Simulating Barrel Cortex Oscillations Using Spiking Neural Networks

Isabel Quinn Garon, BSc
December 2021

Final Report

Host Institution
Technische Universität Graz
Institut für Grundlagen der Informationsverarbeitung

Supervisor
Dr. Robert Legenstein

Table of Contents

Additional Materials

Figures

Tables

Code Fragments

1 Introduction

1.1 Abstract

The following report investigates strategies to emulate and manipulate the delayed oscillatory behavior demonstrated in the barrel cortex in response to a biphasic current injection. This data was gathered as a part of the SYNCH project, and the phenomenon is believed to play a role in short term working memory of somatosensory stimuli. The primary reference for this target behavior was a series of Powerpoint slides of composite depth electrode recordings, and later raw data from the recordings themselves. The first phase of this work explores the conditions in which network-wide oscillations could be captured in simulation in response to the same stimuli used *in vivo*, and confirms earlier suspicions that they are mediated by inhibitory neurons. The behavior is demonstrated in two models - a network of 500 generalized integrate and fire neurons (*microscopic model*), and a mean-field population model (*mesoscopic model*), in which activity in the excitatory and inhibitory populations is captured by a population firing rate [26, 35]. The second phase of this project explores which features of these oscillations can be manipulated, within the constraints of current injections possible *in vivo*. While oscillatory frequency and duration varies widely in response to a single current injection, ultimately it was possible to fix the duration of the oscillatory period using multiple current injections.

1.2 SYNCH Project

This work falls under the umbrella of the European Union's Horizon 2020 project SYNCH (A **SY**naptically connected brain-silicon **N**eural **C**losed-loop **H**ybrid system). The goal of this project is to produce a closed-loop neural-hybrid system, which reads information out from the mouse cortex, interprets it on a neuromorphic chip, and projects a response back onto the cortex. Eventually, it would explore the coadaptation of the artificial and biological networks, and their clinical applications to neurological disorders [14]. My work considers the challenge of projecting behavior onto the mouse cortex, and is performed in collaboration with the Vassanelli Lab at the University of Padua. The model system investigated in the biological data is whisker somatosensation in the mouse barrel cortex.

2 Background

2.1 Biological Networks

The network of the brain comprises nearly 100 billion individual units called neurons. The interactions between these units, across the orders of magnitude from individual cells, local

populations, and brain regions, mediate all thought, learning, and environmental interactions over the lifetime of an animal. While our capacity to understand these interactions is dwarfed by the sheer size and complexity of even simple model organisms, our understanding builds from a few fundamental principles.

2.1.1 Anatomy of a Neuron

Neurons have three sections - the dendrites, cell body, and axon. Electrical signals are received at the dendrites: projections from the cell body with a complex morphology resembling the arbor of a tree. These signals are integrated as they pass down the "branches" of this tree towards the body of the cell, or soma. Once a critical voltage is achieved at the soma, the neuron produces an action potential or "spike," a signal that propagates down the neuron's axon, synapsing on to the dendrites of other cells, and perpetuating the spread of information through the brain [36]. According to the fundamental Hebbian principle, the more often one neuron synapses onto another, the increased synaptic reliability of that connection [19]. Thus, experience mediates changes in the anatomical structure of the brain, in order to support learning and memory.

2.1.2 Neuron Subtypes

These 100 billion neurons are diverse in their electrophysiological and synaptic properties. Broadly categorized, there are two types of neurons: excitatory and inhibitory. Excitatory neurons have an amplifying effect, increasing the likelihood of a spike in the neurons they synapse on to. Excitatory pyramidal neurons are the dominant neuron type in the cortex, and are relatively homogenous when compared with the other major category of cortical neuron - the inhibitory neurons. Inhibitory neurons are more opaque, as their function is instead a subtle dampening of behavior, a decrease in the likelihood of activity in neurons they synapse on to. These inhibitory neurons can be further subdivided into somatostatin, parvalbumin, and ionotropic receptor expressing subtypes, with diverse electrophysiology and connectivity profiles [38].

In addition to these cortex-wide categories, 90% of the mammalian cortex is striated into six layers, from the outermost Layer 1, which is primarily inhibitory, to the innermost Layer 6 [31]. Each of these layers has distinct combinations of cell types and densities, and play different roles in the routing of information through the brain [3, 23, 38].

2.2 Simulated Networks

Simulation is an essential technique in Neuroscience. There are three principal components to a neural network that is intended to simulate a biological behavior (this is distinct from artificial neural networks, which have the added nuance of a learning rule). First, there are the biophysical properties of a neuron. The dynamics of an individual neuron's behavior were first reduced to a system of equations by Hodgkin and Huxley in 1952. These equations have

been used to describe and simulate neurons with complex morphologies in an accurate manner, however these models are also computationally intensive [15, 20]. These equations have also been reduced to more efficient systems, in which a neuron is abstracted to a single point, and its behavior captured by the fluctuation of the membrane voltage over time and a spiking threshold (see equations). These equations have been demonstrated to be surprisingly accurate to the more detailed biophysical models [5, 13].

The second part of a network simulation is the connectivity - the specificity of synapses occurring between axons and dendrites. Layers within the brain have different connection probabilities, as do inhibitory and excitatory populations [21]. The connectivity can be scaled according to the spatial arrangement of neurons [33] and can also be adjusted within a moderate region in order to achieve balanced network activity that resembles activity seen in vivo [7].

The final component is the input that drives network activity. The brain is not a silent system waiting for a sensory stimulus. There is a significant amount of ongoing background activity into which new information is integrated. Network simulations typically have some underlying drive, such as a poisson pool of randomly firing neurons or a base depolarizing current to emulate this sustained activity. It is also possible to stimulate the system with injected current or patterned stimuli.

In addition to these components, there is the analysis and interpretation of system behavior - in a complex model, there is a spectrum of features that can be saved from the network, including the spikes, the membrane voltage of each neuron, or more population-wide features such as spike rate or synchrony level. What is primarily depicted in the following research are raster plots - these represent the spikes or action potentials released by cells over the course of the simulation. The y-axis is the neuron number (in this case, the population is generic and there is no special importance to the order of these neurons) and the dot on the x-axis represents a single event for the neuron. These are qualitative observations, and cannot be compared perfectly to the biological data, which is recorded from a depth electrode (spikes are deduced from changes in conductance).

2.2.1 Benefits of Simulation

There are numerous benefits to using simulation over *in vivo* experiments. While all models are imperfect descriptions of the phenomenon they emulate, when analyzing a system as complex as the brain, a perfect simulation would be just as opaque to researchers as data gathered *in vivo*. Simulation allows scientists to limit the scope and scale of their experiments so that causality can be more reliably determined. These simulations also offer an omniscient perspective, an ability to see and monitor the properties of each individual neuron over the course of a trial, a possibility that is in the distant experimental future. In a sense, it is an opportunity to tap into the disparate data sets generated in experimental neuroscience, and bind them together to form a cohesive model that can be further explored [15].

Alongside this omniscient perspective is complete control over simulation features, that is unbounded by biological constraints. Simulations offer a space where networks which could not

exist in biology can be probed in order to establish a more complete spectrum of behavior, for example a population of neurons with all-to-all connectivity, unreasonably fast integration times, or no noise. Even randomized features of the network can be held constant from trial to trial, by the power of random seeding. Additionally, the navigation of simulations between levels of resolution, from biophysically detailed reconstructions of neurons [20, 18], to point neuron simulations [4,5,7,9,10,13,33,37], to mean-field approximations of population behavior [35, 22, 17], have allowed scientists to explore the level of detail necessary to capture the behavior they study, and dramatically increase the speed with which experiments can be conducted. In this report, the typical trial with a population of 500 neurons can be conducted at about 4x the speed of *in vivo* behavior. This is further improved by the use of computer clusters and parallel processing, and opens up the data gathered to optimization and machine learning procedures.

A final benefit of simulation to highlight here is the comparatively low barrier of entry and cost in biological resources of simulation techniques. Many simulation paradigms, particularly at the resolution of point neuron and population models, can be performed on a laptop computer, and require no more complex experimental construction. These experiments can also source from existing data, and do not require expensive optical setups or additional mouse lives.

2.3 Mouse Barrel Cortex

The recordings emulated here (Figure 1) are captured in the somatosensory cortex of mice, in a region integrating tactile signals from whisking activity known as the barrel cortex. The barrel cortex offers a well-investigated model of cortical column and receptive field organization [16]. Within this region, each whisker is associated with a densely locally connected and discernible area known as a barrel. Cortical Layer 4 of this somatosensory region receives feedforward information from the thalamus, as is typical of early sensory regions, where it is processed by dense local circuitry that has been explored extensively [16]. This region also has a number of afferent projections to the secondary somatosensory cortex, motor cortex, thalamus, pons, and brain stem [2].

Connectivity and neuron-subtype distribution within the barrel cortex is in many ways distinct. While the widely accepted average ratio of excitatory to inhibitory neurons in the cortex is 4:1, only about 11% of neurons in the barrel cortex are inhibitory [24, 16]. These inhibitory neurons receive intense feedforward excitation from the thalamus, primarily targeted towards Layer 4, where projections from the thalamus are densest, a feature which likely plays a critical role in the formation of receptive fields, sparse coding, and the resetting of network balance after a strong sensory stimulus [3, 23, 12, 10]. This intense inhibition is facilitated by dense connectivity between the fast-spiking interneurons and excitatory cells within each barrel, and dense afferent connectivity from L4 to the surrounding layers [23, 24].

The qualitative behavior that is the focus of this report is the rapid oscillations seen in response to whisker stimulation in the barrel cortex in anesthetized mice. There are several notable differences between the oscillatory response generated from the whisker stimulation (via thalamus) and electrode stimulation (direct to barrel cortex) in the same recorded region. While

electrode stimulation and whisker stimulation both produce oscillatory behavior, for a similar duration of about one second following stimulus presentation, there are some qualitative differences in the recordings (Figure 1). The electrode stimulus results in larger oscillation intensities, perhaps resulting from a larger spatial spread of the current injection. The frequency of these oscillations is also much slower, in the range of 10-12Hz. There is also a peak in spiking activity at onset of the whisker stimulus that is not seen with the current injection.

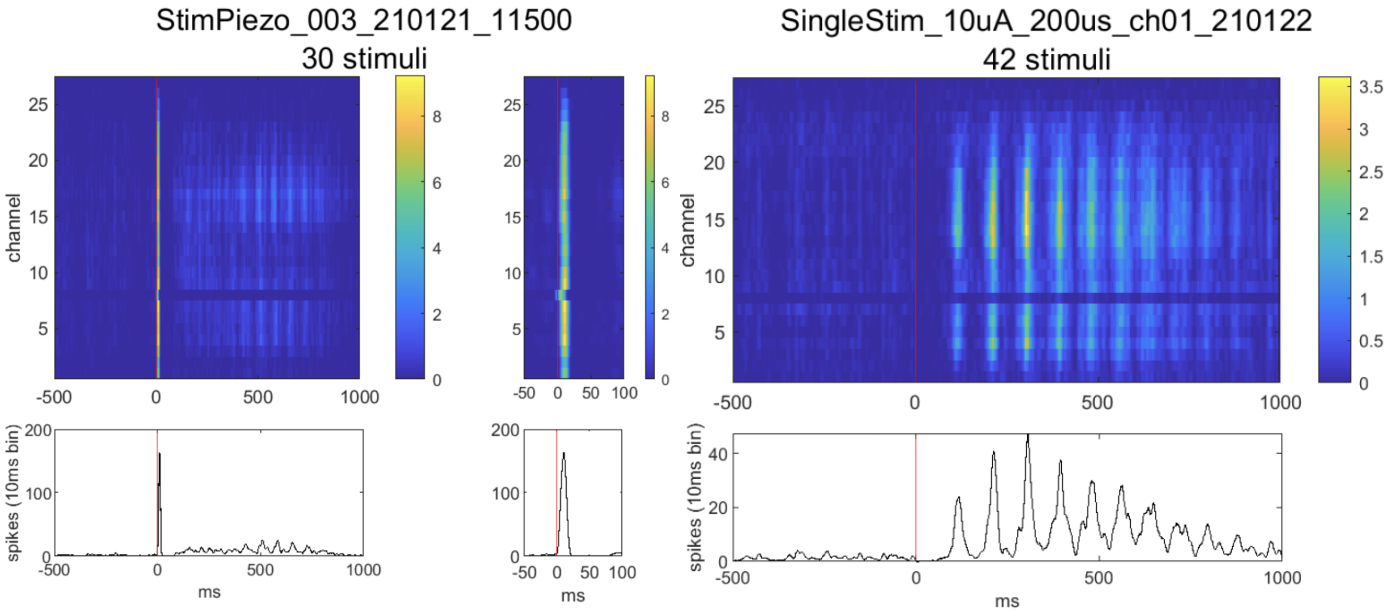

Figure 1: *Recordings from barrel cortex. Channels 14-20 approximately align with Layer 4. (Left) Oscillations in response to whisker stimulation. Top is a composite of 30 stimuli, recorded across all 6 layers. Bottom represents a spike histogram over the same time course. (Center) Zoom of the spike peak at stimulation of the whisker stimuli. (Right) Electrode biphasic stimulation composite.*

---

3 Modeling Strategies

Three types of modeling strategies were explored over the course of this project. The Allen Institute visual cortex model [5] was investigated for its hyper-realistic connectivity, although this strategy was ultimately eliminated, as it was computationally expensive and too slow for optimization procedures. The remaining models explored were microscopic simulations of individual point neurons, and mesoscopic simulations, in which the activity of the excitatory and inhibitory populations was abstracted to a single node [35]. The mesoscopic model performs marginally faster than the microscopic model (5x versus 4x real time). Oscillations were first elicited from this model, as the connectivity is necessarily all-to-all, and one of the constraints

considered in the results as required to produce oscillations is dense recurrent connectivity in a population. After this initial breakthrough, this model was primarily used as a control, as its structure eliminates topological concerns of clustering, in-degree and out-degree confounding factors. It also could be expanded in future work to incorporate multiple layers or diverse inhibition in an efficient fashion.

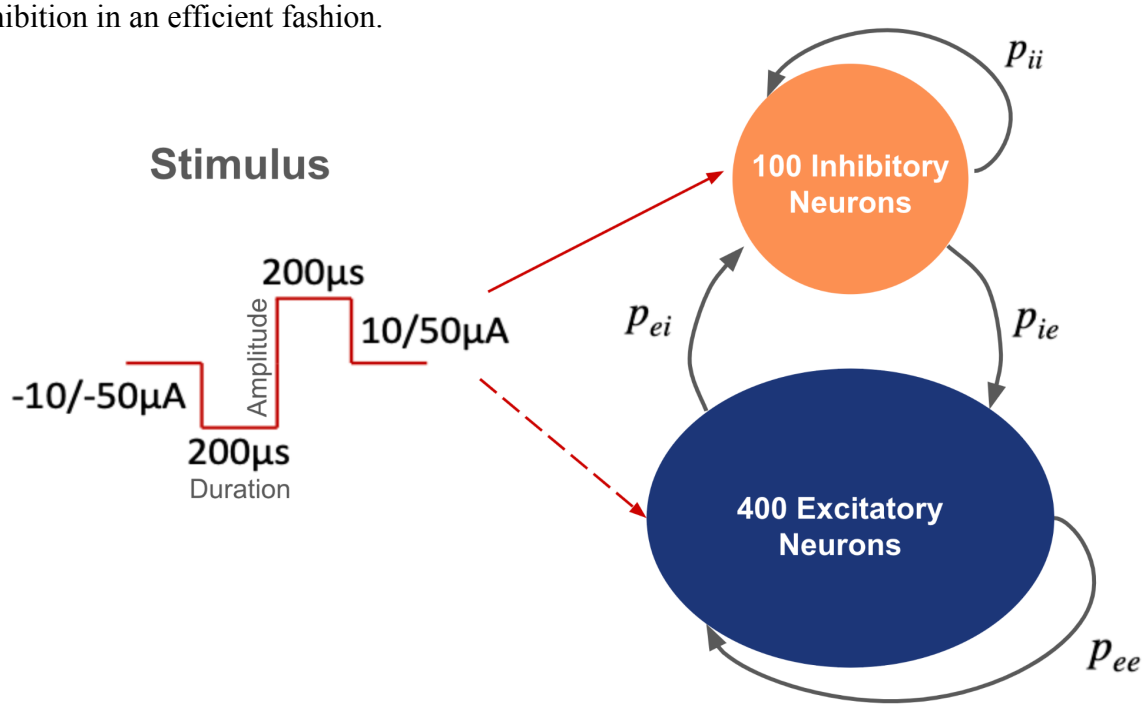

Figure 2: *Diagram of the network, including correct connection probability notation and in vivo stimulus details.*

These simulations were performed on networks of 500 neurons, and for a duration of two seconds, with stimulus onset at one second, unless otherwise noted (Figure 2). The stimulus was identical to that used in the *in vivo* experiments, a biphasic current injection, and the duration and amplitude of this injection was varied within a reasonable range (amplitude $.01\ \mu A$ to $500\ \mu A$, duration $200\ \mu S$, $1\ ms$, $5\ ms$, or $10\ ms$). The time step of the simulation is 0.5 ms, unless the duration of current injection is smaller than that, in which case the timestep is 0.2 ms - this change has no apparent impact on network dynamics.

3.1 Microscopic Model

The microscopic population activity is governed by current-based generalized-integrate-and-fire equations [26]. The membrane voltage of each point neuron evolves over time according to the equation:

$$dV(t)/dt = [-g_L * (V(t) - E_L) + I(t)]/C$$

The firing intensity λ of each neuron determines stochastic spike emission given the firing threshold $V_T(t)$ and the threshold softness Δ_V:

$$V_T(t) = V_{T_{star}} + \gamma_1(t) + \gamma_2(t)$$

$$\lambda(t) = \lambda_0 * exp(V(t) - V_T(t))/\Delta_V$$

The firing threshold incorporates spike-frequency-adaptation kernels γ_i, which decay according to:

$$d\gamma_i/dt = -\gamma_i / \tau_{\gamma i}$$

$$\gamma_i = \gamma_i + q_i$$

Individual parameter descriptions and values are available in Table 1. Although possible using this neuron type in NEST simulation software, no spike-triggered currents are used in these simulations [25]. Each of these neurons is randomly connected based on a fixed indegree reflective of the connection probability.

3.2 Mesoscopic Model

The mesoscopic model is governed by equations describing interacting populations of neurons [25, 35]. The key value in the mesoscopic population model is that it accounts for the effects of the finite size of a population of neurons. The population firing rate is corrected, such that it demonstrates the random rate fluctuations seen in populations smaller than infinite size: a more accurate model of the brain. The first integral in the equation below represents the expected firing rate in the case of an infinite population of neurons, defined by the conditional firing intensity given the time of the neuron's last spike, and the density of last spike times in the population. The second term corrects for finite population size using the population optimal firing rate:

$$\overline{A}(t) = \int_{-\infty}^{t} \lambda(t|\hat{t}\,)\, S(t|\hat{t}\,)\, A_N(\hat{t})\, d\hat{t} + \Lambda(t)(1 - \int_{-\infty}^{t} S(t|\hat{t}\,)\, A_N(\hat{t})\, d\hat{t}\,)$$

The behavior of a single neuron is determined by the hazard function:

$$h(t) = \lambda_0\, exp(\frac{V_m(t) - E_{sfa}(t)}{\Delta_V})$$

Since the mesoscopic model describes populations of neurons based on the GIF equations used in the microscopic model, the same parameters are used in each situation [25].

Table 1 - *Network parameters used in both mesoscopic and microscopic models [30, 35]*

Network Parameters		
C_m	Membrane Capacitance	250 pF
g_L	Membrane Capacitance	C_m / τ_m
E_L	Reversal Potential	0 mV
T_{ref}	Refractory period	4 ms
V_{reset}	Voltage reset potential	0 mV
μ	Depolarizing voltage, accounts for all external currents $\mu = R(I_{ext} + V_{rest})$	24 mV
q_{sfa}	Increase in spike frequency adaptation after each emission	1 mV, 10 mV
τ_{sfa}	Spike frequency adaptation time constant	100 ms, 1000 ms
Δ_u	Softness of threshold (noise level)	2.5 mV
τ_m	Membrane time constant	20 ms
τ_e	Excitatory synaptic time constant	3 ms
τ_i	Inhibitory synaptic time constant	6 ms
V_{th}	Base firing threshold	15 mV
c	Escape rate at firing threshold	10 Hz

3.3 Allen Institute Visual Cortex Model

Initial investigations for this project considered the large-scale visual cortex model described in [5]. This model was appealing for several reasons. Foremost, it is currently the largest publicly available cortical model, simulating nearly a millimeter of cortex. It contains over 230,000 neurons, including one excitatory and three inhibitory subtypes with electrophysiologically realistic features for 4 of the cortical layers. The model also is available at two resolutions, a biophysically detailed model (i.e. containing spatially detailed dendritic arbors) and a point neuron model, which uses the same connectivity as the more detailed model, but abstracts each neuron to a single point, which behaves according to the generalized leaky-integrate-and-fire modeling paradigm outlined in [37]. It is a proof of concept of the accuracy of point-neuron moes, as the point neuron simulations were extremely accurate to the biophysical simulations, but were conducted nearly 8,000x as fast as the more detailed model [5,

13]. The incorporation of layers was an interesting feature to capture, in an effort to stimulate the SYNCH project recordings, as the depth electrodes did span the cortical layers, and spatial spread throughout those layers may have played a role in the ramp and decay of the oscillations. While a different sensory system (vision instead of somatosensation), early sensory systems follow a similar structure, with information traveling from the sensory organ, through the thalamus, and the Layer 4 of the primary sensory cortex. While a somatosensory cortex model does exist, it is not publicly available [32].

Ultimately, this model was not a good fit for the goals of this project. The runtime was not optimal for exploring a large parameter space - while the simulation itself was reasonable to perform on the computing resources available, constructing a network of that size in the Nest simulation software took over an hour, even distributed over several computer nodes. Additionally, the complexity of the model could have obscured the underlying cause of the oscillations, as there were several sources of noise and randomization at each level of the simulation. Finally, the base drive to the Allen Institute model is also complex, consisting of a background source of shared activity, as well as a simulated Lateral Geniculate Nucleus. The simulated LGN is a linear-non-linear poisson process which creates spike trains in response to a time-dependant visual stimulus [5, 13]. It is difficult to disentangle the roles of each of these network stimuli from the stimulus investigated in this project, the biphasic current injection, even when using a simple “grey-screen” or poisson LGN stimulus.

4 Parameter Selection

4.1 Parameter sweeps

Because models are imperfect simulacrums of biological networks, the parameters governing network activity cannot be directly pulled from biological data. Parameter sweeps are a valuable tool in determining what values should be used in a simulation that are within a biologically feasible range, but will still produce balanced, sustained network activity. In this case, a grid-zoom parameter search is employed. First, plausible parameter ranges along two axes of investigation were identified: network connectivity, and input stimulus. For each investigation, a high granularity grid was constructed, considering all possible combinations of the target parameters (in this case inter- and intra- neuron population connectivity, and input amplitude and duration). Then multiple simulations were performed at each parameter combination, scoring the performance of the simulation using a series of scoring functions. Using these scores, a subrange of parameters that performs well was identified, and the sweep was constructed and rerun using a lower-granularity grid.

Function 1: *Grid-Zoom Parameter Sweep - performs grid zoom parameter sweep across all given variables [28]. Format for **kwargs: simulation_key = [lower_bound, upper_bound, step]*

```
from sklearn.utils.extmath import cartesian
def grid_zoom(**kwargs):
    keys = []
    values = []
    num_trials= 1
    # unpack arguments, establish ranges
    for key, value in kwargs.items():
        keys.append(key)
        values.append(np.arange(value[0], value[1],
                                value[2]))
        num_trials = num_trials * len(values[-1])
    run_params = cartesian(values) # Construct grid
    # run grid zoom
    current = 1
    for run in run_params:
        new_args = {}
        filename = ""
        for i, num in enumerate(run):
            new_args[keys[i]]=num
            filename = filename+"_"+str(keys[i])
                                +"_"+str(num)
        sim = Simulation(**new_args,save =
                                filename,randomize = True)
        sim.run_micro()
```

4.1.1 Scoring Strategies

Three general scoring metrics were considered in these parameter sweeps: synchrony, oscillatory frequency, and oscillatory duration. Synchrony was measured using the SPIKE-synchronization method from the PySpike python package. This is a pairwise coincident detector, where two perfectly synchronized spike trains produce a score of 1 [27]. During the initial oscillatory period, network average synchrony scores are almost exactly 1, and the variance in pairwise scores is low, making this measurement an effective way of distinguishing the duration of the initial oscillatory period from the residual synchrony that can be detected in some samples. Any relationship between pre- and post- stimulus synchrony, as well as synchrony in base current versus poisson driven networks, is still being investigated.

Distinct oscillations and their frequencies are determined using the scipy signal-analysis Python package, which determines the intensity, width, and interpeak intervals of all peaks at least 20 ms apart and a standard deviation above the average prestimulus network activity. Frequency

is calculated using only the peaks in the initial oscillatory burst, as this frequency sometimes increases as intensity decreases.

Function 2: *Population Synchrony - Calculate synchrony in network behavior. Format: array of spike times for each neuron. Precisely synchronized spikes produce a score of 1 [27].*

```
import pyspike
def SPIKEdist(spike_trains):
    sp = []
    # Format spike trains for PySpike
    for i in spike_trains:
        sp.append(pyspike.SpikeTrain(i,
                    edges = (0.0,400.)))
    # Calculate spike synchrony profile
    spike_profile = pyspike.spike_sync_profile(sp)

    print("SPIKE distance: %.8f" % spike_profile.avrg())
    print("SPIKE profile integral: %.8f" %
          (spike_profile.integral()[0]/2000.))
    # Plot results
    x, y = spike_profile.get_plottable_data()
    plt.figure()
    plt.plot(x, y)
    return True
```

Function 3: *Identifying Oscillation Peaks - given the rate for a single trial, returns the times of the oscillatory peaks, the maximum prominence, the width of the individual peak and the interval between peaks (Figure 5). Calculated using the SciPy signal analysis package [40].*

```
import scipy.signal as sig
def peaks(spike_rate, init_prom=0):
    # Option to set a base prominence or re-calculate for
    the given trial
    if init_prom == 0:
        prom = sig.peak_prominences(spike_rate,
               sig.find_peaks(spike_rate)[0])[0]
        init_prom = np.mean(prom)+np.std(prom)*2.5

    peaks = sig.find_peaks(spike_rate,
              prominence=init_prom, distance = 30)[0]
    maxes = spike_rate[peaks]
    widths = sig.peak_widths(spike_rate, peaks)[0]
    ipi = np.diff(peaks, n=1)

    return [peaks, maxes, widths, ipi, init_prom]
```

Function 4: *Calculating Frequency - calculates the frequency of oscillation in the initial set of oscillatory peaks, based on the standard width of oscillations prior to any gaps, and disregarding peaks that occur after a single oscillation is missed.*

```
def osc_rate(peaks, widths_):
    widths = np.zeros(len(peaks))
    widths[1:] = widths_

    if len(peaks) > 2:
        combined = zip(peaks, widths)
        no_gaps = []

        for p, w in combined:
            if w < (1.5*widths[1]):
                no_gaps.append(p)
            else:
                break
        rate = len(no_gaps)/(max(no_gaps)/1000)

    else:
        rate = 0
        no_gaps = []
    return [rate, len(no_gaps)]
```

4.1.2 Generic parameter sweep

My initial parameter sweep considered inter- and intra- population connectivity. The range of connection probabilities was reflective of the range of values that have been measured in cortical slices ($p_{e\to e} = 20\% - 35\%$, $p_{e\to i} = 30\% - 60\%$, $p_{i\to e} = 40\% - 70\%$, $p_{i\to i} = 30\% - 70\%$) [4]. The expectation, based on previous research where all-to-all connectivity produced the most coherent oscillations, was that denser local connectivity would be favored [6, 8, 41]. The step of the connectivity sweep was 10%, and was later repeated on areas of interest with a decreased granularity of 2%.

The most successful connectivity regimes had higher excitatory to inhibitory connectivity ($p_{e\to i} \approx 55\%$) and inhibitory to inhibitory connectivity ($p_{i\to i} \approx 35\%$) than recurrent excitatory connectivity ($p_{e\to e} \approx 30\%$). Still, the recurrent connectivity within the excitatory population was high. Nearly all regimes with a lower connection probability between excitatory and inhibitory populations ($p_{e\to i}$) than within the excitatory population ($p_{e\to e}$) produced synchronous, regular activity [9]. Oscillations with a similar gain and decay structure and frequency to the ICMS slides could be found at a number of stimulus parameter combinations. In response to an identical stimulus as in the slides ($10\ \mu A$ for $200\ \mu S$) several connectivity regimes produced a very similar

result to the experimental recordings (Figure 3, top right), with a marginally faster frequency (12.5 Hz, versus 10 Hz in the original).

At each connectivity combination, the magnitude and duration of current injection, and base drive to the network was also considered .The amplitude of the biphasic current injected into the inhibitory population spanned a range of orders of magnitude, from. $01\ \mu A$to $500\ \mu A$. The duration of the input was $200\ \mu S$, $1\ ms$, $5\ ms$, or $10\ ms$. The base depolarization for each neuron is calculated according to the equation $\mu = R(I_{external} + V_{rest}) = 24\ mV$, but a range of base depolarizations from $\mu = 16\ mV - 30\ mV$ was tested. Additionally, distinct depolarization values for the inhibitory and excitatory populations were considered, to investigate how this impacted the balance of the network [35]. Beyond the simple binary of producing or not producing oscillations, input had little impact on the performance of the oscillations. Input that was too small (short in duration or small in amplitude) or a weak base drive to the network would only produce a single post-stimulus feature. Otherwise, the input had no discernable effect on oscillation duration or frequency.

4.1.3 Barrel cortex specific parameter sweep

Regardless of which layer current is initially injected in the ICMS slides, increased activity is generally first seen in Layer 4. Although there are some exceptions to this trend, it appears as if the oscillations are somewhat left-skewed in their intensity, and that the L4 region is maximally engaged by the second or third oscillation (Figure 2, bottom left, channels 17-20). This behavior is unsurprising, given the situation of Layer 4 within the barrel cortex. This region has a greater number of recurrent excitatory connections than other layers, as well as a greater number of afferent projections to other layers, with a comparatively small amount of efferent feedback [24]. It is possible that the gain in signal can be partially attributed to the integration of the current injection as it spreads across the cortex. Given these features and the fact that L4 is the primary target of projections conveying whisking information from the thalamus, it is likely that this oscillatory activity is driven primarily by Layer 4.

In light of this, in addition to the initial parameter sweep on networks using generic cortical network properties (4:1 ratio of excitatory to inhibitory neurons), a parameter sweep was conducted with network properties more specific to the barrel cortex. This region of the brain, and in particular Layer 4, has significantly fewer inhibitory neurons than elsewhere - only about 11% of neurons are GABAergic. Additionally, there is evidence that excitatory connectivity is slightly higher than the generic 20% probability of connection [24]. In combination to this higher excitatory connectivity, the connectivity between inhibitory and excitatory pools within a single barrel is dense (45%-60% inhibitory-to-excitatory connectivity, depending on inhibitory subtype, and 25%-60% excitatory-inhibitory connection probability in L2/3, with potentially denser connectivity in L4) [3, 23].

The results of this secondary sweep were very similar to the initial sweep, with two regions of successful connectivity values. One region that is similar to the original

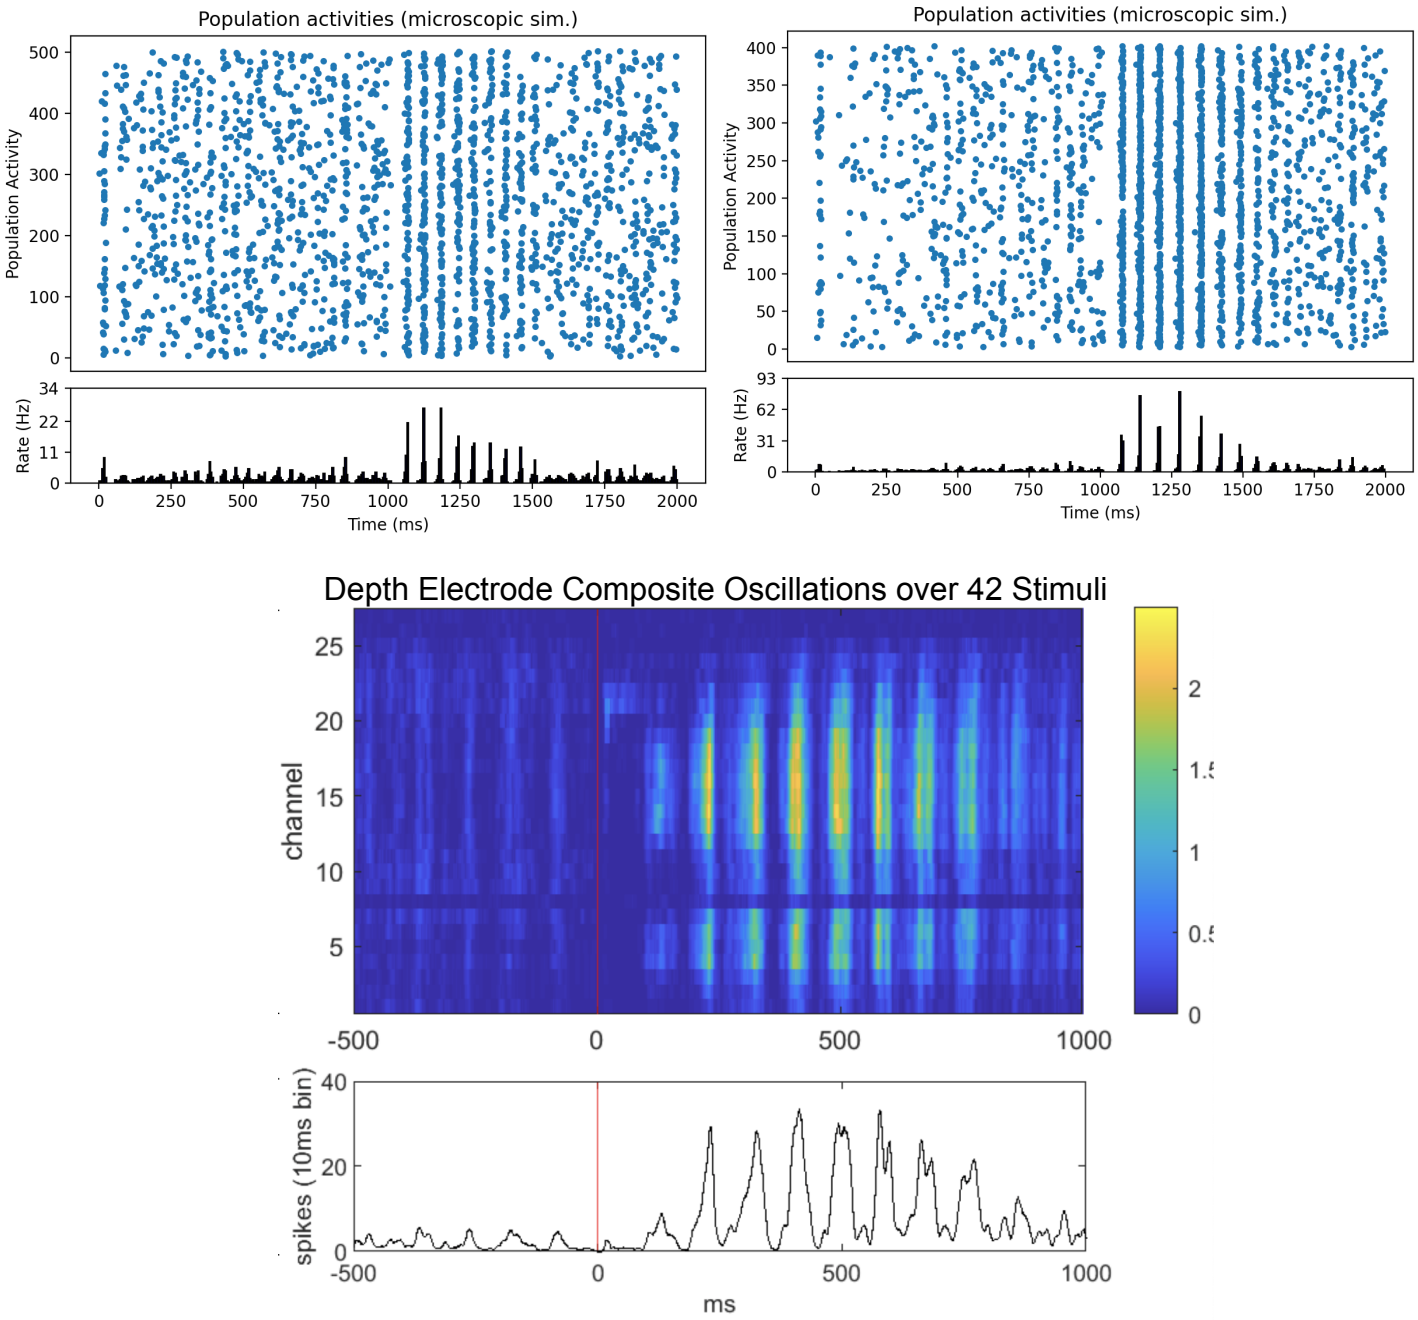

Figure 3 - *Microscopic network activity for the barrel cortex parameter sweep (top left), the general parameter sweep (top right) and a sample from the ICMS slides provided by the Vassanelli Lab (bottom).*

($p_{ee} = 30\%$, $p_{e\to i} = 50\%$, $p_{i\to e} = 40\%$, $p_{i\to i} = 30\%$) and one in which the excitatory to excitatory connectivity was fixed to the value seen *in vivo*, and less realistic inhibitory connectivity ($p_{e\to e} = 24\%$, $p_{e\to i} = 70\%$, $p_{i\to e} = 35\%, p_{i\to i} = 15\%$). This regime was more

prone to synchronous events, with low inhibitory to inhibitory connectivity. It could, however, sustain activity at much lower base currents, meaning that the network fired at lower initial firing rate (2-4 Hz) and produced slower oscillations (10-12 Hz, as low as 8 Hz), that were more reflective of the ICMS data (Figure 7). This region also demonstrated the greatest number of simulations with consistent gain and decay in its signal response (Figure 3, top left).

The requirement for low i-to-i connectivity is not wholly unbased. There is no single review on inhibitory to inhibitory connectivity in barrel cortex, but in general Layer 4 of the cortex has a higher concentration of parvalbumin interneurons, which preferentially synapse onto themselves and excitatory neurons, but have low i-i connection probabilities compared to SST or Htr3a neurons which typically target other inhibitory populations [5, 38].

4.2 Weight scaling

In the original [35] code, the weights were scaled such that they always matched a simulation paradigm in which the population size was 1000 (800 E 200 I) and the probability of connection between each population was 20%. This scaling prevented changes in activity resulting from new connection probabilities, which was a requirement for producing oscillations in the microscopic simulations, and was eliminated from the final simulation procedure.

4.3 Electrophysiological Parameters

Next, a broad exploration of the values governing neuron dynamics (Table 1) was conducted to see if there were any easily identified stiff or sloppy dimensions. This search indicated that these results were generally robust against small changes in network parameters [29]. This was examined by checking whether an increase or decrease of a parameter by 10% caused an immediate state change in network dynamics. While the response to the stimulus of course evolved, dynamics were generally stable. There were two notable exceptions to this rule - the parameter governing colored-noise current, Δ_u, is sensitive to a tenth of a millivolt in the original simulation paradigm. Δ_u determines the softness of the spiking threshold, and describes the range around the set threshold at which the neuron is guaranteed to fire ($V_{th} - \Delta_u$) and unlikely to fire at all ($V_{th} + \Delta_u$). Functionally, this value represents the noise within the population as well as fluctuations in background activity [11, 35]. While this sensitivity is unsurprising, it suggests a limited range of depolarization/membrane voltages within the network. This hypothesis is supported by the comparative response of a network driven by a poisson population - the range of Δ_u values that produce the same balanced network activity in the microscopic simulation are much wider (> 1mV, compared to 0.01mV).

Following this adjustment, the relationship between the oscillations and the Δ_u value was explored. The system appears to act as a damped oscillator, in which the excitatory/inhibitory

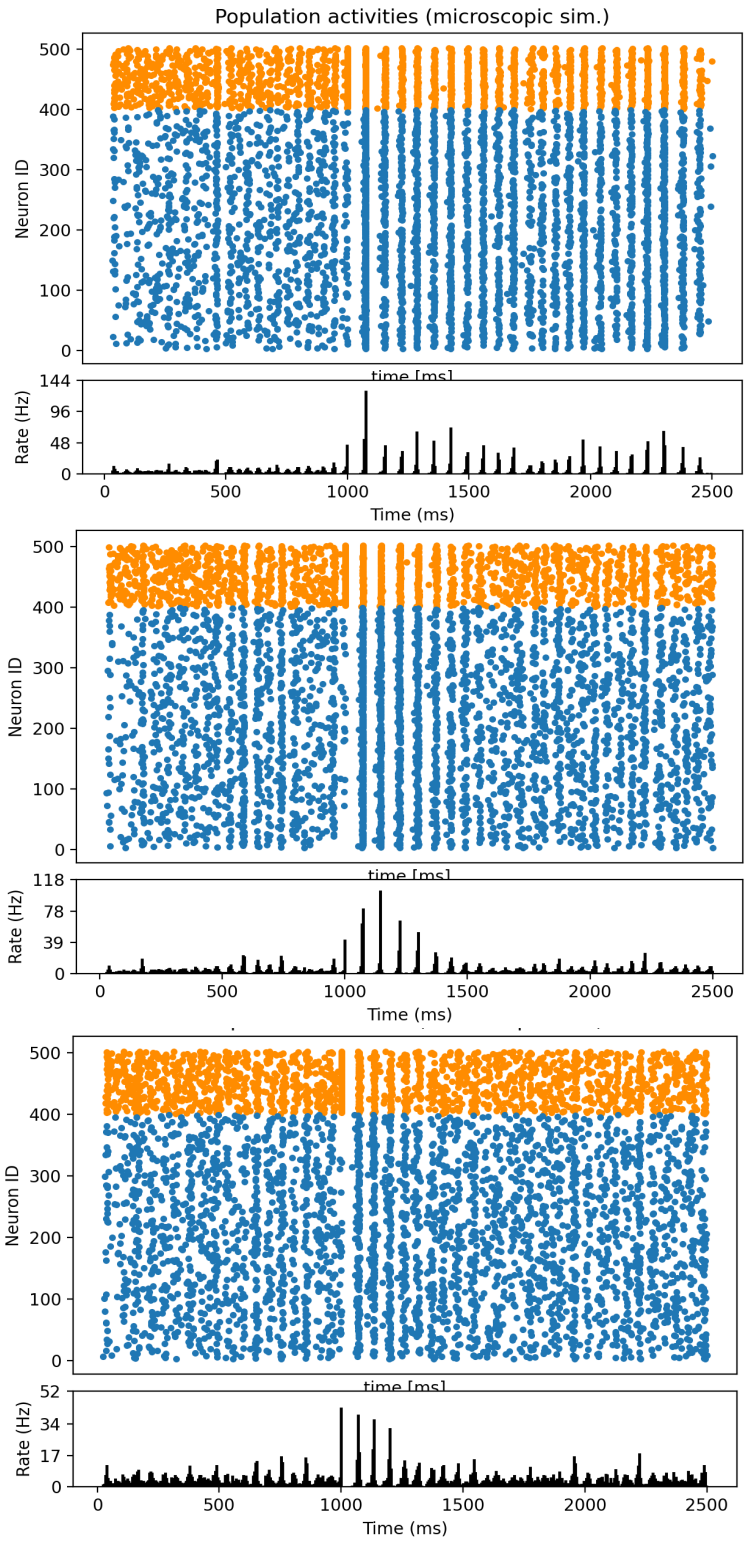

Figure 4 - *Example simulations with Δ_u values of 2.0 mV (top), 2.5 mV (middle), and 3.0 mV (bottom).*

bursts of activity would continue to appear indefinitely in the absence of noise in the system. Δ_u therefore closely determines the duration of oscillations - higher levels of noise in the network eliminates the oscillations sooner, and lower Δ_u values coordinate with longer or indefinite oscillatory phases (Figure 3). Network firing state was also sensitive to changes in the base weight. Further investigation of how different weight scaling methods impact dynamics is warranted [4].

4.4 Poisson Drive and Randomization

Earlier simulation paradigms had limited randomization, and exhibited extremely low trial to trial variability. Next, methods of adding noise to the system were explored. The initial membrane voltage was randomly selected from a normal distribution, with a low of -60mV and a high of $-30\,mV$. Initial weights were also randomized - in the original simulations, excitatory weight was $.3\,mV$ and inhibitory $1.5\,mV$if the number of connections is equal to the fixed target value, and so this value was adjusted to a range $e = .25mV - .35\,mV$ and $e = -1\,mV - -2mV$.

No comprehensive parameter sweeps have been performed using a poisson drive instead of a constant depolarizing input voltage. This report only compared specific parameter combinations of interest. A regime with similar drive to the depolarizing voltage used in the original code ($\mu = 24\,mV$), maintains a base of

$\mu = 10\,mV$, with additional input from 100 neurons firing at a rate of 10 Hz, connected to the excitatory population with a probability of 20%. The poisson input has only been explored for the microscopic model, and significantly slows the runtime (2x realtime, instead of 1.1x).

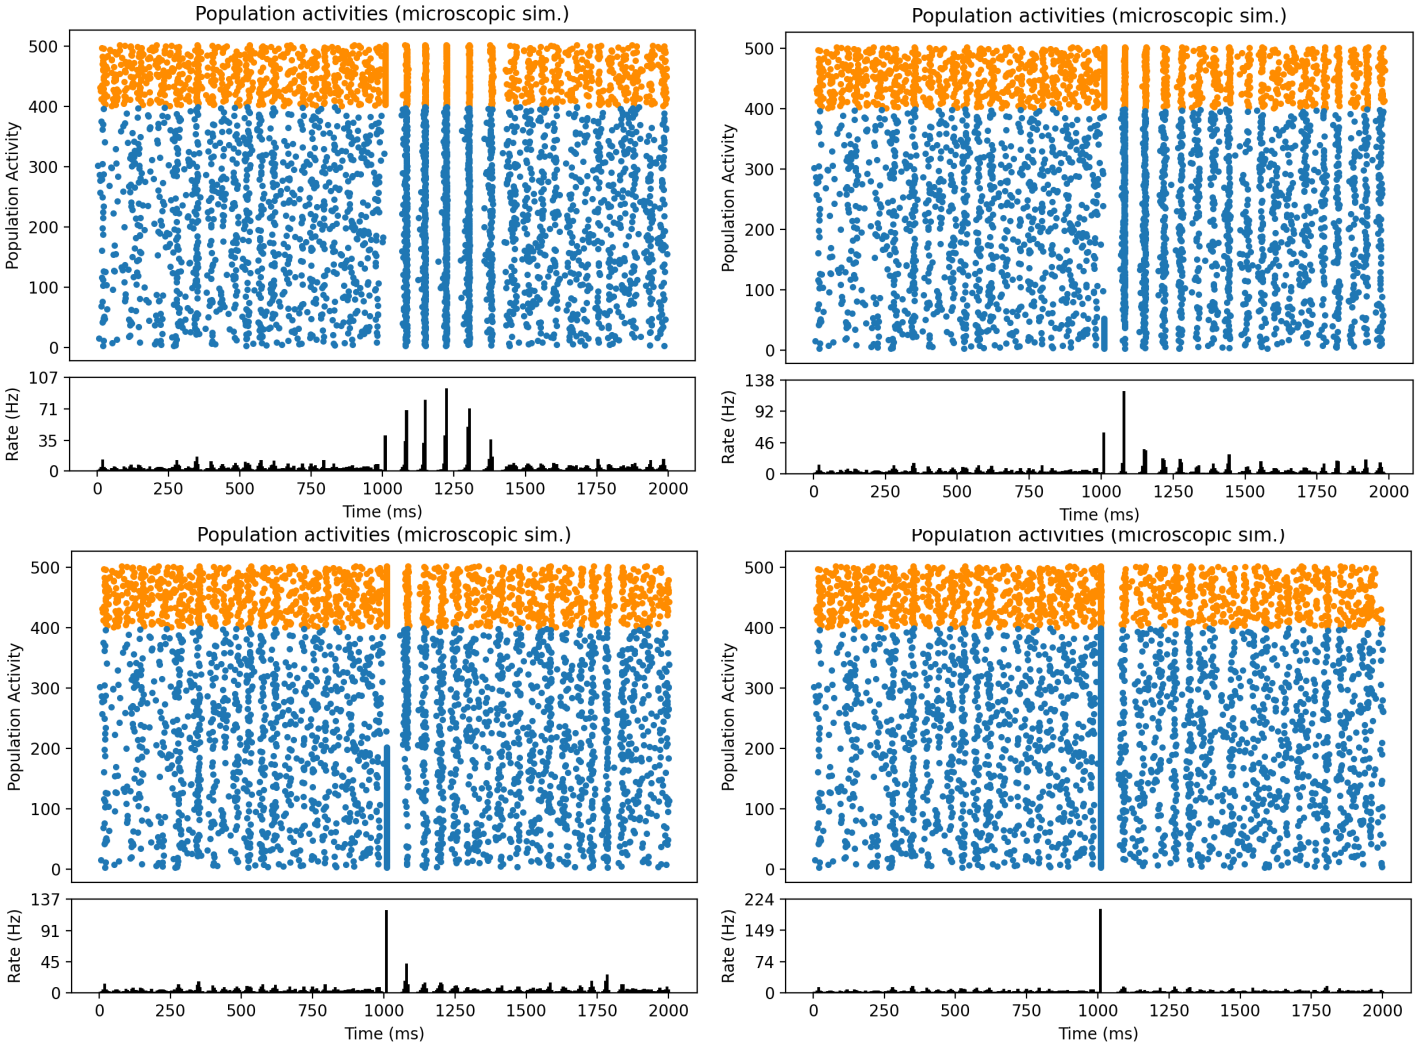

Figure 5: *Oscillatory behavior is inhibition mediated. (Top left) Excitatory Population activity when current is injected to Inhibitory population only ($N_i = 100$). (Top right) Excitatory activity when current is injected to all inhibitory neurons and 50 excitatory neurons. (Bottom left) Current to inhibitory neurons and 200 excitatory neurons. (Bottom right) Current injected to Excitatory neurons only.*

5 Results

5.1 Requirements for oscillation

Two simulation features were mandatory in eliciting oscillatory behavior after an injection of current into the network. First, current injection had to be predominantly targeted to the inhibitory population of neurons (Figure 5, top left). Injecting current to excitatory neurons as well as inhibitory neurons lessened the amplitude of oscillatory response (Figure 5, top right, bottom left), and injecting current to all neurons or only excitatory neurons produced no oscillatory behavior (Figure 5 bottom right). The second requirement was either extremely dense (greater than 70%) connectivity, or independent connection probabilities between neuron populations (i.e. connection probability cannot be fixed at 20% regardless of neuron identity). Given that this second connectivity scheme is more realistic to that seen in vivo, the following focuses on populations with distinct connectivity.

5.1.1 Oscillation qualitative comparison

A majority of the simulations only produced oscillations until around the 500ms mark, before rapidly decaying, consistent with the most intense oscillations in the ICMS slides. After this point, frequency of oscillations in the simulations increases slightly as they decay, a change of no more than 1.5 Hz. This may be consistent with the in vivo data, it is difficult to assess such a small change visually.

As a further metric of comparison, the width and the interpeak intervals were scored. As the intensity of the oscillations increases, the width of the peak decreases, corresponding to an increase in interpeak intervals. These changes are on the order of 1-5ms, so very small, but were seen across simulation paradigms (Figure 6).

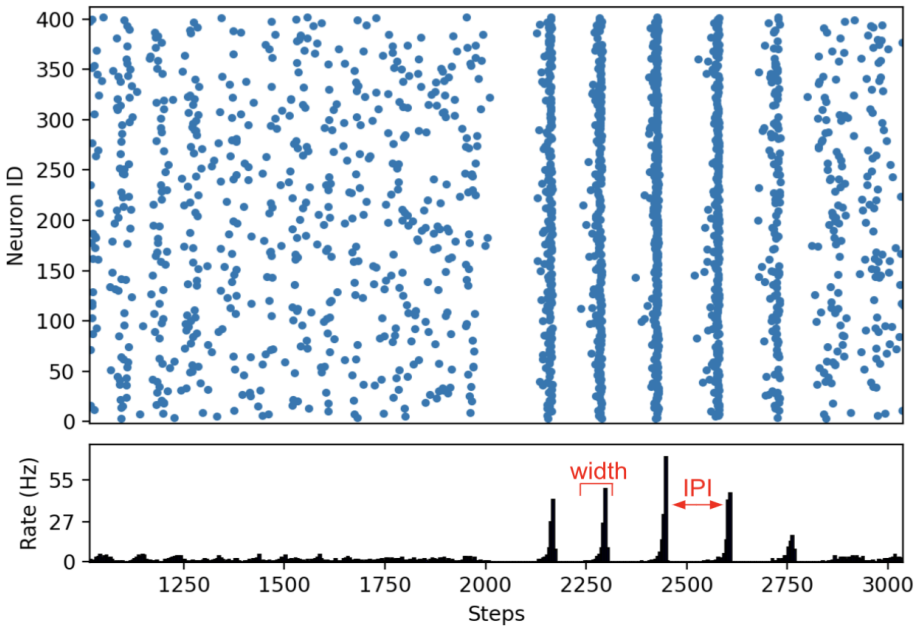

Figure 6 - *Raster plot of oscillating network activity in response to a* 50 µ*A*/200 µ*Sstimulus, and zoomed to show only the 500 ms preceding and following stimulus. Intensity and IPI increase, and a subsequent decrease in peak width can be seen in the first three oscillations.*

5.1.2 Post-stimulus delay

Lower inhibitory to inhibitory connectivity is related to a longer delay before oscillation initiation post-stimulus. These delays last the duration of a single oscillatory period +/-5 ms (Figure 7). These results required no adjustment of the inhibitory synaptic time constant, although [39] suggests that this delay *in vivo* is mediated by $GABA_B$ receptors in response to a volley of spikes. The delay periods seen in these simulations are slightly shorter than in the simulated $GABA_B$mediated delay, however they seem consistent with the delay before the first oscillation in the ICMS slides, especially given the higher base firing rates and oscillation frequencies. Further exploration of the relationship between this delay and the time constant governing inhibitory postsynaptic decay is required [1, 8].

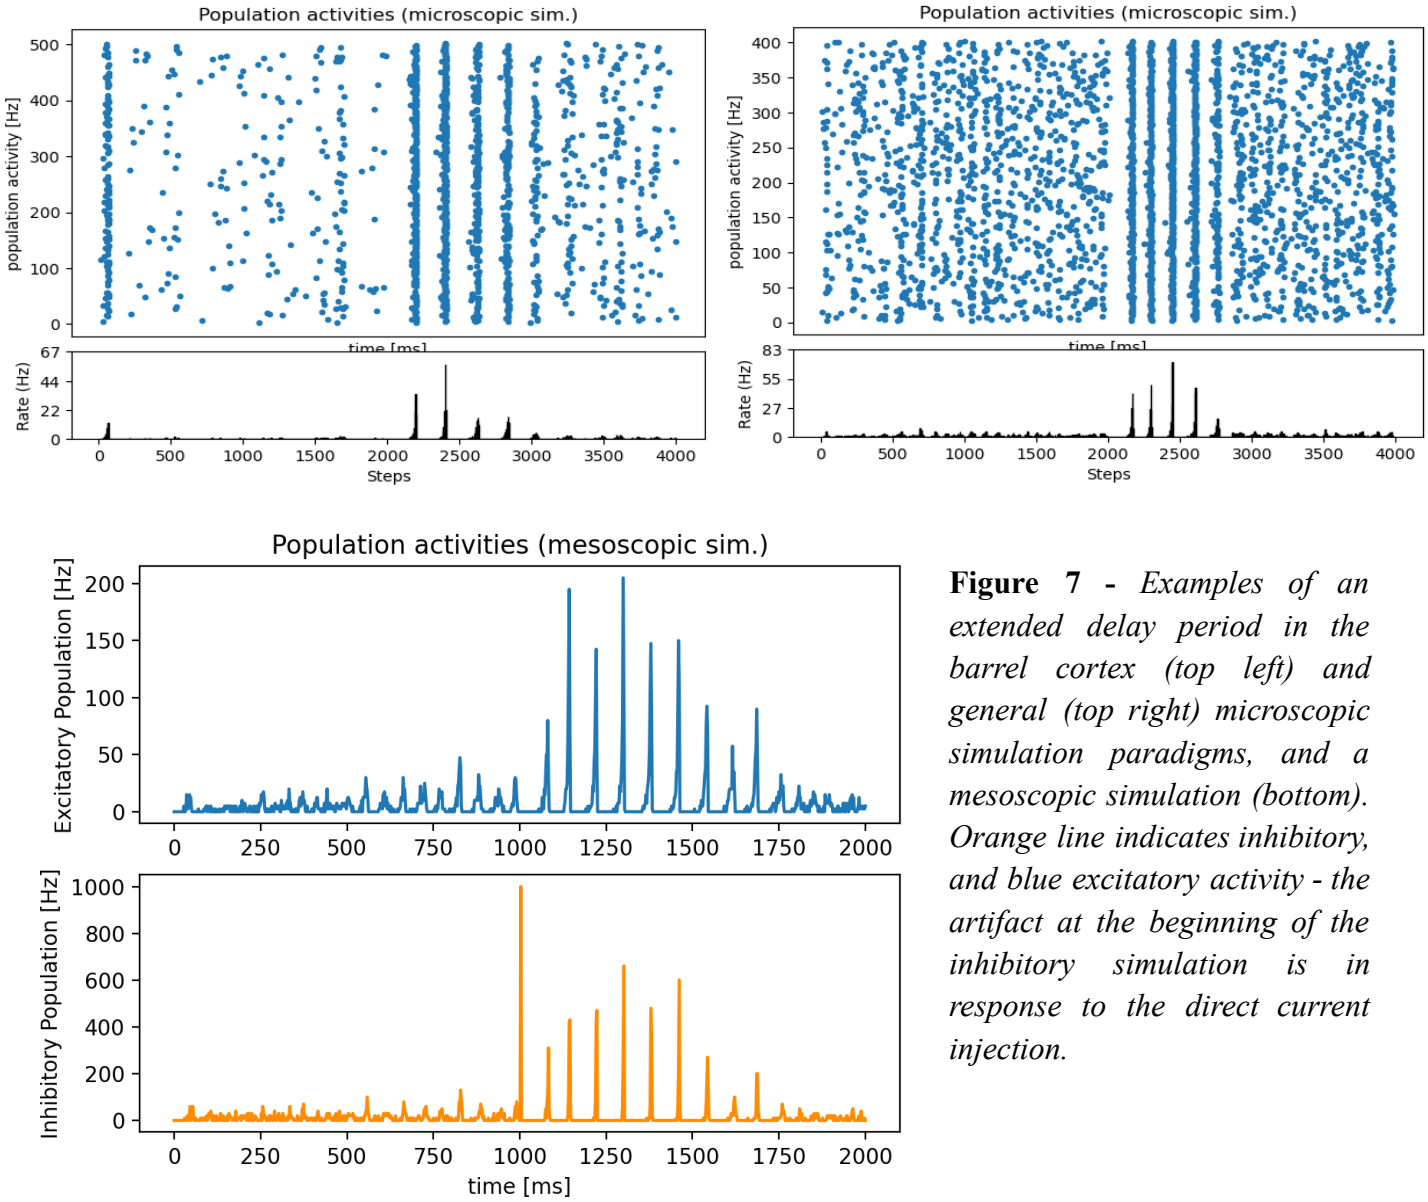

Figure 7 - *Examples of an extended delay period in the barrel cortex (top left) and general (top right) microscopic simulation paradigms, and a mesoscopic simulation (bottom). Orange line indicates inhibitory, and blue excitatory activity - the artifact at the beginning of the inhibitory simulation is in response to the direct current injection.*

5.1.3 Synchrony in the network

At a number of both stimulus and connectivity combinations across those surveyed, the networks would demonstrate large, defined oscillations for 500-800ms, which decay in intensity, but some networks never fully resumed an asynchronous-irregular regime (Figure 8). This is not mitigated by the addition of poisson input into the network. In some cases, the intensity of these sustained oscillations appeared to themselves oscillate. It is not possible to determine the duration of any residual oscillations in the ICMS slides, since the data shows just 1 second after stimulus, but it does appear that in some trials, there is some small oscillations in network activity before stimulus presentation. In some cases it appears there were some small, fast oscillations in these simulations before the stimulus as well (Figure 3, bottom left).

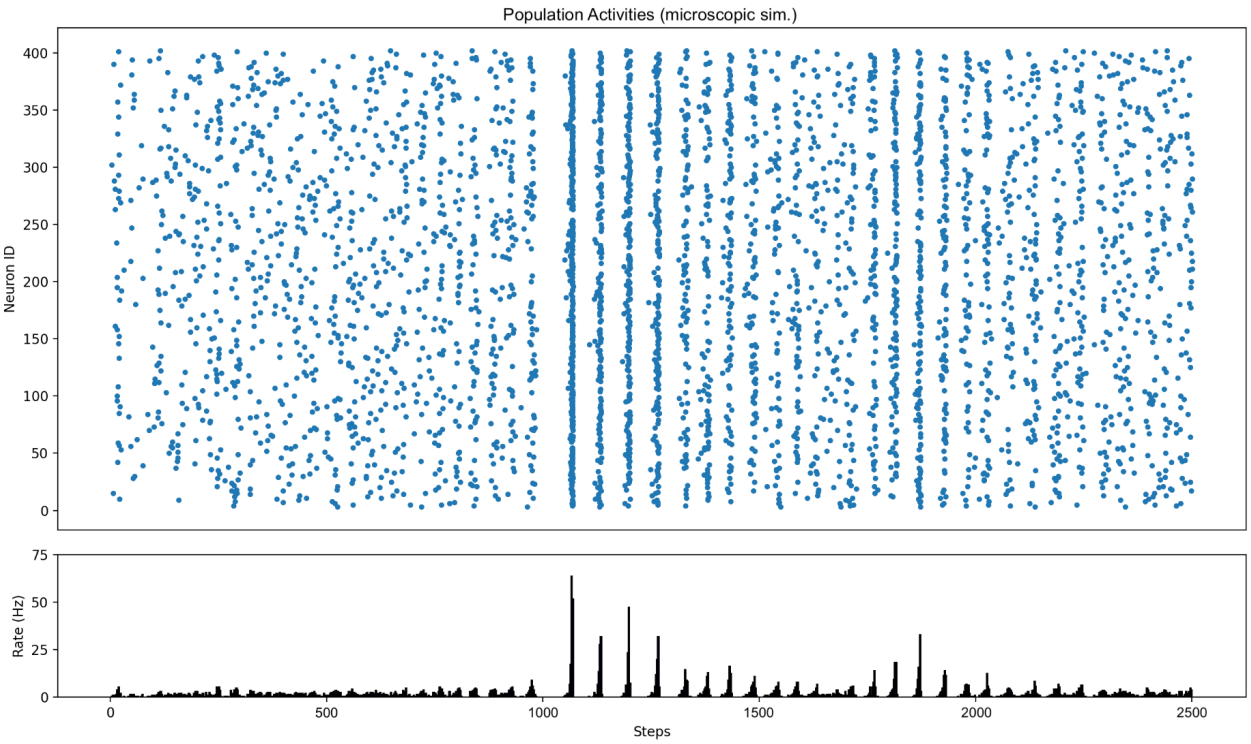

Figure 8 - *Microscopic simulations demonstrating a secondary oscillation. Simulation runtime was extended to 2.5 seconds in order to completely capture this effect.*

5.2 Microscopic & Mesoscopic Model Comparison

It should be noted that connectivity in the mesoscopic model is necessarily all-to-all, given that the excitatory and inhibitory populations are represented by a single node. The weights between those populations are, however, scaled to reflect the different connection probabilities. On average, the behavior of the meso and micro scopic models are similar (Figure 9), and the resulting firing and oscillation rates are within 1.5 Hz for each simulation paradigm, although

identical signal gain is not seen in every simulation, and extreme changes in the parameters governing spiking dynamics affect the simulations differently.

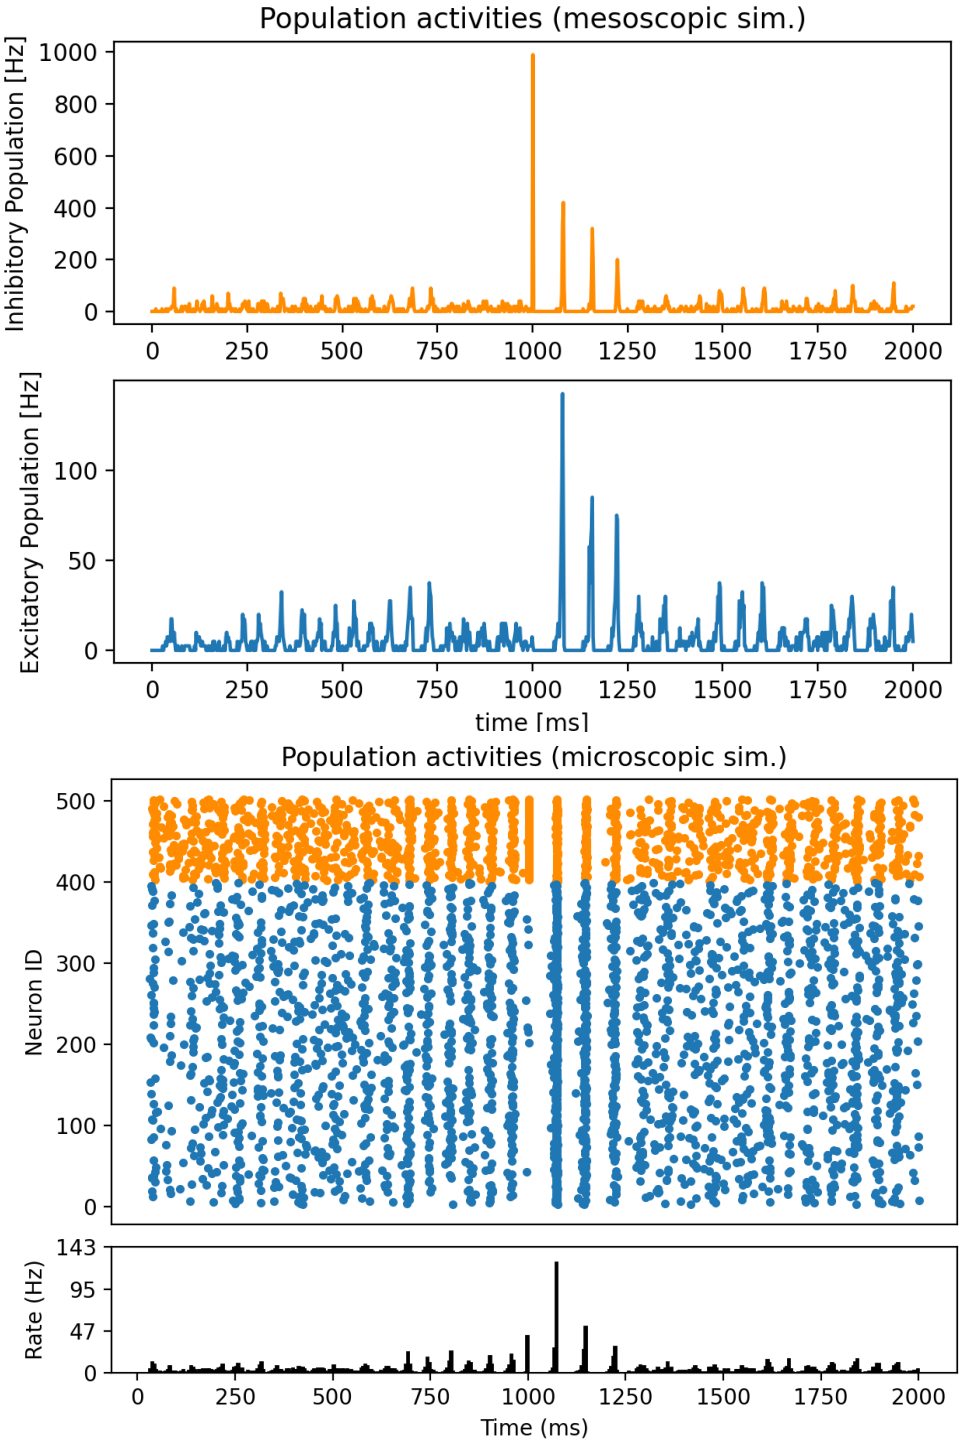

Figure 9 - *Example of a mesoscopic (top) simulation. A_N- population activity, $\overline{A}$- instantaneous population firing rate. (Bottom) Excitatory activity of a microscopic simulation with the same parameters. Spike histogram below.*

5.3 Multiple Stimulation Strategies

Although the trial-to-trial response is highly variable in response to a single stimulus, change in network behavior can be consistently affected with multiple current injections. Next, a routine was developed to demonstrate one approach to controlling the number of oscillations in a single trial using multiple current injections - this procedure required identifying two features of the oscillatory response. First, the ideal time to inject current to perpetuate the oscillations. Given that current injected during the population wide refractory period (the 60 ms following an oscillation) has a very limited impact, considering the range of injection times from 20ms before the oscillation peak to 10ms after. The most effective injection time was the first 2-5 ms after the peak of an oscillation. Given this feature, it was possible to minimize the number of current injections used, only injecting current after oscillations more than two standard deviations below the mean oscillatory prominence. This strategy is therefore not exclusive to computer simulation - because the injection period is after the oscillatory peak, it would be possible to determine when to inject *in vivo* without leveraging the omniscient perspective of the simulation. Each injection does slightly increase the subsequent inter-oscillatory interval by a few milliseconds, but this is still within the range of inter-oscillatory interval variability seen over the course of trials without secondary injections.

Function 5: *NEST Code to Produce a Fixed Number of Oscillations*

```
1  def fixed_number_oscillations(targ_num):

2      sim = Simulation(###Parameters###)

3      first = sim.run_micro()
4      current = len(first)
5      stimuli=[1000.]

6      while current != targ_num:
7          sim = Simulation(###Parameters###)
8          no_gaps = sim.run_micro()
9          current = len(no_gaps[0])

10         if current > targ_num-2:
11             last = no_gaps[0][-1*(current-targ_num)-2]
12             stimuli.append(last+1005.)
13             sim = Simulation(###Parameters###)
14             no_gaps = sim.run_micro()
15             break
16         else:
17             stimuli.append(no_gaps[0][-1]+1002.)

18     sim = Simulation(###Parameters###)
19     sim.run_micro()
```

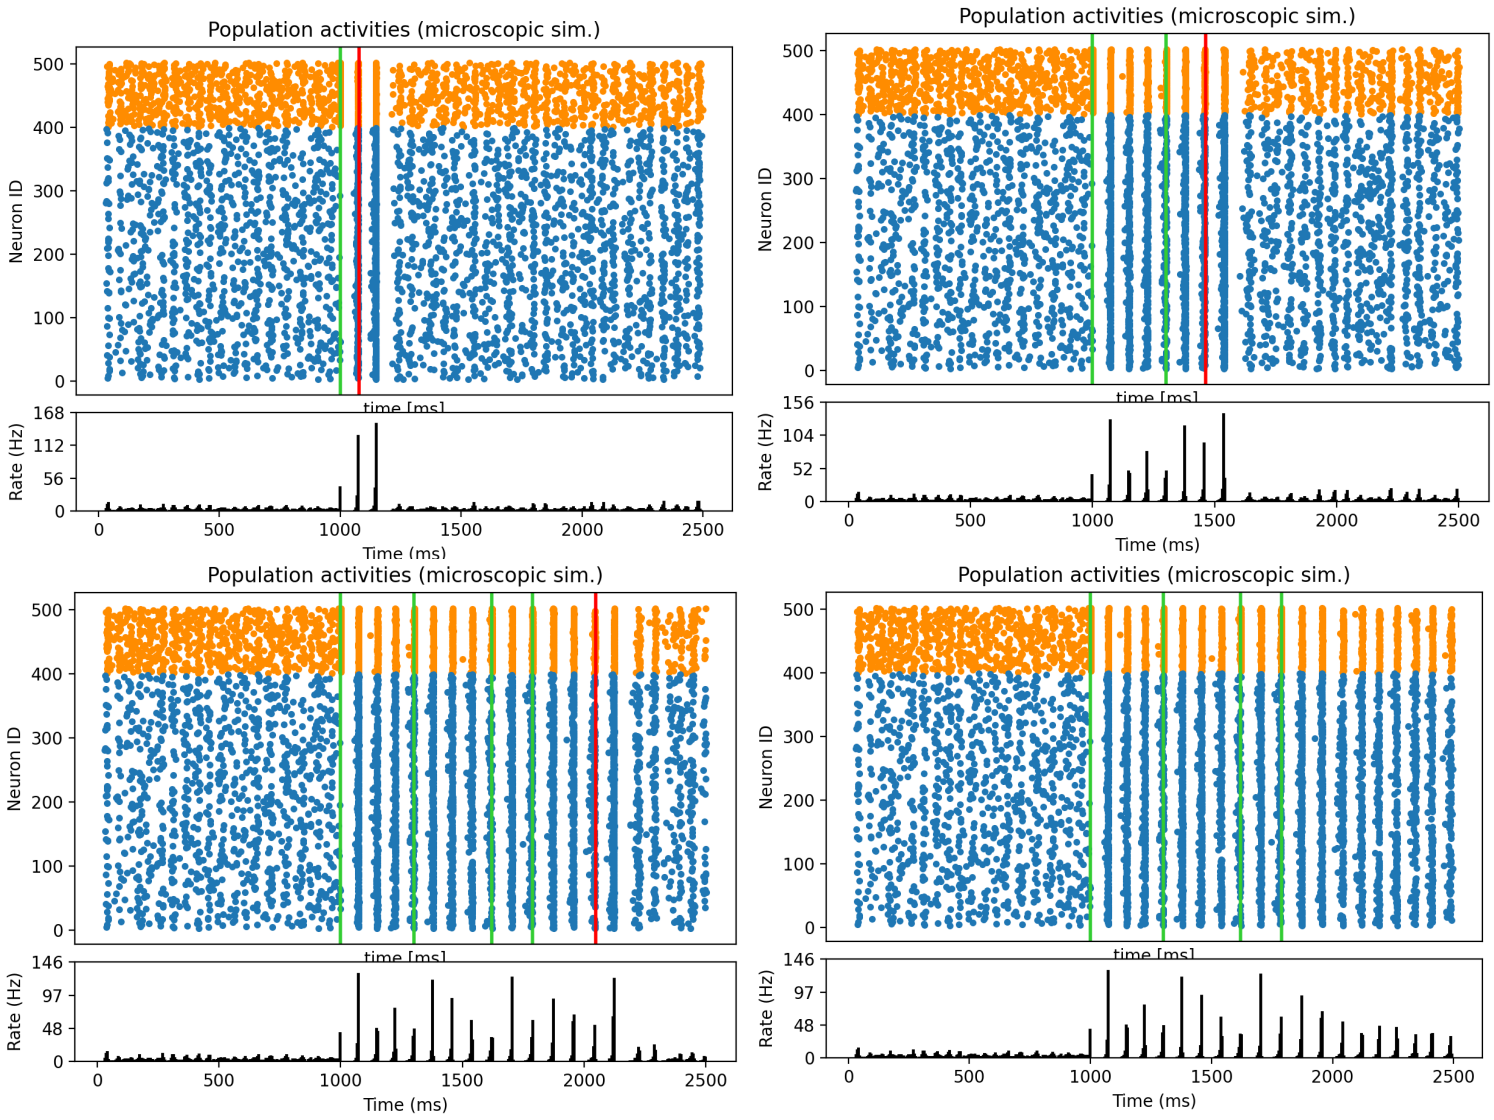

Figure 10 - *example simulations, using the same random seed, of oscillatory responses with a target length of 2 (top left), 7 (top right), and 14 (bottom left) periods. Green lines represent biphasic current injections, and the final red line in each simulation is a hyperpolarizing current. (Bottom right) current stimuli are identical to bottom left, but with no hyperpolarizing current.*

The second feature required was how to exit the oscillatory state once the target duration had been reached. Injection of a hyperpolarizing current to the inhibitory population returns the network to the asynchronous state, although the amplitude of this current also varies from trial to trial (Figure 10). If it is excessively large, it will simply initiate further oscillation; strong hyperpolarizing current to the entire inhibitory population causes a burst in excitation, further perpetuating oscillations. The hyperpolarizing current can be much smaller than the depolarizing used to initiate the oscillations - about 1 nA versus10 μA. This was further scaled by the

prominence of the penultimate oscillation, in order to avoid bursts of excitation. Although this is more effective than scaling by a fixed value, it does not perform perfectly, particularly in longer simulations - other methods of scaling the final current injection should be explored further.

6 Conclusion

Two features were required to achieve sustained oscillations in these spiking neural network simulations. First, the network had to be wired with dense recurrent and inter-population connectivity. Second, the stimulus had to be of sufficient intensity, with a significant amount of background activity, in order to produce oscillating activity. A damped oscillator is an apt metaphor for this behavior - a spring must have some recoil, and a pull to the spring must be of a sufficient strength in order for oscillations to occur. Noise in the network acts like friction to the spring. Increased noise into the system decreases the intensity and duration of the oscillations, however small adjustments to features of the single current injection do not reliably affect the same response. Additionally, there is a plateau in the intensity of network response to stronger stimuli. Stimulus duration of more than 5 ms never produced oscillations with the desired structure. Network response also stopped evolving for stimuli with an intensity greater than $50\ \mu A$, indicating an upper bound in the ability of a network to integrate extremely high amplitude currents.

The most effective way to control oscillations was using multiple network simulations. By injecting current slightly after the oscillatory peak, the behavior could be extended indefinitely, or terminated abruptly using a slightly hyperpolarizing injection. In the future, this strategy can be implemented in vivo, as the timing of these injections can be determined in an online fashion.

6.1 Future Directions

These simulations will be used on two fronts. They will be employed by the University of Padua to explore and optimize stimulation strategies before expending biological resources on the task. They are further being used as a reinforcement learning environment, to explore strategies of changing the stimulus to cause network response to better meet a desired behavior, or to optimize the simulated behavior to be closer to the biological activity.

In addition to these two immediate projects, expanding the model to incorporate more diverse inhibition would be an interesting test of the nature and timescale of this behavior. Layer 4 of the barrel cortex receives intense feedforward inhibition from the Thalamus, targeting fast spiking (likely Parvalbumin) inhibitory cells [12]. The longer response time of other subpopulations (SST or ionotropic) may diminish this behavior, or by a source of further synchronization with their inhibitory-neuron targeted inhibition. It would also be interesting to explore adding new cortical layers to the model, to investigate how the oscillations spread across these distinct populations, and compare to the cross-layer data captured in the depth electrode recordings. Incorporating dynamic synapses, where the strength of the synapse changes based on

recent spiking history, would also be a valuable area to pursue, in an effort to more accurately capture the rise and decay of oscillation intensity [34]. The most efficient approach to answer these questions would be to implement these additional features in the mesoscopic model.

7 Acknowledgements

My research has been made possible by a number of mentors and advisors. First, I would like to thank the Fulbright-Austrian Commission and the Fulbright-Marshall Plan Foundation. Without these programs I would not have been able to pursue any of my work here, or enjoyed the spectacular experience that was living in Austria. The staff at the Fulbright Commission, particularly Darrah Lustig and Herman Agis, persisted through difficult conditions, to provide adapted versions of many quintessential Fulbright experiences during this year of coronavirus, and I am very thankful for their support. I would also like to thank Professor Robert Legenstein, for his continued patience and advice over the course of this project, and for his help and time in making this program possible by acting as my supervisor. Horst Pentsching has offered valuable advice, and generously provided initial code for both my experiments with the Allen Brain Institute cortical model and my efforts in optimization and reinforcement learning on this project. Professor Vassanelli at the University of Padua and his graduate student Mattia Tambaro have generously shared their depth electrode recordings, as well as thoughts and advice on the underlying cause of these oscillations. Additionally, I would like to thank the staff and students at the IGI: Daniela Windisch-Scharler, Oliver Friedl, Ceca Kraisnikovic, Dr. Spela Brglez, Dr. Ozan Özdenizci, Dr. Guozhang Chen, Dr. Michel Müller, Christoph Stöckl, Maximilian Baronig, Romain Ferrand, Titouan Cormier, Florian Unger, and Eben Kadile. Finally, I would like to thank my undergraduate Principal Investigator Dr. Jason MacLean, who encouraged me to apply to the Fulbright program.

8 References

1. Amit, Daniel J., and Nicolas Brunel. "Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex." Cerebral cortex (New York, NY: 1991) 7.3 (1997): 237-252.
2. Aronoff, Rachel, et al. "Long‐range connectivity of mouse primary somatosensory barrel cortex." European Journal of Neuroscience 31.12 (2010): 2221-2233.
3. Avermann, Michael, et al. "Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex." Journal of neurophysiology 107.11 (2012): 3116-3134.
4. Barral, Jérémie, and Alex D. Reyes. "Synaptic scaling rule preserves excitatory–inhibitory balance and salient neuronal network dynamics." Nature neuroscience 19.12 (2016): 1690-1696.
5. Billeh, Yazan N., et al. "Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex." Neuron 106.3 (2020): 388-403.
6. Boccaletti, Stefano, et al. "Complex networks: Structure and dynamics." Physics reports 424.4-5 (2006): 175-308.
7. Bojanek, Kyle, Yuqing Zhu, and Jason MacLean. "Cyclic transitions between higher order motifs underlie sustained asynchronous spiking in sparse recurrent networks." PLoS computational biology 16.9 (2020): e1007409.
8. Börgers, Christoph, and Nancy Kopell. "Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity." Neural computation 15.3 (2003): 509-538.
9. Brunel, Nicolas. "Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons." Journal of computational neuroscience 8.3 (2000): 183-208.
10. Bruno, Randy M., and Daniel J. Simons. "Feedforward mechanisms of excitatory and inhibitory cortical receptive fields." Journal of Neuroscience 22.24 (2002): 10966-10975.
11. Chizhov, Anton V., and Lyle J. Graham. "Efficient evaluation of neuron populations receiving colored-noise current based on a refractory density method." Physical Review E 77.1 (2008): 011910.
12. Cruikshank, Scott J., Timothy J. Lewis, and Barry W. Connors. "Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex." Nature neuroscience 10.4 (2007): 462-468.
13. Dai, Kael, et al. "Brain Modeling ToolKit: An open source software suite for multiscale modeling of brain circuits." PLOS Computational Biology 16.11 (2020): e1008386.
14. Eurocord 2020. “A synaptically connected brain-silicon neural closed-loop hybrid system. Objectives | A SYnaptically connected brain-silicon Neural Closed-loop Hybrid system.” Retrieved December 12, 2021, from https://synch.eucoord2020.com/objectives/.
15. Fan, Xue, and Henry Markram. "A brief history of simulation neuroscience." Frontiers in neuroinformatics 13 (2019): 32.

16. Feldmeyer, Dirk. "Excitatory neuronal connectivity in the barrel cortex." Frontiers in neuroanatomy 6 (2012): 24.
17. Gast, Richard, et al. "PyRates—A Python framework for rate-based neural simulations." PloS one 14.12 (2019): e0225900.
18. Gouwens, Nathan W., et al. "Systematic generation of biophysically detailed models for diverse cortical neuron types." Nature communications 9.1 (2018): 1-13.
19. Hebb, Donald Olding. The organization of behavior: A neuropsychological theory. Psychology Press, 2005.
20. Hodgkin, Alan L., and Andrew F. Huxley. "A quantitative description of membrane current and its application to conduction and excitation in nerve." The Journal of physiology 117.4 (1952): 500-544.
21. Jiang, Xiaolong, et al. "Principles of connectivity among morphologically defined cell types in adult neocortex." Science 350.6264 (2015).
22. Kamps, M. de. "A simple and stable numerical solution for the population density equation." Neural computation 15.9 (2003): 2129-2146.
23. Koelbl, Christian, et al. "A barrel-related interneuron in layer 4 of rat somatosensory cortex with a high intrabarrel connectivity." Cerebral Cortex 25.3 (2015): 713-725.
24. Lefort, Sandrine, et al. "The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex." Neuron 61.2 (2009): 301-316.
25. Marc-Oliver Gewaltig and Markus Diesmann (2007) “NEST.” Scholarpedia, 2(4):1430.
26. Mensi, Skander, et al. "Parameter extraction and classification of three cortical neuron types reveals two distinct adaptation mechanisms." Journal of neurophysiology 107.6 (2012): 1756-1775.
27. Mulansky, Mario, and Thomas Kreuz. "PySpike—A Python library for analyzing spike train synchrony." SoftwareX 5 (2016): 183-189.
28. Pedregosa, Fabian, et al. "Scikit-learn: Machine learning in Python." the Journal of machine Learning research 12 (2011): 2825-2830.
29. Ponce-Alvarez, Adrian, et al. "Cortical state transitions and stimulus response evolve along stiff and sloppy parameter dimensions, respectively." Elife 9 (2020): e53268.
30. Pozzorini, Christian, et al. "Automated high-throughput characterization of single neurons by means of simplified spiking models." PLoS computational biology 11.6 (2015): e1004275.
31. Purves, D., et al. (2018) “Neuroscience.” 6th Edition, Sinauer Associates, New York.
32. Reimann, Michael W., et al. "An algorithm to predict the connectome of neural microcircuits." Frontiers in computational neuroscience 9 (2015): 28.
33. Rosenbaum, Robert, et al. "The spatial structure of correlated neuronal variability." Nature neuroscience 20.1 (2017): 107-114.
34. Schmutz, Valentin, Wulfram Gerstner, and Tilo Schwalger. "Mesoscopic population equations for spiking neural networks with synaptic short-term plasticity." The Journal of Mathematical Neuroscience 10.1 (2020): 1-32.
35. Schwalger, Tilo, Moritz Deger, and Wulfram Gerstner. "Towards a theory of cortical columns: From spiking neurons to interacting neural populations of finite size." PLoS computational biology 13.4 (2017): e1005507.

36. Spruston, Nelson. "Pyramidal neurons: dendritic structure and synaptic integration." Nature Reviews Neuroscience 9.3 (2008): 206-221.
37. Teeter, Corinne, et al. "Generalized leaky integrate-and-fire models classify multiple neuron types." Nature communications 9.1 (2018): 1-15.
38. Tremblay, Robin, Soohyun Lee, and Bernardo Rudy. "GABAergic interneurons in the neocortex: from cellular properties to circuits." Neuron 91.2 (2016): 260-292.
39. Venkatraman, Subramaniam, and Jose M. Carmena. "Behavioral modulation of stimulus-evoked oscillations in barrel cortex of alert rats." Frontiers in integrative neuroscience 3 (2009): 10.
40. Virtanen, Pauli, et al. "SciPy 1.0: fundamental algorithms for scientific computing in Python." Nature methods 17.3 (2020): 261-272.
41. Whittington, Miles A., et al. "Inhibition-based rhythms: experimental and mathematical observations on network dynamics." International journal of psychophysiology 38.3 (2000): 315-336.