-min, Berenyi, Antal, Fujisawa, Shigeyoshi, Hsu, Yun-Wei A., Garcia, Alfredo J., III, Gu, Xuan, Zanella, Sebastien, et al. A toolbox of cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nature Neuroscience*. 2012; 15(5):793–802. [PubMed: 22446880]
4. Pala, Aurélie, Petersen, Carl CH. In vivo measurement of cell-type-specific synaptic connectivity and synaptic transmission in layer 2/3 mouse barrel cortex. *Neuron*. 2015; 85(1):68–75. [PubMed: 25543458]
5. Wimmer, Ralf D., Ian Schmitt, L., Davidson, Thomas J., Nakajima, Miho, Deisseroth, Karl, Halassa, Michael M. Thalamic control of sensory selection in divided attention. *Nature*. 2015; 526(7575): 705–709. [PubMed: 26503050]
6. Grosenick, Logan, Marshel, James H., Deisseroth, Karl. Closed-loop and activity-guided optogenetic control. *Neuron*. 2015; 86(1):106–139. [PubMed: 25856490]
7. Ching, ShiNung, Ritt, Jason T. Control strategies for underactuated neural ensembles driven by optogenetic stimulation. *Frontiers in Neural Circuits*. 2013; 7
8. Ahmadian, Yashar, Packer, Adam M., Yuste, Rafael, Paninski, Liam. Designing optimal stimuli to control neuronal spike timing. *Journal of Neurophysiology*. 2011; 106(2):1038–1053. [PubMed: 21511704]
9. Berényi, Antal, Belluscio, Mariano, Mao, Dun, Buzsáki, György. Closed-loop control of epilepsy by transcranial electrical stimulation. *Science*. 2012; 337(6095):735–737. [PubMed: 22879515]
10. Kozák, Gábor, Berényi, Antal. Sustained efficacy of closed loop electrical stimulation for long-term treatment of absence epilepsy in rats. *Scientific Reports*. 2017; 7
11. Paz, Jeanne T., Davidson, Thomas J., Frechette, Eric S., Delord, Bruno, Parada, Isabel, Peng, Kathy, Deisseroth, Karl, Huguenard, John R. Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nature Neuroscience*. 2013; 16(1):64–70. [PubMed: 23143518]
12. Krook-Magnuson, Esther, Armstrong, Caren, Oijala, Mikko, Soltesz, Ivan. On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy. *Nature Communications*. 2013; 4:1376.
13. Ehrens, Daniel, Sritharan, Duluxan, Sarma, Sridevi V. Closed-loop control of a fragile network: application to seizure-like dynamics of an epilepsy model. *Frontiers in neuroscience*. 2015; 9
14. Santaniello, Sabato, Fiengo, Giovanni, Glielmo, Luigi, Grill, Warren M. Closed-loop control of deep brain stimulation: a simulation study. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2011; 19(1):15–24. [PubMed: 20889437]
15. Choi, John S., Brockmeier, Austin J., McNeil, David B., Von Kraus, Lee M., Príncipe, José C., Francis, Joseph T. Eliciting naturalistic cortical responses with a sensory prosthesis via optimized microstimulation. *Journal of Neural Engineering*. 2016; 13(5):056007. [PubMed: 27518368]
16. O'Connor, Daniel H., Hires, S Andrew, Guo, Zengcai V., Li, Nuo, Yu, Jianing, Sun, Qian-Quan, Huber, Daniel, Svoboda, Karel. Neural coding during active somatosensation revealed using illusory touch. *Nature Neuroscience*. 2013; 16(7):958–965. [PubMed: 23727820]
17. Stark E, Roux L, Eichler R, Senzai Y, Royer S, Buzsáki G. Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. *Neuron*. 2014; 83(2):467–480. [PubMed: 25033186]
18. Latchoumane, Charles-Francois V., Ngo, Hong-Viet V., Born, Jan, Shin, Hee-Sup. Thalamic spindles promote memory formation during sleep through triple phase-locking of cortical, thalamic, and hippocampal rhythms. *Neuron*. 2017
19. Newman, Jonathan P., Fong, Ming-fai, Millard, Daniel C., Whitmire, Clarissa J., Stanley, Garrett B., Potter, Steve M. Optogenetic feedback control of neural activity. *eLife*. 2015; 4:e07192. [PubMed: 26140329]

20. Whitmire, Clarissa J., Waiblinger, Christian, Schwarz, Cornelius, Stanley, Garrett B. Information coding through adaptive gating of synchronized thalamic bursting. *Cell reports*. 2016; 14(4):795–807. [PubMed: 26776512]
21. Whitmire, Clarissa J., Millard, Daniel C., Stanley, Garrett B. Thalamic state control of cortical paired-pulse dynamics. *Journal of neurophysiology*. 2017; 117(1):163–177. [PubMed: 27760816]
22. Waiblinger, Christian, Brugger, Dominik, Schwarz, Cornelius. Vibrotactile discrimination in the rat whisker system is based on neuronal coding of instantaneous kinematic cues. *Cerebral Cortex*. 2015; 25(4):1093–1106. [PubMed: 24169940]
23. Kleinfeld, David, Ahissar, Ehud, Diamond, Mathew E. Active sensation: insights from the rodent vibrissa sensorimotor system. *Current Opinion in Neurobiology*. 2006; 16(4):435–444. [PubMed: 16837190]
24. Hill, Dan N., Bermejo, Roberto, Zeigler, H Philip, Kleinfeld, David. Biomechanics of the vibrissa motor plant in rat: rhythmic whisking consists of triphasic neuromuscular activity. *The Journal of Neuroscience*. 2008; 28(13):3438–3455. [PubMed: 18367610]
25. Petersen, Carl CH. Cortical control of whisker movement. *Annual Review of Neuroscience*. 2014; 37:183–203.
26. Nawrot, Martin, Aertsen, Ad, Rotter, Stefan. Single-trial estimation of neuronal firing rates: from single-neuron spike trains to population activity. *Journal of Neuroscience Methods*. 1999; 94(1): 81–92. [PubMed: 10638817]
27. Koyama, Shinsuke, Shinomoto, Shigeru. Histogram bin width selection for time-dependent poisson processes. *Journal of Physics A: Mathematical and General*. 2004; 37(29):7255.
28. Shimazaki, Hideaki, Shinomoto, Shigeru. A method for selecting the bin size of a time histogram. *Neural Computation*. 2007; 19(6):1503–1527. [PubMed: 17444758]
29. Shimazaki, Hideaki, Shinomoto, Shigeru. Kernel bandwidth optimization in spike rate estimation. *Journal of Computational Neuroscience*. 2010; 29(1–2):171–182. [PubMed: 19655238]
30. McFarland, James M., Cui, Yuwei, Butts, Daniel A. Inferring nonlinear neuronal computation based on physiologically plausible inputs. *PLoS Computational Biology*. 2013; 9(7):e1003143. [PubMed: 23874185]
31. Paninski, Liam. Convergence properties of three spike-triggered analysis techniques. *Network: Computation in Neural Systems*. 2003; 14(3):437–464.
32. Park, Mijung, Pillow, Jonathan W. Receptive field inference with localized priors. *PLoS Comput Biol*. 2011; 7(10):e1002219. [PubMed: 22046110]
33. Stanley, Garrett B. Adaptive spatiotemporal receptive field estimation in the visual pathway. *Neural Computation*. 2002; 14(12):2925–2946. [PubMed: 12487798]
34. Pillow, Jonathan W., Paninski, Liam, Uzzell, Valerie J., Simoncelli, Eero P., Chichilnisky, EJ. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *Journal of Neuroscience*. 2005; 25(47):11003–11013. [PubMed: 16306413]
35. Zhuang, Mang, Atherton, DP. *IEE Proceedings D-Control Theory and Applications*. Vol. 140. IET; 1993. Automatic tuning of optimum pid controllers; p. 216-224.
36. Waiblinger, Christian, Brugger, Dominik, Whitmire, Clarissa J., Stanley, Garrett B., Schwarz, Cornelius. Support for the slip hypothesis from whisker-related tactile perception of rats in a noisy environment. *Frontiers in Integrative Neuroscience*. 2015; 9
37. Borden, Peter Y., Ortiz, Alex D., Waiblinger, Christian, Sederberg, Audrey J., Morrissette, Arthur E., Forest, Craig R., Jaeger, Dieter, Stanley, Garrett B. Genetically expressed voltage sensor arelight for imaging large scale cortical activity in the anesthetized and awake mouse. *Neurophotonics*. 2017; 4(3):031212–031212. [PubMed: 28491905]
38. Churchland, Mark M., Byron, M Yu, Cunningham, John P., Sugrue, Leo P., Cohen, Marlene R., Corrado, Greg S., Newsome, William T., Clark, Andrew M., Hosseini, Paymon, Scott, Benjamin B., et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*. 2010; 13(3):369–378. [PubMed: 20173745]
39. Carandini, Matteo, Demb, Jonathan B., Mante, Valerio, Tolhurst, David J., Dan, Yang, Olshausen, Bruno A., Gallant, Jack L., Rust, Nicole C. Do we know what the early visual system does? *Journal of Neuroscience*. 2005; 25(46):10577–10597. [PubMed: 16291931]

40. Olshausen, Bruno A., Field, David J. How close are we to understanding v1? *Neural Computation*. 2005; 17(8):1665–1699. [PubMed: 15969914]
41. Liu, Jianbo, Khalil, Hassan K., Oweiss, Karim G. Model-based analysis and control of a network of basal ganglia spiking neurons in the normal and parkinsonian states. *Journal of Neural Engineering*. 2011; 8(4):045002. [PubMed: 21775788]
42. Miranda-Dominguez O, Gonía J, Netoff TI. Firing rate control of a neuron using a linear proportional-integral controller. *Journal of Neural Engineering*. 2010; 7(6):066004. [PubMed: 20975212]
43. Nandi, Anirban, Kafashan, Mohammad Mehdi, Ching, ShiNung. American Control Conference (ACC), 2016. IEEE; 2016. Controlling point process generalized linear models of neural spiking; p. 5779–5784.
44. Iolov, Alexandre, Ditlevsen, Susanne, Longtin, André. Stochastic optimal control of single neuron spike trains. *Journal of Neural Engineering*. 2014; 11(4):046004. [PubMed: 24891497]
45. Urbain, Nadia, Salin, Paul A., Libourel, Paul-Antoine, Comte, Jean-Christophe, Gentet, Luc J., Petersen, Carl CH. Whisker-related changes in neuronal firing and membrane potential dynamics in the somatosensory thalamus of awake mice. *Cell Reports*. 2015; 13(4):647–656. [PubMed: 26489463]
46. Gutnisky, Diego Adrian, Yu, Jianing, Hires, Samuel Andrew, To, Minh-Son, Bale, Michael Ross, Svoboda, Karel, Golomb, David. Mechanisms underlying a thalamocortical transformation during active tactile sensation. *PLoS Computational Biology*. 2017; 13(6):e1005576. [PubMed: 28591219]
47. Harris, Kenneth D., Henze, Darrell A., Hirase, Hajime, Leinekugel, Xavier, et al. Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature*. 2002; 417(6890):738. [PubMed: 12066184]
48. Churchland, Mark M., Cunningham, John P., Kaufman, Matthew T., Foster, Justin D., Nuyujukian, Paul, Ryu, Stephen I., Shenoy, Krishna V. Neural population dynamics during reaching. *Nature*. 2012; 487(7405):51. [PubMed: 22722855]
49. Åström, Karl J., Wittenmark, Björn. Adaptive control. Courier Corporation; 2013.
50. Sorokin, Jordan M., Davidson, Thomas J., Frechette, Eric, Abramian, Armen M., Deisseroth, Karl, Huguenard, John R., Paz, Jeanne T. Bidirectional control of generalized epilepsy networks via rapid real-time switching of firing mode. *Neuron*. 2017; 93(1):194–210. [PubMed: 27989462]
51. Shadlen, Michael N., Newsome, William T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience*. 1998; 18(10):3870–3896. [PubMed: 9570816]
52. Goris, Robbe LT., Movshon, J Anthony, Simoncelli, Eero P. Partitioning neuronal variability. *Nature Neuroscience*. 2014; 17(6):858–865. [PubMed: 24777419]
53. Lesica, Nicholas A., Jin, Jianzhong, Weng, Chong, Yeh, Chun-I., Butts, Daniel A., Stanley, Garrett B., Alonso, Jose-Manuel. Adaptation to stimulus contrast and correlations during natural visual stimulation. *Neuron*. 2007; 55(3):479–491. [PubMed: 17678859]
54. Chander, Divya, Chichilnisky, EJ. Adaptation to temporal contrast in primate and salamander retina. *Journal of Neuroscience*. 2001; 21(24):9904–9916. [PubMed: 11739598]
55. Mante, Valerio, Frazor, Robert A., Bonin, Vincent, Geisler, Wilson S., Carandini, Matteo. Independence of luminance and contrast in natural scenes and in the early visual system. *Nature Neuroscience*. 2005; 8(12):1690. [PubMed: 16286933]
56. Maravall, Miguel, Alenda, Andrea, Bale, Michael R., Petersen, Rasmus S. Transformation of adaptation and gain rescaling along the whisker sensory pathway. *PloS One*. 2013; 8(12):e82418. [PubMed: 24349279]
57. Castro-Alamancos, Manuel A., Oldford, Elizabeth. Cortical sensory suppression during arousal is due to the activity-dependent depression of thalamocortical synapses. *The Journal of Physiology*. 2002; 541(1):319–331. [PubMed: 12015438]
58. Patel, Yogi A., George, Ansel, Dorval, Alan D., White, John A., Christini, David J., Butera, Robert J. Hard real-time closed-loop electrophysiology with the real-time experiment interface (rtxi). *PLOS Computational Biology*. 2017; 13(5):e1005430. [PubMed: 28557998]

Figure 1. Closed loop optogenetic control of firing rate
(A) Physical diagram.
(B) System block flow diagram.
(C) Procedure for closed-loop stimulation experiments. The observer was designed for a given reference firing rate pattern previous to experiments. A model was fit to data recorded for system identification during the experiment. Using this model, controller gains were optimized in simulation. These parameters were then used for experimental closed-loop stimulation.

Figure 2. Closed- vs. open-loop optogenetic control of dynamic firing rate trajectories

(A) Closed- and open-loop control of sinusoidally-modulated firing rate. Closed-loop (black) control and pulsatile open-loop (red) stimulation were used to elicit a 1 Hz sinusoidally modulated firing rate. Light lines correspond to single trial firing rates estimated by smoothing spike trains with a Gaussian window (120 ms SD); bold lines are the trial-averaged rate. Average control inputs (*i.e.*, light) are below the corresponding firing rate trajectory.

(B) Closed- and open-loop control in presence of a disturbance. Control was challenged with a whisker disturbance at 2 seconds into the control epoch, as shown in gray (top).

(C) Closed-loop and open-loop control of non-sinusoidal firing rate. Top, firing rates for closed-loop (black) vs. open-loop (red) control: average in bold, while fills represent 95% confidence intervals for smoothed PSTH. Middle, Fano factor calculated in 250-ms sliding

window for closed- and open-loop control ($n = 25$ trials). Bottom, trial-averaged control inputs for closed-loop (black) or open-loop (red).

Figure 3. Observer design: Choosing filter bandwidth

(A) Conceptual diagram. A given sinusoidal firing rate (λ) drove a Poisson spike generator (P). The resulting spike train was multiplied by λ^{-1} (not shown) ahead of filtering. Filters parameterized by a time constant, τ , yielded an estimate of the true rate.

(B) Optimal time constant as a function of expected number of spikes per period and degree of modulation about mean (top) and as a function of z (bottom). For visualization, the optimal filter has been normalized by the frequency of each sinusoid. Fit: $a = 0.423$, $b = 0.389$.

(C) Filter time constants designed for single-trial estimation where $\mu = 20$ spikes/s, $\sigma/\mu = 1$. Bold purple, region of frequencies where the derived design equation was fit. Light purple, frequencies at which the design equation was extrapolated ($z < 12$). Insets, example single-trial estimates (purple) of the ground truth rate (green) at indicated modulation frequencies.

Figure 4. Controller Design: Tuning the controller around an LNP model neuron

(A) Controller design through simulation. The closed-loop system was simulated with a model of the neural system for design purposes.

(B) Example tuning surface for 5 Hz sinusoidal trajectory. In simulation, the controller was tasked with tracking a sinusoidal trajectory (here, 5 Hz), using the observer designed previously for the corresponding reference. The objective was to minimize the squared tracking error, weighted as a function of frequencies important for the control task.

(C) Examples of optimal and suboptimal controller gains. Frequency-domain error (top row) corresponds to amplitude of error between the raw instantaneous rate (n) and the reference at DC and the modulation frequency (here, 5 Hz). For comparison, the square root of the frequency-weighted squared error (J_{fwt}) is also provided. Error bars correspond to +1 SD. Corresponding error spectra are provided (middle row), as compared to a simulated Poisson

process modulated at the reference rate (light grey). Green lines highlight DC and f_{mod} . Finally, time-domain tracking is provided (bottom).

Figure 5. LNP Model Performance: Open-loop vs. closed-loop

(A) Fitting the linear-nonlinear (LN) model.

(B) A typical LN model fit to training data. Left, the linear system estimated using whitened spike triggered averaging; error bars correspond to ± 1 SD for the lagged coefficients of the kernel when fit to 10 subsets of the full training dataset. Right, the static nonlinearity fit by maximum likelihood, given the observed spikes. Gray points indicate the experimental firing rate (PSTH smoothed with 1-ms Gaussian) versus the kernel-filtered stimulus.

(C) LNP prediction of response to open-loop 'replay' of stimulation used experimentally during a 5 Hz sinusoidal control task, using the same LN model shown in B. Top, firing rate predicted by the model (red) as compared to the experimental data for the same cell (black); bottom, experimental optical input.

(D) LNP prediction of response to simulated closed-loop stimulation. Firing rate and light input predicted by the model (blue) as compared to the experimental data for the same cell (black); bottom, simulated (blue) and experimental (black) closed-loop stimulation

Figure 6. Robustness of Control to Model Inaccuracy

(A) Model Perturbation and Simulation. Tire static gain (g) and tire bias (m) of tire linear component of tire LNP model were systematically perturbed and simulated either in closed- or open-loop.

(B) Grid Search Over Gain and Bias. Fractional changes in g and m ranged from five times smaller to five times greater than the original parameter value. Grayscale represents the percentage change in tracking performance (J_{fwt}) between that of the original model and each perturbed version. This tracking error was calculated from 1 second onward for 5-second control epochs.

(C) Changing Bias. Holding all else constant, m was changed, and the tracking performance was quantified for closed-loop control vs. open-loop replay of the light traces used to stimulate the original neuron. Gray circles indicate the estimated biases of all recorded neurons relative to the model used for perturbation study ($n=20$).

(D) Examples of Simulated Control When Bias Estimation Inaccurate. (Top) Outcome when the actual neuron (blue) was 2 times less biased than the model around which controller were tuned (black); (Bottom) the outcome when the actual neuron (blue) was 2 times more negatively biased than the model around which the controller was tuned (black). Scale bar indicates 20 spikes/s.

(E) Changing Gain. Holding all else constant, g was changed and the tracking performance quantified for closed-loop control vs. open-loop replay of the light traces used to stimulate

the original neuron. Gray circles indicate the estimated gains of all recorded neurons relative to the model used for perturbation study ($n=20$).

(F) Examples of Simulated Control When Gain Estimation Inaccurate. (Top) Outcome when the actual neuron (blue) was 2 times less sensitive than the model around which controller were tuned (black); (Bottom) the outcome when the actual neuron (blue) was 2 times more sensitive than the model around which the controller was tuned (black). Scale bar indicates 20 spikes/s

Figure 7. Sinusoidal tracking performance

(A) Example experimental implementation (“Cell 3”) where the controller and observer were tuned for a trajectory modulated at 5 Hz. The third of a five-second control epoch is shown.

(B) Population performance for tracking a 5 Hz sinusoidal trajectory (LNP prediction vs. experimental): average (black bar) and individual cells (colored symbols). 95% confidence intervals for this metric were calculated for simulated Poisson firing at the reference rate and plotted in light grey. Left, results of design procedure predicted by the LNP models fit and tuned around during the experiment. Middle, experimental closed-loop tracking performance. Right, experimental open-loop tracking performance.

(C) Closed- vs. open-loop experimental performance on 1, 5, & 10 Hz sinusoidal control tasks. Closed-loop tracking error is significantly less than open-loop ($p < 0.05$, Wilcoxon signed rank test, $n = 12$ comparisons, 4 different cells).

Figure 8. Non-sinusoidal tracking performance

(A) Example implementation (“Cell 5”) on the more naturalistic, non-sinusoidal trajectory. Controller and observer were tuned for a 10 Hz sinusoidally modulated trajectory. The corresponding error spectrum is also shown (bottom). For comparison, the error spectrum of simulated Poisson firing at the reference rate is plotted in light grey.

(B) Population tracking error for non-sinusoidal trajectory. 95% confidence intervals for this metric are calculated for simulated Poisson firing at the reference rate and plotted in light grey. Left, simulated LNP performance predicted by design procedure. Middle, experimental closed-loop tracking performance. Right, experimental open-loop tracking performance.

(C) Experimental across-trial variability in closed-loop vs. open-loop. Treating 750-ms onward as ‘steady-state’, time-averaged Fano factor calculated in a 250-ms sliding window for the closed- vs. open-loop control cases.