Predictive routing emerges from self-supervised stochastic neural plasticity

Author Names and Affiliations: Hamed Nejat[1], Jason Sherfey[2], André M. Bastos[1,3]

[1]: Department of Psychology, Vanderbilt University, Nashville, TN, USA

[2]: Department of Psychological and Brain Sciences, Boston University, Boston, MA, USA

[3]: Vanderbilt Brain Institute, Vanderbilt University, Nashville, TN, USA

^Correspondence: Hamed Nejat (hamed.nejat@Vanderbilt.Edu) and André M. Bastos (andre.bastos@Vanderbilt.Edu)

Conflict of interest statement: None.

Acknowledgements: RO0MH116100(AMB), VU-startup funds (AMB)

Abstract

Neurophysiology studies propose that predictive coding is implemented via alpha/beta (8-30 Hz) rhythms that prepare specific pathways to process predicted inputs. This leads to a state of relative inhibition, reducing feedforward gamma (40-90 Hz) rhythms and spiking to predictable inputs. We refer to this model as predictive routing. It is currently unclear which circuit mechanisms implement this push-pull interaction between alpha/beta and gamma rhythms. To explore how predictive routing is implemented, we developed a self-supervised learning algorithm we call generalized Stochastic Delta Rule (gSDR). It was necessary to develop this learning rule because manual tuning of parameters (frequently used in computational modeling) is inefficient to search through a non-linear parameter space that defines how neuronal rhythms emerge and interact. We used gSDR to train biophysical neural circuits and validated the algorithm on simple tasks, e.g., tuning membrane potentials and firing rates. We next applied gSDR to model observed neurophysiology. We asked the model to reproduce a shift from baseline oscillatory dynamics (~<20Hz) to stimulus induced gamma (~40-90Hz) dynamics recorded in the macaque monkey visual cortex. This gamma-band oscillation during stimulation emerged by self-modulation of synaptic weights via gSDR. We further showed that the gamma-beta push-pull interactions implied by predictive routing could emerge via stochastic modulation of both local inhibitory circuitry as well as top-down modulatory inputs to a network. To summarize, we implemented gSDR to train biophysical neural circuits based on a series of objectives. gSDR succeeded in implementing these objectives. This revealed the inhibitory neuron mechanisms underlying the gamma-beta push-pull dynamics that are observed during predictive processing tasks in systems and cognitive neuroscience.

Significant Statement

This study contributes to the advancement of self-supervised learning for modeling the behavior of complex neural circuits and specifically, biophysical models based on predictive routing framework. Since gSDR is an evolutionary algorithm and does not rely on specific model-based assumptions, it could improve autonomous approaches both in computational neuroscience and neural network research.

Introduction

Empirical observations of neuronal activity during tasks that involve processing predictable and unpredictable sensory inputs support a predictive routing framework[1]. The framework proposes that predictive coding[2] may be implemented via low-frequency (8-30Hz, Alpha/Beta) rhythms that modulate pathways processing predicted inputs by inhibiting feedforward high-frequency (>40Hz, Gamma) rhythms and spiking activity. This leads to enhanced bottom-up gamma and suppressed top-down alpha/beta during unpredictable states, and enhanced top-down alpha/beta suppressing bottom-up gamma during predictable stimuli. Further, this top-down beta-rhythmic predictive suppression of gamma

has been proposed to be highly specific to the cortical locations that represent the predicted stimulus.

The predictive routing model has recently received causal support via a study that investigated auditory oddball responses (stimuli that deviate from a predictable pattern) during propofol-induced anesthesia[3–6], which eliminated neuronal responses to stimuli in prefrontal cortex but retained stimulus processing in sensory areas. During propofol-induced anesthesia, oddball responses in auditory cortex were no longer modulated by beta suppression. As a result of the loss of prefrontal inputs and beta modulation during propofol, neuronal activity in sensory cortex became disinhibited[3]. These data support the proposal of top-down beta acting in a suppressive manner, as the removal of this activity resulted in enhanced sensory oddball-related gamma and spiking.

It remains an open question how this predictive routing mechanism is achieved in cortical circuits. One possibility is that top-down inputs target circuitry in the receiving areas that generate beta-rhythmic activity[7,8] that suppresses activity locally. Another possibility is that alpha/beta rhythms modulate gamma activity via phase-amplitude coupling[9–11] to implement predictive suppression at specific phases of the lower frequency. These proposals rely on the interaction between specific populations of inhibitory interneurons[12–14] across layers with pyramidal excitatory cells, but the exact biophysical architecture that allows one rhythm (e.g., beta) to act in a suppressive way on another (e.g., gamma) remain largely unknown.

Previous computational models have implemented certain aspects of the predictive routing model. A model of V1 and V4 oscillatory coherence implemented with Dynamic Causal Modeling was able to fit observed electrophysiological recordings[15]. However, this model did not perform predictive coding as it was not trained on a task, and also did not simulate a push-pull dynamic. A subsequent model was able to accurately model attention-dependent changes in coherence between these areas based on modulation of local inhibitory interneuron connectivity in V1[16]. Another modeling study simulated gamma and beta by incorporating different inhibitory time constants in superficial vs. deep layers, respectively[13]. This model was able to accurately capture the push-pull between top-down beta and bottom-up gamma. But once again, prediction as such was not explicitly modeled. A more recent model extended these previous mass models by incorporating a physical dimension, enabling the model to both express aspects of predictive routing (cross-frequency coupling and push-pull relationships between alpha/beta and gamma) and infer locations of synaptic inputs to the gamma/beta generating circuits[17]. This model also did not explicitly consider predictive processing and task modulation elements.

The neuronal mass models reviewed above make the assumption that the necessary biophysics to model cognitive processes exist at the level of neuronal populations, which does not enable these models to express more detailed aspects of biophysics, such as neuron to neuron connectivity. To overcome this, research using a detailed laminar circuit implementation studied attentional modulation, which was assumed to come from a top-down beta-rhythmic source[18]. This top-down beta was targeted to specific cell types that could enhance gamma at the attended location and weaken it at the unattended location, which replicated previous experimental observations[19]. However, this model was

fit by hand-tuning of parameters, making it difficult to search the full parameter space and precluding fitting to specific tasks. The model also did not consider predictive processing, in which a push-pull relationship between beta and gamma has been more commonly observed [1,3]

To summarize, the modeling literature has investigated aspects of beta-gamma predictive routing like computations, but it has either used these to investigate attention and have not explicitly considered prediction in the modeling framework using more detailed biophysical models.

To overcome these challenges, in this work we introduce a learning algorithm we call generalized stochastic delta rule (gSDR, Fig. 1) and apply it to train biophysically-detailed neural models to capture electrophysiologically observed cortical dynamics. The goal of this learning algorithm is to train biophysically detailed models (models consisting of single cells and their synaptic connections) in order to gain insight into complex circuit dynamics, such as beta-gamma push-pull, and how they arise. In this work, we begin by establishing that this learning rule can be applied to modulate both simple (e.g., membrane potentials and firing rates) and more complex network properties (e.g., modulate population activity and oscillations) of biophysical neurons modeled with Hodgkin-Huxley equations. We then apply this modeling framework to understand several questions pertinent to predictive routing. We first consider how a cortical network flexibly switches between oscillatory motifs (e.g, from beta to gamma) by modulating local (within-area) synaptic connectivity between inhibitory and excitatory interneurons. Second, we consider how modulation of top-down synaptic connections to a circuit can induce similar spectral shifts between beta and gamma. Finally, we apply the model to observed brain dynamics by training a model of biophysical neurons on high-density electrophysiology data (dense, in-vivo, spiking recordings performed in area MT/MST during presence of stimulus) acquired from awake behaving non-human primates (NHPs). We were able to train the model to match the network transition from low-frequency activity during baseline to gamma oscillations during the presence of stimuli. Together, these modeling results advance our understanding of the biophysical mechanisms that could implement predictive routing computations in cortex.

Materials and Methods

One of the challenges in biophysically detailed models is to develop automated methods for training these neural models to help generate insights into the neural circuit dynamics that underlie cognition. We used the DynaSim toolbox in Matlab to implement the neural models[20]. Our model consists of a network of neurons with distinct cell types modeled by corresponding Hodgkin-Huxley circuit equations. In addition, some other mechanisms (e.g, receptors & ion channels) have been added to the Hodgkin-Huxley equation model. The system of equations that govern how voltage changes over time in neurons of a neural population is given in equation 1:

$$C_m \frac{dV}{dt} = -I_{inp}(t, V) + -\sum I_{int} - \sum I_{syn} \quad [1]$$

Where t is time (ms), C_m is the membrane capacitance, I_{int} denotes the intrinsic membrane currents (such as I_{Na}, I_K, I_{Leak}), $I_{inp}(t, V)$is the current reflecting inputs from external sources and I_{syn} denotes synaptic currents from the other neurons driving this neuron. Using this framework, we are able to define detailed neuronal models with multiple populations of similar or different cell types and various synaptic connection mechanisms (such as I_{AMPA}, I_{NMDA}, I_{GABA}, …) between them.

gSDR's unique contribution is making DynaSim models trainable, specially to reproduce complex observed in-vivo dynamics with the least dependence on supervision. To accomplish this, we implemented an extension to DynaSim called DynaLearn. DynaLearn takes any DynaSim model, and designs a training environment for it. DynaLearn efficiently explores a DynaSim's model parameters in order to achieve a single or multiple objectives. These objectives are user-defined metrics that are derived from the model (e.g, membrane potential, firing rates and spectral population dynamics). There are various optimization or learning methods in DynaLearn but here we focus on the generalized Stochastic Delta-Rule (gSDR) which is inspired from the stochastic delta rule[21], spike-timing dependent plasticity[22] and genetic-optimization learning[23].The general form of our algorithm is shown in equation 2:

$$U_t = U_{t-\Delta t} + (L(t)\,\delta(\lambda)(1 - \alpha) + \alpha)\,.\,R \quad [2]$$

Where U_t represents all variables at time t that gSDR allows to change, $\delta(\lambda) \sim uniform(-\lambda, \lambda)$ (λ : exploration factor) is a random sample from a uniform distribution, $L(t)$ is the output of the evaluation function (see methods, defining objectives in gSDR) at that time, α is the unsupervised factor (if $\alpha = 0$, the model will no longer have unsupervised changes and if $\alpha = 1$, the model will be fully unsupervised) and R is the mutual-correlation dependent plasticity (MCDP) factor around the step time ($t - \Delta t$) based on $V_{t \times N}$ which is a matrix containing the membrane potential of all N neurons. This function acts as an approximation of spike-timing dependent plasticity (STDP) depending on the time length Δt. Details are in this equation 3:

$$R = MCDP(V_t, \Delta t) = \frac{r - \mu_r}{\sigma_r}, \; r_{N \times N} = Corr(V_{t-\Delta t}, V_t) \quad [3]$$

Consider that there are two components in gSDR; one is the value of loss or reward that affects how the model will change its parameters (semi-supervised, loss/reward feedback is provided yet parameter changes are not directly controlled by supervision) and the other is the MCDP part that is completely unsupervised. In addition, the value of R will be 1 for any variable in U_t that is not a synaptic weight, in case if other model parameters such as channel properties are included in the simulation. This means that only synaptic weights can be changed by MCDP (Fig.1).

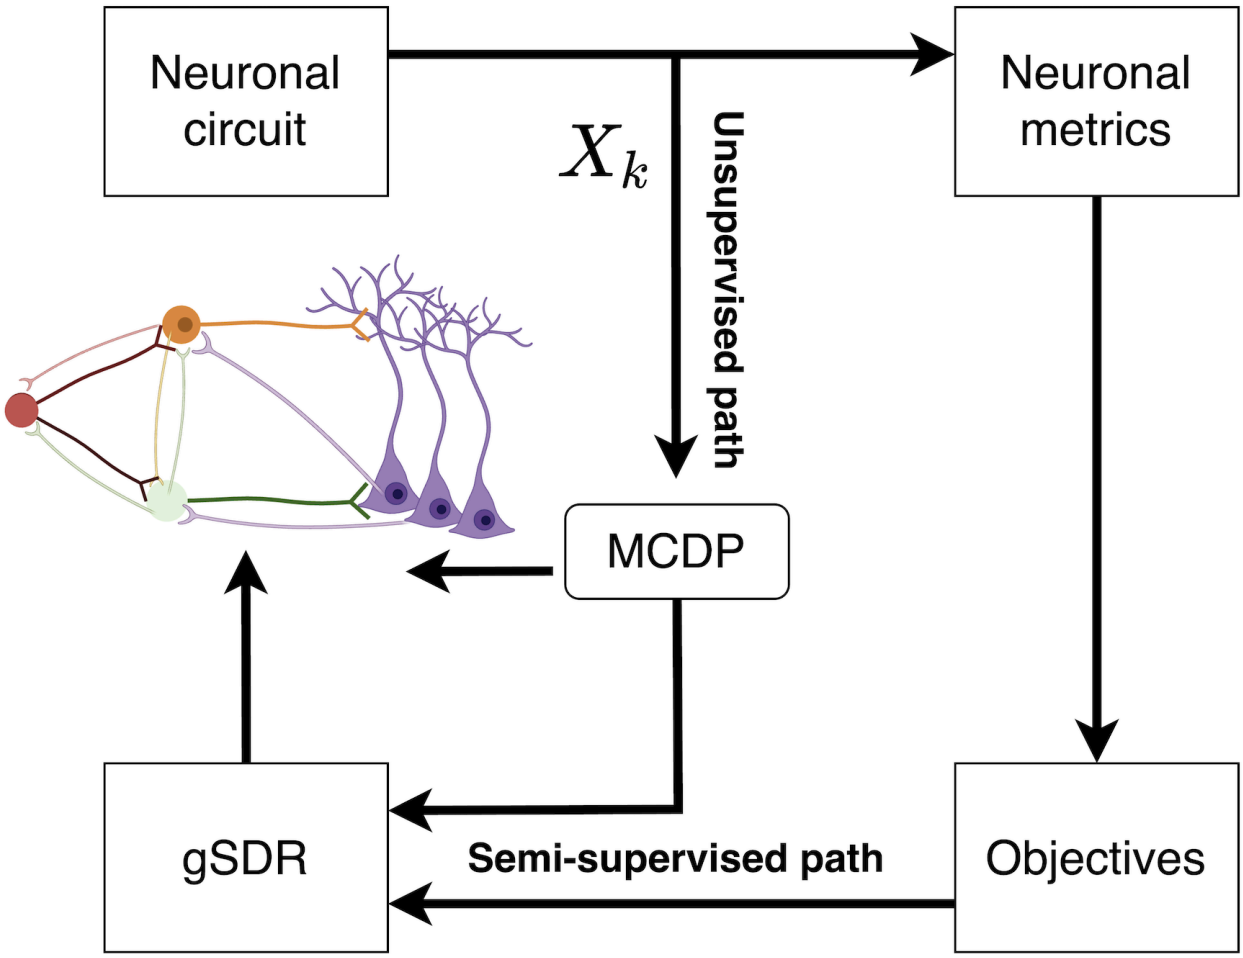

Fig1: Simplified gSDR-DynaSim interaction flowchart. A DynaSim model (Neuronal circuit) generates a series of neural responses represented by X[k]. In the unsupervised path these voltages are used to derive synaptic co-activation patterns calculated by MCDP. This specifies a learning rule independent from any external objective that enhances connections between co-active neurons. In addition, the semi-supervised path calculates the neuronal metrics and their distance to prespecified objectives, which defines the loss or reward value for the current trial. Generalized stochastic delta rule (gSDR) uses this loss or reward value(s) to change the model's parameters U(t). Over time, the model explores many possible parameter spaces and keeps the most optimal parameters minimizing the loss or maximizing the reward and thereby achieving the objectives.

Algorithm 1 Generalized stochastic delta rule

$U, \lambda, \alpha, L_{opt} \leftarrow Init(.)$ ▷ Model construction/initialization
while $Training$ **do**
  $M \leftarrow Metrics(U)$ ▷ Measurements from the model
  $X \leftarrow Reponse(U)$ ▷ Membrane potential responses $(Neurons \times Time)$
  $L(t) \leftarrow Evaluate(M, \hat{Targets})$ ▷ Model evaluation
  $R \leftarrow Corr(X_t, X_{t-\tau})$ ▷ Pairwise temporal correlation
  if $L \leq L_{opt}$ **then** ▷ Update optimal state
    $L_{opt} \leftarrow L$
    $U_{opt} \leftarrow U$
  else if $L \geq C \times L_{opt}$ **then** ▷ Divergence condition
    $U \leftarrow U_{opt}$ ▷ Return to the last optimal state
  else
    for $W_{pre,post} \in U$ **do** ▷ Synaptic weights
      $\Delta \leftarrow (L.\delta(\lambda)(1-\alpha)+\alpha).R$ ▷ $\delta(\lambda) \sim u(-\lambda, \lambda)_{n \times m}$
      $W_{new} \leftarrow W + \Delta$
      $W \leftarrow Adj(W_{new})_{[0,1]}$ ▷ Adjust synaptic weights
    end for
  end if
end while

gSDR pseudocode

Reinforcing objectives such as specific firing rates, spectral or temporal patterns requires defining the corresponding metric for each one. A metric is a direct or indirect measurement from the model, similar to a probe recording in-vivo neural signals. gSDR metrics are derived from any parameter that exists in a DynaSim model's responses. A metric generally consists of three arguments; spatial annotation (which part of the circuit to be measured), temporal annotation (when in time points or intervals of simulation should be selected for this measurement) and the measurement function (what aspect of the annotated model output should be measured). In our simulations we have utilized both metrics that have a biophysical unit (current and voltage) as well as metrics that are derived from the model's response (firing rate and spectral response). Algorithm 1 describes how gSDR uses these metrics to train models.

Objective(s) consisted of one or multiple components. The loss function in the single neuron simulations (Fig 2) and the initial population simulations (Fig 3-4) were calculated by the mean square error (MSE) between measured metric and the target. In the neuronal population firing rate tuning simulation (Fig 4) two objective components were used. First, MSE between the excitatory population's firing rate and the target (70 Spike/sec) and second, MSE between the inhibitory population's firing rate and the target (20 Spike/sec). For the population natural frequency tuning (Fig 5), the loss was calculated by the MSE between the measured natural frequency (the frequency with the peak power in the power spectrum response) and the assigned target (30Hz). In the population spectral relative power tuning (Fig 6) the error was calculated using binary error (0-1) for the model which

in this case means that a reward value of 1 was assigned when beta was greater than gamma and otherwise, 0.

In the last simulations comparing the spectral response of the model to the spectral response of the data (Fig 10-12), we defined a loss function (eqn. 4) by calculating the log ratio of the model's spectral response to the target spectral response, after normalization. Both the model and target spectral response was computed by the taking the FFT of each individual neuron's spiking timecourse and then averaging across neurons in the population.

$$L = LogRatio(Metric_t, Target_t) = \sum_{i \in Metric} Log(\frac{Metric_{i,t}}{Target_{i,t}}) \quad [4]$$

In biophysical neural models, it is important to consider the plausibility of objectives, as some target states for this type of neural circuits may violate biological assumptions (such as non-biological membrane potentials or very high firing rates). In such cases, the model may not be able to achieve certain targets. It is also possible that gSDR brings the model out of biological constraints in case if there are no biologically plausible boundary conditions for training or tuning variables. Thus, we have made it possible to define boundary conditions for model variables. Boundary conditions have been used to keep neuronal circuits biologically plausible and stable in gSDR.

Here, we illustrate some basic examples of gSDR utilization. These models are designed in order to illustrate the gSDR and the overall logic of the method (Fig.1). In the first single neuron simulation example, the model was trained to modulate the membrane potential from an initial starting point of -64.9mV (model's initial trial response is shown in Fig. 2C) to a target of -60.0mV. The model was only allowed to tune the reversal potential of the leak channel (E-leak, Fig. 2A) during training. As can be seen in Fig. 2A, comparing pre to post training, the membrane potential shifted to the desired target. In Fig. 2B it can be seen that the model's loss (mean square error between target and measured membrane potential) decreased over trials despite transient divergence. At trial 151, the model was manually stopped due to a relatively negligible loss.

A

Variables(*V*) / Metrics(*M*)	Initial	Final
V: ***E***(Leak, mV)	-54.4	-26.0
M: Average membrane voltage (mV)	-64.9	-60.0

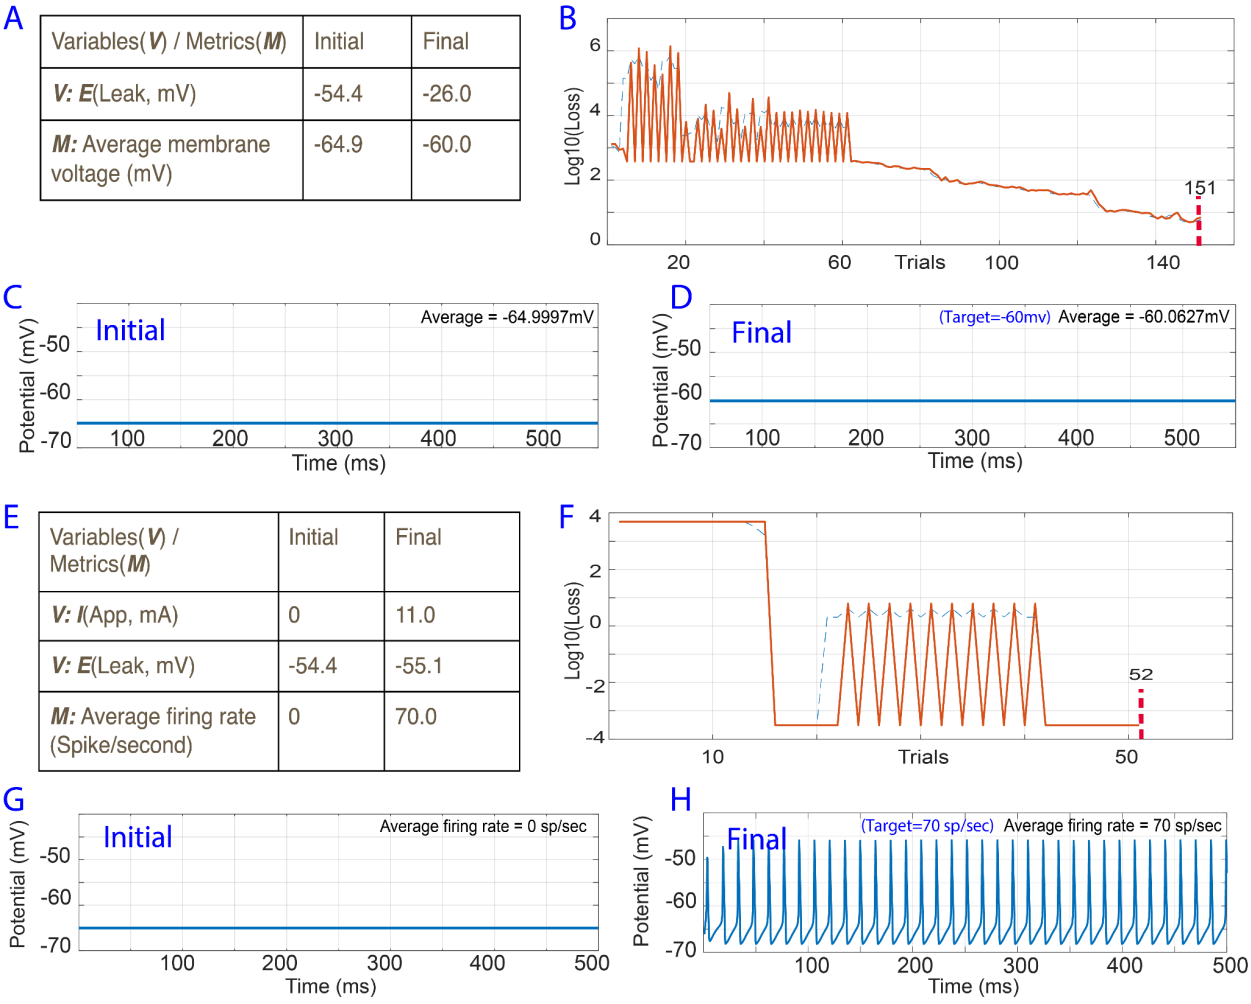

E

Variables(*V*) / Metrics(*M*)	Initial	Final
V: ***I***(App, mA)	0	11.0
V: ***E***(Leak, mV)	-54.4	-55.1
M: Average firing rate (Spike/second)	0	70.0

Fig2: Single cell basic optimizations. (A) Summary of initial and final training variables and metrics (B) Loss in trials corresponding to the target metrics. (C) Initial membrane potential response. (D) Final membrane potential response. (E) Summary of initial and final training variables and metrics for firing rate optimization. In this simulation both the reversal potential of the leak channel (Eleak) and the internal drive (Iapp) are variables that change in training. (F) Loss in trials corresponding to the target metrics. (G) Neuron response (Initial) (H) Neuron response (final)

In the second single neuron simulation, the model was trained to modulate the firing rate from an initial starting point of 0 spike/sec to a target of 70 spike/sec. The model was only allowed to tune the reversal potential of the leak channel (E-leak) and an external constant current drive during training. As can be seen in Fig. 2E, comparing pre to post training, the firing rate approached the desired target. In Fig. 2F it can be seen that the model's loss (mean square error between target and measured firing rate) reduced after exploratory behavior (Fig. 2F). At trial 52, the model was manually stopped due to a relatively negligible loss.

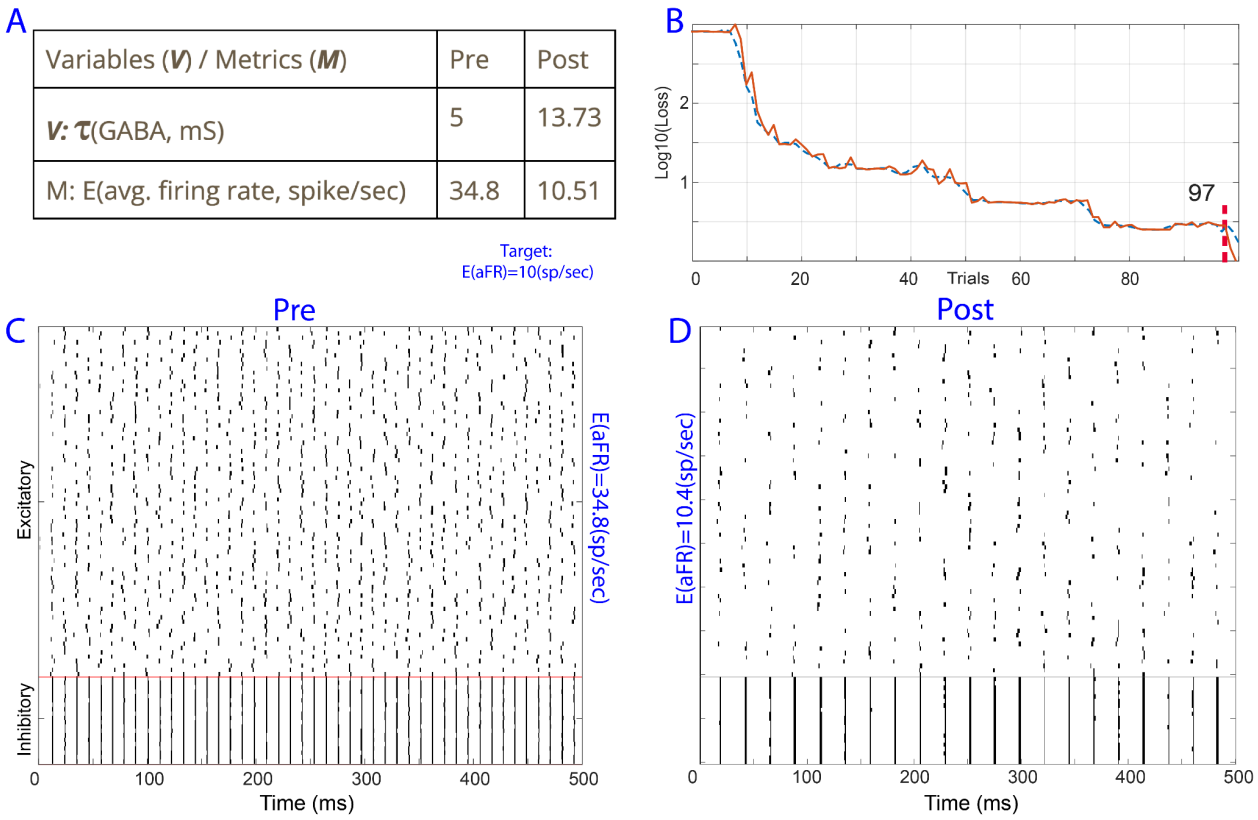

Fig3: E-I (Excitatory-Inhibitory) neuronal optimization of population firing rate and membrane potential. (A) Summary of initial and final training variables and metrics. (B) Loss in trials corresponding to the target metrics, firing rate and membrane potential. (C) Initial spiking response as a raster plot. Each row corresponds to a single neuron's spiking during one trial (500ms). (D) Raster plot after training.

Next, we show how to simulate neuronal populations. For this we introduced a model containing 100 neurons (80 excitatory, 20 inhibitory) and additional mechanisms (Glutamate, GABA) and parameters for their connectivity (synaptic connections, synaptic time constants, E-noise and I-noise). The model was initialized with a uniform connectivity matrix between all populations. The first objective was to train the model to achieve a specific average firing rate in the excitatory neuron population. The initial firing rate was 34.8 spike/sec (Fig. 3C) and the target was 10 spike/sec. Only inhibitory synaptic time constants were allowed to be changed by the model. As seen in fig. 3, the model achieved the target average firing rate in the excitatory population by increasing inhibitory time constants which effectively decreases the overall spiking activity.

Neural population model (dual firing rate tuning): In this example (Fig.4) we used the same model architecture as the first E-I simulation (Fig.3). We trained this model using two parallel objectives to demonstrate the flexibility of gSDR to explore a multi-target parameter space. The equally weighted objectives were to achieve an average firing rate of 20 spike/sec and 70 spike/sec for excitatory and inhibitory populations, respectively. The initial firing rates of the excitatory and inhibitory populations were 34.8 spike/sec and 91.81 spike/sec respectively (Fig. 4A). Similar to the previous simulation, only inhibitory

synaptic time constants were allowed to be changed by the model. As seen in fig. 4, the model achieved the target average firing rates by increasing inhibitory time constants

A

Variables(***V***) / Metrics(***M***)	Pre	Post
V: τ(GABA, ms)	5	8.05
M: Avg. firing rate *(Exc., spike/sec)*	34.83	19.61
M: Avg. firing rate *(Inh., spike/sec)*	91.81	71.41

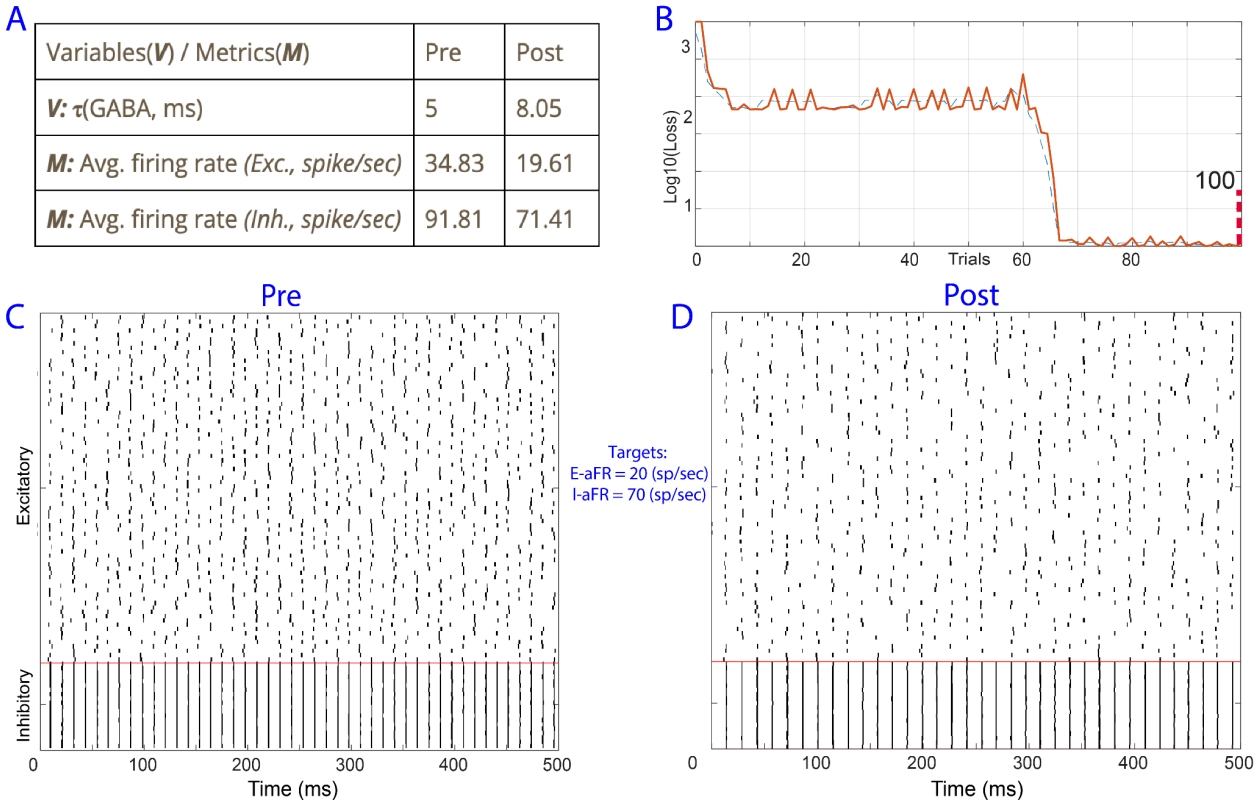

Fig4: E-I (Excitatory-Inhibitory) neuronal population multitarget firing rate optimization on the same model in Fig4. (A) Summary of initial and final training variables and metrics. In this simulation, the firing rates of E and I populations should converge to their respective objectives (Target firing rate of 70 spikes per second for I population and 20 for E population. (B) Loss in trials corresponding to the target metrics. Each population's firing rate converges to the desired target. (C) Initial raster plot. (D) Raster plot after about 100 trials.

Although it is trivial for the model to decrease the firing rate of both populations by increasing the inhibitory synaptic time constant, the relation between the time constant and firing rate is negatively correlated but not linearly (i.e, multiplying the time constant by 2 will not necessarily divide the frequency by 2). Indeed, it can be seen (Fig4-A) that the post to pre time constant ratio is ~1.6 but pre to post average firing rate of each population is different (1.77 for the excitatory population, 1.28 for the inhibitory population). This shows that model parameters can non-linearly interact making it difficult to manually tune biophysical models to reproduce desired responses.

Neural population natural frequency tuning: Natural frequency of a circuit is the dominant frequency peak in a circuit. Consequently, we call the highest peak in the spectral response of the model as its natural frequency. Pre-training natural frequency of the excitatory population were 91Hz (Fig. 5A). For training, inhibitory synaptic time constants and baseline Gaussian noise gain were allowed to vary in the model. As seen in fig. 5, the model converged to the target natural frequency (~30Hz) by increasing inhibitory time constants. Similar to the model in Fig. 4, the relation between synaptic time constants and

spectral response of the model is negatively correlated yet pre to post time constant ratio and the natural frequency peak are not linearly related (Fig. 5A,C).

A

Variables(***V***) / Metrics(***M***)	Pre	Post
V: τ(GABA, ms)	5	29.98
V: *Noise (Exc.)*	40	50.00
V: *Noise(Inh.)*	40	41.86
M: *Nat. freq. (Hz)*	91.27	30.60

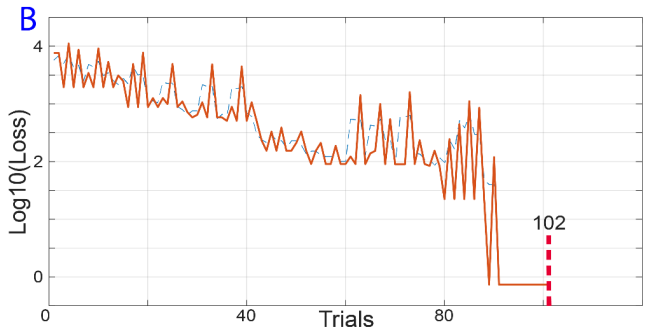

C

Parameter log	Trial 1	Trial 25	Trial 50	Trial 75	Trial 102
V: τ(GABA, ms)	5	17.48	22.00	24.12	29.98
V: *Noise (Exc.)*	40	39.58	47.03	50.00	50.00
V: *Noise(Inh.)*	40	43.32	50.00	44.56	41.86
M: *Nat. freq. (Hz)*	91.27	75.65	64.27	54.85	30.60

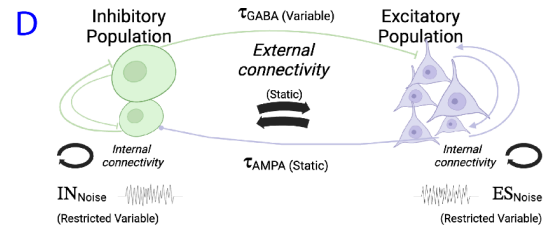

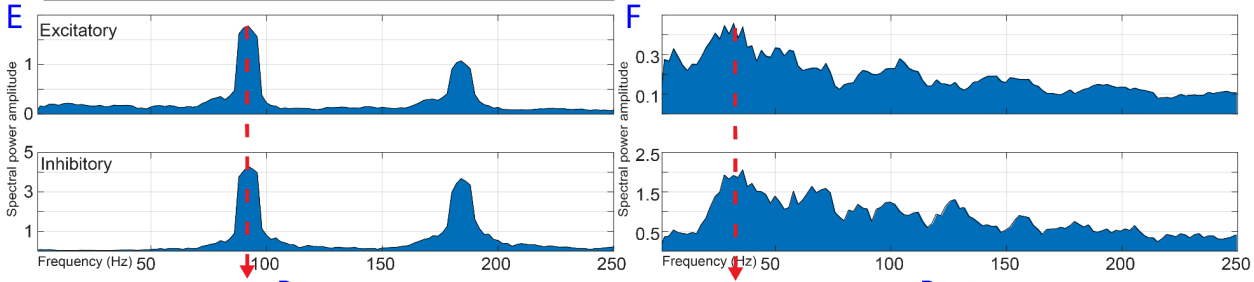

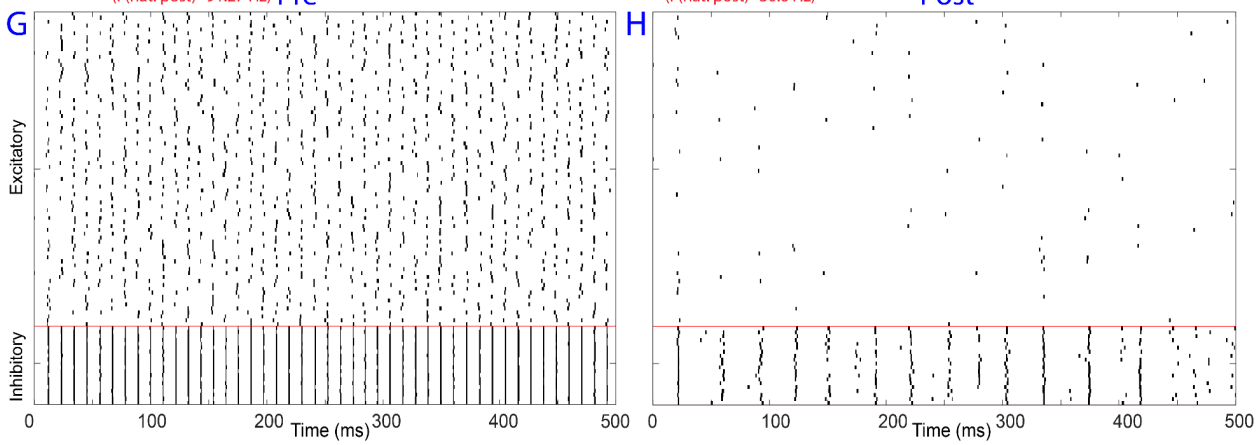

Fig5: E-I natural frequency tuning. (A) Summary of initial and final training variables and metrics. In this simulation time constant of inhibitory synaptic transmission and internal noise are optimized to achieve a natural (dominant) spectral peak at a particular frequency (30Hz). (B) Loss in trials corresponding to the target metrics. (C) Summary of training variables and metrics during training. (D) bANN circuit schematic. (E) Initial frequency response. (F) Frequency response after training showing that bANN response has achieved the target metrics. (G) Initial raster plot. (H) Final raster plot.

Experimental Design

In-vivo neural signal recordings: One adult bonnet macaque (Macaca radiata) aged 18 years (7.4kg) was used in this study. The monkey was first implanted with a 18mm recording chamber (Christ) placed over visual/temporal cortex and a headpost. To insert electrodes into the brain we used a mechanical Microdrive (Narishige, Tokyo, Japan) that was mounted onto the chambers. A guide tube was lowered through a recording grid (Christ) to penetrate the dura mater and granulation tissue. We then acutely introduced a 128-channel deep array linear into area MT/MST. We used a linear 128-channel recording array where the inter-contact spacing was 40 microns (Diagnostic Biochips, Glen Burnie, MD). Recordings were acquired using a RHD System (Intan technologies, CA) sampling at 30 kHz. Recordings were electrically grounded to the guide tube.

Grid positions were pre-determined based on a pre-recording MRI scan with the chambers and grid in place, with water-saline to mark the trajectories of each grid hole position. We advanced the electrode until there was visually responsive neuronal activity on most channels and until the local field potentials showed a distinct spectral signature of the cortical sheet, characterized by large amplitude alpha/beta (10-30 Hz) oscillations in deep channels and gamma (40-150 Hz) oscillations in superficial channels[24].

We used monkeylogic (developed and maintained at the National Institute for Mental Health) to control the behavioral task. The monkey was trained to fixate its eyes around a central fixation dot (radius of fixation window: 1.5 visual degrees). Eye position was monitored using an Eyelink 1000 (SR Research, Ottawa, Canada). A task began with an isoluminant gray screen. Once the monkey maintained fixation for 500ms, a sequence of drifting gratings appeared (radius = 12 visual degrees, drift rate: 2Hz, 1 cycle/visual degree, 0.8 Michelson contrast, angle of drifting grating was 45 or 135 degrees from the vertical midline) for 500ms, and was replaced by a blank screen with the fixation dot remaining for 500ms. This sequence was repeated 4 times in a trial, for a total trial length of 4500ms (500ms pre-sequence fixation + 4x1000ms stimulus presentations). The monkey was rewarded for maintaining fixation throughout the trial with a small juice reward. Only the first 1000ms of data was used for this analysis.

Analysis

Single units were sorted using Kilosort-2.0[25]. We only selected neurons which satisfied both quality requirements (Signal-to-noise ratio > 1.2, Presence ratio > 97.5%). We convolved all single unit spike times with a post-synaptic potential kernel with a rise time constant of 0.5ms and a decay time constant of 4.5ms. This convolution was applied to the model's response and to the neurophysiological spike time data. This step was necessary in order to facilitate spectral analysis, which is more well-behaved on continuous data as opposed to point-process data[26].

We estimated each neuron's spectrogram (time-frequency response) using Matlab's multitaper spectral estimation with Kaiser window of length 400ms and 90% overlapping between adjacent time windows. For both the model and the electrophysiology data,

performed this analysis starting from 500ms before the stimulus onset to 1000ms after the stimulus onset. Time-frequency responses of these neurons in two different conditions, stimuli-on and no-stimuli, became gSDR's objective in order to generate the corresponding neural response similar to it.

Results

Establishing the gSDR learning rule: We begin by describing the generalized Stochastic Delta-Rule (gSDR) and how it is applied to biophysical models (Fig. 1, Algorithm 1, see Materials and Methods). The first step is to use a biophysical circuit model to generate neuronal dynamics such as membrane potentials that change over time. In this work, we have used DynaSim[20] to model individual excitatory/inhibitory neurons's membrane fluctuations using Hodgkin-Huxley equations. For single neuron parameter optimization tasks (e.g., see Fig. 2), gSDR uses only a supervised path for training. For population (more than one neuron) parameter optimization tasks, gSDR uses both a supervised and unsupervised path (Fig. 1). In the unsupervised path the membrane voltages are used to derive synaptic co-activation patterns calculated by Mutual Correlation Dependent Plasticity (MCDP, essentially the correlation between membrane time courses over time, see Materials and Methods). This unsupervised learning rule is independent from any external objective and enhances connections between co-active neurons. In addition, the semi-supervised path calculates the neuronal metrics and their distance to prespecified objectives. An objective can be defined on any behavior derived from the neuronal model (e.g., firing rates, membrane potentials, synchronization of neurons, etc.). The difference between the model's performance and the objective defines the loss value for the current trial. The generalized stochastic delta rule (gSDR) uses this loss value, together with the unsupervised MCDP rule to change the model's parameters. Over time, the model explores many possible parameter spaces and keeps the most optimal parameters minimizing the loss thereby iterates until the objectives are achieved.

Single Neuron Membrane, Firing Rate, Population Firing Rate and Population spectral power tuning: We begin by demonstrating that the novel gSDR Rule is able to automatically train models on a number of simple tasks (Figs. 2-5, see Materials and Methods for additional details). These included membrane time constant optimization and firing rate optimization to a pre-defined target in single neurons, where the model could change only the reversal potential of the E_leak channel (Fig. 2). Next, we showed that the model could tune the population firing rate to a target rate for a neuronal population. This could be achieved by varying the time constants of inhibitory neurotranmission, and the firing rate of either a single (E population only, Fig 3) or multiple (simultaneously both E and I populations, Fig. 4) groups of neurons could be tuned to achieve a desired population firing rate. We further showed that the model could tune its natural frequency (the peak in the power spectra of the population spiking activity) by adjusting the time constants of inhibitory neurotransmission (Fig. 5). These relatively simple modeling tasks could in principle have been accomplished with manual tuning of parameters, but they allowed us to establish the validity of the gSDR algorithm in various scenarios where intuition of the model's and their parameters can be used as a sanity check for the model's behavior.

Using gSDR to model gamma-beta push-pull dynamics and Predictive Routing: Having established the utility of the gSDR algorithm for simple tasks, we next set the task of modulating a push-pull interaction. We trained the model to implement the flexible switch between oscillatory motifs, particularly the observed gamma-beta push pull seen in predictive routing[1,9,27]. To do this, we first simulated a model that could switch between two network states. In one network state, gamma dominates (is more powerful than) beta and in the other network beta dominates gamma. In this simulation, the model architecture is the same as the model in Fig. 5 but synaptic connections are included as training variables, making the model's dimensions (free parameters) comparable to ANNs. The model's objective is to match a spectral ratio to each of the two conditions. In condition 1 (simulation of a sensory stimulus processing, Fig. 6C), the objective is to have higher gamma band (40-100Hz) power than beta (10-40Hz), and in condition 2 (simulation of no stimulus processing, Fig.6 D), higher relative beta band power than gamma.

After 74 trials, the model achieves the conditional objective (more gamma than beta on condition 1 and vice versa). Notice that the population average gain which is equal to the average of synaptic connectivity weights is not affected significantly (Fig. 6A, C). However, the variance induced by gSDR to the distribution of synaptic connection weights is carrying the learned information to match the objective. This shows that synaptic plasticity is capable of modulating the relative spectral response of a neural circuit, without any change to the time constants or population average gain.

In addition, the transition from a gamma-dominated state (condition 1, Fig. 6C) to a beta-dominated state (condition 2, Fig 6D) resulted in substantially suppressed population activity, especially amongst the excitatory cells, consistent with a suppressive influence of beta on overall spiking activity[28].

A

Variables(*V*) / Metrics(*M*)	Initial	Final
V: Avg/SD Netcon	0.50/0.0	0.49/0.22
V: $ES_{Noise\ (mv)}$	40	39.96
V: $IN_{Noise\ (mv)}$	40	39.96

Parameter log	Trial 1 (initial)	Trial 20	Trial 40	Trial 60	Trial 74 (final)
V: Avg/SD Netcon (IE)	.5/.0	.47/.06	.44/.14	.51/.19	.49/.22
V: $ES_{Noise\ (mv)}$	40	40.28	40.09	39.92	39.96
V: $IN_{Noise\ (mv)}$	40	40.10	40.17	39.86	39.96
$\frac{\beta[10-40Hz]}{\gamma[40-90Hz]}$ (1)	0.986	1.062	1.092	1.074	1.127
$\frac{\beta[10-40Hz]}{\gamma[40-90Hz]}$ (2)	0.971	0.845	0.852	0.847	0.837

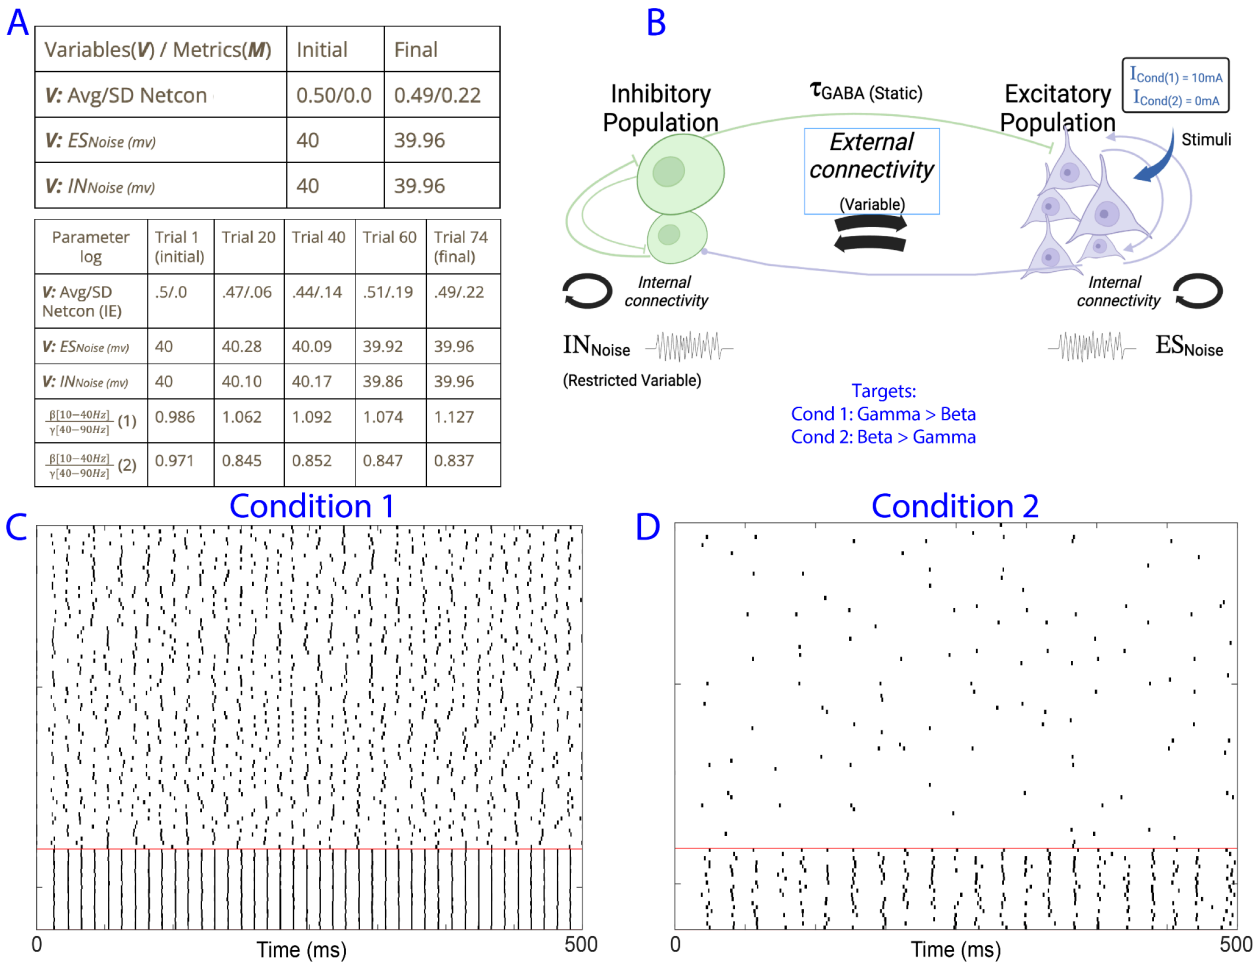

Fig6: E-I population spectral power ratio tuning. (broad Beta [10-40Hz] average power to broad Gamma [40-90Hz] average power). (A) Variables and metrics pre, during and post training (B) Network schematic. (C) Post-training raster plot on condition 1. (D) Post-training raster plot on condition 2.

Neural population contextual spectral relative power tuning: In predictive routing, we have proposed that top-down inputs carrying predictions (via beta) serve to suppress bottom-up processing (via gamma) in a selective manner[1]. We therefore next modeled how top-down inputs can modulate the network state to shift dynamics from a gamma to a beta dominated spectral state. In this simulation, the model's goal is to tune its spectral relative power between beta and gamma(Fig.7). This is similar to the previous model (Fig.6) but in this example local synaptic weights were static and could not be modulated by the model. Instead, top-down contextual inputs were allowed to change. This change was modeled by a gain parameter, which could modulate the strength of the top-down synaptic connections to each of the neuronal populations included in this model (Fig. 7B): E cells, interneurons with slow synaptic time constant (SST-cell like) and fast inhibitory population (PV-cell like).

Top-down input gain to all populations was initially set to 0 for both conditions. Next, gSDR was allowed to modulate this gain to achieve a gamma-dominated state in condition 1 and a beta-dominated state in condition 2. Interestingly, the model found a solution with positive

gain on slower inhibitory population but a negative gain for the fast inhibitory population. Post-training responses show a context-dependent spectral power shift (Sup. Fig. 2 C-F) with respect to beta (20-40Hz) to gamma (55-75Hz) power ratio objectives. This solution shows the benefit of gSDR-based exploration for biophysical models and is consistent with the predictive routing mechanism.

A

Variables(*V*) / Metrics(*M*)	Pre	Post
V: $Gain_{ES}$	0	-0.32
V: $Gain_{INfast}$	0	-2.86
V: $Gain_{INslow}$	0	6.75
M: B/G ratio(condition 1)	0.85	0.90
M: B/G ratio(condition 2)	0.85	1.79

B

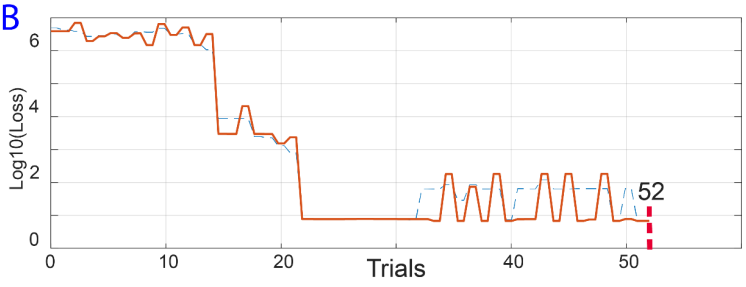

C

Parameter log	Trial 1	Trial 10	Trial 20	Trial 30	Trial 52
V: $Gain_{ES}$	0	2.522	1.546	0.292	-0.326
V: $Gain_{INfast}$	0	-4.291	0.557	-2.293	-2.868
V: $Gain_{INslow}$	0	2.776	1.441	5.9821	6.759
M: B/G ratio(condition1)	0.858	0.923	0.886	0.899	0.901
M: B/G ratio(condition2)	0.858	0.979	1.441	1.742	1.796

D

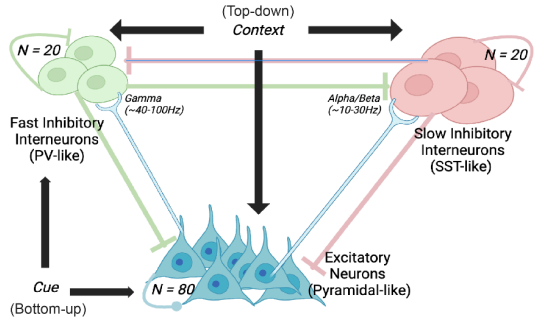

E

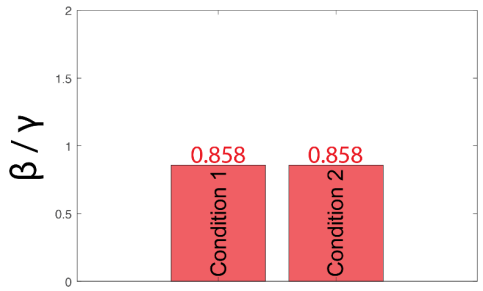

F

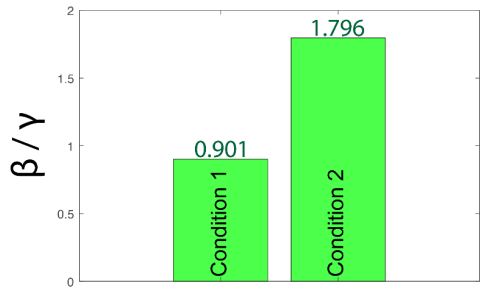

Pre-training → Post-training

Fig7: E-I contextual spectral relative power tuning. (High Beta [20-40Hz] average power to middle Gamma [55-75Hz] average power). (A) Summary of pre and post training variables and metrics. (B) Loss in trials corresponding to the target metrics. (C) Summary of training variables and metrics during training. (D) network schematic. (E) Spectral response (pre) (F) Raster plot (pre)

Predictive routing, lower order local cortical circuit based on electrophysiological recordings: In this simulation, we used gSDR to replicate electrophysiological responses observed in in-vivo recordings. Our objective here is to ask gSDR to modulate synaptic connectivities in order to recapitulate observed gamma band neuronal dynamics across experimental conditions. Current models of oscillatory dynamics (e.g, Pyramidal - Interneuron Network Gamma or PING) are capable of explaining observed gamma band neuronal dynamics, distinguishing between “strong” and “weak” versions[29]. In strong PING, gamma band oscillations (35-100Hz) require a strong input drive and feedback inhibition with fast inhibitory time constants (e.g 5-15ms), yielding highly synchronous, coherent gamma oscillations that dominate the response of most individual neurons (see Fig. 4-5).

In weak PING models, noise is added to the input signal or synaptic weights to generate a gamma oscillation with sparse, irregular firing in the pyramidal cells.

Before applying gSDR, we first determined whether real neurons display a strong or weak PING phenotype in-vivo. To do this, we performed electrophysiological recordings in the visual cortex (areas MT and MST) of awake macaque monkeys during presentation of drifting grating stimuli known to induce gamma oscillations (see methods for details about neurophysiology).

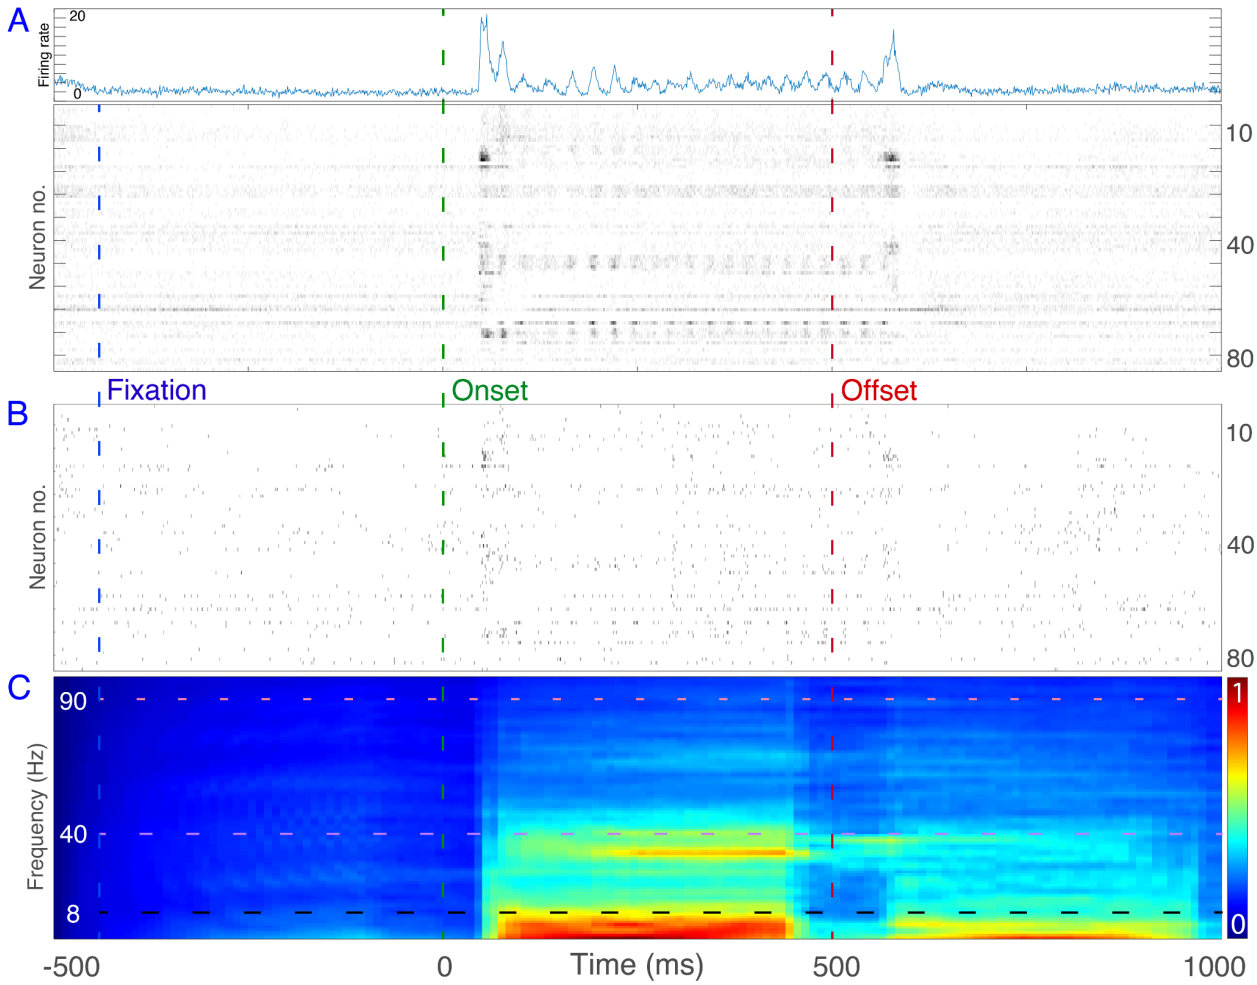

Fig8: Single unit neural response from visual area MST (NHP) (A) PSTH and all-trials cumulative raster plot (100 trials) (B) Single-trial raster plot of neurons (blue line: fixation, green line: onset of stimuli, red line: offset). (C) Corresponding spectrogram, showing a stimulus dependent evoked spectral response around ~38Hz.

We simultaneously recorded ~80 neurons in areas MT/MST to study a large neuronal population participating in gamma oscillation. During pre-stimulus baseline the neuronal population had an oscillatory resonance around beta (~20Hz, Sup. Fig.4). During presentation of the stimuli, strong narrow-band oscillatory gamma was present in the population spiking raster (Fig 8A) and single trial raster (Fig 8B). In addition, time-frequency analysis of the population spiking activity indicated the presence of a sustained narrow-band gamma (~37.5Hz) with the same duration as the duration of stimulus (500ms). These observations are consistent with the known properties of gamma

oscillations. These neuronal dynamics appear to be phenomenologically distinct from strong PING models and more similar to weak PING models in the sense that on single trials individual neurons only sparsely participate in the oscillation (Fig. 8B).

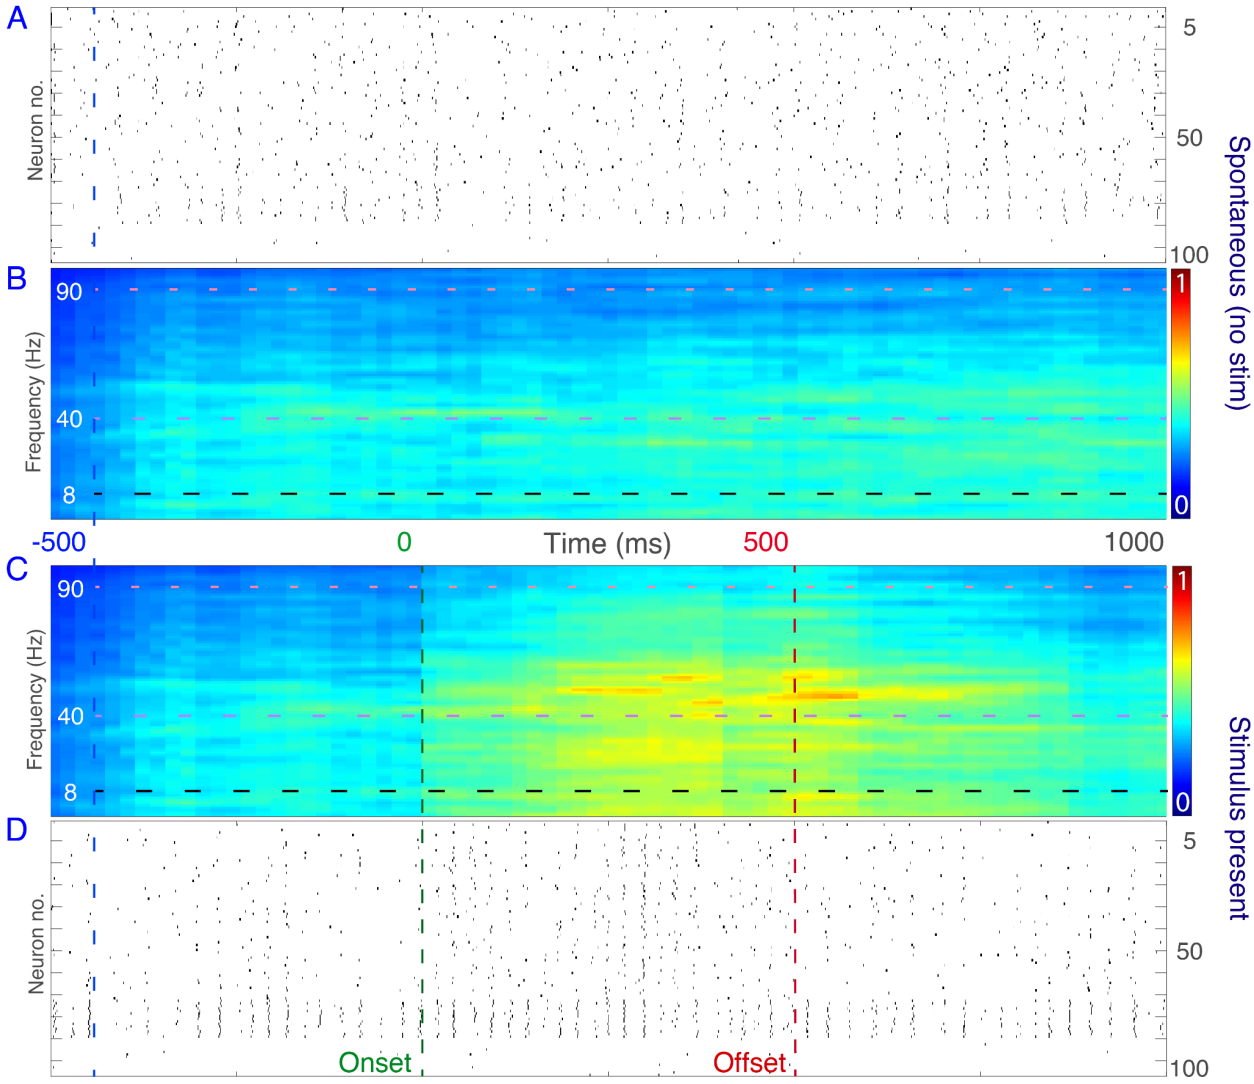

Fig9: Pre-training, single unit neural responses from the model. (A) Raster plot of the model during spontaneous activity (no stimuli) (B) Corresponding spectrogram, showing non-evoked wide-band spectral response. (C) Spectrogram of the model's response during stimulation. (D) Raster plot during stimulation.

To study the circuit mechanisms that produce this gamma oscillation, we implemented a model with 100 neurons (70 excitatory, 30 inhibitory). For this model, instead of relying on synaptic time constants, inhibitory neurons were defined based on two classes of connectivity[30]. One class had a pattern of locally-focused connectivity (restricted to establish local synaptic connections, to maximum 10% of total neurons, Sup. Fig. 1) and another class had a global widespread connectivity (no restrictions on number of synapses to other neurons). Unlike the classic PING model that relies on fast inhibitory time constants for producing gamma, we did not predefine or modulate the synaptic time constants in our model (time constants were set to a random distribution with a mean of 50ms and standard deviation of 5ms for all connections). Instead, we allowed gSDR to

modulate the synaptic connections. The objective of the model was to approximate the spectrotemporal features present in the observed neurophysiology (Fig. 8) in both the pre-stimulus fixation and stimulation time period.

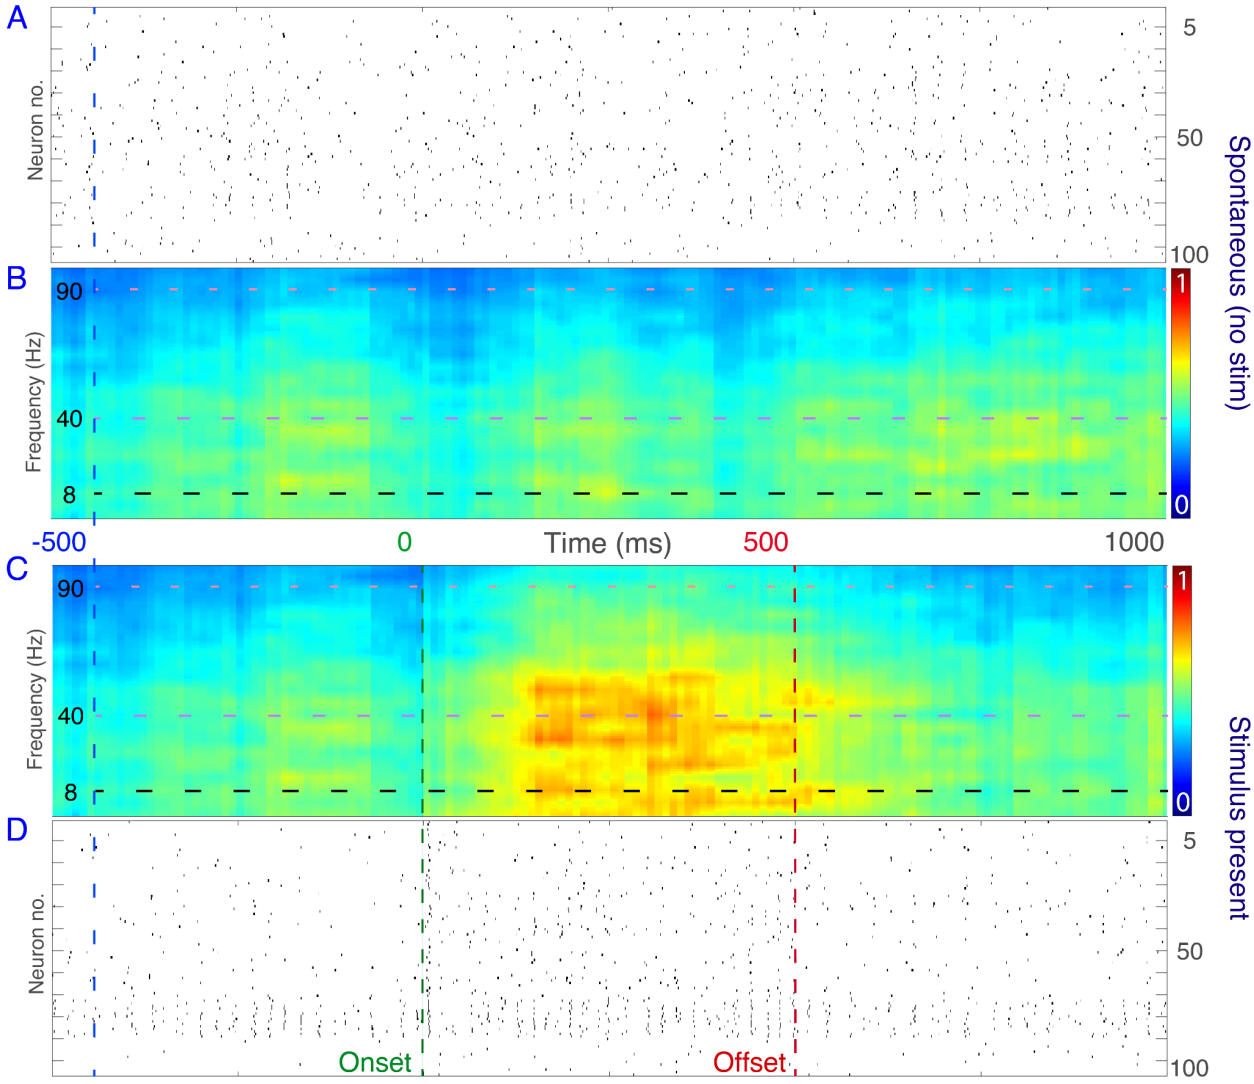

Fig10: Mid-training (trial 200) (A) No-stim raster plot. (B) No-stim spectrogram. (C) Spectrogram (during stim.). (D) Stimulus-present raster plot.

The spectrotemporal response of both the model and the data was quantified by taking each neuron's power spectrum and averaging these power spectra across all neurons in the model/data (see methods). The model simultaneously optimized its spectral response to the pre-stimulus data (From 500ms pre-stimulus to stimulus onset) as well as its spectral response to the post-stimulus data (From stimulus onset to 500ms post-stimulus). The loss (used to train the model) was calculated using the log-ratio (see methods) between the model and the data's spectral response. To simulate the stimulus, we assumed in the model that at time 0 (onset of the stimuli) there was an increase in the external input (modeled as a direct current-DC step function with amplitude of 5mA) with a 500ms duration.

Background noise was a voltage-based stationary random process from a uniform distribution in the interval of [-10mV,+10mV], which is always present independently.

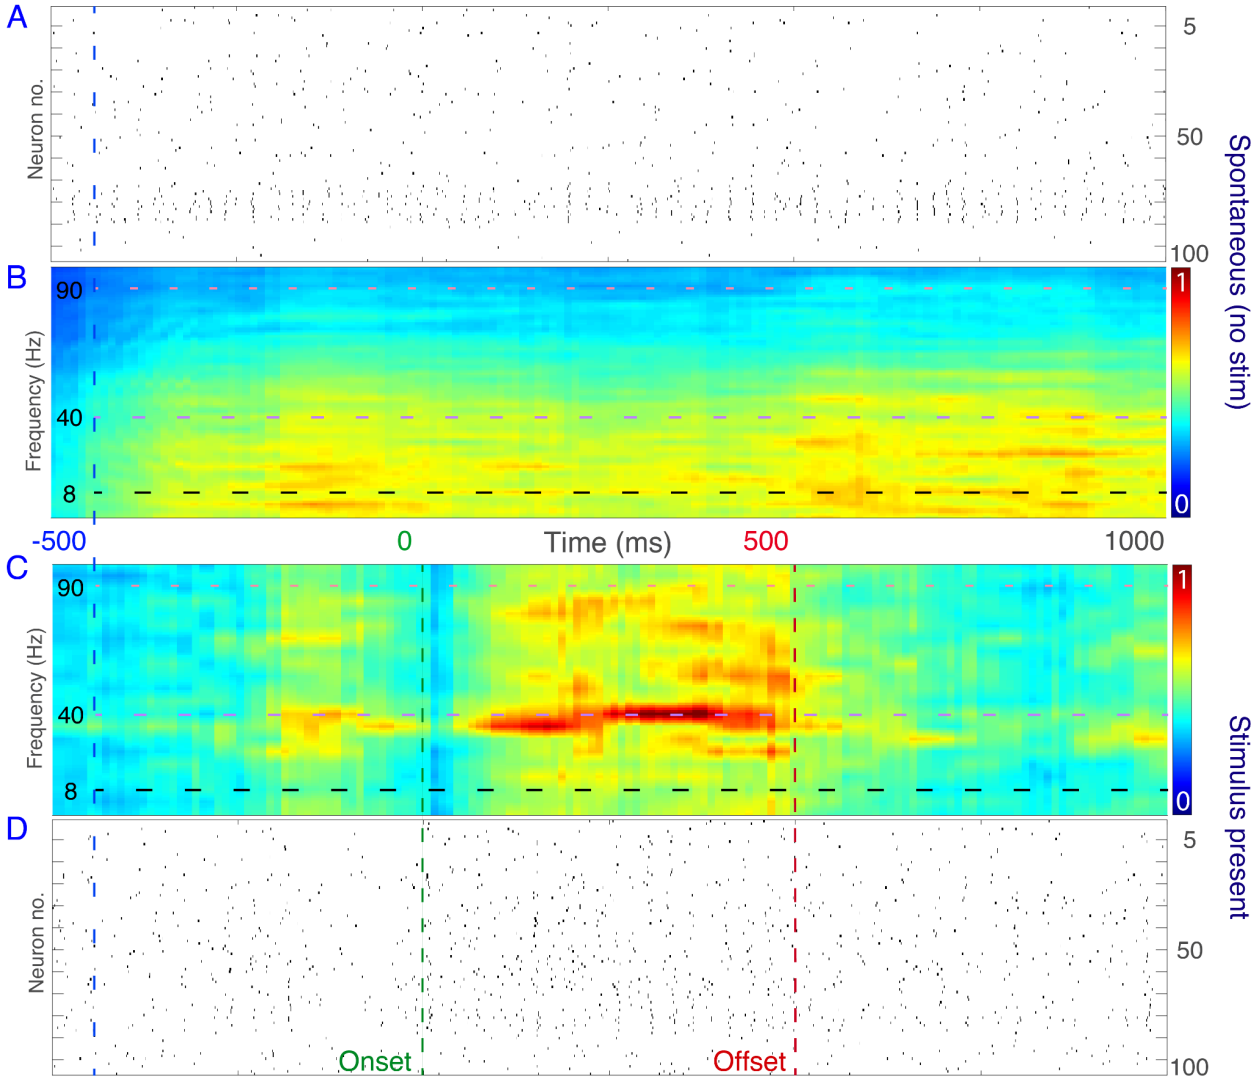

Fig11: Post-training, single unit neural responses from the model. (A) No-stim raster plot. (B) No-stim spectrogram. (C) Corresponding spectrogram, showing that the evoked spectral response is shifted from high-gamma to low-gamma and beta. (D) Stimulus-present raster plot.

During training, both conditions (pre-stim and stim) were equally presented to the model and the loss function penalized both conditions equally. The model's initial synaptic weights were set-up to be uniform across all connections (Sup. Fig. 1 A,B), leading to the pre-training responses (Fig. 9). As can be seen, the spectral response of the model increases during stimulus presentation (Fig. 9C) but in a broad-band manner that did not capture the narrow-band gamma response present in the data (Fig. 8C). After 200 trials (Fig. 10) the model showed enhanced gamma amplitude during stimulation but still not in a narrow band regime.

Eventually, after 370 trials (Sup. Fig. 1 C), the model achieved a plausible stimulus dependent narrow band gamma response and pre-stimulus spontaneous dynamics (Fig. 11,

Sup. Fig. 4). This demonstrates that our stochastic learning rule is capable of training a biophysical model to achieve observed neuronal dynamics without direct supervision or manual dictation of parameters. Specifically, it did so by modulating the synaptic weight matrix in a highly sparse and non-uniform manner (Sup. Fig. 1 A,B). Moreover, recapitulating the observed gamma band oscillations did not require manual adjustment of time constants or strong and uniform I to E synaptic connections, instead, via sparse synaptic connections bidirectionally (Sup. Fig. 1 A-B). This further suggests that the biophysically plausible models when trained on real data implements oscillatory dynamics with weak PING characteristics, e.g. the oscillation is mostly observed at a population level, not at the single neuron (Fig. 9-11, Sup. Fig. 4). This is because the post-training connectivity weights were highly variable and therefore not uniform.

We next quantified which connections were mostly changed in the model to account for these shift in spectral dynamics from beta to gamma. To do this, we quantified the average synaptic connection change from the model in the pre-training state to the post-training state, separately for synaptic connections between the excitatory neurons, the local inhibitory interneurons (which had a connectivity structure that limited to make connections only with 10% of local neurons) and the global inhibitory interneurons (which had no restrictions on number of neurons they connect to). Although the synaptic connections between all classes of neurons changed, we found that the greatest magnitude of change occurred between the local inhibitory interneurons and the other two populations (excitatory and global inhibitory interneurons, Sup. Fig. 1&4). These results allow us to understand that a network can transition from a beta dominated state to a gamma dominated state, both via local synaptic modulation and top-down modulation of the circuit, during the pre-stimulus to stimulation time (Figs. 6, 7, 9, 10, 11, Sup. Figs. 1, 2, 4).

Discussion

Our work extends efforts to train biophysical neuronal models. We applied a novel learning algorithm, gSDR, to biophysical neural circuits, which made it possible to explore unseen solutions for model parameters. gSDR lets these models decide how to find these solutions by adjusting the trade-off between unsupervised and supervised learning (see methods, Fig.1 and Alg.1). Thus, this framework facilitates computational studies on neural dynamics seen during cognitive tasks, such as predictive routing[1,3,31], whereby predictable stimuli are associated with more beta than gamma and this spectral shift is specific to predicted stimulus representations. The gSDR framework also enabled biophysical networks to mimic neuronal gamma oscillations, allowing us to study which parameters and network architectures are critical for their formation.

Relationship between gSDR and other computational modeling frameworks

Another approach that has emerged to train biophysical models is Dynamic Causal Modeling (DCM)[15,16,32]. DCM has been used to simulate observed brain dynamics by training model parameters using gradient descent with convex assumptions. Current DCM implementations assume the population scale (e.g., summarizing the behavior of hundreds to thousands of neurons as a neuronal mass) is a sufficient level of complexity to explain

brain dynamics. In our view, it is important to also consider and model the microscopic elements of neural circuits (single cells, single synapses and neuro-chemical action at the micro-scale, meaning between neurons). Additionally, current implementations of DCM make the assumption that the equations describing the model are solvable by gradient descent, implying a convex and continuous relation between parameter space and the model's behavior[33]. In contrast, previous modeling has shown that population behavior can be highly-non convex and non-continuous across subtle parameter changes[34]. This was one of the motivations for incorporating a more stochastic approach to learning with gSDR, with fewer assumptions on how to explore the parameter space.

Another modeling framework that has been inspired by neuronal computational modeling is Artificial Neural Network (ANN) and Deep Neural Network (DNN) models, which have revolutionized artificial intelligence (AI) and deep learning[35]. In general, these models use more abstracted “neurons” and prioritize performance and high-level behavior over biological plausibility[35]. ANNs and DNNs generally do not consider neural dynamics observed in electrophysiology [15], specific cell types [33], excitatory-inhibitory (E-I) interactions and neurochemicals (e.g neurotransmitters)[34]. While these models are of course useful for learning complex tasks, their non-biophysically realistic architectures and non-physiological learning assumptions make their results difficult to map to neuroscience, due to structural disparities and non-equivalent computations[36–38].

With this in mind, one might ask what the best level of detail is in computational modeling to understand the brain? In our view, it is critical to include different neuronal cell types, with specific neuromodulators, that inhibit or excite other neurons, as these interactions will generate rich spectro-temporal neuronal dynamics as observed both in-vivo[1,3,9,27] and in-silico[1818,50,51]. Since our goal is to test theoretical hypotheses about the brain at the single neuron, microcircuit and/or neuronal population scales, we cannot rely on models that are not as consistent as possible with the neurobiology.

Many works have utilized biophysically inspired modeling to mirror in-vivo brain dynamics[17]. For example. spiking neural networks (SNNs) have been used to study cognition[41,42] and to model findings from human electrophysiology during a mismatch negativity task. [43,44] In one modeling study, an SNN learned predictive-coding computations by changing synaptic connections in response to changing stimulus patterns and sequences. Other studies have proposed various hypotheses about many other mechanisms contributing to the oscillatory and temporal dynamics of mismatch negativity and predictive coding[45–49]. In our view, these detailed biophysical models are important because they can also be used to explain, in a computational framework, how observed brain dynamics may change from a healthy to a disordered state and lead to mental disorders[50–52]. Computational explanations for these disorders should be able to link specific low-level mechanisms (e.g., specifical cell types, their connectivity parameters) to higher-level behaviors (e.g., observed decreases in mismatch responses in schizophrenia[53]). This type of modeling should therefore also generate novel targets for new treatments and therapeutics, to re-balance networks into a healthier regime. Therefore, it is imperative to build accurate biophysical models that can be trained on large scale neurophysiological recordings to better understand these circuit dysfunctions.

What has gSDR taught us about training biophysical models?

Many computational studies have utilized biophysical details of neural microcircuits to perform computations[43,49,54]. However, most of these models require manual tuning and optimization prior to the simulation due to their complexity[37,38] and biophysical constraints[18,39,40,55]. One of the major advances of gSDR is to overcome manual tuning with an efficient search algorithm capable of finding modeling solutions for both simple and more complex tasks. For example, the highly non-uniform connectivity matrices found by gSDR (Supp Fig. 1B) would have been difficult and highly computationally costly to find via manual tuning, but could be efficiently discovered in a few hundred trials with gSDR. A future goal will be to use the gSDR framework for not only exploring biophysical solutions to well-defined problems but also for task learning, to then examine whether the mechanisms of gSDR-tuned models are similar to observed brain networks. This technique could also be applied in a recurrent manner, to drive experiments based on insight from modeling, and also to use experimental data to create more realistic models.

What has gSDR taught us about gamma oscillations?

We applied this biophysically detailed modeling framework to study gamma oscillations that we also recorded in vivo. One of our motivations was to use gSDR-guided modeling to study whether the observed neuronal oscillations were better explained by so-called "weak" vs. "strong" Pyramidal-Interneuron Network Gamma (PING). Current models of oscillatory dynamics are capable of explaining observed gamma band neuronal dynamics using either mechanism. In strong PING models, gamma band oscillations are result of a strong input drive and feedback inhibition with fast inhibitory time constants[18,40]. In weak PING models, internal noise generates a gamma oscillation with sparse, irregular firing in the pyramidal cells[29,34,56].

Previous studies[18,57,58] regarding biophysical mechanisms contributing to the oscillatory dynamics in predictive routing have suggested that the gamma oscillation can act as a winner-take-all mechanism. This mechanism can bias the network to process one stimulus over another, when both are present but top-down modulatory drive is present onto one population. This winner-take-all mechanism has been used to explain stimulus specific attention and distractor suppression. These studies also show that desynchronizing effect of random connectivity can be reduced by strengthening the E→I synapses[59]. More surprising, it cannot be reduced by strengthening the I→E synapses. However, the decay time constant of inhibition plays an important role. Faster decay yields tighter synchrony. In particular, in models in which the inhibitory synapses are assumed to be instantaneous, the effects of sparse, random connectivity cannot be seen. This is in accordance to our modeling results (Supp. Fig. 3). The synaptic connectivity from I focal to E and I focal to other I cells was most strongly modulated by the model to achieve a dynamic gamma rhythm. The model therefore converged onto mechanisms[57,58] whereby inhibitory synapses were strong enough to enable PING rhythms, but "weak enough" not to suppress all noise-driven spiking in the presence of the PING rhythm, implying the role of a weak-PING like circuit.

In our model, we were able to replicate the in-vivo visually evoked narrow-band gamma (Fig. 10-11) without specific time constants nor uniform inhibitory feedback connections.

We adjusted the model's architecture by assuming that fast-spiking inhibitory interneurons make a limited number of locally-focused and strong synapses based on a cortical circuit model[30] with similar assumptions based on neurophysiology[60,61]. We also included alternative mechanisms for inhibitory interneuron populations, which we called the locally-focused connectivity interneurons and globally-widespread connectivity interneurons. The reason for including both inhibitory interneuron populations was to evaluate both time-constant based interneuron models and connectivity based interneuron models proposed in a computational model of primary visual cortex in rodents[30].

gSDR was trained to match the power spectra of neurons in this model to the electrophysiologically observed neuronal spectral response. In the post-training state, the model exhibited several features which were more realistic compared to previous models. First, unlike previous models this narrow band oscillation did not dominate the response during spontaneous activity, but was highly dependent on strong, excitatory synaptic inputs which switched the network from a beta- to gamma-dominated state. This flexible transition from pre-stimulus beta to stimulus-induced gamma was learned via gSDR exploration of synaptic connectivity, with notable contribution of strong synapses from locally-focused interneurons. Second, the gamma oscillation did not orchestrate the entire neuronal population with high synchrony. Instead, sparse activity on a single trial level characterized the oscillation, which could be best explained as a population phenomenon. These signatures are more in line with weak PING than strong PING models.

What has gSDR taught us about predictive routing?

Our modeling framework allowed us to establish several neurobiologically plausible mechanisms for predictive routing. In previous empirical work, we demonstrated that when unpredictable inputs (such as an auditory oddball stimulus) reach a cortex lacking top-down input -such as during propofol-induced anesthesia[3]- gamma activity becomes disinhibited due to the absence of top-down beta control. In the awake state, however, beta oscillations (associated with top-down processing[1,7,27]) exert control over gamma activity. Our modeling results show that top-down contextual weights, which we assumed act via a gain parameter onto specific excitatory and inhibitory interneurons, can be learned during training (Fig. 7) to shift network dynamics between gamma and beta. This supports the predictive routing model in which top-down, preparatory signals are used to prime cortical networks for processing predicted stimuli.

We also considered a simpler possibility, namely that changes in local synaptic connectivity are sufficient to induce a spectral shift between beta and gamma (Fig. 6, Figs. 8-11). This may be sufficient to explain deviant responses to stimuli that are so-called local oddballs (e.g., in the stimulus sequence x-x-x-Y, Y is an unrepeated stimulus that deviates from the repeated context of stimulus x). Local oddball responses do not necessitate top-down connections carrying predictions, because they persist, or are even enhanced, during deep states of anesthesia when frontal to sensory top-down inputs are functionally disconnected[3]. Therefore, this opens the possibility that local connectivity changes are sufficient to carry these more basic forms of prediction that also induce beta/gamma shifts in population activity. Indeed, we found that modulations to these local synaptic connectivity parameters could model a spectral shift from beta to gamma. We found this to be the case both in a pure modeling framework, where gSDR operated on a metric (to

increase or decrease the beta/gamma ratio) as well as when gSDR was trained to maximize the fit between neurophysiologically observed and modeled power spectra of neuronal activity. In both cases, we identified that a critical parameter to transition between beta to gamma and vice-versa is a limited number of strong synapses from local inhibitory interneurons. This is accordance to previous studies[49] which show that the loss of evoked gamma by antagonizing or blocking NMDA receptors, leads to decreased activity in the GABAergic interneurons and increased pyramidal excitation[62], because the GABAergic interneurons are tenfold more sensitive to the NMDA-r antagonists than the pyramidal neurons[63]. In fact, the proposed theory of NMDA-r hypofunction[64,65] causing reduced evoked gamma via dampening local cortical PV-interneuron activity in schizophrenia and related psychosis disorders is supported by significance of strong-local synaptic inhibition[66,67].

Therefore, gSDR has taught us that both of these network implementations (top-down inputs, as well as local intrinsic circuitry) are capable of inducing beta-gamma push-pull dynamics. Meanwhile, more empirical neurophysiology experiments are required to causally control for such mechanisms (such as the modulation of local inhibition) in order to unveil the trade off between local circuitry mechanisms and global top-down modulation. Further empirical, causal, and modeling studies will be necessary to reveal the contributions of local vs. top-down connections for oscillatory push-pull dynamics observed during predictive routing.

Future extensions of the gSDR modeling framework.

The gSDR algorithm was inspired by evolutionary processes and the broader family of genetic optimizations[23,68,69], because of the stochastic nature of search through parameter space. Currently, gSDR with biophysically plausible neuronal networks might not surpass other learning algorithms (e.g., ANNs and DNNs trained with gradient descent) in task performance but offer a more natural learning framework, mimicking the innate mechanisms observed in the brain. The brain itself relies on exploration and model-based learning, such as Spike-Timing-Dependent Plasticity (STDP), chemical neuromodulators and genetic expression. Furthermore, the unsupervised nature of predictive coding and broadly the Bayesian brain concept, emphasize the importance of self-guided learning processes[70,71]. We performed simulations from the simple to more complicated models. We validated gSDR on single neuron optimizations, population response and spectral interactions. We then implemented a neuronal circuit in order to replicate the spectrotemporal response observed in electrophysiology and predictive routing framework.

In future work, we are interested in the role of different cell types and their laminar composition on neuronal dynamics. Inspired from knock-out studies in neurobiology, first we will implement multiple groups of models with different compositions of inhibitory interneurons. The objective of these models would be the same but consisting of three objectives; maintaining similarity to the electrophysiology, maintaining anatomical properties such as between layer connections, and performing a basic task via the predictive routing concept. After tuning, task performance and bioplausibility will be compared to evaluate the role of layer-specific cell type composition. In addition to this,

similar to the recent studies involving rodent biophysical models[37,38] we will aim to estimate the local-field potential (LFP) and current-source density (CSD).

This will enable more detailed study of the coherence between spiking activity and LFP (spike-field coherence, SFC[72]). Since we are able to estimate SFC from these models, gSDR would be utilized in order to penalize or enhance SFC as an objective. Using this approach, we could disrupt or enhance SFC to examine the effects it could have on model's performance, neural dynamics and bioplausibility. This is important because causal studies on functional and anatomical connectivity are difficult to perform in-vivo, unlike in-silico studies. Since we are able to use gSDR to define specific connectivity restrictions, one direction of studies could be to restrict connectivity patterns via the objective function to perform in-silico knock-out studies of particular circuit elements.

Conclusion

We implemented gSDR, a novel hybrid learning algorithm that naturally explores the distribution of the parameters to achieve user-specified objectives. This hybrid learning approach made it possible for models to learn using both supervised methods and also an unsupervised, trial-and-error approach. We demonstrated how this algorithm can be used to train highly non-convex, non-linear and stochastic models of biophysical neural networks. This made it possible to study how neural networks produce neurophysiologically realistic gamma oscillations. The network converged onto dynamics that were more consistent with weak PING, as opposed to strong PING dynamics. We also used this method to show that predictive routing can be implemented using multiple mechanisms. The shift from beta to gamma (and vice-versa) network dynamics implied by predictive routing could be achieved via changes to local circuitry as well as top-down connectivity.

Acknowledgement

This work was supported by NIH-R00MH116100 (AMB) and Vanderbilt startup funds (AMB). Electrophysiology, neural recording and procedures were approved by the Vanderbilt University Institutional Animal Care and Use Committee. We thank Nancy Kopell and Homero Esmeraldo for critical comments to an earlier version of this manuscript.

Code availability

Github reporsitory: https://github.com/DynaSim/DynaSim/tree/devDL

Supporting information

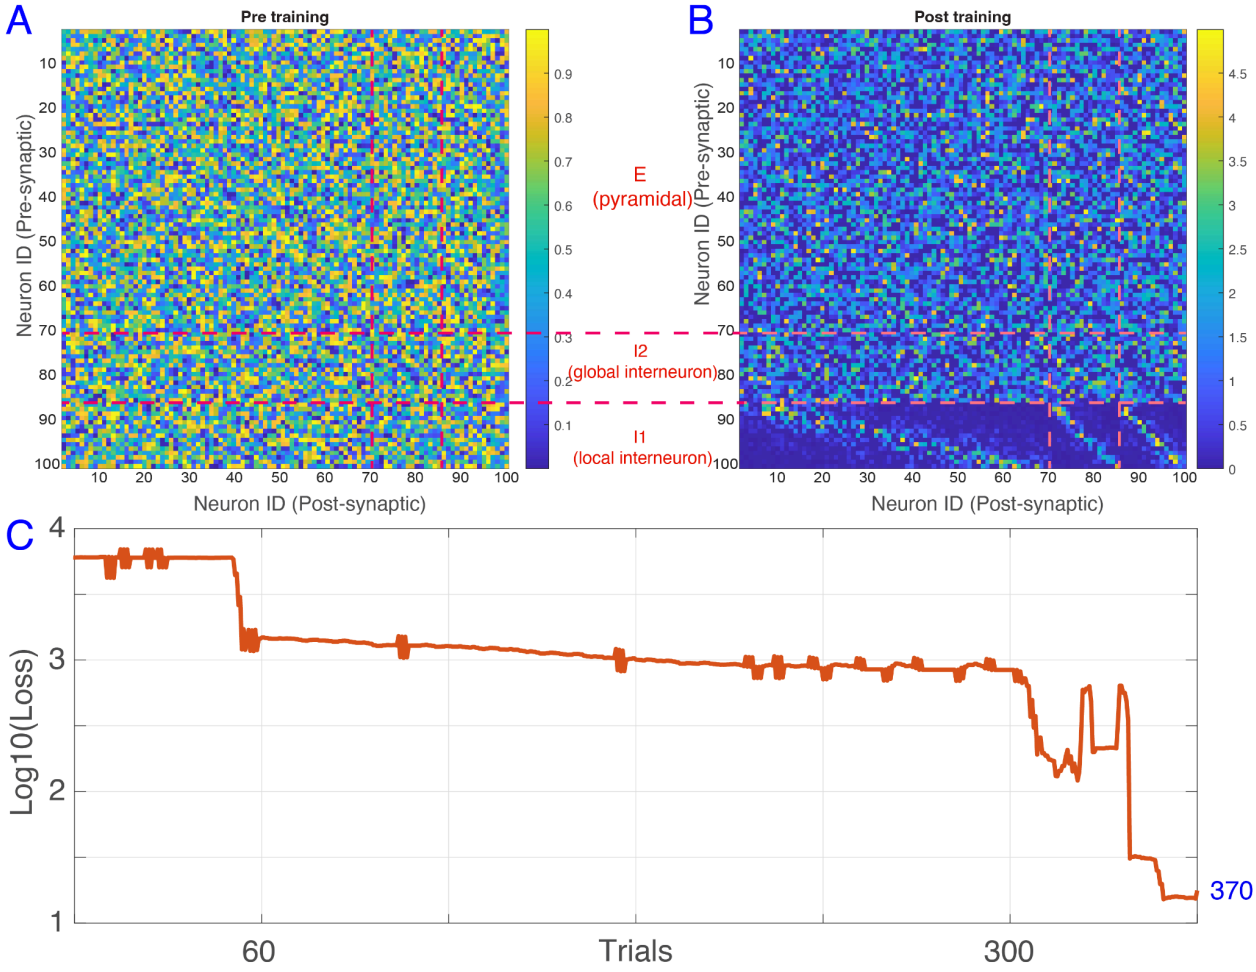

Sup. Fig 1: *(A-B) Pre-Post training synaptic connectivity matrices, corresponding to figures 9-11 (C) Loss across trials during training for the model in Fig 10-12, up to trial 370 (last trial)*

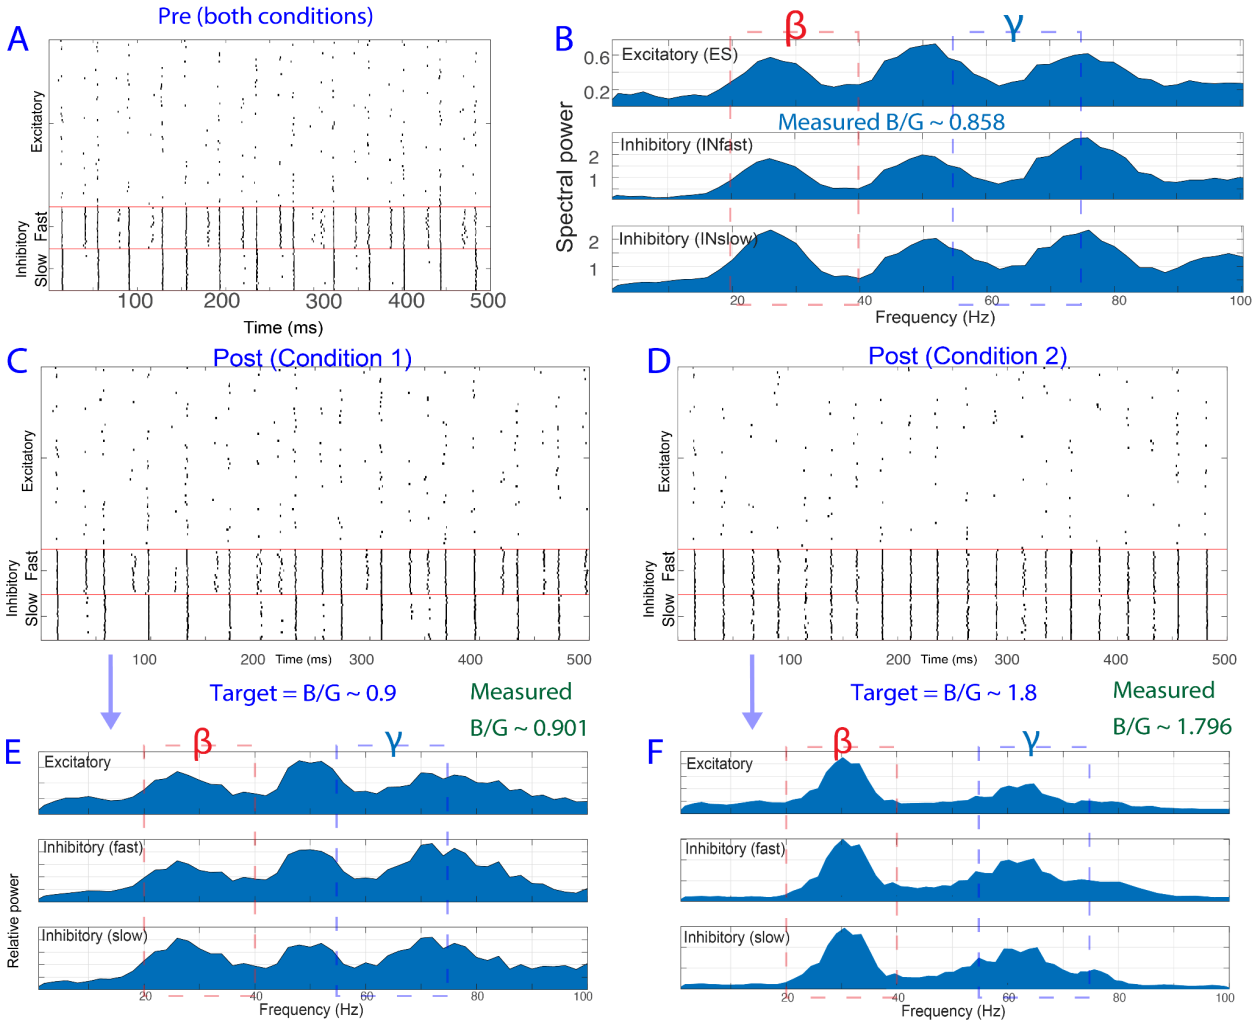

***Sup. Fig 2**: E-I dual-context specific band frequency ratio optimization responses after training, corresponding to the Fig. 7. (A) Post-training raster plot, condition 1 (B) Post-training raster plot, condition 2 (C) Spectral response, condition 1. (D) Spectral response, condition 2.*

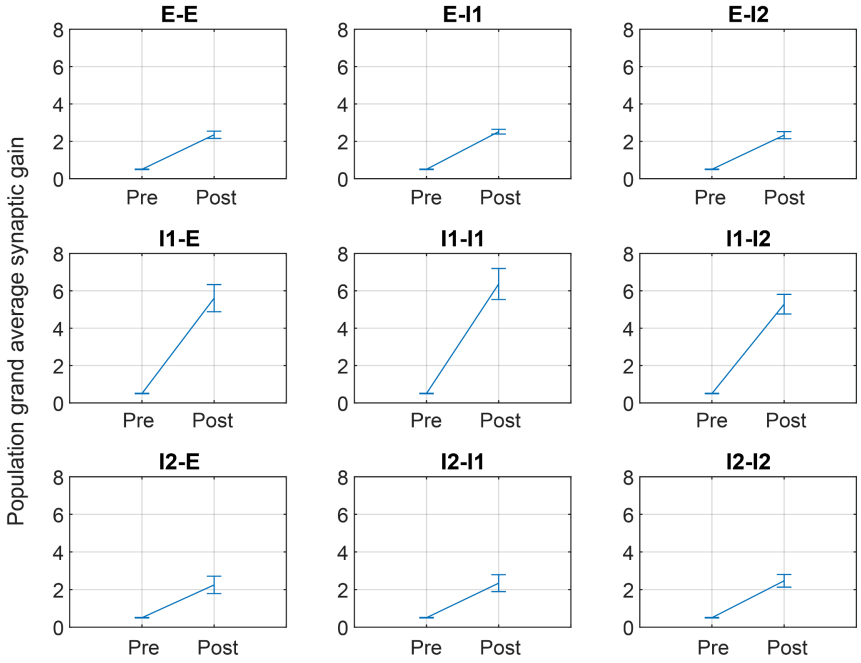

***Sup. Fig 3**: Pre and post-training grand average population synaptic gain between different populations, in format of Source-Target (e.g, E-I1 means synaptic connections from E population to the I1). 'E': excitatory pyramidal neurons, 'I1': locally-focused interneurons. 'I2': globally widespread connectivity interneurons.*

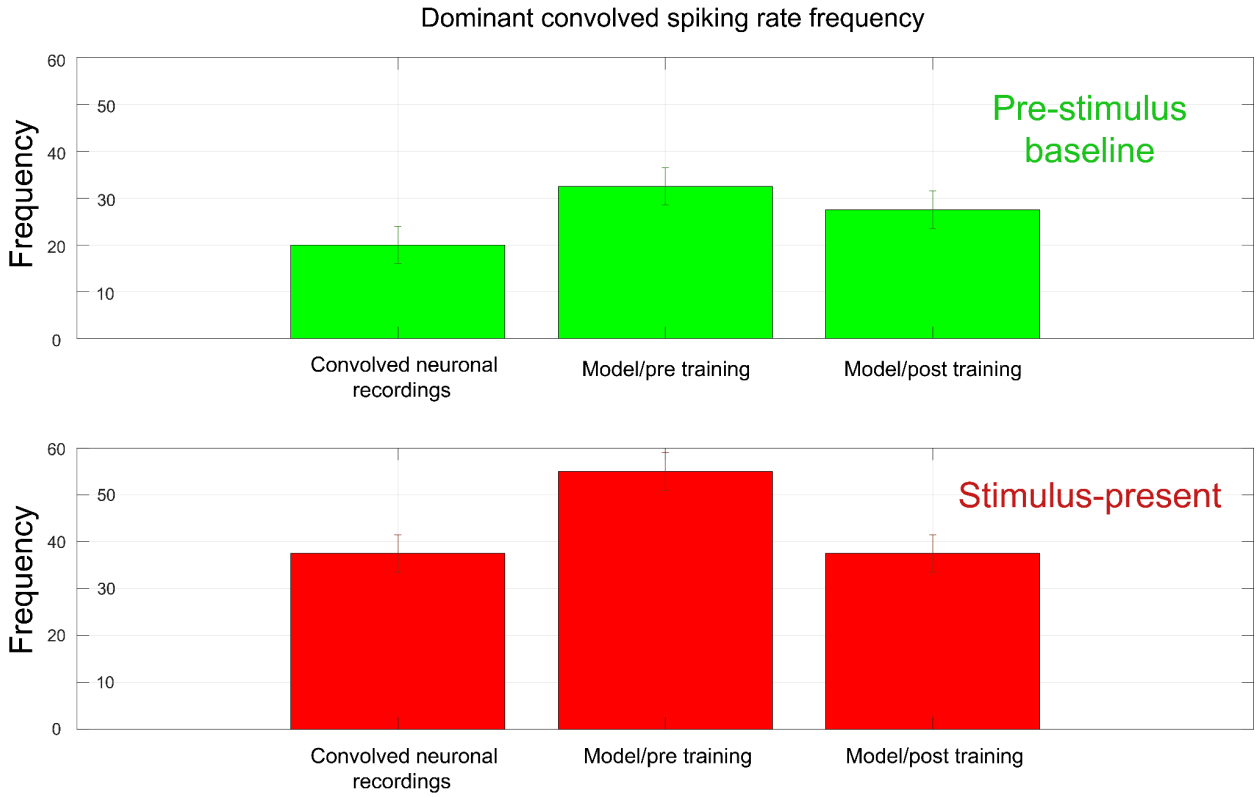

***Sup. Fig 4**: Dominant frequency in the recorded spiking and model (pre training/post training, pre-stim/post-stim). Upper subpanel corresponds to the pre-stimulus baseline and bottom subpanel corresponds to the stimulus-present time.*

References

1. Bastos, A. M., Lundqvist, M., Waite, A. S., Kopell, N. & Miller, E. K. Layer and rhythm specificity for predictive routing. *Proc. Natl. Acad. Sci.* **117**, 31459–31469 (2020).
2. Bastos, A. M. *et al.* Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
3. Xiong, Y. (Sophy) *et al.* Propofol-mediated loss of consciousness disrupts predictive routing and local field phase modulation of neural activity. *Proc. Natl. Acad. Sci.* **121**, e2315160121 (2024).
4. Xiao, J., Chen, Z. & Yu, B. A Potential Mechanism of Sodium Channel Mediating the General Anesthesia Induced by Propofol. *Front. Cell. Neurosci.* **14**, 593050 (2020).
5. Bastos, A. M. *et al.* Neural effects of propofol-induced unconsciousness and its reversal using thalamic stimulation. *eLife* **10**, e60824 (2021).
6. Eisen, A. J. *et al.* Propofol anesthesia destabilizes neural dynamics across cortex. *Neuron* **112**, 2799-2813.e9 (2024).
7. Bastos, A. M. *et al.* Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron* **85**, 390–401 (2015).
8. van Kerkoerle, T. *et al.* Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proc. Natl. Acad. Sci. U. S. A.* **111**, 14332–14341 (2014).
9. Bastos, A. M., Loonis, R., Kornblith, S., Lundqvist, M. & Miller, E. K. Laminar recordings in frontal cortex suggest distinct layers for maintenance and control of working memory. *Proc. Natl. Acad. Sci.* **115**, 1117–1122 (2018).
10. Haegens, S. *et al.* Laminar Profile and Physiology of the α Rhythm in Primary Visual,

Auditory, and Somatosensory Regions of Neocortex. *J. Neurosci.* **35**, 14341–14352 (2015).

11. Spaak, E., Bonnefond, M., Maier, A., Leopold, D. A. & Jensen, O. Layer-specific entrainment of γ-band neural activity by the α rhythm in monkey visual cortex. *Curr. Biol. CB* **22**, 2313–2318 (2012).

12. Pfeffer, C. K., Xue, M., He, M., Huang, Z. J. & Scanziani, M. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nat. Neurosci.* **16**, 1068–1076 (2013).

13. Mejias, J. F., Murray, J. D., Kennedy, H. & Wang, X.-J. Feedforward and feedback frequency-dependent interactions in a large-scale laminar network of the primate cortex. *Sci. Adv.* **2**, e1601335–e1601335 (2016).

14. Wang, X.-J. & Yang, G. R. A disinhibitory circuit motif and flexible information routing in the brain. *Curr. Opin. Neurobiol.* **49**, 75–83 (2018).

15. Bastos, A. M. *et al.* A DCM study of spectral asymmetries in feedforward and feedback connections between visual areas V1 and V4 in the monkey. *NeuroImage* **108**, 460–475 (2015).

16. Katsanevaki, C. *et al.* Attentional effects on local V1 microcircuits explain selective V1-V4 communication. *NeuroImage* **281**, 120375 (2023).

17. Sanchez-Todo, R. *et al.* A physical neural mass model framework for the analysis of oscillatory generators from laminar electrophysiological recordings. *NeuroImage* **270**, 119938 (2023).

18. Lee, J. H., Whittington, M. A. & Kopell, N. J. Top-Down Beta Rhythms Support Selective Attention via Interlaminar Interaction: A Model. *PLoS Comput. Biol.* **9**,

e1003164 (2013).

19. Richter, C. G., Thompson, W. H., Bosman, C. A. & Fries, P. Top-Down Beta Enhances Bottom-Up Gamma. *J. Neurosci. Off. J. Soc. Neurosci.* **37**, 6698–6711 (2017).

20. Sherfey, J. S. *et al.* DynaSim: A MATLAB Toolbox for Neural Modeling and Simulation. *Front. Neuroinformatics* **12**, 10 (2018).

21. Frazier-Logue, N. & Hanson, S. J. The Stochastic Delta Rule: Faster and More Accurate Deep Learning Through Adaptive Weight Noise. *Neural Comput.* **32**, 1018–1032 (2020).

22. Markram, H., Gerstner, W. & Sjöström, P. J. Spike-Timing-Dependent Plasticity: A Comprehensive Overview. *Front. Synaptic Neurosci.* **4**, (2012).

23. Vikhar, P. A. Evolutionary algorithms: A critical review and its future prospects. in *2016 International Conference on Global Trends in Signal Processing, Information Computing and Communication (ICGTSPICC)* 261–265 (IEEE, Jalgaon, India, 2016). doi:10.1109/ICGTSPICC.2016.7955308.

24. Mendoza-Halliday, D. *et al.* A ubiquitous spectrolaminar motif of local field potential power across the primate cortex. *Nat. Neurosci.* **27**, 547–560 (2024).

25. Pachitariu, M., Sridhar, S., Pennington, J. & Stringer, C. Spike sorting with Kilosort4. *Nat. Methods* **21**, 914–921 (2024).

26. Kass, R. E., Ventura, V. & Brown, E. N. Statistical Issues in the Analysis of Neuronal Data. *J. Neurophysiol.* **94**, 8–25 (2005).

27. Miller, E. K., Lundqvist, M. & Bastos, A. M. Working Memory 2.0. *Neuron* **100**, 463–475 (2018).

28. Lundqvist, M., Bastos, A. M. & Miller, E. K. Preservation and Changes in Oscillatory Dynamics across the Cortical Hierarchy. *J. Cogn. Neurosci.* **32**, 2024–2035 (2020).

29. Lee, S. & Jones, S. R. Distinguishing mechanisms of gamma frequency oscillations in human current source signals using a computational model of a laminar neocortical network. *Front. Hum. Neurosci.* **7**, (2013).

30. Del Rosario, J. *et al.* Lateral inhibition in V1 controls neural & perceptual contrast sensitivity. Preprint at https://doi.org/10.1101/2023.11.10.566605 (2023).

31. Gabhart, K., Xiong, Y. (Sophy) & Bastos, A. Where are global oddballs? A predictive routing theory. (2024) doi:10.31234/osf.io/7sz3w.

32. Friston, K. J., Harrison, L. & Penny, W. Dynamic causal modelling. *NeuroImage* **19**, 1273–1302 (2003).

33. Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B. & LeCun, Y. The Loss Surfaces of Multilayer Networks.

34. Sherfey, J. S., Ardid, S., Hass, J., Hasselmo, M. E. & Kopell, N. J. Flexible resonance in prefrontal networks with strong feedback inhibition. *PLOS Comput. Biol.* **14**, e1006357 (2018).

35. Dong, S., Wang, P. & Abbas, K. A survey on deep learning and its applications. *Comput. Sci. Rev.* **40**, 100379 (2021).

36. Soo, W. W. M., Goudar, V. & Wang, X.-J. Training biologically plausible recurrent neural networks on cognitive tasks with long-term dependencies. Preprint at https://doi.org/10.1101/2023.10.10.561588 (2023).

37. Rimehaug, A. E. *et al.* Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex. *eLife* **12**, e87169 (2023).

38. Billeh, Y. N. *et al.* Systematic Integration of Structural and Functional Data into Multi-scale Models of Mouse Primary Visual Cortex. *Neuron* **106**, 388-403.e18 (2020).

39. Roopun, A. K. concatenation underlies interactions between gamma and beta rhythms in neocortex. *Front. Cell. Neurosci.* **2**, (2008).

40. Börgers, C., Epstein, S. & Kopell, N. J. Gamma oscillations mediate stimulus competition and attentional selection in a cortical network model. *Proc. Natl. Acad. Sci.* **105**, 18023–18028 (2008).

41. Zhao, D., Zeng, Y. & Li, Y. BackEISNN: A deep spiking neural network with adaptive self-feedback and balanced excitatory–inhibitory neurons. *Neural Netw.* **154**, 68–77 (2022).

42. Bellec, G., Wang, S., Modirshanechi, A., Brea, J. & Gerstner, W. Fitting summary statistics of neural data with a differentiable spiking network simulator.

43. Wacongne, C., Changeux, J.-P. & Dehaene, S. A Neuronal Model of Predictive Coding Accounting for the Mismatch Negativity. *J. Neurosci.* **32**, 3665–3678 (2012).

44. Wacongne, C. *et al.* Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proc. Natl. Acad. Sci.* **108**, 20754–20759 (2011).

45. Erickson, M. A., Ruffle, A. & Gold, J. M. A Meta-Analysis of Mismatch Negativity in Schizophrenia: From Clinical Risk to Disease Specificity and Progression. *Biol. Psychiatry* **79**, 980–987 (2016).

46. Hamilton, H. K. *et al.* Mismatch Negativity in Response to Auditory Deviance and Risk for Future Psychosis in Youth at Clinical High Risk for Psychosis. *JAMA Psychiatry* **79**, 780 (2022).

47. Kirihara, K. *et al.* A Predictive Coding Perspective on Mismatch Negativity Impairment in Schizophrenia. *Front. Psychiatry* **11**, 660 (2020).

48. Fitzgerald, K. & Todd, J. Making Sense of Mismatch Negativity. *Front. Psychiatry* **11**,

468 (2020).

49. Chien, V. S. C., Maess, B. & Knösche, T. R. A generic deviance detection principle for cortical On/Off responses, omission response, and mismatch negativity. *Biol. Cybern.* **113**, 475–494 (2019).

50. Rappel, P. *et al.* Subthalamic theta activity: a novel human subcortical biomarker for obsessive compulsive disorder. *Transl. Psychiatry* **8**, 118 (2018).

51. Ostlund, B. D., Alperin, B. R., Drew, T. & Karalunas, S. L. Behavioral and cognitive correlates of the aperiodic (1/f-like) exponent of the EEG power spectrum in adolescents with and without ADHD. *Dev. Cogn. Neurosci.* **48**, 100931 (2021).

52. Peterson, E. J., Rosen, B. Q., Belger, A., Voytek, B. & Campbell, A. M. Aperiodic Neural Activity is a Better Predictor of Schizophrenia than Neural Oscillations. *Clin. EEG Neurosci.* **54**, 434–445 (2023).

53. Adams, R. A. *et al.* Computational Modeling of Electroencephalography and Functional Magnetic Resonance Imaging Paradigms Indicates a Consistent Loss of Pyramidal Cell Synaptic Gain in Schizophrenia. *Biol. Psychiatry* **91**, 202–215 (2022).

54. Sherfey, J., Ardid, S., Miller, E. K., Hasselmo, M. E. & Kopell, N. J. Prefrontal oscillations modulate the propagation of neuronal activity required for working memory. *Neurobiol. Learn. Mem.* **173**, 107228 (2020).

55. Ma, J. Biophysical neurons, energy, and synapse controllability: a review. *J. Zhejiang Univ.-Sci. A* **24**, 109–129 (2023).

56. Ardid, S. *et al.* Biased competition in the absence of input bias revealed through corticostriatal computation. *Proc. Natl. Acad. Sci.* **116**, 8564–8569 (2019).

57. Börgers, C., Epstein, S. & Kopell, N. J. Gamma oscillations mediate stimulus

competition and attentional selection in a cortical network model. *Proc. Natl. Acad. Sci.* **105**, 18023–18028 (2008).

58. Börgers, C. & Kopell, N. J. Gamma Oscillations and Stimulus Selection. *Neural Comput.* **20**, 383–414 (2008).

59. Christoph Börgers, Nancy Kopell. Synchronization in Networks of Excitatory and Inhibitory Neurons with Sparse, Random Connectivity. *Neural Comput.* (2003) doi:https://doi.org/10.1162/089976603321192059.

60. Adesnik, H., Bruns, W., Taniguchi, H., Huang, Z. J. & Scanziani, M. A neural circuit for spatial summation in visual cortex. *Nature* **490**, 226–231 (2012).

61. Nienborg, H. *et al.* Contrast Dependence and Differential Contributions from Somatostatin- and Parvalbumin-Expressing Neurons to Spatial Integration in Mouse V1. *J. Neurosci.* **33**, 11145–11154 (2013).

62. Grunze, H. *et al.* NMDA-dependent modulation of CA1 local circuit inhibition. *J. Neurosci.* **16**, 2034–2043 (1996).

63. Rujescu, D. *et al.* A Pharmacological Model for Psychosis Based on N-methyl-D-aspartate Receptor Hypofunction: Molecular, Cellular, Functional and Behavioral Abnormalities. *Biol. Psychiatry* **59**, 721–729 (2006).

64. Nakazawa, K. & Sapkota, K. The origin of NMDA receptor hypofunction in schizophrenia. *Pharmacol. Ther.* **205**, 107426 (2020).

65. Dallérac, G. *et al.* Dopaminergic neuromodulation of prefrontal cortex activity requires the NMDA receptor coagonist D -serine. *Proc. Natl. Acad. Sci.* **118**, e2023750118 (2021).

66. Lewis, D. A., Curley, A. A., Glausier, J. R. & Volk, D. W. Cortical parvalbumin

interneurons and cognitive dysfunction in schizophrenia. *Trends Neurosci.* **35**, 57–67 (2012).

67. Konradi, C. *et al.* Hippocampal interneurons are abnormal in schizophrenia. *Schizophr. Res.* **131**, 165–173 (2011).

68. Slowik, A. & Kwasnicka, H. Evolutionary algorithms and their applications to engineering problems. *Neural Comput. Appl.* **32**, 12363–12379 (2020).

69. Whitley, D. An overview of evolutionary algorithms: practical issues and common pitfalls. *Inf. Softw. Technol.* **43**, 817–831 (2001).

70. Yuille, A. & Kersten, D. Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* **10**, 301–308 (2006).

71. Friston, K. & Kiebel, S. Predictive coding under the free-energy principle. *Philos. Trans. R. Soc. B Biol. Sci.* **364**, 1211–1221 (2009).

72. Buffalo, E. A., Fries, P., Landman, R., Buschman, T. J. & Desimone, R. Laminar differences in gamma and alpha coherence in the ventral stream. *Proc. Natl. Acad. Sci.* **108**, 11262–11267 (2011).