

**TOWARDS OPTOGENETIC FEEDBACK CONTROL OF NEURAL POPULATION
DYNAMICS**

A Dissertation Proposal
Presented to
The Academic Faculty

By

Kyle Johnsen

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
College of Engineering
Department of Biomedical Engineering

Georgia Institute of Technology and Emory University

Jan 2023

© Kyle Johnsen 2023

TOWARDS OPTOGENETIC FEEDBACK CONTROL OF NEURAL POPULATION DYNAMICS

Thesis committee:

Dr. Nabil Imam
Computational Science and Engineering
Georgia Institute of Technology

Dr. Garrett Stanley
Biomedical Engineering
Georgia Institute of Technology & Emory University

Dr. Chethan Pandarinath
Biomedical Engineering
Emory University & Georgia Institute of Technology

Dr. Patricio Vela
Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Christopher Rozell (advisor)
Electrical and Computer Engineering
Georgia Institute of Technology

Date approved:

TABLE OF CONTENTS

Chapter 1: Specific Aims	1
1.1 Rationale	1
1.2 Aim 1: A CLOC experiment simulation testbed	1
1.3 Aim 2: Multi-input CLOC	2
1.4 Aim 3: Using CLOC to manipulate latent neural dynamics	2
Chapter 2: Background	4
2.1 Closed-loop control in neuroscience	4
2.2 Various scales and tools for closed-loop control	5
2.3 Previous work	6
2.4 Unexplored applications for network-level control	7
Chapter 3: Aim 1 - A CLOC simulation testbed	9
3.1 Rationale	9
3.1.1 Innovation	9
3.2 Approach	10
3.2.1 Guiding principles	10
3.2.2 Closed-loop simulation architecture	10
3.2.3 Electrode recording	11

3.2.4	Optogenetic stimulation	13
3.3	Results	13
3.3.1	Model validation	13
3.3.2	Example experiments	15
3.4	Limitations	17
3.5	Summary of contributions	17
Chapter 4: Aim 2 - Multi-input CLOC		19
4.1	Rationale	19
4.1.1	Innovation	20
4.2	Approach	21
4.2.1	System and controller formulation	21
4.2.2	Control method comparison <i>in silico</i>	23
4.3	Expected results	24
4.4	Preliminary results	24
4.5	Potential pitfalls & alternative strategies	25
Chapter 5: Aim 3 - Control of latent population dynamics		27
5.1	Rationale	27
5.1.1	Innovation	28
5.2	Approach	29
5.2.1	Formulate task with latent dynamics hypothesis	29
5.2.2	Train RSNN models	30
5.2.3	Fit dynamical systems model	31

5.2.4	Control of latent factors	32
5.2.5	Exploration of experiment parameters and expected effects	33
5.2.6	Exploration of model realism and expected effects	34
5.3	Expected results	35
5.4	Potential pitfalls & alternative strategies	36
Chapter 6: Proposed timeline		38
References		39

SUMMARY

As the importance of causal inference becomes increasingly apparent in neuroscience, so does the need for technology enabling precise manipulation of neural activity. Feedback control is a prime example—widely used throughout the engineering disciplines, it has had a significant impact via techniques such as the voltage clamp and the dynamic clamp on cellular neuroscience. However, it has yet to be widely applied at the mesoscale of circuits and networks despite recent improvements in interfacing technology, such as optogenetics. Challenges to such applications include the complexity of implementing fast closed-loop experiments, the need to adapt the mature methods of control theory to the idiosyncratic constraints of systems neuroscience experiments, and the lack of established technical guidelines for applying feedback control to address complex scientific questions.

In this work I propose to begin to address these challenges in three aims. In Aim 1, I develop a simulation framework for easily prototyping closed-loop optogenetic control (CLOC) experiments *in silico*, thus enabling CLOC experiment design and method development without the costs of *in-vivo* experiments or up-front investments in hardware-software systems. In Aim 2, I will translate sophisticated model-based feedback control algorithms to the experimental setting of multi-input CLOC—the simultaneous use of multiple light sources and/or opsins—and test the virtues of these algorithms compared to the limited, simpler ones previously used. Finally, in Aim 3, I will explore how control quality varies with experimental parameters in a promising future application of CLOC—controlling the latent dynamics of neural population activity—and test the hypothesis that per-neuron actuation will not be needed. I will do this *in silico* with recurrent spiking neural networks trained using biologically plausible methods and differing degrees of brain-like structure. This work will thus equip the systems neuroscience community to more fully take advantage of CLOC with an accessible testing and development environment, multi-input actuation, and a point of reference for designing experiments capable of answering complex scientific

questions.

CHAPTER 1

SPECIFIC AIMS

1.1 Rationale

As the importance of causal inference becomes increasingly recognized in neuroscience, the need for technology enabling precise manipulation of neural variables becomes apparent. Feedback control is an important class of such manipulations for its ability to increase inference power by reducing response variability. Widely used throughout the engineering disciplines, it has had a significant impact through a variety of techniques (e.g., voltage clamp, dynamic clamp) on cellular neuroscience. However, feedback control also has great potential at the mesoscale/systems level, potentially enabling researchers to *unambiguously infer the downstream effects of circuit/population-level neural activity*.

For a number of reasons, though, *this potential has not been widely realized*. I posit that the main challenges to adoption rather include the **complexity of implementing** fast closed-loop experiments, the need to **adapt the mature methods of control theory** to the idiosyncratic constraints of systems neuroscience experiments, and the **lack of established technical guidelines** for applying feedback control to address complex scientific questions. **The proposed work aims to begin to address these challenges, and thus strengthen the set of causal tools available to probe neural systems.**

1.2 Aim 1: A CLOC experiment simulation testbed

One significant obstacle to closed-loop optogenetic control (CLOC) experiments is the cost of acquiring and configuring compatible hardware-software systems. Moreover, the

maintenance of animals or cell cultures inherent in lab experiments can slow the pace of developing novel techniques. In Aim 1, I attempt to address these obstacles by developing a simulation framework for easily prototyping CLOC experiments *in silico*, thus enabling faster, cheaper CLOC experiment design and method development. We demonstrate the software's utility in different virtual experiments and provide it to the public as open-source software with thorough documentation.

1.3 Aim 2: Multi-input CLOC

Multi-input CLOC—the simultaneous use of multiple light sources and/or opsins—is necessary for precise manipulation of neural systems. However, the basic control theory methods previously used for CLOC, while fast, do not take actuator constraints into account and are thus likely to be inadequate for multi-actuator problems. The field of control theory provides elegant, powerful—if slower—solutions to this class of problems, but applying them requires interdisciplinary expertise. In this aim I will translate more sophisticated model-based feedback control algorithms to the multi-input CLOC setting and assess their merits compared to limited, simpler methods *in silico* for real-time control of brain networks.

1.4 Aim 3: Using CLOC to manipulate latent neural dynamics

To our knowledge, CLOC has yet to be applied in answering complex systems neuroscience questions. In this aim, to pave the way for future *in-vivo* experiments that accomplish this, I propose to develop technical and conceptual guidelines as I control the latent dynamics of simulated neural populations. First, I will produce these virtual models by training recurrent spiking neural networks with state-of-the-art, biologically plausible methods—each differing in their degrees of brain-like architecture and training procedure complexity. I will then use the simulation testbed of Aim 1 to explore how control quality varies with recording, stimulation, and control parameters and the complexity and size of the system—thus giving researchers a tentative idea of the relative importance of each factor of CLOC. I

will test the hypothesis that low-dimensional dynamics can be manipulated with similarly low-dimensional, rather than per-neuron, actuation.

CHAPTER 2

BACKGROUND

2.1 Closed-loop control in neuroscience

Mesoscale neuroscience—on the level of populations of neurons, rather than the whole brain or individual cells—is currently undergoing a revolution fueled by advances in neural manipulation [1, 2, 3, 4, 5, 6, 7, 8] and measurement [9, 10, 11, 12, 13, 14, 15, 16] technologies as well as data analysis methods [17, 18, 19, 20, 21, 22]. These have yielded unprecedented datasets [23, 24] and insights into network activity and plasticity [25, 26, 27, 28, 29]. Moreover, they enable novel experimental paradigms such as direct closed-loop control of neural activity [30, 32, 33, 34, 35, 36, 37, 38, 39, 31]. This closed-loop stimulation offers exciting prospects of intervention in processes that are too fast or unpredictable to control manually or with pre-defined stimulation, such as sensory information processing, motor planning, and oscillatory activity. Unlike other forms of closed-loop control using environmental/stimulus input [40], behavioral output [41] or neurofeedback training [42, 8], the direct control of neural activity itself provides more opportunities for revealing the downstream effects of that activity.

Closed-loop control of neural activity can be implemented in an event-triggered sense [38, 39, 31]—enabling the experimenter to respond to discrete events of interest, such as the arrival of a traveling wave [43] or sharp wave-ripple [44]—or in a feedback sense [34, 35, 36, 37], driving the system towards a reference state or trajectory. The latter has multiple advantages over open-loop control (delivery of a pre-defined stimulus): by rejecting exogenous inputs, noise, and disturbances, it reduces variability across time and across trials,

allowing for finer-scale inference. Additionally, it can compensate for model mismatch as it actively reduces error, whereas an open-loop stimulus is defined and limited by imperfect models. Moreover, whereas traditional perturbation methods often include lesioning [45], unnatural silencing, or extreme stimulation, feedback control poses a more naturalistic alternative. Clamping the activity of a population of neurons to their baseline, for example, effectively shuts down information flow from those neurons without departing from typical firing rates.

2.2 Various scales and tools for closed-loop control

Closed-loop control of neural activity can be performed at multiple scales and with different sets of tools. At the smallest, sub-neuron scale, feedback control has yielded decades of fruitful research in the form of tools such as the voltage clamp [46] and dynamic clamp [47, 48], controlling electrical properties of the cell membrane. The frontiers of this small-scale neuroscience often involve scaling up to many neurons and scaling down to subcellular structures such as dendrites, but multi-electrode, *in-vivo*, intracellular recording methods are challenging [49, 50, 51, 52, 53]. Optical tools—e.g., optogenetics and fluorescence microscopy—can circumvent the difficulties of working with electrodes at such small scales, but an optical approach is not yet feasible for this purpose. The obstacles lie mainly in recording technology: the kinetics of both voltage indicators [16] and intracellular calcium [54] are too slow to capture phenomena faster than a typical action potential.

By contrast, the current state of technology is ripe for innovating closed-loop control methods at larger scales of neural activity, from single neurons to populations and circuits. Several promising combinations of recording and stimulation modalities are possible and still relatively novel: electrode recording with optogenetic stimulation [55, 34, 35, 36], fluorescence microscopy with electrical stimulation [16], fluorescence microscopy with photostimulation (all-optical control) [56, 57, 58, 37, 59], and fMRI with optionally transcranial

photostimulation [60, 61]. Each of these tool combinations has pros and cons in terms of spatial and temporal resolution, crosstalk [62], and degrees of freedom. A natural starting point for many neuroscientists is the first of these tool sets—electrode recording combined with optogenetics—since the two methods are so widely used, interfere little with each other (as long as metal electrodes are not directly illuminated [55, 62]), and allow for genetically targeted stimulation. I will henceforth refer to this combination as CLOC, following the convention established by previous works [35, 36].

2.3 Previous work

The proposed work builds on the work my lab and collaborators have done previously in developing CLOC. This includes pioneering work taking feedback control from the cellular to the network level by controlling population firing rates *in vitro*—first by Wagenaar *et al.* [63], who reduce unnatural global bursting and conclude as a result that such bursting results from the lack of the exogenous input present in the live, awake cortex. Later, Newman *et al.* [64] used the same proportional error-driven controller in a demonstration of user-friendly real-time electrophysiology software.

These model-free control methods later saw further development and the introduction of optogenetic actuation. In the first known demonstration of optogenetic feedback control, Newman *et al.* [34] used bidirectional stimulation for fixed and slowly varying firing rate targets *in vitro* (Figure 2.1a). The authors improved control methods by introducing an integral term to the model-free controller and, in another first, demonstrated feedback control of firing rates *in vivo* by bringing CLOC to the anesthetized rat. Bolus *et al.* [35] used PI control again, but developed a more principled approach to set estimation and control parameters and tracked dynamic firing rate trajectories down to a ~100-ms timescale (Figure 2.1b).

Later, in a direct precursor to the proposed work, Bolus *et al.* [36] employed more so-

phisticated and scalable optimal feedback control methods that can be characterized as an adaptive linear quadratic regulator (LQR). Built on a state-space, linear dynamical system model, the controller estimates the latent state using Kalman filtering [65] and calculates in real-time the optimal control input based on the current error. The authors enable adaptivity by estimating a second latent state, added to the first, that represents a slowly varying bias or disturbance to the system—important especially in awake animal experiments, where dynamic brain state changes contribute to high per-trial variability (Figure 2.1c).

2.4 Unexplored applications for network-level control

While preliminary technological foundations for feedback control of neural activity have already been laid, CLOC has yet to be applied to network-level variables beyond population firing rates to further causal hypothesis testing. Examples of these potential targets for control include the activity of different cell types [66, 67, 68, 69, 70, 71]; the type [72, 73, 74], frequency [75], amplitude [75], spike coherence [76, 77] and interactions [78, 79] of different oscillatory patterns [80]; discrete phenomena such as bursts [81], sharp wave ripples [44], oscillatory bursts [82, 83, 84, 85, 86], traveling waves [87, 88, 89, 37, 90, 43], or sleep spindles [91]; and latent states describing neural dynamics [92, 93, 94, 95, 96, 97, 98, 99], including those most predictive of behavior [20, 100, 101]. In the proposed work, I focus on the latter for their success in generating computational hypotheses for brain function and in predicting behavior. In the following chapters, I describe plans to move towards *in-vivo* control of latent dynamics by developing CLOC infrastructure and algorithms and finally performing simulations and analysis *in silico*.

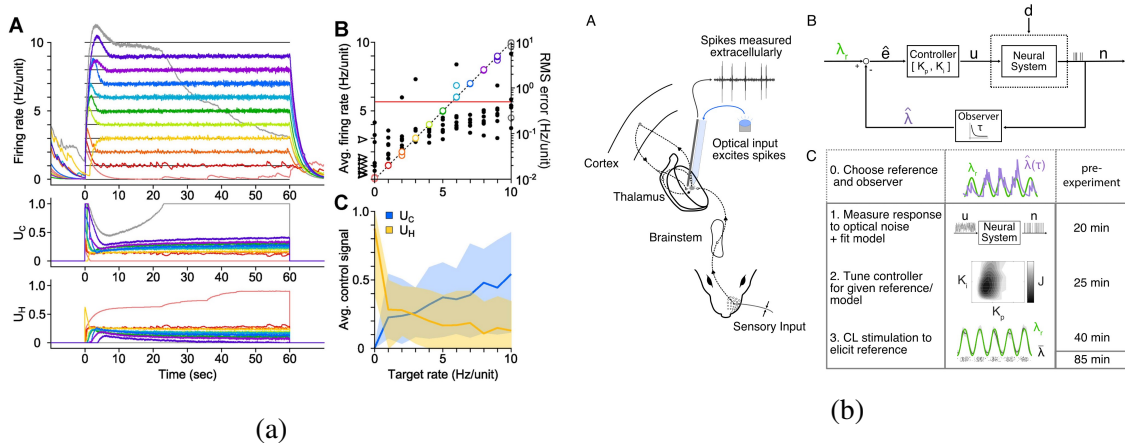


Figure 2.1: Previous CLOC experiments. (A) Figure 2 from Newman *et al.* [34], demonstrating control clamping a cultured neuron to different firing rates. U_C refers to the control signal for channelrhodopsin2(H134R) (ChR2_R), parametrizing 470-nm light delivery. U_H likewise parametrizes 590-nm light delivery to activate enhanced halorhodopsin-3.0 (eNpHR3.0). (B) Figure 2 from Bolus *et al.* [35], outlining an *in-vivo* experiment setup, a control diagram showing how optical input is determined from the firing rate estimated in real time, and a strategy for tuning the controller. (C) Figure 4b from Bolus *et al.* [36], showing how CLOC clamps the firing rate of a single thalamic neuron in the awake mouse over multiple trials, reducing response variability (as measured by the Fano factor).

CHAPTER 3

AIM 1 - A CLOC SIMULATION TESTBED

3.1 Rationale

CLOC experiments can be difficult and costly, posing a barrier to entry for neuroscientists that find CLOC's advantages attractive. Their lab might lack the funds or time to invest in hardware and software infrastructure or the confidence that such an investment would yield successful experiments. Moreover, when the proposed experiment requires signal processing/control method development, iterating on designs *in-vivo* may be cumbersome, given the additional cost of animal care and training. However, these costs and risks can be mitigated through *in-silico* prototyping. Given a reliable model of the system of interest, one can simulate an experiment and assess its effectiveness. Alternate models of a system can be tested to assess robustness of a method to unknown model properties, or that method can be validated on a variety of systems to determine its general applicability. This strategy not only allows for a researcher to evaluate and optimize methods before committing significant resources to them, but also accelerates the development cycle.

3.1.1 Innovation

For these reasons, I have developed Cleo: Closed Loop, Electrophysiology, and Optogenetics experiment simulation testbed. Unlike existing software, Cleo simultaneously provides a high-level interface to fast and flexible neural network simulations; easy, model-independent injection of electrode recording and optogenetic perturbations; and a real-time, closed-loop processing module capable of modeling communication and processing delays inherent in real experiments. I thus provide a low-cost environment to design and develop CLOC

methods, as well as for researchers to see if CLOC may serve their research agenda.

3.2 Approach

3.2.1 Guiding principles

Two factors drove our choice of recording and stimulation models to integrate into Cleo. First, because a main purpose of Cleo is to enable prototyping of experiments, we focused on models at the level of the parameters an experimenter is able to alter—this led us to choose models embedded in physical, 3D space. Second, we assumed that highly realistic models of individual neurons are not needed to capture most mesoscale-level phenomena. Accordingly, Cleo was developed with the intention of using point neurons models, rather than multi-compartment neuron models with realistic morphologies.

In addition to our modeling priorities, the goals of usability and extensibility guided our choices in software dependencies and infrastructure. Ease of use is important to make Cleo as accessible as possible to researchers with varied backgrounds, and motivated Cleo’s modular design, which allows users to add different recording or stimulation devices with little or no modification to the underlying network model. Additionally, extensibility is important for the testbed to remain relevant under changing needs in the future, allowing for new functionality to be easily added in a “plug-in,” modular architecture.

3.2.2 Closed-loop simulation architecture

We chose Brian 2 [102] as the spiking neural network simulator to build Cleo around. Brian has the advantages of being flexible and intuitive, where models are defined with user-provided mathematical equations. Moreover, it is written in Python, an open-source, intuitive language [103] widely used in computational neuroscience [104, 105]—users do not need to use any other languages to use Brian and Cleo. Brian is also relatively fast (especially since it is developed primarily for point neuron simulations), as shown in

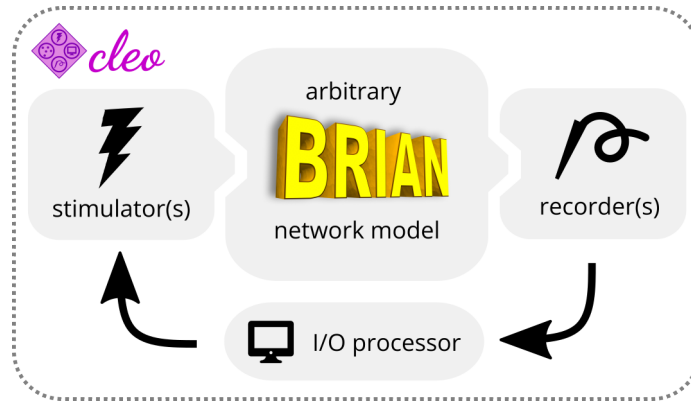


Figure 3.1: A conceptual diagram of Cleo’s functionality. Cleo wraps a Brian network model, injects stimulation and recording devices, and interfaces with the network in real time through a simulated “I/O processor”.

benchmarks [102].

Cleo provides three modules—recording, stimulation, and closed-loop processing—for integration with an existing Brian model (see Figure 3.1). Cleo’s functionality centers around a `CLSimulator` object that orchestrates the interactions between these three modules and the Brian simulation by injecting devices, running the simulation, and communicating with an `IOProcessor` object at each time step. The `IOProcessor` receives recorder measurements according to a specified sampling rate and returns any updates to stimulator devices. To simulate the effects of latency in closed-loop control, Cleo provides a `LatencyIOProcessor` class which stores control signals in a buffer to later deliver them after an arbitrary delay.

3.2.3 Electrode recording

Because we have prioritized point neuron simulations, the electrode functionality currently implemented in Cleo does not rely on raw extracellular potentials, which can only be computed from multi-compartment neurons [106, 107]. This biophysical forward modeling approach has been taken in other software [108, 109, 110, 111]. To approximate sorted and multi-spike recording without filtering and thresholding of extracellular potentials,

Cleo simply takes ground-truth spikes and stochastically determines which to report with probability inversely proportional to distance to the electrode [112, 113] (see Figure 3.2).

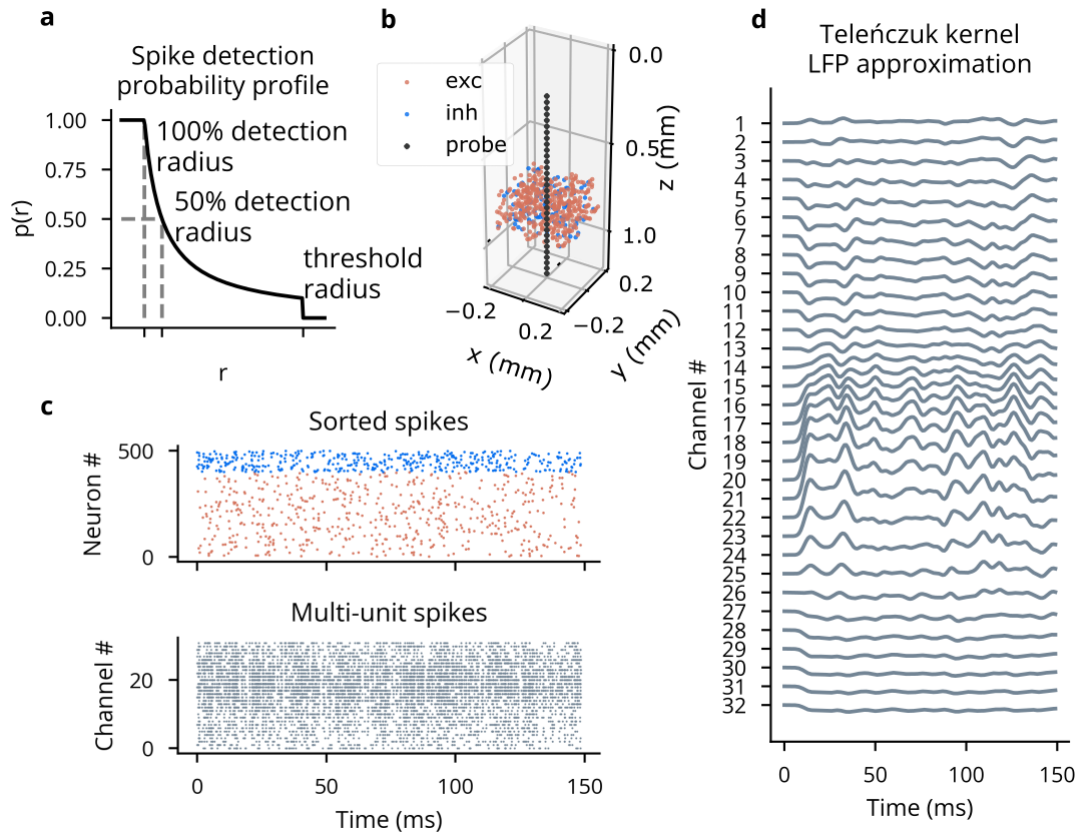


Figure 3.2: Illustration of LFP and spiking from Cleo’s electrophysiology module. (A) The probabilistic spike detection model. All spikes within the 100% detection radius, 50% of spikes at the 50% detection radius, and none of those outside the threshold radius are recorded. The detection probability decays with $1/r$. (B) An example plot generated by Cleo showing the positions of neurons and electrode contacts. (C) Randomly generated spikes for the neurons shown in B. Top: the sorted spike signal, which gives the ground truth source neuron for every spike as a perfect proxy for spike sorting. Bottom: multi-unit spikes, where spikes are reported on every channel they are detected on, regardless of the source neuron. (D) The TKLFP signal generated from the spikes in C for each channel. Y-axis units are not shown.

In order to approximate cortical LFP without recurring to morphological neurons and biophysical forward modeling, we implemented the kernel LFP approximation from Teleńczuk *et al.* [114], which we term TKLFP (Teleńczuk kernel LFP). This method approximates the per-spike contribution to LFP (termed uLFP: unitary LFP) with a delayed Gaussian

function, where amplitude and delay depend on the position of the neuron relative to the electrode. This implementation is available as a standalone Python package on PyPI [115]. Accuracy of this implementation is verified in automated test suites in both TKLFP and Cleo packages.

3.2.4 Optogenetic stimulation

Cleo simulates optogenetic stimulation by combining a model of light propagation with an opsin model relating light to current. The light model is based on Kubelka-Munk light propagation, operating on the assumption that the medium is optically homogeneous and that particles are larger than the light wavelength [116, 117]. Cleo includes absorbance, scattering, and refraction parameters for 473-nm (blue) light as given in Foutz *et al.* [116].

Independent of the light propagation model, Cleo provides two different opsin models. One is a four-state Markov model as presented in [118], which captures rise, peak, plateau, and fall dynamics as the opsin is activated and deactivated through a Markov process. Additionally, by defining conductance rather than current directly, the model is able to reproduce the photocurrent's dependence on the membrane potential (see Figure 3.3). Because the Markov model depends on somewhat realistic membrane potential and resistance values, however, it is not well suited for many pre-existing models that do not. To that end, Cleo also provides a simplified model that delivers current (of arbitrary units) proportional to light at the neuron. Users can also specify both the probability that cells of a target population express an opsin and the per-cell expression level.

3.3 Results

3.3.1 Model validation

To verify that light and opsin models behaved as expected, we reproduced figures illustrating the optic fiber light transmission profile found in Figure 2a,b of Foutz *et al.* [116] and the

photocurrent in response to the ramping light intensity from Figure 4c of Evans *et al.* [118]. Light intensity-firing rate relationship was qualitatively similar to that originally reported, though differing in some respects due to the use of point neurons rather than morphological neurons. See Figure 3.3 for details. Additionally, preliminary experiments show that the simplified, proportional current opsin model is able to produce a firing response qualitatively similar to that of the Markov model.

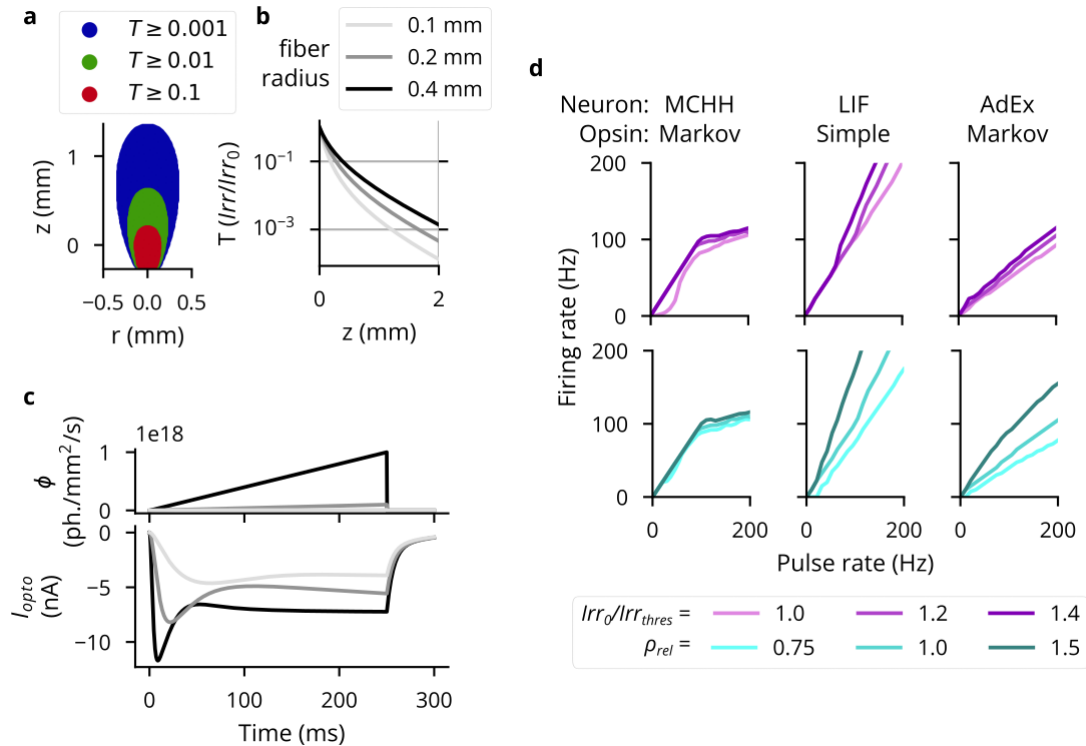


Figure 3.3: Validation of the optogenetics module. (A) Light transmittance T as a function of radius and axial distance from the optic fiber tip. Transmittance refers to the irradiance Irr as a proportion of the irradiance at the fiber tip Irr₀. After Figure 2a from Foutz *et al.* [116]. (B) Light transmittance T as a function of distance z straight out from the fiber tip for different optic fiber sizes. After Figure 2b from Foutz *et al.* [116]. (C) Photocurrent I_{opto} for ramping light of different intensities. After the figure produced by the “ramp” protocol from the default PyRhO simulator Evans *et al.* [118]. (D) Neuron firing rates in response to optical stimulation with 5-ms pulse frequencies ranging from 1 to 200 Hz. The left column re-plots data from Foutz *et al.* [116]. The middle column shows results for an LIF neuron with a simple opsin, and the right column for a tonic AdEx neuron [81] with a Markov opsin model. The top row shows results for different light intensities: 100%, 120%, and 140% of the threshold for producing a single spike with a 5-ms pulse. The bottom row shows results for different expression levels relative to the default, ρ_{rel} . The irradiance used for these simulations was 120% of the single-spike threshold.

3.3.2 Example experiments

In order to demonstrate Cleo’s utility to the public, we implemented three example experiments to feature in the upcoming publication:

1. Closed-loop inhibition of a traveling wave in a rodent somatosensory cortex model [90]. This experiment features a simple on/off control scheme where light activating an inhibitory opsin is turned on whenever the number of spikes detected at the optrode exceeds a threshold. The result is that while a wave of activity spreads outward through sensory cortex in response to a stimulus, the area around the optrode remains unperturbed. An experiment of this sort could explore the consequences of traveling waves in sensory cortex on perception.
2. Feedback control of layer 2/3 interneurons, disrupting plasticity in a model of primary visual cortex [119]. In this model, a top-down reward signal to vasoactive intestinal peptide (VIP) interneurons inhibits a population of parvalbumin (PV) interneurons, in turn disinhibiting a population of pyramidal cells (PC). However, using a proportional-integral (PI) feedback controller, we clamp the firing rate of interneurons, blocking this disinhibition. Without increased PC firing, synaptic weights corresponding to the rewarded stimulus no longer increase, thus demonstrating how feedback control can disrupt normal circuit function.
3. Optogenetic evocation of sharp wave-ripples in an anatomically detailed model of hippocampus [120, 121]. This uses linear quadratic regulator control to evoke a reference sharp wave-ripple LFP signature. See [Figure 3.4](#) for details, and note that feedback control, not looking ahead, so to speak, fails to evoke the reference signal at the desired time. The strategy I propose in Aim 2 should remedy this.

Additionally, we have endeavored to make Cleo accessible to the public to aid in the exploration and adoption of CLOC methods. Cleo is open-source and can be installed from the Python Package Index under the name `cleosim`. The code can be found on

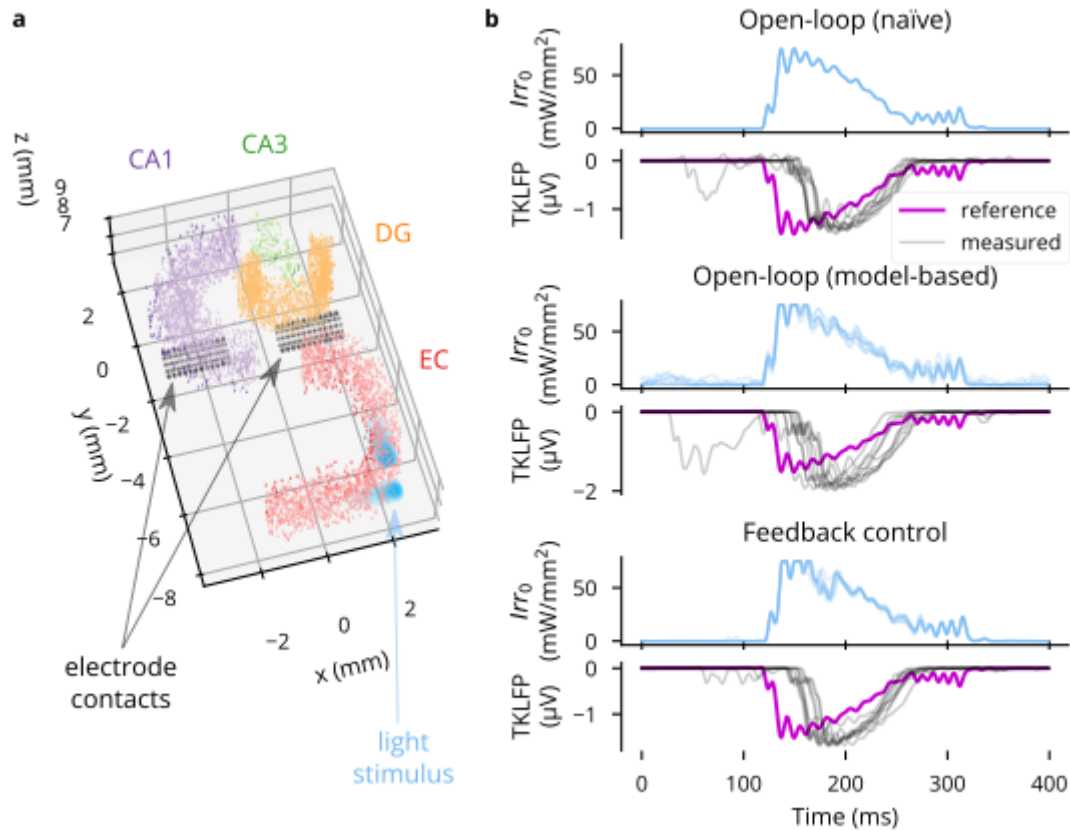


Figure 3.4: An example application of Cleo to the anatomical hippocampus model from Aussel et al. [120, 121]. (A) A 2.5 mm-thick slice of the 15 mm-tall model is shown. The model consists of four regions, entorhinal cortex (EC), dentate gyrus (DG), CA3, and CA1. Electrode contacts are represented as black dots and are in the same location as in the original model. Two light sources are shown in EC. Nine other such pairs (for a total of 20 light sources) not pictured here were spaced regularly parallel to the z axis. (B) Results are shown for ten trials each of three stimulation paradigms: naïve open-loop, where the input is simply a mirrored, rectified version of the reference signal; model-based open-loop, where a controller is run offline on a simulated model; and feedback control, where a controller is run using online measurements. Input Irr_0 is the light intensity at the tip of each optic fiber. The system output TKLFP refers to the Teleńczuk kernel LFP approximation.

[GitHub](#). Documentation, including an overview, tutorials, and API reference, can be found at <https://cleosim.readthedocs.io>.

3.4 Limitations

Because we (and Brian) prioritize point neuron models, Cleo does not currently support advanced, morphology-dependent features of neural dynamics as well as recording and stimulation. This would preclude realistically simulating extracellular potentials as well as the effects of spike sorting, for example. However, if features such as these are needed, Cleo could be developed further to integrate with Brian’s morphological neuron features. And in the case of LFP, at least one other method exists for point-neuron LFP approximation which could be implemented to help compensate for the limitations of TKLFP [122].

Both a strength and a limitation of Cleo’s design is that it depends on whatever model the user provides. This avoids the pitfall of painstakingly developing stock models that may or may not prove useful to researchers and instead lets them identify or develop a model that adequately describes the phenomena being studied. If a sufficiently realistic model for the studied system does not exist, however, developing one may be prohibitively costly, becoming a computational project in its own right as opposed to simply a stepping-stone towards an experiment. In these cases we suggest that a workaround could be to test a variety of potential models to identify which experimental configurations would be robust to unknown properties of the system. Indeed, the desired experiment in this case could be one that best adjudicates between these hypotheses.

3.5 Summary of contributions

Cleo thus provides a simulation environment facilitating and accelerating the adoption and development of CLOC methods. Multiple features were designed to reflect this need, for example, neural simulation and experimental configuration at an intermediate level

of realism, relevant recording and stimulation modalities, modular software architecture, flexible interfacing that can be easily adapted to existing Brian models, and arbitrary closed-loop processing enabling realistic processing delays. Moreover, we have developed virtual experiments to serve as examples as well as thorough online documentation so Cleo can help experimental labs test and potentially adopt CLOC methods.

CHAPTER 4

AIM 2 - MULTI-INPUT CLOC

4.1 Rationale

The power of closed-loop optogenetic control (CLOC, henceforth referring specifically to feedback control; see [section 2.1](#)) is limited by the degrees of freedom provided by the optogenetic stimulation. Naturally, we would want at least as many actuators as degrees of freedom in the system to control it effectively—for example, we may want to stimulate different layers, cell types, or columns separately in the cortex. Moreover, actuation/stimulation can be unidirectional or bidirectional, referring to whether a single opsin type or both excitatory and inhibitory opsins are used simultaneously. Unidirectional control has obvious shortcomings: for example, an excitatory opsin alone could not lower the firing rate of transfected cells, making it unsuitable to clamp baseline activity or follow a time-varying reference with steep drops.

While a previous study [34] has implemented bidirectional CLOC, it does not feature the generalizability and scalability of the model-based, optimal control algorithms introduced by later work [36] (see [section 2.3](#)) for unidirectional actuation. This adaptive linear-quadratic regulator (LQR) approach is more robust to disturbances and can scale to multi-input multi-output (MIMO) systems. Moreover, its behavior can be easily configured by setting penalties on state error, the control signal, and even the derivative of the control signal to encourage smooth actuation.

Thus, a natural goal for furthering CLOC is to combine the advantages of multi-input/bidirectional actuation and model-based optimal control—however, this poses

additional challenges and opportunities. The adaptive LQR method previously developed has limited application for multi-input actuation because it does not model the constraint that the input (light intensity) must be nonnegative. This would cause problems in the case of bidirectional (or any spatially overlapping) actuation, since the controller might call for negative excitatory input rather than positive excitatory input. A heuristic workaround to this is to place light sources in spatially non-overlapping pairs (e.g., blue and red) and treat them as the positive and negative directions of a single actuator. This allows the continued use of simple and fast LQR methods but fails to model kinetic differences or spectral crosstalk between inhibitory and excitatory opsins and precludes alternate light configurations. While I hypothesize that methods for optimal control with constraints will outperform this heuristic LQR, such methods are more costly computationally. It is thus unclear which method(s) are preferable for real-time control on the timescale of network-level variables of neural activity using compute resources typically available to an experimental neuroscience lab.

4.1.1 Innovation

I propose addressing this problem by comparing LQR to model predictive control (MPC), which is widely used for its flexibility in implementing optimal control with constraints. Rather than computing the control signal from the current error signal at each step, MPC looks ahead, optimizing over the predicted trajectory some finite number of steps into the future, in what is known as “receding horizon control.” The quadratic program optimization required at every control step, however, introduces latency which could harm control performance compared to LQR. I thus plan on testing my hypothesis that MPC will be able to better optimize multi-input optogenetic stimulation while accommodating experimental constraints and considerations during real-time control on timescales relevant to network-level descriptions of neural activity. I will do this by assessing control quality of MPC compared to the heuristic LQR approach previously described as I simulate multi-output feedback control of firing rates and gamma oscillations observed in local field potentials

(LFP).




Figure 4.1: An illustration of how MPC optimizes the system input over a receding horizon. By [Martin Behrendt](#), licensed under [CC BY-SA 3.0](#).

4.2 Approach

4.2.1 System and controller formulation

Naturally, the model is a vital element of MPC. I will use a previously developed Gaussian linear dynamical system (GLDS) model [36], which has been shown to reliably capture firing rate dynamics in a light-driven spiking neuron system. The discrete-time GLDS is governed by the following equations:

$$x_{t+1} = Ax_t + Bu_t + w_t$$

$$y_t = Cx_t + d$$

$$z_t = y_t + v_t$$

where $x_t \in \mathbb{R}^n$ is the n -dimensional state, $u_t \in \mathbb{R}^k$ is the k -dimensional stimulus (i.e., $k = 2$ for two opsins, one light source each), $y_t \in \mathbb{R}^m$ is the firing rate in spikes/timestep (for each of m measured neurons), and $z_t \in \mathbb{R}^m$ is the number of binned spikes observed at time t . $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times k}$, and $C \in \mathbb{R}^{m \times n}$ are the state transition, input, and output matrices, respectively.

$w_t \sim \mathcal{N}(0, Q)$ and $v_t \sim \mathcal{N}(0, R)$ are Gaussian-distributed process and measurement noise, respectively, and $d \in \mathbb{R}^{m \times 1}$ represents baseline firing rates. Model order (n) and horizon length ($T \in \mathbb{N}$) will be chosen to balance complexity and prediction error for noise-driven fitting data generated from the test network. The latent state x_t will be estimated online using the Kalman filter [65], driven by the prediction error $z_t - \hat{y}_t|_{t-1}$.

I will set hard non-negative constraints on the light input as well as a ceiling determined by hardware limitations (i.e., the maximum voltage deliverable to the LED driver). To design an appropriate cost function, I will use a conventional per-time step quadratic form

$$\ell(x_t, r_t, u_t) = (x_t - r_t)^T Q^{\text{ctrl}}(x_t - r_t) + u^T R^{\text{ctrl}} u,$$

where $r_t \in \mathbb{R}^n$ is the reference trajectory at time t . Q^{ctrl} and R^{ctrl} are real $n \times n$ and $k \times k$ matrices chosen to appropriately penalize tracking error and stimulus size, respectively. This quadratic cost function formulation lends the problem well to standard optimization techniques—combined with a linear dynamical system, it constitutes the classical linear-quadratic-Gaussian (LQG) control problem.

Then, at every time step t the controller solves the following quadratic program:

$$\begin{aligned} & \text{minimize} && \sum_{\tau=t}^{t+T} \ell(x_\tau, u_\tau) \\ & \text{subject to} && u_\tau \succeq 0 \\ & && x_{\tau+1} = Ax_\tau + Bu_\tau \end{aligned}$$

where $T \in \mathbb{N}$ is the number of steps in the prediction/control horizon and \succeq indicates an inequality for each element of $u_\tau \in \mathbb{R}^k$. This yields the solution $\tilde{u}_t, \dots, \tilde{u}_{t+T-1}$, of which we take just the first step \tilde{u}_t to apply to the system.

4.2.2 Control method comparison *in silico*

To confirm our assumption that bidirectional and multi-channel configurations will improve control quality, I will test both unidirectional and bidirectional actuation as well as one-channel and multi-channel configurations for each experiment. (In the bidirectional actuation case, one “channel” includes both an inhibition-triggering and an excitation-triggering light source). Another condition of interest will be to penalize a low-pass-filtered version of the stimulus, which could reflect overheating or ion imbalances caused by prolonged stimulation [123, 124, 125]. This could be added to the linear dynamics and quadratic cost functions without changing the optimization methods, but, as with opsin and channel count, the increased size of the problem increases latency which could affect which method performs best. Most importantly, I will compare open-loop, heuristic LQR, and MPC approaches to see whether MPC attains better performance despite a longer computational delay, as hypothesized. Control algorithm computation time will be measured for each method and used during simulations as a realistic delay. To evaluate controller performance, I will use metrics such as the mean-squared error (MSE).

In the first experiment, I will simulate multi-output control of firing rates in an attempt to clamp population activity. Wagenaar *et al.* [63] and Newman *et al.* [64] laid the foundation for this by clamping aggregate, population firing rates, and Bolus *et al.* [36] took this further by treating firing rates of individual neurons separately, though with a single optic fiber input. This is an obvious case where multiple inputs should provide more control, allowing us to manipulate neurons (or groups of neurons in the case of unsorted threshold crossings) more individually, each to its own baseline. The experiment will be performed on a Poisson linear dynamical system (PLDS) model [126] fit to optogenetic input/spike output data from a spiking neural network (SNN) model. The SNN model will contain necessary features, such as cell types, time-varying exogenous input, and connectivity profiles, to produce stochastic firing patterns with unpredictable disturbances and will be simulated together

with the recording and stimulation facilities of Aim 1.

As another experiment, I propose another form of network-level control—manipulating the oscillatory signatures found in LFP signals. Gamma oscillations (30-90 Hz) are one such example, which have been shown to have time-varying properties and interact with other oscillatory bands. Interestingly, gamma rhythms can even show phase coherence in distant brain regions such as visual cortex in opposite hemispheres, which has been hypothesized to play a role in information integration [127]. I plan to test control methods on a virtual experiment of this phenomenon by simulating two populations of neurons exhibiting gamma oscillations [128, 129] with long-range connections mediating gamma frequency coherence, again using the framework from Aim 1. I will then perform feedback control on the LFP signal to counter the natural phase locking that arises, providing an opportunity to test how well LFP can be controlled on a <100-ms timescale of a single gamma cycle.

4.3 Expected results

I expect that bidirectional, multi-channel, and model predictive control will perform better than unidirectional, single-channel, and LQR control, respectively for both experiments, despite the added computational cost these methods require. I also expect that restricting the controller through hard constraints or prolonged stimulation penalties will also be possible with the more sophisticated methods without greatly increasing the error.

4.4 Preliminary results

Basic simulations controlling a linear dynamical system model fit to experimental data show the advantages of bidirectional control and of MPC (see [Figure 4.2](#)). Bidirectional actuation allows the system to avoid overshooting the reference, in the case of LQR, or to minimize error faster by first exciting then inhibiting, in the case of MPC. MPC's advantages in looking ahead also clearly allow it to follow the reference more closely than the heuristic

LQR controller (assigning negative inputs to the second light source). However, these results do not yet account for control signal computation latency.

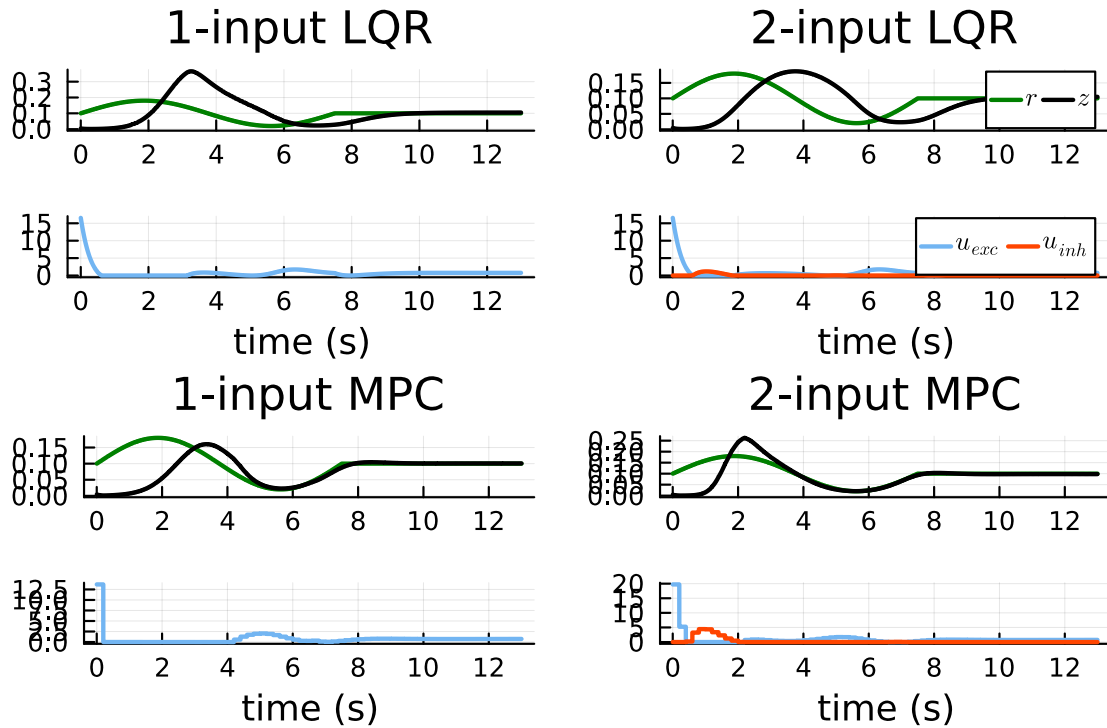


Figure 4.2: Simulated control of an linear dynamical system with 1- and 2-input control, using LQR and MPC controllers. The top panel of each contains the reference and the actual firing rate, in spikes/second. The bottom contains the light intensity, in terms of mW/mm^2 , where blue represents light for an excitatory opsin (such as Chr2) and red-orange that for an inhibitory opsin (such as Jaws).

4.5 Potential pitfalls & alternative strategies

There are some limitations in the proposed GLDS model that may need to be addressed. While it was adequate for the experiments in Bolus *et al.* [36], a Poisson [126] or other output nonlinearity may be needed in the case that a standard GLDS does not fit the data well. While this would likely make the estimation of x more expensive, the underlying dynamics could remain linear, leaving the same underlying quadratic program for the controller to solve.

This touches another potential concern: the speed of MPC. If the optimization problem

solution is slow, there are a few options to explore. One is that some variations in the control scheme can help balance speed and performance, such as letting the control horizon be shorter than the prediction horizon, which shrinks the optimization problem. Likewise, the control period can be longer than the time step of the system, reducing how often the control signal is computed. If conventional methods such as these are unsuccessful, I may turn to approximate methods such as that described in Wang and Boyd [130] or training an artificial neural network. Another potential solution is explicit MPC [131], which finds a piecewise-affine explicit solution to the quadratic program which can be faster than obtaining the implicit solution for small-enough problems.

CHAPTER 5

AIM 3 - CONTROL OF LATENT POPULATION DYNAMICS

5.1 Rationale

As technology for recording from the brain has improved, systems neuroscience has shifted increasingly towards a reduced-dimensionality, population coding perspective in many brain areas [132, 92, 94, 133]. This reflects the observation that while the activity of any single neuron can vary greatly across trials where external variables are controlled, underlying latent variables can be decoded from the population which are much more predictable and reproducible. Moreover, formulating this latent variable as a dynamical state whose evolution can be predicted has been shown to improve inference and enable state-of-the-art “decoding” of downstream variables of interest such as movement [134, 135, 20, 100], mood [136], speech [137], and decision states [138, 139]. Not only does this allow us to infer the output of a given brain region, but it allows us to form hypotheses about how it produces that output by analyzing the dynamical landscape of fixed points [140, 141, 142]—a view formalized in the Computation through Dynamics (CTD) framework [97].

However, while these latent variables have been used to successfully decode other variables of interest, this is a necessary, but insufficient, condition to demonstrate a causal relationship. That is, an association between neural activity variable a and some other variable b may reflect the causal relationship $a \rightarrow b$, but could also reflect $b \rightarrow a$ or $a \leftarrow c \rightarrow b$. This may be adequate for brain-computer interface (BCI) applications, but verification of that causal relationship is necessary for neuroscience’s goal of deepening our understanding of the architecture and algorithms of brain computations. This requires experimental control [143]

on the level of neural populations, but, as stated by Vyas *et al.* [97],

The challenge is nontrivial; testing CTD models requires a high degree of control over neural activity. The experimenter must be able to manipulate neural states arbitrarily in state space and then measure behavioral and neural consequences of such perturbations[.]

CLOC is a natural candidate for this kind of experimental control of latent states for various reasons. The optimal state-space methods already formulated map directly to the latent dynamics models which have generated so much interest in recent systems neuroscience. High-dimensional optogenetic actuation is possible through micro-LED devices [144, 147, 148, 149, 150, 151, 152, 153, 154, 145, 146, 8], two-photon targeting of individual neurons [155, 156, 157, 37, 3], and genetic targeting. Moreover, real-time feedback can drive a variable neural system towards complex latent state targets which would be attainable with low accuracy at best and not at all at worst with open-loop stimulation ¹. However, while it is clear a high degree of control will be needed, it is unknown how this translates to recording, stimulation, and control parameters. One significant open question is whether this kind of control will require actuation on the level of individual neurons, or whether lower-dimensional actuation will suffice. I hypothesize that the latter is the case, but testing this hypothesis will likely require extensive trial and error, as well actuation hardware that does not exist or is not readily available, making it costly if not infeasible to do *in vivo*.

5.1.1 Innovation

Thus, to give a preliminary answer to this question, I propose to develop technical and conceptual guidelines as I control the latent dynamics of simulated neural populations. First, I will produce virtual models by training recurrent spiking neural networks with state-of-the-art, biologically plausible methods—each differing in their degrees of brain-like

¹Low accuracy because open-loop stimulation would not counteract spontaneous variability. Inevitable model mismatch would result in open-loop stimulation potentially entirely missing the target state.

architecture and training procedure complexity. I will then use the simulation testbed of Aim 1 and the multi-input control methods of Aim 2 to explore how control quality varies with both experimental parameters (such as recording and stimulation channel counts or control algorithms) and system characteristics (such as the size, complexity of the network model). This will both give researchers a tentative idea of the relative importance of each factor of CLOC and, as parameters approach experimental realism, allow us to conjecture as to the number of independent actuators needed to control the latent dynamics of a brain system. My hypothesis is that control quality will plateau with a number of actuators on or near the same level of magnitude as the latent factors used to describe the neural activity—that per-neuron actuation limit not be needed.

5.2 Approach

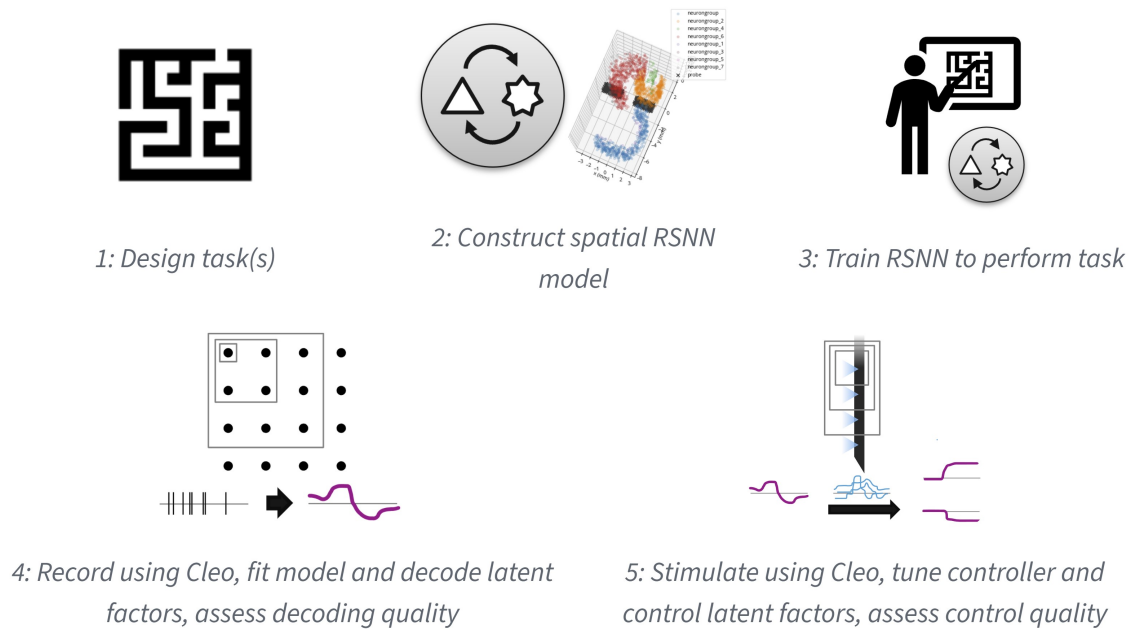


Figure 5.1: Overview of latent factor control experiments.

5.2.1 Formulate task with latent dynamics hypothesis

To model an experiment of interest to neuroscientists, I will first choose a task where latent dynamic analyses have yielded testable hypotheses in animal research. For example, it is

hypothesized that motor cortex works as a dynamical system where movement planning essentially sets an initial condition during a preparatory phase. This is demonstrated specifically by decoding reach direction in a delayed reach task from dorsal premotor cortex before movement onset [158, reviewed in 96] (see Figure 5.2). One way to causally test this hypothesis would be to manipulate the latent factors corresponding to a certain reach direction despite an absent or contradictory cue and verify whether the subject reaches in the predicted direction.




Figure 5.2: Illustration of a delayed center-out reach experiment. From Santhanam *et al.* [159].

5.2.2 Train RSNN models

After identifying a task and latent dynamics hypothesis to test, I will train recurrent spiking neural network (RSNN) models to perform the task. These RSNN models will be defined in 3D space to make them compatible with the electrode and optogenetics models from Aim 1. To avoid painstaking implementation of every known anatomical detail of the brain region(s) involved in the task—an approach which is not guaranteed to include every important detail or capture true circuitry—I propose using a simple, abstract form such as a rectangular prism of cortical tissue and measuring the effect of adding progressively more brain-like structural constraints (see subsection 5.2.6).

I will then train these models to perform the given task. While a few different methods for training RSNNs exist, I plan on using e-prop [160], a biologically plausible approximation to

backpropagation through time (BPTT), the method typically used to train artificial recurrent neural networks (RNNs). This biological plausibility consists in updating synaptic weights using only local information about past activity—“eligibility traces”—and a top-down learning signal and enables learning sparsely firing spike-coding—as opposed to only rate-coding—solutions. While e-prop has been shown to learn complex tasks such as Atari games via reinforcement learning, I propose a simpler supervised learning approach. This is both more data-efficient and would directly allow for the learning of auxiliary variables, which has been shown in one study to make model dynamics more brain-like [161] (see [section 5.4](#)).

To give a concrete example of inputs and outputs for the case of the delayed center-out reaching task before mentioned, inputs might include the task phase (stop or go), the target position, and the current hand position. The model output might be hand velocity, and the learning signal could be computed from the x/y distance between the current position and the center (target) before (after) the go cue is given. A regularizer term on acceleration could be added to ensure trajectories are smooth.

5.2.3 Fit dynamical systems model

After a model has been trained to some threshold performance level, I will record spikes and/or LFP from the network while it performs the task. The resulting data will then be used to fit a low-dimensional dynamical systems model. Seeing that the goal of the virtual experiment is to test the causal effect of a latent factor on model output, I propose using system identification methods that prioritize the discovery of those latent factors that are most relevant to behavior. The preferential subspace identification (PSID) method presented by Sani *et al.* [20] is a good candidate, having been shown to predict behavior well with few dimensions. The linear system it produces is also ideal for the control theory methods I develop in Aim 2. In a follow-up study, Sani *et al.* [100] introduce RNN-PSID supporting nonlinearities as well and find empirically that linear dynamics with only a nonlinear

mapping from latent to behavior are often sufficient to describe experimental data well. Using a potentially nonlinear output mapping in this way would increase the expressiveness of my models while maintaining underlying linear dynamics for fast optimization.

After fitting a model to data, I will identify latent factors corresponding to variables of interest. These could be individual components of the latent state x directly or simple (e.g., linear) transformations of x . In the delayed center-out reaching task, for example, I would expect two factors with which the state at the end of the preparatory period can be used to predict the cued target position.

5.2.4 Control of latent factors

The next step is then to test whether these latent factors do in fact cause the behavior we observe. Importantly, the model and hypothesis from the previous step were formed without stimulation, since our stated goal is to causally test the latent factors we deduce *from passive observation*. Thus, I will first need to expand the model by simulating random optogenetic stimulation and fitting an input model—e.g., an input matrix B for the typical LDS case. If this is insufficient to fit the data well, I may need to increase the dimensionality of the latent state x to account for the actuator state (i.e., opsin kinetics) or dimensions of neural activity not arising during unperturbed activity.

With an input model, I will then be able to run control experiments. The methods from Aim 2 will calculate optimal actuation strategies as the simulation progresses to drive the identified latent factors to the desired state. Control quality will be assessed for each experiment using metrics such as mean-squared error (MSE) and will serve as the criterion for how successful the experiment was, seeing that low error serves the larger goal of testing the causal relationship with model behavior.

5.2.5 Exploration of experiment parameters and expected effects

To guide experimental design, I propose repeating the above process many times to explore how different factors affect control quality, thus informing which investments might be most fruitful. These factors might include:

- *Control type.* Open-loop control is faster and easier but cannot counteract variability inherent to the system and in the presence of model mismatch will fail to reach the target. Closed-loop methods will almost certainly perform better, but are harder to implement. Moreover, closed-loop control offers a range of options spanning the cost/quality spectrum, with the fast linear quadratic regulator (LQR) on one end and the expensive long-horizon, high-resolution model predictive control (MPC) on the other.
- *Total recording/stimulation channel count ($m + k$).* It may be helpful to think of these two parameters together, since they both require limited space in a device that is placed in or on the brain. I expect control quality to increase with channel count up to a point where the computational cost required for closed-loop control outweighs the benefits.
- *Recording/stimulation channel ratio (m/k).* Related to the previous point, when space for interface is hardware, it is an open question whether the current landscape of neuroscience technology, which has prioritized high-channel-count recording over stimulation, is the most effective for causal perturbations of neural dynamics.
- *Data collection budget.* Data collection is not free, and thus must be considered especially when considering the data-efficiency advantages of feedback control. This is also relevant for choosing models to control the system—with little training data, for example, a simple, easy-to-fit model may outperform a more expressive, data-hungry one.
- *Optogenetic stimulation parameters.* These could include the number and type of

opsins and genetic targeting. Using at least two opsins could be expected to perform better than just one, as shown in Aim 2, though many more might not be useful due to overlapping spectral sensitivity and increased computation time. The effect of genetic targeting will likely depend on how well the targeted population correlates with the latent factor of interest.

5.2.6 Exploration of model realism and expected effects

As explained above, I propose using abstract RSNN models with varying degrees of brain-like realism rather than extremely detailed models. This is because highly detailed models are difficult to develop and expensive to simulate—moreover, it is hard to say exactly when a model has enough detail to behave similarly to a real brain. On the other hand, by assessing the effect of a handful of brain-like features on how well we can control latent neural dynamics *in silico*, I hope to identify trends that give a rough outlook of doing so *in vivo*. For example, if control quality tends to drop with the addition of each feature, we might infer control of a real brain will be considerably more challenging than in simulations. On the other hand, if control quality does not decrease considerably with brain similarity, we might have reason to be more optimistic.

The brain-like features I propose assessing include structural characteristics such as:

- *Variability*. Brains are far from deterministic, with unpredictability on the level of synaptic transmission all the way up to firing rates across a population. Small-scale stochasticity encourages redundancy and robustness, which may make the model harder to “hack,” or perturb unnaturally. I expect that large-scale variability, such as an unmodeled input to the region, will highlight the advantages of feedback control.
- *Network size (how many neurons are in the model)*. More neurons would likely mean more dimensions along which activity can vary, complicating the prospect of arbitrary control.

- *Subpopulation structure.* This could include constraining connectivity to circuits of brain regions, cortical layers, and cell types. This could also make control harder if the latent factor to control operates in cells that are multiple synapses away from those stimulated or if dynamics become driven by recurrent connections more than stimulation.
- *Connectivity profiles.* Cell pairs within a given population might have a uniform or a spatially dependent connection probability. I expect this, like subpopulation structure, to encourage segregated dynamics which could be either easier or harder to perturb.
- *Cell diversity.* Cells might have uniform or a random distribution of parameters such as membrane resistance, synaptic transmission delays, resting potentials, adaptivity, etc. Again, I expect this would encourage the segregation of neural dynamics.
- *Neuron model complexity.* Individual cell models could range from simple leaky integrate-and-fire (LIF) to Hodgkin-Huxley dynamics. I do not expect this to considerably impact control quality.

They also include characteristics about how the model is trained. A real brain regions is capable of performing not only one, simple, stereotyped task, but many tasks in many contexts. Therefore, another way to make models more brain-like is to diversify their experience—to train them on other tasks. It has been shown that this can result in an RNN reusing computational motifs learned in one task to solve others [162, 163].

5.3 Expected results

In addition to the specific expectations for experimental and model parameters described above, I expect that control quality will be relatively low (error will be high) especially for low actuator counts k . However, as previously stated, I do not anticipate that k will need to approach the per-neuron limit either; rather, that low-error control will be possible with k close to the same order of magnitude as the latent state dimensionality. While

expectations for each brain-like model feature are explained above, I predict that on the whole, control quality will decrease as models become more realistic. However, I predict that that that decrease will plateau, providing some confidence in my results and assurance that the complexity of the real brain may not be an insurmountable barrier to future *in-vivo* experiments.

5.4 Potential pitfalls & alternative strategies

One potential problem is that the RSNN models could learn dynamics that are not very brain-like. For example, they might learn a complicated feedforward function instead of tracking latent variables in a more natural way, in which case explicitly teaching the models to track variables of interest has been shown to make dynamics more brain-like [161]. Alternatively, the latent variables decoded could be dominated by very few neurons, thus producing “latent variables” which are not actually latent and distributed, as in the brain. This can be avoided by regularizing the state inference step to discourage sparsity, thus inferring the state from a more diverse set of neurons.

Another caveat to generalizing conclusions is that the proposed perturbations may be unnatural—by perturbing the latent factors decoded from passive observation, I may unwittingly be pushing neurons “off-manifold,” causing them to fire in unnatural combinations [98]. While it is beyond the scope of the proposed work to avoid this, I can at least quantify the phenomenon via the proxy of neural activity *unexplained* by our latent dynamical, assuming our intrinsic manifold is a subspace rather than a lower-dimensional attractor structure embedded therein [97, 164]. For example, if under passive observation we can predict 95% of neural activity, and that drops to 50% under perturbation, we could conclude that 45% of the activity is now unnatural, as it must be explained by dimensions of activity introduced by stimulation.


Besides conceptual obstacles, the proposed work also has practical challenges, such as its

large scale. To adequately assess control prospects across the whole space of experiment and structural factors described could represent a prohibitively high computational cost. This can be mitigated by avoiding a grid search of this space, analyzing random combinations of features or just one at a time rather than every single combination. Another factor in the computational cost is the speed of the training, model fitting, and control simulations themselves. In the case that simulations are too slow with the Brian/Cleo framework described in Aim 1, it may be worth exploring simpler simulations for the training step which does not require Cleo. This could involve using fast differential equation solvers [165] for the training step or even a less biological learning algorithm such as backpropagation through time with surrogate gradients for the nondifferentiable spike threshold function [166]. Training might also be accelerated using a form of “warm start” where weights are initialized from a pre-trained model, though care would need to be taken that the effects of the warm start on learned dynamics not overshadow those of the structural features being tested.

Additionally, there is a possibility that RNN-PSID requires nonlinear terms to fit model data well, despite the authors’ findings that linear dynamics are often sufficient [100]. In the case that nonlinearities preclude the formulation of MPC as a fast quadratic program, I may need to turn to neural networks, which, after training, could quickly produce approximate solutions for nonlinear optimization problems (see [section 4.5](#)).

CHAPTER 6

PROPOSED TIMELINE



REFERENCES

- [1] L. Fenno, O. Yizhar, and K. Deisseroth, “The development and application of optogenetics,” *Annual Review of Neuroscience*, vol. 34, pp. 389–412, Jul. 2011.
- [2] J. S. Wiegert, M. Mahn, M. Prigge, Y. Printz, and O. Yizhar, “Silencing Neurons: Tools, Applications, and Experimental Constraints,” *Neuron*, vol. 95, no. 3, pp. 504–529, 2017.
- [3] S. Sridharan *et al.*, “High-performance microbial opsins for spatially and temporally precise perturbations of large neuronal networks,” *Neuron*, vol. 110, no. 7, 1139–1155.e6, Feb. 2022.
- [4] J. Vierock *et al.*, “BiPOLES is an optogenetic tool developed for bidirectional dual-color control of neurons,” *Nature Communications*, vol. 12, no. 1, pp. 1–20, Jul. 2021.
- [5] H. Adesnik and L. Abdeladim, “Probing neural codes with two-photon holographic optogenetics,” *Nature Neuroscience*, vol. 24, no. 10, pp. 1356–1366, Oct. 2021.
- [6] G. Faini *et al.*, “Ultrafast Light Targeting for High-Throughput Precise Control of Neuronal Networks,” *bioRxiv*, p. 2021.06.14.448315, Jun. 2021.
- [7] B. L. Roth, “DREADDs for Neuroscientists,” *Neuron*, vol. 89, no. 4, pp. 683–694, Feb. 2016.
- [8] D. Eriksson *et al.*, “Multichannel optogenetics combined with laminar recordings for ultra-controlled neuronal interrogation,” *Nature Communications*, vol. 13, no. 1, p. 985, Feb. 2022.
- [9] N. A. Steinmetz *et al.*, “Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings,” *Science*, vol. 372, no. 6539, Apr. 2021.
- [10] J. H. Siegle, A. C. López, Y. A. Patel, K. Abramov, S. Ohayon, and J. Voigts, “Open Ephys: An open-source, plugin-based platform for multichannel electrophysiology,” *Journal of Neural Engineering*, vol. 14, no. 4, p. 045 003, Jun. 2017.
- [11] P. Gutruf and J. A. Rogers, “Implantable, wireless device platforms for neuroscience research,” *Current Opinion in Neurobiology*, Neurotechnologies, vol. 50, pp. 42–49, Jun. 2018.

- [12] W. Göbel and F. Helmchen, “In Vivo Calcium Imaging of Neural Network Function,” *Physiology*, vol. 22, no. 6, pp. 358–365, Dec. 2007.
- [13] T. Knöpfel and C. Song, “Optical voltage imaging in neurons: Moving from technology development to practical tool,” *Nature Reviews Neuroscience*, vol. 20, no. 12, pp. 719–727, Dec. 2019.
- [14] K. Svoboda and R. Yasuda, “Principles of Two-Photon Excitation Microscopy and Its Applications to Neuroscience,” *Neuron*, vol. 50, no. 6, pp. 823–839, Jun. 2006.
- [15] A. Kazemipour *et al.*, “Kilohertz frame-rate two-photon tomography,” *Nature Methods*, vol. 16, no. 8, pp. 778–786, Jul. 2019.
- [16] J. Wu *et al.*, “Kilohertz two-photon fluorescence microscopy imaging of neural activity in vivo,” *Nature Methods*, vol. 17, no. 3, pp. 287–290, Mar. 2020.
- [17] L. van der Maaten and G. Hinton, “Visualizing Data using t-SNE,” *Journal of Machine Learning Research*, vol. 9, no. 86, pp. 2579–2605, 2008.
- [18] G. J. Berman, D. M. Choi, W. Bialek, and J. W. Shaevitz, “Mapping the stereotyped behaviour of freely moving fruit flies,” *Journal of The Royal Society Interface*, vol. 11, no. 99, p. 20140672, Oct. 2014.
- [19] A. Mathis *et al.*, “DeepLabCut: Markerless pose estimation of user-defined body parts with deep learning,” *Nature Neuroscience*, vol. 21, no. 9, pp. 1281–1289, Sep. 2018.
- [20] O. G. Sani, H. Abbaspourazad, Y. T. Wong, B. Pesaran, and M. M. Shanechi, “Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification,” *Nature Neuroscience*, vol. 24, no. 1, pp. 140–149, 2021.
- [21] O. Sporns, “Graph theory methods: Applications in brain networks,” *Dialogues in Clinical Neuroscience*, vol. 20, no. 2, pp. 111–121, Jun. 2018.
- [22] S. Schneider, J. H. Lee, and M. W. Mathis, *Learnable latent embeddings for joint behavioral and neural analysis*, Apr. 2022. arXiv: [2204.00673](https://arxiv.org/abs/2204.00673) [cs, q-bio].
- [23] L. K. Scheffer *et al.*, “A connectome and analysis of the adult *Drosophila* central brain,” *eLife*, vol. 9, E. Marder, M. B. Eisen, J. Pipkin, and C. Q. Doe, Eds., e57443, Sep. 2020.
- [24] A. L. Juavinett, G. Bekheet, and A. K. Churchland, “Chronically implanted Neuropixels probes enable high-yield recordings in freely moving mice,” *eLife*, vol. 8, L. L. Colgin and N. Steinmetz, Eds., e47188, Aug. 2019.

- [25] E. R. Oby *et al.*, “New neural activity patterns emerge with long-term learning,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 30, pp. 15 210–15 215, Jul. 2019.
- [26] Y. Yang *et al.*, “Modelling and prediction of the dynamic responses of large-scale brain networks during direct electrical stimulation,” *Nature Biomedical Engineering*, vol. 5, no. 4, pp. 324–345, Feb. 2021.
- [27] B. R. Cowley, A. C. Snyder, K. Acar, R. C. Williamson, B. M. Yu, and M. A. Smith, “Slow Drift of Neural Activity as a Signature of Impulsivity in Macaque Visual and Prefrontal Cortex,” *Neuron*, vol. 108, no. 3, 551–567.e8, Nov. 2020.
- [28] L. Avitan and C. Stringer, “Not so spontaneous: Multi-dimensional representations of behaviors and context in sensory areas,” *Neuron*, Jul. 2022.
- [29] M. Jazayeri and S. Ostojic, “Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity,” *Current Opinion in Neurobiology*, vol. 70, pp. 113–120, Oct. 2021. arXiv: [2107.04084](https://arxiv.org/abs/2107.04084).
- [30] L. Grosenick, J. H. Marshel, and K. Deisseroth, “Review Closed-Loop and Activity-Guided Optogenetic Control,” *Neuron*, vol. 86, pp. 106–139, 2015.
- [31] S. Dutta, E. Ackermann, and C. Kemere, “Analysis of an open source, closed-loop, realtime system for hippocampal sharp-wave ripple disruption,” *Journal of Neural Engineering*, vol. 16, no. 1, p. 016 009, Dec. 2019.
- [32] A. Kumar, I. Vlachos, A. Aertsen, and C. Boucsein, “Challenges of understanding brain function by selective modulation of neuronal subpopulations,” *Trends in Neurosciences*, vol. 36, no. 10, pp. 579–586, Oct. 2013.
- [33] S. M. Potter, A. El Hady, and E. E. Fetz, “Closed-loop neuroscience and neuroengineering,” *Frontiers in Neural Circuits*, vol. 0, p. 115, 2014.
- [34] J. P. Newman, M. F. Fong, D. C. Millard, C. J. Whitmire, G. B. Stanley, and S. M. Potter, “Optogenetic feedback control of neural activity,” *eLife*, 2015.
- [35] M. F. Bolus, A. A. Willats, C. J. Whitmire, C. J. Rozell, and G. B. Stanley, “Design strategies for dynamic closed-loop optogenetic neurocontrol in vivo,” *Journal of Neural Engineering*, vol. 15, no. 2, p. 026 011, Apr. 2018.
- [36] M. F. Bolus, A. A. Willats, C. J. Rozell, and G. B. Stanley, “State-space optimal feedback control of optogenetically driven neural activity,” *Journal of neural engineering*, vol. 18, no. 3, p. 036 006, Jun. 2021.

- [37] Z. Zhang, L. E. Russell, A. M. Packer, O. M. Gauld, and M. Häusser, “Closed-loop all-optical interrogation of neural circuits in vivo,” *Nature Methods*, vol. 15, no. 12, pp. 1037–1040, Dec. 2018.
- [38] E. Krook-Magnuson, C. Armstrong, M. Oijala, and I. Soltesz, “On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy,” *Nature Communications*, vol. 4, no. 1, pp. 1–8, Jan. 2013.
- [39] A. Witt, A. Palmigiano, A. Neef, A. El Hady, F. Wolf, and D. Battaglia, “Controlling the oscillation phase through precisely timed closed-loop optogenetic stimulation: A computational study,” *Frontiers in Neural Circuits*, vol. 7, no. April, pp. 1–17, 2013.
- [40] S. Tafazoli, C. J. MacDowell, Z. Che, K. C. Letai, C. R. Steinhardt, and T. J. Buschman, “Learning to control the brain through adaptive closed-loop patterned stimulation,” *Journal of Neural Engineering*, vol. 17, no. 5, p. 056007, Oct. 2020.
- [41] S. S. Srinivasan, B. E. Maimon, M. Diaz, H. Song, and H. M. Herr, “Closed-loop functional optogenetic stimulation,” *Nature Communications*, vol. 9, no. 1, pp. 1–10, Dec. 2018.
- [42] M. Prsa, G. L. Galiñanes, and D. Huber, “Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons,” *Neuron*, vol. 93, no. 4, 929–939.e6, Feb. 2017.
- [43] Z. W. Davis, L. Muller, J. Martinez-Trujillo, T. Sejnowski, and J. H. Reynolds, “Spontaneous travelling cortical waves gate perception in behaving primates,” *Nature*, vol. 587, no. 7834, pp. 432–436, Nov. 2020.
- [44] G. Buzsáki, “Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning,” *Hippocampus*, vol. 25, no. 10, pp. 1073–1188, Oct. 2015.
- [45] A. R. Vaidya, M. S. Pujara, M. Petrides, E. A. Murray, and L. K. Fellows, “Lesion Studies in Contemporary Neuroscience,” *Trends in Cognitive Sciences*, vol. 23, no. 8, pp. 653–671, Aug. 2019.
- [46] A. L. Hodgkin, A. F. Huxley, and B. Katz, “Measurement of current-voltage relations in the membrane of the giant axon of Loligo,” *The Journal of Physiology*, vol. 116, no. 4, pp. 424–448, Apr. 1952.
- [47] A. A. Sharp, M. B. O’Neil, L. F. Abbott, and E. Marder, “The dynamic clamp: Artificial conductances in biological neurons,” *Trends in Neurosciences*, vol. 16, no. 10, pp. 389–394, Oct. 1993.
- [48] A. A. Prinz, L. F. Abbott, and E. Marder, “The dynamic clamp comes of age,” *Trends in Neurosciences*, vol. 27, no. 4, pp. 218–224, Apr. 2004.

- [49] J. T. Davie *et al.*, “Dendritic patch-clamp recording,” *Nature Protocols*, vol. 1, no. 3, pp. 1235–1247, Sep. 2006.
- [50] G. Wang *et al.*, “An optogenetics- and imaging-assisted simultaneous multiple patch-clamp recording system for decoding complex neural circuits,” *Nature Protocols*, vol. 10, no. 3, pp. 397–412, Feb. 2015.
- [51] D. Engel, “Subcellular patch-clamp recordings from the somatodendritic domain of nigral dopamine neurons,” *Journal of Visualized Experiments*, vol. 2016, no. 117, e54601, Nov. 2016.
- [52] Y. Peng, F. X. Mittermaier, H. Planert, U. C. Schneider, H. Alle, and J. R. P. Geiger, “High-throughput microcircuit analysis of individual human brains through next-generation multineuron patch-clamp,” *eLife*, vol. 8, e48178, Nov. 2019.
- [53] A. Alegria, A. Joshi, J. O’Brien, and S. B. Kodandaramaiah, “Single neuron recording: Progress towards high-throughput analysis,” *Bioelectronics in Medicine*, vol. 3, no. 3, pp. 33–36, 2020.
- [54] M. Inoue, “Genetically encoded calcium indicators to probe complex brain circuit dynamics in vivo,” *Neuroscience Research*, vol. 169, pp. 2–8, Aug. 2021.
- [55] J. A. Cardin *et al.*, “Targeted optogenetic stimulation and recording of neurons in vivo using cell-type-specific expression of Channelrhodopsin-2,” *Nature Protocols*, vol. 5, no. 2, pp. 247–254, Jan. 2010.
- [56] N. C. Flytzanis *et al.*, “Archaerhodopsin variants with enhanced voltage-sensitive fluorescence in mammalian and *Caenorhabditis elegans* neurons,” *Nature Communications*, vol. 5, no. 1, p. 4894, Sep. 2014.
- [57] D. R. Hochbaum *et al.*, “All-optical electrophysiology in mammalian neurons using engineered microbial rhodopsins,” *Nature Methods*, vol. 11, no. 8, pp. 825–833, Jun. 2014.
- [58] V. Emiliani, A. E. Cohen, K. Deisseroth, and M. Häusser, “All-optical interrogation of neural circuits,” *Journal of Neuroscience*, vol. 35, no. 41, pp. 13 917–13 926, Oct. 2015.
- [59] K. E. Kishi *et al.*, “Structural basis for channel conduction in the pump-like channel-rhodopsin ChRmine,” *Cell*, vol. 185, no. 4, pp. 672–689.e23, Feb. 2022.
- [60] J. Y. Lin, P. M. Knutsen, A. Muller, D. Kleinfeld, and R. Y. Tsien, “ReaChR: A red-shifted variant of channelrhodopsin enables deep transcranial optogenetic excitation,” *Nature Neuroscience*, vol. 16, no. 10, pp. 1499–1508, Sep. 2013.

- [61] S. Chen *et al.*, “Near-infrared deep brain stimulation via upconversion nanoparticle-mediated optogenetics,” *Science*, vol. 359, no. 6376, pp. 679–684, Feb. 2018.
- [62] A. M. Packer, B. Roska, and M. Häuser, “Targeting neurons and photons for optogenetics,” *Nature Neuroscience*, vol. 16, no. 7, pp. 805–815, Jun. 2013.
- [63] D. A. Wagenaar, R. Madhavan, J. Pine, and S. M. Potter, “Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation,” *Journal of Neuroscience*, vol. 25, no. 3, pp. 680–688, Jan. 2005.
- [64] J. Newman, R. Zeller-Townson, M.-f. Fong, S. Arcot Desai, R. Gross, and S. Potter, “Closed-Loop, Multichannel Experimentation Using the Open-Source NeuroRighter Electrophysiology Platform,” *Frontiers in Neural Circuits*, vol. 6, 2013.
- [65] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Journal of Fluids Engineering, Transactions of the ASME*, vol. 82, no. 1, pp. 35–45, Mar. 1960.
- [66] E. A. Mukamel and J. Ngai, “Perspectives on defining cell types in the brain,” *Current Opinion in Neurobiology, Neuronal Identity*, vol. 56, pp. 61–68, Jun. 2019.
- [67] P. Němec and P. Osten, “The evolution of brain structure captured in stereotyped cell count and cell type distributions,” *Current Opinion in Neurobiology, Neurobiology of Behavior*, vol. 60, pp. 176–183, Feb. 2020.
- [68] R. I. Martinez-Garcia *et al.*, “Two dynamically distinct circuits drive inhibition in the sensory thalamus,” *Nature*, vol. 583, no. 7818, pp. 813–818, Jul. 2020.
- [69] C. Lee, A. Lavoie, J. Liu, S. X. Chen, and B.-h. Liu, “Light Up the Brain: The Application of Optogenetics in Cell-Type Specific Dissection of Mouse Brain Circuits,” *Frontiers in Neural Circuits*, vol. 14, 2020.
- [70] A. Joglekar *et al.*, “A spatially resolved brain region- and cell type-specific isoform atlas of the postnatal mouse brain,” *Nature Communications*, vol. 12, no. 1, p. 463, Jan. 2021.
- [71] H. Zeng, “What is a cell type and how to define it?” *Cell*, vol. 185, no. 15, pp. 2739–2755, Jul. 2022.
- [72] S. R. Cole and B. Voytek, “Brain Oscillations and the Importance of Waveform Shape,” *Trends in Cognitive Sciences*, vol. 21, no. 2, pp. 137–149, Feb. 2017.
- [73] S. Cole and B. Voytek, “Cycle-by-cycle analysis of neural oscillations,” *Journal of Neurophysiology*, vol. 122, no. 2, pp. 849–861, Aug. 2019.

- [74] M. S. Fabus, A. J. Quinn, C. E. Warnaby, and M. W. Woolrich, “Automatic decomposition of electrophysiological data into distinct nonsinusoidal oscillatory modes,” *Journal of Neurophysiology*, vol. 126, no. 5, pp. 1670–1684, Nov. 2021.
- [75] A. B. Saleem *et al.*, “Subcortical Source and Modulation of the Narrowband Gamma Oscillation in Mouse Visual Cortex,” *Neuron*, vol. 93, no. 2, pp. 315–322, Jan. 2017.
- [76] E. A. Buffalo, P. Fries, R. Landman, T. J. Buschman, and R. Desimone, “Laminar differences in gamma and alpha coherence in the ventral stream,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 108, no. 27, pp. 11 262–11 267, Jul. 2011.
- [77] T. J. Buschman, E. L. Denovellis, C. Diogo, D. Bullock, and E. K. Miller, “Synchronous Oscillatory Neural Ensembles for Rules in the Prefrontal Cortex,” *Neuron*, vol. 76, no. 4, pp. 838–846, Nov. 2012.
- [78] J. Aru *et al.*, “Untangling cross-frequency coupling in neuroscience,” *Current Opinion in Neurobiology*, vol. 31, pp. 51–61, Apr. 2015.
- [79] L. Zhang, J. Lee, C. Rozell, and A. C. Singer, “Sub-second dynamics of theta-gamma coupling in hippocampal CA1,” *eLife*, vol. 8, Jul. 2019.
- [80] G. Buzsáki and A. Draguhn, “Neuronal oscillations in cortical networks,” *Science*, vol. 304, no. 5679, pp. 1926–1929, Jun. 2004.
- [81] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, 2014, ISBN: 978-1-107-44761-5.
- [82] T. Akam and D. M. Kullmann, “Oscillatory multiplexing of population codes for selective communication in the mammalian brain,” *Nature Reviews Neuroscience*, vol. 15, no. 2, pp. 111–122, Feb. 2014.
- [83] I. Tal, S. Neymotin, S. Bickel, P. Lakatos, and C. E. Schroeder, “Oscillatory Bursting as a Mechanism for Temporal Coupling and Information Coding,” *Frontiers in Computational Neuroscience*, vol. 14, 2020.
- [84] M. Lundqvist, J. Rose, P. Herman, S. L. L. Brincat, T. J. J. Buschman, and E. K. K. Miller, “Gamma and Beta Bursts Underlie Working Memory,” *Neuron*, vol. 90, no. 1, pp. 152–164, Apr. 2016.
- [85] M. Lundqvist *et al.*, “Reduced variability of bursting activity during working memory,” *Scientific Reports*, vol. 12, no. 1, p. 15 050, Sep. 2022.

- [86] G. Karvat, A. Schneider, M. Alyahyay, F. Steenbergen, M. Tangermann, and I. Diester, “Real-time detection of neural oscillation bursts allows behaviourally relevant neurofeedback,” *Communications Biology*, vol. 3, no. 1, pp. 1–10, Feb. 2020.
- [87] T. K. Sato, I. Nauhaus, and M. Carandini, “Traveling Waves in Visual Cortex,” *Neuron*, vol. 75, no. 2, pp. 218–229, Jul. 2012.
- [88] L. Muller, F. Chavane, J. Reynolds, and T. J. Sejnowski, “Cortical travelling waves: Mechanisms and computational principles,” *Nature Reviews Neuroscience*, vol. 19, no. 5, pp. 255–268, Apr. 2018.
- [89] M. E. Rule, C. Vargas-Irwin, J. P. Donoghue, and W. Truccolo, “Phase reorganization leads to transient β -LFP spatial wave patterns in motor cortex during steady-state movement preparation,” *Journal of Neurophysiology*, vol. 119, no. 6, pp. 2212–2228, Jun. 2018.
- [90] S. Moldakarimov, M. Bazhenov, D. E. Feldman, and T. J. Sejnowski, “Structured networks support sparse traveling waves in rodent somatosensory cortex,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 115, no. 20, pp. 5277–5282, May 2018.
- [91] L. M. J. Fernandez and A. Lüthi, “Sleep Spindles: Mechanisms and Functions,” *Physiological Reviews*, vol. 100, no. 2, pp. 805–868, Apr. 2020.
- [92] M. M. Churchland *et al.*, “Neural population dynamics during reaching,” *Nature*, vol. 487, no. 7405, pp. 51–56, 2012.
- [93] K. V. Shenoy, M. Sahani, and M. M. Churchland, “Cortical Control of Arm Movements: A Dynamical Systems Perspective,” *Annual Review of Neuroscience*, vol. 36, pp. 337–359, 2013.
- [94] J. P. Cunningham and B. M. Yu, “Dimensionality reduction for large-scale neural recordings,” *Nature neuroscience*, vol. 17, no. 11, pp. 1500–9, Nov. 2014.
- [95] M. T. Kaufman, M. M. Churchland, S. I. Ryu, and K. V. Shenoy, “Cortical activity in the null space: Permitting preparation without movement,” *Nature Neuroscience*, vol. 17, no. 3, pp. 440–448, Feb. 2014.
- [96] J. A. Gallego, M. G. Perich, L. E. Miller, and S. A. Solla, “Neural Manifolds for the Control of Movement,” *Neuron*, vol. 94, no. 5, pp. 978–984, Jun. 2017.
- [97] S. Vyas, M. D. Golub, D. Sussillo, and K. V. Shenoy, “Computation through Neural Population Dynamics,” *Annual Review of Neuroscience*, vol. 43, pp. 249–275, Jul. 2020.

- [98] K. V. Shenoy and J. C. Kao, “Measurement, manipulation and modeling of brain-wide neural population dynamics,” *Nature Communications*, vol. 12, no. 1, pp. 1–5, Dec. 2021.
- [99] D. Peixoto *et al.*, “Decoding and perturbing decision states in real time,” *Nature*, vol. 591, no. 7851, pp. 604–609, 2021.
- [100] O. G. Sani, B. Pesaran, M. M. Shanechi, and M. Hsieh, “Where is all the nonlinearity: Flexible nonlinear modeling of behaviorally relevant neural dynamics using recurrent neural networks,” *bioRxiv*, p. 2021.09.03.458628, Sep. 2021.
- [101] C. Hurwitz *et al.*, “Targeted Neural Dynamical Modeling,” in *Advances in Neural Information Processing Systems*, vol. 35, 2021, pp. 29 379–29 392, ISBN: 978-1-71384-539-3. arXiv: [2110.14853](https://arxiv.org/abs/2110.14853).
- [102] M. Stimberg, R. Brette, and D. F. Goodman, “Brian 2, an intuitive and efficient neural simulator,” *eLife*, vol. 8, Aug. 2019.
- [103] A. Bogdanchikov, M. Zhaparov, and R. Suliyev, “Python to learn programming,” in *Journal of Physics: Conference Series*, vol. 423, IOP Publishing, Apr. 2013, p. 012 027.
- [104] A. P. Davison, M. L. Hines, and E. Muller, “Trends in programming languages for neuroscience simulations,” *Frontiers in Neuroscience*, vol. 3, no. DEC, pp. 374–380, 2009.
- [105] E. Muller, J. A. Bednar, M. Diesmann, M. O. Gewaltig, M. Hines, and A. P. Davison, “Python in neuroscience,” *Frontiers in Neuroinformatics*, vol. 9, no. APR, p. 11, Apr. 2015.
- [106] K. H. Pettersen, H. Lindén, A. M. Dale, and G. T. Einevoll, “Extracellular spikes and CSD,” *Handbook of neural activity measurement*, vol. 1, pp. 92–135, 2012.
- [107] G. Buzsáki, C. A. Anastassiou, and C. Koch, “The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes,” *Nature Reviews Neuroscience* 2012 13:6, vol. 13, no. 6, pp. 407–420, May 2012.
- [108] E. Hagen, S. Næss, T. V. Ness, and G. T. Einevoll, “Multimodal modeling of neural network activity: Computing LFP, ECoG, EEG, and MEG signals with LFPy 2.0,” *Frontiers in Neuroinformatics*, vol. 12, p. 92, Dec. 2018.
- [109] H. Parasuram, B. Nair, E. D’Angelo, M. Hines, G. Naldi, and S. Diwakar, “Computational modeling of single neuron extracellular electric potentials and network local field potentials using LFPsim,” *Frontiers in Computational Neuroscience*, vol. 10, no. Jun, p. 65, Jun. 2016.

- [110] R. J. Tomsett *et al.*, “Virtual Electrode Recording Tool for EXtracellular potentials (VERTEX): Comparing multi-electrode recordings from simulated and biological mammalian cortical tissue,” *Brain Structure and Function*, vol. 220, no. 4, pp. 2333–2353, Jul. 2015.
- [111] C. Thornton, F. Hutchings, and M. Kaiser, “The virtual electrode recording tool for extracellular potentials (VERTEX) Version 2.0: Modelling in vitro electrical stimulation of brain tissue,” *Wellcome Open Research*, vol. 4, 2019.
- [112] S. R. Nason *et al.*, “A low-power band of neuronal spiking activity dominated by local single units improves the performance of brain–machine interfaces,” *Nature Biomedical Engineering* 2020, pp. 1–11, Jul. 2020.
- [113] G. R. Holt and C. Koch, “Electrical interactions via the extracellular potential near cell bodies,” *Journal of Computational Neuroscience*, vol. 6, no. 2, pp. 169–184, 1999.
- [114] B. Telenczuk, M. Telenczuk, and A. Destexhe, “A kernel-based method to calculate local field potentials from networks of spiking neurons,” *Journal of Neuroscience Methods*, vol. 344, p. 108 871, Oct. 2020.
- [115] K. Johnsen, *Kjohnsen/tklfp: V0.2.0*, Zenodo, Jul. 2022.
- [116] T. J. Foutz, R. L. Arlow, and C. C. McIntyre, “Theoretical principles underlying optical stimulation of a channelrhodopsin-2 positive pyramidal neuron,” *J Neurophysiol*, vol. 107, pp. 3235–3245, 2012.
- [117] T. Vo-Dinh, *Biomedical Photonics: Handbook*. CRC Press, 2003, ISBN: 978-0-203-00899-7.
- [118] B. D. Evans, S. Jarvis, S. R. Schultz, and K. Nikolic, “PyRhO: A Multiscale Optogenetics Simulation Platform,” *Frontiers in Neuroinformatics*, vol. 10, no. MAR, p. 8, Mar. 2016.
- [119] K. A. Wilmes and C. Clopath, “Inhibitory microcircuits for top-down plasticity of sensory representations,” *Nature Communications*, vol. 10, no. 1, p. 5055, Nov. 2019.
- [120] A. Aussen, L. Buhry, L. Tyvaert, and R. Ranta, “A detailed anatomical and mathematical model of the hippocampal formation for the generation of sharp-wave ripples and theta-nested gamma oscillations,” *Journal of Computational Neuroscience*, vol. 45, no. 3, pp. 207–221, Dec. 2018.
- [121] A. Aussen, R. Ranta, O. Aron, S. Colnat-Coulbois, L. Maillard, and L. Buhry, “Cell to network computational model of the epileptic human hippocampus suggests specific

roles of network and channel dysfunctions in the ictal and interictal oscillations,” *Journal of Computational Neuroscience*, Aug. 2022.

- [122] A. Mazzoni, H. Lindén, H. Cuntz, A. Lansner, S. Panzeri, and G. T. Einevoll, “Computing the Local Field Potential (LFP) from Integrate-and-Fire Network Models,” *PLOS Computational Biology*, vol. 11, no. 12, e1004584, 2015.
- [123] O. Yizhar, L. E. Fenno, T. J. Davidson, M. Mogri, and K. Deisseroth, “Optogenetics in Neural Systems,” *Neuron*, vol. 71, no. 1, pp. 9–34, Jul. 2011.
- [124] M. Kokaia, M. Andersson, and M. Ledri, “An optogenetic approach in epilepsy,” *Neuropharmacology*, vol. 69, pp. 89–95, Jun. 2013.
- [125] J. M. Stujenske, T. Spellman, and J. A. Gordon, “Modeling the Spatiotemporal Dynamics of Light and Heat Propagation for InVivo Optogenetics,” *Cell Reports*, vol. 12, no. 3, pp. 525–534, Jul. 2015.
- [126] J. H. Macke, L. Buesing, J. P. Cunningham, B. M. Yu, K. V. Shenoy, and M. Sahani, “Empirical models of spiking in neural populations,” in *Advances in Neural Information Processing Systems*, vol. 24, Curran Associates, Inc., 2011.
- [127] G. Buzsáki and X.-J. Wang, “Mechanisms of Gamma Oscillations,” *Annual Review of Neuroscience*, vol. 35, no. 1, pp. 203–225, 2012.
- [128] X.-J. Wang and G. Buzsáki, “Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model,” *Journal of neuroscience*, vol. 16, no. 20, pp. 6402–6413, 1996.
- [129] N. Brunel and X.-J. Wang, “What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance,” *Journal of neurophysiology*, vol. 90, no. 1, pp. 415–430, 2003.
- [130] Y. Wang and S. Boyd, “Fast Model Predictive Control Using Online Optimization,” *IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY*, vol. 18, no. 2, p. 267, 2010.
- [131] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, “The explicit linear quadratic regulator for constrained systems,” *Automatica*, vol. 38, no. 1, pp. 3–20, Jan. 2002.
- [132] J. F. Kalaska, R. Caminiti, and A. P. Georgopoulos, “Cortical mechanisms related to the direction of two-dimensional arm movements: Relations in parietal area 5 and comparison with motor cortex,” *Experimental Brain Research*, vol. 51, no. 2, pp. 247–260, Jul. 1983.

- [133] D. L. Barack and J. W. Krakauer, “Two views on the cognitive brain,” *Nature Reviews Neuroscience*, vol. 22, no. 6, pp. 359–371, Jun. 2021.
- [134] C. Pandarinath *et al.*, “Inferring single-trial neural population dynamics using sequential auto-encoders,” *Nature Methods*, vol. 15, no. 10, pp. 805–815, Oct. 2018.
- [135] F. R. Willett, D. T. Avansino, L. R. Hochberg, J. M. Henderson, and K. V. Shenoy, “High-performance brain-to-text communication via handwriting,” *Nature*, vol. 593, no. 7858, pp. 249–254, May 2021.
- [136] O. G. Sani, Y. Yang, M. B. Lee, H. E. Dawes, E. F. Chang, and M. M. Shanechi, “Mood variations decoded from multi-site intracranial human brain activity,” *Nature Biotechnology*, vol. 36, no. 10, pp. 954–961, Nov. 2018.
- [137] G. K. Anumanchipalli, J. Chartier, and E. F. Chang, “Speech synthesis from neural decoding of spoken sentences,” *Nature*, vol. 568, no. 7753, pp. 493–498, Apr. 2019.
- [138] A. S. Morcos and C. D. Harvey, “History-dependent variability in population dynamics during evidence accumulation in cortex,” *Nature Neuroscience*, vol. 19, no. 12, pp. 1672–1681, Dec. 2016.
- [139] T. D. Kim, T. Z. Luo, J. W. Pillow, and C. D. Brody, “Inferring Latent Dynamics Underlying Neural Population Activity via Neural Differential Equations,” in *Proceedings of the 38th International Conference on Machine Learning*, PMLR, Jul. 2021, pp. 5551–5561.
- [140] D. Sussillo and O. Barak, “Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks,” *Neural Computation*, vol. 25, no. 3, pp. 626–649, Mar. 2013.
- [141] D. Sussillo, “Neural circuits as computational dynamical systems,” *Current Opinion in Neurobiology*, Theoretical and Computational Neuroscience, vol. 25, pp. 156–163, Apr. 2014.
- [142] J. T. Smith, S. W. Linderman, and D. Sussillo, “Reverse engineering recurrent neural networks with Jacobian switching linear dynamical systems,” in *Advances in Neural Information Processing Systems*, vol. 20, Dec. 2021, pp. 16 700–16 713, ISBN: 978-1-71384-539-3. arXiv: [2111.01256](https://arxiv.org/abs/2111.01256).
- [143] J. Pearl, *Causality*. New York: Cambridge University Press, 2009, ISBN: 978-0-521-89560-6.
- [144] S. Dufour and Y. D. Koninck, “Optrodes for combined optogenetics and electrophysiology in live animals,” *Neurophotonics*, vol. 2, no. 3, p. 031 205, Jul. 2015.

- [145] S. Jeon *et al.*, “Implantable Optrode Array for Optogenetic Modulation and Electrical Neural Recording,” *Micromachines*, vol. 12, no. 6, p. 725, Jun. 2021.
- [146] C. Kathe *et al.*, “Wireless closed-loop optogenetics across the entire dorsoventral spinal cord in mice,” *Nature Biotechnology*, vol. 40, no. 2, pp. 198–208, Feb. 2022.
- [147] K. Y. Kwon, H.-M. Lee, M. Ghovanloo, A. Weber, and W. Li, “Design, fabrication, and packaging of an integrated, wirelessly-powered optrode array for optogenetics application,” *Frontiers in Systems Neuroscience*, vol. 9, 2015.
- [148] M. Welkenhuysen *et al.*, “An integrated multi-electrode-optrode array for in vitro optogenetics,” *Scientific Reports*, vol. 6, no. 1, p. 20 353, Feb. 2016.
- [149] L. Wang, K. Huang, C. Zhong, L. Wang, and Y. Lu, “Fabrication and modification of implantable optrode arrays for in vivo optogenetic applications,” *Biophysics Reports*, vol. 4, no. 2, pp. 82–93, Apr. 2018.
- [150] N. McAlinden *et al.*, “Multisite microLED optrode array for neural interfacing,” *Neurophotonics*, vol. 6, no. 3, p. 035 010, Aug. 2019.
- [151] D. Mao, N. Li, Z. Xiong, Y. Sun, and G. Xu, “Single-Cell Optogenetic Control of Calcium Signaling with a High-Density Micro-LED Array,” *iScience*, vol. 21, pp. 403–412, Nov. 2019.
- [152] D. Mao, Z. Xiong, M. Donnelly, and G. Xu, “Brushing-Assisted Two-Color Quantum-Dot Micro-LED Array Towards Bi-Directional Optogenetics,” *IEEE Electron Device Letters*, vol. 42, no. 10, pp. 1504–1507, Oct. 2021.
- [153] Y. Ohta *et al.*, “Micro-LED Array-Based Photo-Stimulation Devices for Optogenetics in Rat and Macaque Monkey Brains,” *IEEE Access*, vol. 9, pp. 127 937–127 949, 2021.
- [154] J. Antolik, Q. Sabatier, C. Galle, Y. Frégnac, and R. Benosman, “Assessment of optogenetically-driven strategies for prosthetic restoration of cortical vision in large-scale neural simulation of V1,” *Scientific Reports*, vol. 11, no. 1, pp. 1–18, May 2021.
- [155] A. M. Packer, L. E. Russell, H. W. P. Dalglish, and M. Häusser, “Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo,” *Nature Methods*, vol. 12, no. 2, pp. 140–146, Feb. 2015.
- [156] E. Ronzitti *et al.*, “Submillisecond Optogenetic Control of Neuronal Firing with Two-Photon Holographic Photoactivation of Chronos,” *Journal of Neuroscience*, vol. 37, no. 44, pp. 10 679–10 689, Nov. 2017.

- [157] I.-W. Chen, E. Papagiakoumou, and V. Emiliani, “Towards circuit optogenetics,” *Current Opinion in Neurobiology*, Neurotechnologies, vol. 50, pp. 179–189, Jun. 2018.
- [158] M. M. Churchland, J. P. Cunningham, M. T. Kaufman, S. I. Ryu, and K. V. Shenoy, “Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine?” *Neuron*, vol. 68, no. 3, pp. 387–400, Nov. 2010.
- [159] G. Santhanam *et al.*, “Factor-Analysis Methods for Higher-Performance Neural Prostheses,” *Journal of Neurophysiology*, vol. 102, no. 2, pp. 1315–1330, Aug. 2009.
- [160] G. Bellec *et al.*, “A solution to the learning dilemma for recurrent networks of spiking neurons,” *Nature Communications*, vol. 11, no. 3625, 2020.
- [161] R. Rajalingham, A. Piccato, and M. Jazayeri, “Recurrent neural networks with explicit representation of dynamic latent variables can mimic behavioral patterns in a physical inference task,” *Nature Communications*, vol. 13, no. 1, pp. 1–15, Oct. 2022.
- [162] G. R. Yang, M. R. Joglekar, H. F. Song, W. T. Newsome, and X.-J. Wang, “Task representations in neural networks trained to perform many cognitive tasks,” *Nature Neuroscience*, vol. 22, no. 2, pp. 297–306, Feb. 2019.
- [163] L. Driscoll, K. Shenoy, and D. Sussillo, *Flexible multitask computation in recurrent networks utilizes shared dynamical motifs*, Aug. 2022.
- [164] L. Duncker and M. Sahani, “Dynamics on the manifold: Identifying computational dynamical activity from neural population recordings,” *Current Opinion in Neurobiology*, Computational Neuroscience, vol. 70, pp. 163–170, Oct. 2021.
- [165] C. Rackauckas and Q. Nie, “DifferentialEquations.jl – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia,” *Journal of Open Research Software*, vol. 5, no. 1, p. 15, May 2017.
- [166] F. Zenke and T. P. Vogels, “The Remarkable Robustness of Surrogate Gradient Learning for Instilling Complex Function in Spiking Neural Networks,” *Neural Computation*, vol. 33, no. 4, pp. 899–925, Mar. 2021.