A genetic algorithm for self-supervised models of oscillatory neurodynamics

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)

Predictive processing theories propose that the brain supervises itself, to build an internal model of its environment. This internal model emerges by minimizing the prediction error, the discrepancy between internally generated predictions and external sensory signals. Prior work has proposed that the neurobiological implementation of predictive processing involves neuronal oscillations in the gamma (40-100 Hz) and alpha/beta (10-30 Hz) frequency range. In current computational modeling approaches to predictive processing, there exists a trade-off between algorithmic and implementational aspects. One group of predictive processing models perform self-supervised computations but miss neurobiologic details and oscillatory neurodynamics. In the other group, biophysical models implement neural network models with maximal brain-like structure and oscillatory dynamics but require manual supervision. Here, we propose an evolutionary algorithm, the genetic stochastic delta rule (GSDR), to conduct simulations with biophysical neural networks that can inform predictive processing and other theories by linking the algorithmic and implementation levels. We first evaluate GSDR in a simplified and minimal optimization problem. Then, we simulate commonly observed neural dynamics such as firing rate and modulation of neuronal oscillations. We show that GSDR is capable of replicating oscillatory dynamics in the gamma and alpha/beta frequency bands observed through in-vivo electrophysiology, such that they emerge from synaptic plasticity. This methodology broadens the scope for biology-plausible, automated, large-scale and multi-objective simulations within computational neuroscience. With this virtuous cycle between data and models established by GSDR, we suggest that the search for the circuits underlying predictive processing can be grounded to neuronal data, improving the neurobiological basis of theories.

Author summary

In predictive processing theories, it is hypothesized that the brain creates an internal model of its environment. This is supported both by empirical and theoretical studies in neuroscience, which also emphasize the importance of neuronal oscillations in this process. Importantly, brain and cognitive disorders such as schizophrenia are associated with abnormal neuronal oscillations and impaired predictive processing. To gain further insights about these functions, it is important to build detailed biophysical models

where the underlying mechanisms can be explored and understood. Current models either perform computations without oscillatory rhythms or are biologically detailed but require manual tuning, which limits an unbiased search through parameter space. To address this, we developed a self-supervised evolutionary algorithm, the genetic stochastic delta rule (GSDR). We show that GSDR automatically tunes biophysical neural networks to find the optimal parameters to generate brain-like oscillatory dynamics, which are key features of predictive processing theories. Therefore, GSDR can help us find the real brain circuits that perform predictive processing.

Introduction

Understanding the brain requires linking theoretical mechanisms of neurons to empirical observations in-vivo. Biophysical models in neuroscience [1–6] fill in the gap between fundamental properties of neurons, neurophysiology, cognition and behavior. This provides more insights compared to models at a more abstract or phenomenological level, that while useful for computational applications, often lack explanatory power regarding the underlying mechanisms [7–10]. Biophysical neuronal network models, particularly those based on Hodgkin-Huxley equations [11,12], explicitly represent the ionic currents flowing across the neuronal membrane, the conductances of specific ion channels, and their voltage and time dependencies. This makes it possible to understand how changes in specific ion channel properties [13–15], cell types [16–18], cortical layers [19–21], neurotransmitters [22–24], neuronal metabolism [25–28], neuronal oscillations [29–32], and synaptic integration within neurons [32–34] affects perception, cognition and behavior.

Biophysical models of neuronal oscillations are particularly important for predictive processing theories [35–41] which propose top-down vs. bottom-up processing via distinct oscillations. Top-down processing is defined by the internal or cognitive state of the brain, such as the state of attention or prediction [42,43]. Bottom-up processing refers to how external sensory states are signaled [44–46]. Predictive processing theories propose that the brain constructs an internal model of the external world, and these oscillatory dynamics reflect the underlying computation [46–49]. One of theories that emphasizes the role of neuronal oscillations is predictive routing (PR) [47]. PR proposes how predictive processing is implemented in the brain given the empirical neurophysiologic observations during cognitive tasks. PR states that high-frequency oscillations, gamma (>40Hz) and higher spiking activity represent unpredictable and thus, informative bottom-up signals from lower order sensory cortex (e.g., V1/V2/V4, auditory cortex) towards the higher order cortical areas (e.g., higher-order sensory areas, prefrontal cortex). PR also states that lower-frequency oscillations represent top-down internal predictions [50], affecting the neural activity in the lower order areas by selective inhibition or disinhibition of bottom-up pathways. It is important to note that currently, PR describes and links two out of three levels of analysis in the brain [51,52], the implementational and computational levels, but has not yet specified the representational (algorithmic) level. Here, inspired by the evolutionary nature of neurobiology, we propose an algorithm to add the representational level, thereby connecting all three levels.

A mechanistic account of how these oscillatory dynamics emerge from neuronal circuits and how they perform predictive routing is currently unclear. Numerous studies have modeled oscillatory interactions with computational approaches [17,18,42,53–57]. These biophysical models are capable of generating oscillatory dynamics, including stimulus-evoked gamma oscillations. These models generally utilize Pyramidal-Interneuron networks (PIN), consisting of a strong input drive to pyramidal neurons and feedback inhibition from interneurons. The oscillation frequency depends on the inhibitory synaptic time constant and connectivity, such that a network can be manually-tuned to express PIN-gamma (PING) or PIN-beta (PINB). Note that networks of interconnected inhibitory neurons can also generate gamma oscillations without the participation of Pyramidal neurons, an alternative mechanism referred to as Interneuron Network Gamma [33]. In both PING and ING, the pattern of synaptic connectivity, internal noise and bottom-up inputs can generate oscillations which are either highly synchronous (strong PING, [1,2,17,53]) or sparse and irregular across neurons (weak PING, [58–61]). Biologically observed dynamics in healthy brains are more consistent with the latter, irregular-sparse rhythms.

One limitation of these detailed HH models is that they currently require manual-tuning of parameters. This makes it difficult to search the full parameter space when model building, as an exhaustive parameter sweep is often computationally intractable, and precludes flexible context-dependent optimization. Although some studies employ optimization methods, the problem is that they diverge from basic neurobiological principles [62–64]. Conversely, the brain functions as a self-supervised system. Self-supervised systems, such as the brain, exploit the inherent structure within the data without reliance on external supervision [65–68]. We propose that computational models should emphasize consistency with basic properties of neurobiology, including self-supervision.

In this work, we introduce a parameter optimization algorithm to train networks of Hodgkin-Huxley (HH) neurons with biological principles [69–72]. We call this method the Genetic Stochastic Delta Rule (GSDR), as it is inspired by stochastic delta rule [73–75] and genetic algorithms [76–78]. The method enables the exploration of complex and non-linear parameter spaces to optimize an objective or multiple objectives. Objectives can be any measurable feature of the model, even statistical features of responses generated by the model (such as the strength or synchrony of a neuronal population oscillation). Using simulations, we show that GSDR is capable of training both simple and complex models using extended HH equation models. In particular, models with GSDR were capable of shifting their dominant frequency between gamma and beta, and resembling stimulus-induced gamma oscillations observed in neurophysiological recordings from macaque monkey visual cortical neurons (area MT/MST). Comparative benchmarking revealed the advantages of GSDR compared to exhaustive parameter sweeping and to purely supervised methods. Our work shows that evolutionary algorithms like GSDR are broadly applicable to study complex dynamics in neurobiological systems. Together, these results show that biological self-supervised networks of neurons are capable of linking all three levels of analysis; computational, representational and implementational [51,52]. We discuss the results relative to predictive routing.

Results

Evolutionary exploration with GSDR: The overall strategy of Genetic Stochastic Delta-Rule (GSDR), is to give models the freedom to choose the trade-off between unsupervised and supervised paths (Fig. 1, see also Materials and Methods, Alg. 1, Eqn. 2, parameter α). The neuronal network model generates a series of neural responses (X[t], membrane potential of all neurons across time on trial t). In the unsupervised path, these voltages derive structural representations (R[t], i.e., mutual correlation dependent plasticity matrix) which is independent from external objectives (R[t] only depends on the model's activity, i.e., neuronal membrane potentials). In the semi-supervised path, secondary neuronal metric functions are defined, and applied on the neuronal response to calculate metrics at the current time step (M[t] := F(X(t)), set of metrics). Then, the calculated distance between each metric and its target value (T[t], set of all target values for metrics) defines the loss or reward value for the current trial (L[t] := Eval(M[t] , T[t])). This loss/reward value(s) could change the model's parameters U[t]. Over time, GSDR explores the parameter space of the model, changes the supervision, keeps the most optimal parameters that minimize the loss or maximize the reward and converges to the objective. A key advantage of the GSDR algorithm is that the parameter α, which controls the degree to which the model weighs the supervised vs. the unsupervised path, is itself dynamic - changing over time as models traverse parameter space. More specifically, when $\alpha = 0$, the model is fully supervised, and relies only on an externally defined objective. When $\alpha = 1$, the model is fully unsupervised, and explores the parameter space randomly and via its plasticity rule (Mutual Correlation Dependent Plasticity, see Materials and Methods).

Fig1: Flowchart of evolutionary optimization with GSDR. At each iteration, a simulation is performed to get signals (neuronal membrane potential in time) from the biophysical model. A metric or a set of metrics get measured based on signals (such as, power spectrum density or average firing rate). The difference between target and measured metrics is used to calculate the model's current loss (such as the similarity value between two signals) and passes it to the supervisor function (the delta rule, supervised path). In parallel, representational factors (here, mutual correlation between pairs of neurons) are calculated independent from

objectives (unsupervised path). The model gets updated by a linear average of both update values, weighted by the self-supervision parameter alpha. When alpha is 0, the algorithm becomes fully supervised and when alpha is 1, it becomes fully unsupervised. This parameter can be static or dynamic across iterations.

Advantage of stochastic exploration in noisy, non-linear spaces: We start by showing the advantage of evolutionary strategies in a modified version of the Rastrigin function [77,78]. Note that this simulation will showcase the behavior of GSDR only in the semi-supervised path (the self-supervised path is introduced in the models shown in Fig. 5,6,8). Since neurobiological circuits comprise elements with chaotic and stochastic behavior, we modified the Rastrigin function by adding a uniform noise process to it. We simulated a 1000-iteration run on the noisy Rastrigin function for GSDR, Gradient Descent (GD), and Stochastic Gradient Descent (SGD). The objective of each of these optimization functions was to find the global minimum. The results (Fig. 2) indicate GSDR as an evolutionary algorithm, is particularly effective for these types of parameter spaces, similar to other methods from the same family of algorithms [69,76–80]. Note that although this is a simple 2-parameter search space, it already shows that GSDR can be advantageous for non-linear parameter spaces. In later sections of the manuscript, we also benchmark different methods to GSDR for searching Hodgkin-Huxley biophysical network models, which are much more complex (see Supplemental Fig. 7,8,9).[69,76–80]

$$F(x_1, x_2) = 2(\eta . x_1 . x_2) + (x_1^2 - 2cos(2\pi x_1)) + (x_2^2 - 2cos(2\pi x_2)) ;$$
$$\eta \sim u(0, 1)$$

Fig2: Optimization in a noisy Rastrigin function. Three separate runs (A,B,C) are shown for 1000 iterations for each optimization method, with their corresponding progress. Notice the advantage of gradient-free exploration (red, GSDR) versus gradient-based methods (blue, Gradient Descent - GD, green, Stochastic Gradient Descent - SGD). Stochastic exploration (GSDR) functioned well despite a parameter-loss space that is highly non-convex and noisy. Gradient -based methods tended to get stuck in local minima.

Synaptic time constants in excitatory-inhibitory (E-I) networks: After establishing good performance with this toy optimization problem and other more basic scenarios in neuronal circuits (see Methods and Supplemental Fig. 1 for simple single neuron optimization problems and Supplemental Fig. 2 for optimizing population firing rates), we next tested our approach in a multi-objective optimization problem by setting two firing rate objectives for the population of excitatory (E) and inhibitory (I) cells. We implemented an excitatory-inhibitory (E-I) model comprising 100 neurons, including 80 excitatory and 20 inhibitory neurons (Fig. 3). This model incorporated specific mechanisms: AMPA

receptors for excitatory synapses and GABA receptors for inhibitory synapses, along with other parameters influencing connectivity such as synaptic connections, synaptic time constants, and both excitatory (E-noise) and inhibitory (I-noise) noise. Initial synaptic connectivity was established using a uniform matrix across all neurons. We trained the model to pursue two parallel objectives with equal weight, aiming for average firing rates of 20 spikes/sec for excitatory neurons and 70 spikes/sec for inhibitory neurons. Initially, the excitatory and inhibitory populations exhibited firing rates of 34.8 spikes/sec and 91.81 spikes/sec, respectively (Fig. 3A). Only the inhibitory synaptic time constants were permitted to change. As illustrated in Fig. 3, the model successfully met the target firing rates by increasing the inhibitory time constants, showcasing the capacity of GSDR in modulating mechanistic parameters to simultaneously optimize multiple objectives.

Fig3: E-I (Excitatory-Inhibitory) neuronal population multitarget firing rate optimization. (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 synaptic time constants and firing rate are inversely correlated, their relation is not perfectly linear (i.e, multiplying the time constant by 2 will not necessarily divide the frequency by 2). In the model (Fig. 3), the ratio of time constant (post to pre) is ~1.6 but the corresponding ratio for average firing rates are different (1.77 for the excitatory population, 1.28 for the inhibitory population). This sort of nonlinear interactions makes it difficult to predict solutions and manually tune biophysical models to reproduce desired responses in multi-objective cases.

Fig4: E-I natural (dominant) 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) neural circuit schematic. (E) Initial frequency response. (F) Frequency response after training showing that the neural circuit has achieved the target metrics. (G) Initial raster plot. (H) Final raster plot.

Population natural frequency with noise and synaptic time constant: Next, we used the same neuronal model implemented in Fig. 3 to tune the natural frequency of its average neuronal potential. The natural frequency is defined as the frequency with the maximum oscillatory power in a neural circuit (Fig. 4E-F). The pre-training natural frequency of the excitatory population was 91Hz (Fig. 4A). For training, inhibitory synaptic

time constants and baseline Gaussian noise gain were allowed to vary in the model. As seen in Fig. 4, the model converged to the target natural frequency (30Hz) by increasing inhibitory time constants. The results indicate that 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. 4A,C). We also note that the model achieved the objective by creating a highly inhibited state for the excitatory E cells (Fig. 4H), consistent with an ING mechanism, where a neuronal oscillation is sustained by the I cells. Note that no other objective other than the peak frequency of the model was specified, but in principle, other objectives can be explicitly stated (e.g., to maintain E cell firing above a particular threshold) if desired. It is interesting that the model has converged to a beta rhythm that leads to overall more network inhibition (amongst the E cells), a result consistent with the suggestion of Predictive Routing that beta states are associated with prediction-driven inhibition.

Fig5: *E-I population spectral power ratio tuning. (broad Beta [15-30Hz] 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.*

Spectral push-pull through synaptic self-modulation: Next, we simulated the oscillatory push-pull interaction in an E-I neuronal population. We trained the model to implement the flexible switch between oscillatory motifs, particularly the observed gamma-beta push pull

seen in predictive routing [47,81,82]. 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 Figs. 3-4 but synaptic connections are included as training variables, increasing the number of free parameters in the model. 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. 5C), the objective is to have higher gamma band (35-100Hz) power than beta (15-30Hz), and in condition 2 (simulation of no stimulus processing, Fig. 5D), higher relative beta band power than gamma.

The model achieves the conditional objective (Fig. 5, more gamma than beta on condition 1, and more beta than gamma on condition 2). The change in total synaptic gain (sum of all synaptic connectivity weights) is negligible (Fig. 5A, C), yet the variance of synaptic connection weights has been modulated to match the objective. This shows that synaptic plasticity is capable of modulating the relative spectral response of a neural circuit, without changing time constants or population average/total gain.

In addition, the transition from a gamma-dominated state (condition 1, Fig. 5C) to a beta-dominated state (condition 2, Fig. 5D) resulted in substantially suppressed population activity, especially amongst the excitatory cells, consistent with a suppressive influence of beta on overall spiking activity [83] and with predictive routing models [47,48].

Contextual spectral push-pull with top-down modulation: Predictive routing proposes that top-down inputs carrying predictions (via beta) serve to suppress bottom-up processing (via gamma) in a selective manner [47]. Therefore, we modeled how top-down inputs can modulate the network state to shift dynamics from a gamma to a beta dominated spectral state. Note that we did not model two regions simultaneously but chose instead to model top-down inputs to a hypothetical lower-order area, similar to previous work modeling the top-down effects of attention on sensory processing [42]. In this simulation, the model's goal is to tune its spectral relative power between beta and gamma (Fig. 6). This is similar to the previous model (Fig. 5) 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 inputs to each of the neuronal populations included in this model (Fig. 6B): 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. In condition 1, top-down inputs were set to 0 such that the model was at the baseline (gamma dominated state, i.e., a beta to gamma ratio below 1, Fig. 6E). Next, GSDR was allowed to modulate these top-down gain parameters in the model with the goal of shifting the network to a beta-dominated state in condition 2. Interestingly, GSDR found a solution with positive gain on the slower inhibitory population but a negative gain for the fast inhibitory population (Fig. 6C, red box). Post-training responses showed a context-dependent spectral power shift (Sup. Fig. 3C-F) with respect to the beta (~10-30Hz) to gamma (>35Hz) power ratio objective, which ended at a ratio of ~1.8 (Fig. 6F). The post-training gain parameters are consistent with studies suggesting that top-down feedback drives slow inhibitory interneurons [18,84] and shifts the network towards a beta dominant state [47,85].

Fig6: E-I contextual spectral relative power tuning. (Beta [~10-30Hz] 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). Consider that due to stochasticity of the model caused by factors such as the internal noise and receptor-synaptic interactions in the HH equation, the beta to gamma ratio on condition 1 of the pre vs. post model slightly varies.

Narrow-band gamma oscillations are observed in electrophysiology recordings from the primate visual cortex : The previous model (Fig. 6) generates overly synchronous oscillations (Sup. Fig. 3D) that do not resemble normal in-vivo electrophysiology. To address this problem, we used synaptic plasticity to model electrophysiologically recorded, visually induced gamma oscillations (~38Hz) from the visual cortex of awake macaques. The objective was to let GSDR modulate synaptic connectivities in order to resemble observed gamma band neuronal dynamics across stimulus on/off 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 [42,58–60]. In strong PING, gamma band oscillations (>35Hz) 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. 5C). In weak PING models[58–60], noise is added to the input signal or synaptic weights to generate a gamma oscillation with sparse, irregular firing in the pyramidal cells.

Fig7: Single unit neural response from visual areas MT/MST (NHP) (A) PSTH and all-trials cumulative raster plot (100 trials, 80 neurons) (B) Trial-average cumulative raster plot of neurons (blue line: fixation, green line: onset of stimuli, red line: offset). (C) Corresponding time-frequency response (scaled to 0 as the lowest power and 1 for the highest power across the entire duration of the signal shown above), showing a stimulus dependent evoked spectral response around ~37.5Hz.

Prior to the simulations, we examined whether real neurons display a strong or weak PING phenotype in-vivo. We performed electrophysiological recordings in the visual cortex of awake macaque monkeys during presentation of drifting grating stimuli, known to induce gamma oscillations (see methods for details about neurophysiology). We used neuronal data of the area MT/MST since it was highly sensitive toward moving grating stimuli, and we consistently observed a strong narrow-band gamma response during presentation of the visual stimuli (Fig. 7).

During the pre-stimulus baseline the neuronal population had an oscillatory resonance around beta (~20Hz, Sup. Fig. 5&10). During presentation of the stimuli, strong narrow-band oscillatory gamma was present in the population spiking raster (Fig. 7A) and population raster (Fig. 7B). In addition, time-frequency analysis of the population spiking activity indicated the presence of a sustained narrow-band gamma (~37.5Hz, Fig. 7C) 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.

Self-supervision guides parameter space exploration: To study the circuit mechanisms that produce this gamma oscillation, we implemented a model with 50 neurons (36 excitatory, 14 inhibitory). For this model, instead of relying on synaptic time constants, inhibitory neurons were defined based on two classes of connectivity [21,86]. One class had a pattern of locally-focused connectivity (Parvalbumin-like, restricted to establish local synaptic connections, to ~10% of total neurons, Sup. Fig. 4) and another class had a global widespread connectivity (Somatostatin-like, with 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 change or modulate the synaptic time constants in our model. The decay and rise (tauD and tauR) time constants were static in synaptic mechanisms [87,88]. Instead, the model was allowed to modulate the synaptic connection conductances. The objective of the model was to approximate the spectrotemporal features present in the observed neurophysiology (based on maximizing the correlation between model and data power spectra(Fig. 7A and Supp. Fig. 10 f) in both the pre-stimulus fixation and stimulation time period. We refer to this as the spectral similarity task.

Note that to perform the spectral similarity task while allowing only synaptic weights to be trainable is an objective which is more complex than the objectives used in Figs. 1-6. For this task, we used the self-supervision component of GSDR (see Algorithm 1). Self-supervision refers to the ability of a model to modulate its update rule based on both unsupervised and supervised learning terms. The unsupervised term is calculated by the correlation between synaptic potentials using MCDP (see Methods, equation 3 and Algorithm 2). The supervised term is the output of the inner optimizer, which requires the loss and a model's parameters for calculating the supervised changes (see Methods, equation 2.1 & 2.2 and Algorithm 1). The tradeoff between these terms is controlled by the parameter alpha (ranging between 0 and 1), with values of alpha = 1 corresponding to a fully unsupervised algorithm and values of alpha = 0 leading to a fully supervised algorithm. Also, the algorithm always included a genetic component, meaning that optimal states were kept throughout training, and sub-optimal exploration could trigger a deselection criteria which would revert the model back to the last optimal state (see Algorithm 1). Comparative benchmarking on the spectrotemporal similarity task demonstrated that self-supervision (especially with alpha between 0.6 to 0.8) led to better average model performance compared to fully supervised models (Supplemental Fig. 7&8) and to more traditional parameter sweep approaches (Supplemental Fig. 9), given similar computational time. Note that fully unsupervised models were qualitatively similar on average to self-supervised models, but the best-performing models included self-supervision (Supplemental Fig. 8A).

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 optimized its spectral response to regenerate the stimulus-evoked gamma without replicating the same spectral responses in the absence of the stimuli. The loss (used to train the model) was calculated using the log-ratio (see methods) between the model and the data's spectral response (Supp. Fig. 10), and mean exponential error (MXE, see eqn. 5 in methods section) between spike count of the model and minimum spike counts allowed (to avoid sub-optimal spiking patterns, such

as spike-less responses). 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 noisy direct current-DC step function with amplitude of 1nA convolved to a 120Hz pulse train) with a 500ms duration. Background noise was a stationary random process current [89], which is always present for every neuron independently.

Fig8: Single unit neural responses from the model. (A) Time-frequency response (scaled to 0 as the lowest power

and 1 for the highest power, 1/f adjusted) of the model , pre-training (B) Corresponding raster plot, reflecting highly synchronized spiking. (C) Time-frequency response of the model , mid-training (trial ~200) (D) Corresponding raster plot, mid-training. (E) Time-frequency response of the model , post-training (F) Corresponding raster plot, mid-training. (G) PSTH of all neurons in panel F. During stim-on condition, stimulation was randomly provided to half of the excitatory neurons in the model, from 0ms (onset, green dashed line) to 500ms (offset, red dashed line).

The model's initial synaptic weights were set-up to be uniform across all connections coming from excitatory and global-inhibitory interneurons. Both the excitatory and the global interneuron population had no synaptic connectivity constraints, yet for the local-interneuron population, connectivity was limited to neighboring neurons (Sup. Fig. 4 A,B). This initial state of the model generated the pre-training responses (Fig. 8, A-B). As can be seen, the spectral response of the model increases during stimulus presentation (Fig. 8, A-B) but in a broad-band manner that did not capture the narrow-band gamma response present in the data (Fig. 7C,E). After 200 trials (Fig. 8, C-D) the model showed enhanced gamma amplitude during stimulation but still not in a narrow band regime. After 500 trials , the model achieved a plausible stimulus dependent narrow band gamma response and pre-stimulus spontaneous dynamics (Fig. 8, E-F). Namely, model dynamics and neurophysiology data showed a spectral peak at beta (~15-20Hz, Sup. Fig. 6A) in the pre-stimulus baseline activity. This shifted spectral peaks of the model's dynamics towards gamma (~35-40Hz, Sup. Fig. 5B). The spectral response of the empirical data and the model early vs. late in training is shown in Supplemental Fig. 10, separately for the pre-stimulus period and the stimulus presentation period. The corresponding loss values are shown, indicating a major improvement as the spectral similarity between model and data improves with learning.

This demonstrates that biophysical models trained with GSDR are able to perform well on the spectral similarity task without solely relying on external supervision or manual dictation of parameters. Specifically, recapitulating the observed gamma band oscillations did not require manual adjustment of time constants or strong and uniform I to E synaptic connections (as in previous models). Instead, it achieved the post-training spectral state via sparse and non-uniform synaptic connections (Sup. Fig. 4A-B). This further suggests that the biophysically plausible models when trained on real data implement oscillatory dynamics with weak PING characteristics, e.g. the oscillation is mostly observed at a population level, not at the single neuron (Fig. 7&8).

Discussion

Our work established a novel learning rule, GSDR, for training biophysically realistic neuronal networks. This learning rule combined aspects of supervised learning (the Stochastic Delta Rule) with unsupervised learning (using Mutual Correlation Dependent Plasticity, MCDP) and a stochastic, genetic component (trial and error learning, with optimal solutions kept and sub-optimal solutions de-selected). We established that GSDR worked well for both simple and complex problems. We showed that GSDR can be used to flexibly shift a network between beta and gamma oscillatory states, which was motivated by the predictive routing model. In previous work, the microcircuitry of this flexible

state-dependent spectral shift was unknown. We used GSDR to implement this shift in a neurobiologically plausible manner via three distinct approaches. First, by manipulating local E-I connections (Fig. 5), second, by manipulating the top-down contextual inputs to a network with E, slow (SST-like) I, and fast (PV-like) neurons (Fig. 6), and third, by fitting the power spectra of neurons in the model to an observed time-frequency response in electrophysiology with synaptic weight changes (Fig. 8).

What has GSDR and evolutionary approaches taught us about biophysical models?

Many biophysical models perform complex computations [90–92], yet their complexity and biophysical constraints [20,90,91] often[90–92] require extensive manual tuning. Self-supervised learning, exploration, and genetic algorithms offer compelling alternatives for optimizing parameters in neuronal circuits, as they alleviate the burden of this manual process. Building on the principles of evolutionary learning [69,79,80], the GSDR algorithm is capable of optimizing detailed biophysical models. The primary novelty of GSDR lies in its dynamic self-modulation factor (see Methods Eqn. 2, parameter alpha), which allows the model to dynamically adjust the trade-off between an exploratory, unsupervised path and a reinforced-supervised learning path. This dynamic mixture enabled the model to discover complex connectivity patterns (Sup Fig. 4B) that would have been difficult and costly to find through manual tuning or via exhaustive parameter sweep (Supp. Fig. 9). The choice of specific supervised components (e.g., the Stochastic Delta Rule, Stochastic Gradient Descent, and others) within the GSDR framework opens up future questions regarding the optimal pairing of unsupervised and supervised methods (Supp. Figs 7&8), which will be critical for detailed benchmarking compared to the other semi-supervised approaches.

What has GSDR taught us about oscillatory dynamics?

Previous studies [42,55,56,59,92] regarding biophysical mechanisms contributing to oscillatory dynamics have suggested that the gamma oscillation can act as a winner-take-all mechanism via Pyramidal-Interneuron Network Gamma dynamics [53,55]. 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 [55,59]. In previous models, the PING mechanism could generate either highly synchronous (strong PING [53,55]) or sparse and irregular oscillations [58–61]. Our empirically recorded neurophysiology data was more consistent with a sparse and irregular weak-PING mechanism (Fig. 7). Interestingly, in the spectral similarity task, in the post-training state, the model exhibited a brain-like weak-PING mechanism [58–60], as it was sparse across individual neurons (Fig. 8). In addition, the narrow-band gamma oscillation did not dominate the response during spontaneous activity, which was instead dominated by beta (Supp. Fig. 10). The flexible transition from pre-stimulus beta to stimulus-induced gamma was learned via exploration of synaptic connectivity. Therefore, the same neurons within a network are capable of expressing either beta or gamma depending on the synaptic inputs to the model[55,96].

What has GSDR taught us about predictive routing and oscillatory shifts?

Our results also showed 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. 6) to shift network dynamics between gamma and beta despite strong

and continuous synaptic drive. This supports the predictive routing model in which top-down, preparatory signals are used to prime cortical networks for processing predicted stimuli and shift processing of predictable inputs from gamma to beta.

We also hypothesized that changes in local synaptic connectivity are sufficient to induce a spectral shift between beta and gamma (Fig. 7). 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 [93]. Therefore, GSDR has taught us that multiple network implementations (bottom-up inputs, top-down inputs, as well as local intrinsic circuitry) are capable of inducing beta-gamma push-pull dynamics in realistic biophysical models. Notably, in both scenarios in which we explored the beta-gamma push-pull dynamic (Figs. 5, 6, Supp. Fig. 3), the beta state resulted in a network dominated by dynamics of the I neurons, where the E neurons were selectively suppressed, consistent with an ING mechanism [33]. This is consistent with predictive routing, which hypothesizes that inhibitory states associated with highly predictable stimuli are implemented via preparatory rhythms in the beta frequency range [47,48].

Connections to other computational modeling frameworks

Out of several classes of biophysical computational models, Dynamic Causal Modeling (DCM) [94–98] has been used to train models to simulate observed brain dynamics. DCM implementations generally assume that the population scale (summarizing the behavior of large neuronal masses, or populations of neurons) is a sufficient level of complexity to explain brain dynamics. In contrast, in our view it is important to also consider the microscopic elements of neural circuits (single cells, the synaptic weight matrix and their neuro-chemical interactions) to resolve current questions about the circuitry that underlies predictive processing and predictive routing. Additionally, current DCMs assume that the equations describing the model are solvable by optimization methods such as gradient descent, implying a convex and locally linear parameter space explaining the model's behavior [99]. In contrast, both our own simulations (Fig. 2, Supp. Fig. 7 & 8) and previous studies have shown that neuronal dynamics can be highly-non convex and non-linear across subtle parameter changes [59,92,100].

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 and neuromodulators, that inhibit or excite other neurons, leading to spectro-temporal dynamics observed both in-vivo [81–83,93,101,102,102–104] and in-silico [16,18–20,55,59,92,100,105]. In addition, studies suggest that microscale mechanisms might directly or indirectly perform predictive processing [106–108]. Thus, we emphasize the importance of making models as similar as possible to neurobiology. Such models necessitate a complex, non-linear parameter space which may require genetic, self-supervised training techniques such as GSDR to train (see Fig. 2 and Supp. Figs. 7&8). More exhaustive comparative benchmarking of different techniques should be performed in future work to establish this definitively.

Future extensions of the evolutionary modeling framework.

Evolutionary strategies open an alternative class of representational algorithms for computational modeling in neuroscience. This could be used to test which class of predictive coding (PC) algorithms is the best candidate to match neuronal data [109]. In parallel, the multi-objective strategy can be used to combine the representation proposed by an algorithm (e.g., linear PC, Rao-Ballard PC, PC biased competition (PC-BC)) with neuronal data. This strategy could be useful in determining which algorithm is more likely to be implemented in the brain and compatible with the neuronal data.

In principle, any observed neurophysiology data can be investigated and modeled with this framework. GSDR can be extended to evaluate LFP-spiking relations, properties of spike timing and rate coding, spike-field coherence (SFC) [110], neuronal metabolism and neurochemical-neuromodulator interactions [23,105,111], the observation of multiple co-existing gamma oscillations [112–114], distinct mechanisms for beta vs. alpha oscillations [115–118], and any other property that is measurable from a model and can be compared to data. We note that the models we have trained in the present manuscript contain relatively few neurons (~100 neurons). However, we also showed that these smaller models can be augmented to simulate the dynamics of larger models (for a 1,000 neuron simulation based on augmentation, see Supp. Fig. 11). In future work, we plan to extend these current models to multiple areas and with larger numbers of neurons/parameters to study large-scale network interactions [118]. This is especially important because modern neurophysiological techniques also enable thousands of neurons to be studied during task performance [119]. GSDR enables the relationship between empirically observed neurophysiology and biophysical model parameters to be discovered in a self-supervised manner. Until now, this has not been achieved at the neural circuit level, where model fitting is typically performed using manual tuning. Automation of this tuning is a significant advance of this work.

We propose that GSDR be applied using different objectives on various neurophysiological datasets. In an iterative approach, the models would be trained to satisfy multiple objectives under varying datasets and tasks. The end goal would be to produce powerful, domain-general models covering a large parameter space capable of explaining large amounts of neuronal data. This would allow for large-scale in-silico exploration of parameters that lead to dysfunctional states underlying brain disorders as well as for testing various model-based strategies for treatment. Advanced and detailed models could also be applied to study which implementation of predictive processing [39,68,85,105,109,120–124] is the most consistent with recorded neurophysiology. In this way, models and data can mutually inform one another, and GSDR can help close the gap between the algorithmic and implementation levels. The outputs of such in-silico explorations could open up new questions, which could then be verified using a new round of focused experiments, which feed back onto modeling. This framework in which data and models in neuroscience reinforce each other would create a virtuous cycle of continual refinement and discovery.

Materials and Methods

Here we used the DynaSim toolbox (Matlab) and Jaxley toolbox (Python) to implement biophysical neural circuit models [125,126]. These models consist of one or multiple networks of neurons with different cell types modeled by corresponding Hodgkin-Huxley (HH) equations [11,12]. In addition, equations can be extended by including additional mechanisms in (e.g., equations that describe receptor properties, ion-channels and synaptic mechanisms). The system of equations that govern how voltage changes over time in neurons of a neural population is given in equation 1:

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

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}), $I_{input}(t,\ V)$ 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.

Genetic Stochastic Delta-Rule (GSDR), which is inspired from the stochastic delta rule [74] and genetic-optimization learning [69], explores the free parameter space(s) and deselects the sub-optimal states in order to achieve a single or multiple objectives. The objectives are target values, for metrics that are derived from the model (e.g, membrane potential, firing rates and spectral population dynamics). The general form of GSDR update step is shown in equations 2.1 and 2.2:

$$U_t = U_{t-\Delta t} + (L(t)\ \delta(\lambda)(1 - \alpha) + \alpha)\ .\ R \quad [2.1]$$

$$D(t) = (U_t - U_{t-\Delta t})\ .\ (L(t) - L(t - \Delta t)) \quad [2.2]$$

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 that determines how quickly parameters can change, D(t) is the temporal difference that is the core logic of the delta rule (i.e, the rule that supervises the optimization) $L(t)$ is the output of the evaluation function (see methods, defining objectives in GSDR) at that time, α is the supervision factor ($1 \geq \alpha \geq 0$, if $\alpha = 1$, the model will be fully unsupervised) and R is the model's structural representation. In this work, we used two methods for supervision. One, which we implemented in Jax is the Stochastic Delta Rule (equation 2.3). As a comparative benchmark, we also used the existing Stochastic Gradient Descent (SGD [127,128]) in Jax [126,129] (equation 2.4). We note that by implementing GSDR in Jax, any alternative method of supervision is also available to be used as the supervisor. R was defined using a simplified estimation of spike-time dependent plasticity (STDP) [130], which we call mutual-correlation dependent plasticity (MCDP, eqn.3, alg.2). Note that MCDP only applies for synaptic connectivity properties, since it requires mutual activity between two distinct

elements (such as a synaptic connection or the conductance of a receptor). This will in effect enhance connections between neurons that are mutually correlated (i.e., that have similar activity levels over time). In this work, MCDP-based self supervision was utilized in the simulations with synaptic plasticity, presented in Fig. 5&8.

Algorithm 1 Genetic stochastic delta rule

$U, \lambda, \alpha, C, L_{opt} \leftarrow Init(.)$          $\triangleright$ Initialization

while $Training$ **do**

 $M \leftarrow Metrics(U)$        $\triangleright$ Measurements from the model

 $X \leftarrow Reponse(U)$    $\triangleright$ Membrane potential responses $(Neurons \times Time)$

 $L(t) \leftarrow Evaluate(M, \hat{Targets})$        $\triangleright$ Eqn.4-5

 $R \leftarrow MCDP(X(t))$        $\triangleright$ Eqn.3 and Alg.2

 if $L \leq L_{opt}$ **then**

  $L_{opt} \leftarrow L$

  $U_{opt} \leftarrow U$

  $\alpha \leftarrow \alpha_{opt}$

 else if $L \geq C \times L_{opt}$ **then**       $\triangleright$ Deselection condition

  $U \leftarrow U_{opt}$

 else

  for $W_{pre,post} \in U$ **do**        $\triangleright$ Synaptic weights

   $D \leftarrow SDR(U, T)$        $\triangleright$ Eqn.2.2

   $\Delta \leftarrow (L.D.\delta(\lambda)(1 - \alpha) + \alpha).R_{pre,post}$    $\triangleright$ Eqn.2.1

   $W_{new} \leftarrow W + \Delta$

   $\alpha_{new} \leftarrow \alpha + \eta$       $\triangleright \eta \sim u(-0.1, 0.1)$

   $W \leftarrow Clamp(W_{new})_{[0,max_g]}$    $\triangleright$ Clamp synaptic weights

   $\alpha \leftarrow Clamp(\alpha_{new})_{[0,1]}$    $\triangleright$ Clamp self-supervision

  end for

 end if

end while

---

*Alg1: GSDR pseudocode. Model parameters that are trainable are included in U (eqn. 2-3). After initializing model and optimization parameters (α: self-supervision factor, **λ**: exploration/learning factor, **C**: deselection limit criteria) the training loop begins. At each trial, a set of simulations given current parameters are performed while recording the neuronal membrane potentials **X**. Several metrics are calculated from the model responses X to be used in order to calculate the loss (such as eqn. 4-5). Representational factors **R** are calculated solely based on X. If the batch loss is better than previous optimal loss (initially, inf), the parameters are selected as the optimal state (this is the genetic component). If the batch loss hits the deselection condition (loss becomes C times greater than the optimal loss) the model parameters get reverted to the last optimal state. For each set of parameters in U, update step is performed (eqn. 2.1, 2.2). MCDP values are also calculated but for parameters that are definable (see eqn. 3 and alg.2), otherwise MCDP is 1 (identical for other parameters). Alpha parameter also changes with uniform random msteps, but gets reverted to the optimal value the same way as other parameters. Parameters that are required to be within a certain limit are clamped (to ensure certain variables are non-negative and do not exceed highly non-physiological values, such as synaptic conductances).*

MCDP gets the membrane potential responses X as input, calculates the correlation between each pair of neuronal populations, V_N (pre-synaptic potentials) and V_M (post-synaptic potentials, shifted in time based on the corresponding synaptic time constant, such as τ_{AMPA} for the AMPA synapses, τ_{GABA} for GABA and etc …). The output of MCDP is a matrix with size of N (pre-synaptic neurons) and M (post-synaptic neurons). This serves as a structural representation factor, similar to the STDP but using membrane potential of single neurons with emphasis on synaptic timing. In fact, MCDP increases if two neurons have similar membrane potentials, which also contains their action potentials, subthreshold dynamics and spiking waveform. For each synaptic connection (that exists), post-synaptic potentials are shifted by the corresponding time constant of the synapse in MCDP calculation in order to take timing parameters into account. the Details are in this equation 3:

$$ R \; = \; MCDP(X(t)) \; = \; \frac{r' - \mu_r}{\sigma_r}, \; r_{N \times M} \; = \; Corr(V_{t \times N}, V_{(t-\tau) \times M}) \;\; [3] $$

The value of R will be 1 for any variable in U_t, if MCDP is not valid for it, or in case of model parameters such as certain ion-channel properties (if included in the simulation).

Algorithm 2 Mutual correlation dependent plasticity

$N \leftarrow rows(X(t))$ ▷ $X_{N \times N}(t)$: Membrane potential matrices
$R \leftarrow 0_{N \times N}$
for $W_{pre,post} \in U$ **do**   ▷ Synaptic connection parameters
     $\tau' \leftarrow \tau_{mech}$ ▷ $mechanism \leftarrow AMPA, GABA, ...$
     $\hat{V}_m \leftarrow X_{post}(t - \tau')$ ▷ Post-synaptic potentials (shifted)
     $V_n \leftarrow X_{pre}(t)$ ▷ Pre-synaptic potentials
     $R'_{n \times m} \leftarrow corr(V_n, \hat{V}_m)$ ▷ Eqn.3
     $R_{[n \in pre, m \in post]} \leftarrow R'_{n \times m}$ ▷ Append to representations
end for
$R \leftarrow (R - \mu_R)\sigma_R^{-1}$ ▷ Eqn.2.1, R

Alg2: MCDP pseudocode. For parameters that are shared between a pre-synaptic and post-synaptic neuron, the model potential matrix X is used to calculate the cross correlation. For the post-synaptic matrix, signals are shifted in accordance with the time constant of that pair to take the timing of the reference signal (pre-synaptic) into account. After acquiring all MCDP factors, the normalization step (eqn. 3) is performed and the representational factor matrix R is returned to be used in the unsupervised term of GSDR (alg. 1, eqn. 2.1)

GSDR can include objectives, such as specific firing rates, spectral or temporal patterns. This requires defining the corresponding metric-target for each objective. A metric (M) can be a direct or indirect measurement from the model, a function that is applied on the model's response. 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) consist of one or multiple target values (T). The loss function in the single neuron simulations (Sup. Fig. 1) and the initial population simulations (Fig. 3-4) were calculated by the mean square error (MSE) between the 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 (Sup. Fig. 1), 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. Any spectral analysis involving the power spectrum density (PSD) of model and electrophysiology data is performed on neuronal membrane potentials (Hanning window) and then averaged each spectrum bin across neurons.

$$L_{LR} = LogRatio(M_t, T_t) = \sum_{i \in Metrics} |Log(\frac{Metric_i}{Target_i})| \quad [4]$$

In addition, we defined a one-sided loss function (eqn. 5, mean exponential error, MXE) penalizing models with a very low spike count. This was necessary because without this penalty, some models produced little to no spiking.

$$L_{MXE} = MXE(M_t, T_t) = \frac{1}{N} \sum_{i \in Metrics} e^{-(Metric_i - Target_i)} \quad [5]$$

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 the model falls out of the biological constraints in case there are no biologically plausible boundary conditions. Boundary conditions can be used to keep neuronal circuits biologically plausible and stable, such that the model parameters stay in a limited parameter subspace. For example in this work, parameters such as synaptic connectivity are bound to be non-negative and scaled to maximum strength of 1 (a.u).

Integrating GSDR into biophysical neural network models: As the basic step, we use a minimal biophysical circuit model to generate neuronal dynamics such as membrane potentials that change over time. We model individual excitatory/inhibitory neurons's

membrane fluctuations. For single neuron parameter optimization tasks (e.g., see Sup. Fig. 1), GSDR is simply the SDR since only the supervised path is included 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 potential time courses, 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 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 iterating until the objectives are achieved.

We have included some basic examples of GSDR utilization in supplements (see sup. fig 1-3) 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 Sup. Fig. 1A) to a target of -60.0mV. The trainable parameter was the reversal potential of the leak channel (E-leak, Sup. Fig. 1A) during training. Comparing pre to post training, the membrane potential shifted to the desired target (Sup. Fig. 1A). In the next simulation (Sup. Fig. 1B) 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,we manually stopped early due to a relatively negligible difference between the measured metric and the target (average membrane potential=-60.0mv, target=-60mv) .

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 Sup. Fig. 1E, comparing pre to post training, the firing rate approached the desired target. In Sup. Fig. 1F it can be seen that the model's loss (mean square error between target and measured firing rate) reduced after exploratory behavior (Sup. Fig. 1F). At trial 52, the model was manually stopped due to a relatively negligible loss.

We have also included simulations of neuronal populations. For this, we implemented a model containing 100 neurons (80 excitatory, 20 inhibitory) based on the channel mechanisms (AMPA, GABA [87,88]) and parameters for their connectivity (synaptic connection conductances, synaptic time constants, intrinsic E-noise and I-noise [89]). 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 (Sup. Fig. 2C) and the target was 10 spikes/sec. Only inhibitory synaptic time constants were allowed to be changed by the model. As seen in Sup. Fig. 2, the model achieved the target average

firing rate in the excitatory population by increasing inhibitory time constants which effectively decreases the overall spiking activity.

Experimental Design

In-vivo neural signal recordings: All experimental procedures were approved by the Vanderbilt University Institutional Animal Care and Use Committee. 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 the 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[131].

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[132]. 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 [133].

We estimated each neuron's spectrogram (time-frequency response) using multitaper spectral estimation with a Kaiser window of length 400ms and 98% 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 during stimuli-on (Fig. 8) became GSDR's objective in order to generate neural responses similar to it.

Acknowledgements

This work was supported by NIH-R00MH116100 (AMB), the NSF CAREER grant 2449277 and Vanderbilt startup funds (AMB). We thank Nancy Kopell (and the Cognitive Rhythms and Cognition - CRC team), Homero Esmeraldo, Eli Sennesh for their critical comments to an initial version of this manuscript.

Code availability: Github repository:

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

https://github.com/HNXJ/GSDR

References

1. White JA, Chow CC, Ritt J, No CS-T, Kopell N. Synchronization and Oscillatory Dynamics in Heterogeneous, Mutually Inhibited Neurons.

2. Whittington MA, Traub RD, Kopell N, Ermentrout B, Buhl EH. Inhibition-based rhythms: experimental and mathematical observations on network dynamics. Int J Psychophysiol. 2000;38: 315–336. doi:10.1016/S0167-8760(00)00173-2

3. Lee S, Sen K, Kopell N. Cortical Gamma Rhythms Modulate NMDAR-Mediated Spike Timing Dependent Plasticity in a Biophysical Model. Graham LJ, editor. PLoS Comput Biol. 2009;5: e1000602. doi:10.1371/journal.pcbi.1000602

4. Petousakis K, Apostolopoulou AA, Poirazi P. The impact of Hodgkin–Huxley models on dendritic research. J Physiol. 2023;601: 3091–3102. doi:10.1113/JP282756

5. Ma J. Biophysical neurons, energy, and synapse controllability: a review. J Zhejiang Univ-Sci A. 2023;24: 109–129. doi:10.1631/jzus.A2200469

6. Cattani A, Arnold DB, McCarthy M, Kopell N. Basolateral amygdala oscillations enable fear learning in a biophysical model.

7. Song HF, Yang GR, Wang X-J. Training Excitatory-Inhibitory Recurrent Neural Networks for Cognitive Tasks: A Simple and Flexible Framework. Sporns O, editor. PLOS Comput Biol. 2016;12: e1004792. doi:10.1371/journal.pcbi.1004792

8. Zhao D, Zeng Y, Li Y. BackEISNN: A deep spiking neural network with adaptive self-feedback and balanced excitatory–inhibitory neurons. Neural Netw. 2022;154: 68–77. doi:10.1016/j.neunet.2022.06.036

9. Zhu R-J, Zhao Q, Li G, Eshraghian JK. SpikeGPT: Generative Pre-trained Language Model with Spiking Neural Networks. arXiv; 2024. doi:10.48550/arXiv.2302.13939

10. Yao M, Hu J, Hu T, Xu Y, Zhou Z, Tian Y, et al. SPIKE-DRIVEN TRANSFORMER V2: META SPIKING NEURAL NETWORK ARCHITECTURE INSPIRING THE DESIGN OF NEXT-GENERATION NEUROMORPHIC CHIPS. 2024.

11. Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. J Physiol. 1952;117: 500–544. doi:10.1113/jphysiol.1952.sp004764

12. Hodgkin AL, Huxley AF. Currents carried by sodium and potassium ions through the membrane of the giant axon of Loligo. J Physiol. 1952;116: 449–472. doi:10.1113/jphysiol.1952.sp004717

13. Rabuffo G, Bandyopadhyay A, Calabrese C, Gudibanda K, Depannemaecker D, Takarabe LM, et al. Biophysically inspired mean-field model of neuronal populations driven by ion exchange mechanisms. 2025. doi:10.7554/eLife.104249.1

14. Xiao J, Chen Z, Yu B. A Potential Mechanism of Sodium Channel Mediating the General Anesthesia Induced by Propofol. Front Cell Neurosci. 2020;14: 593050. doi:10.3389/fncel.2020.593050

15. Alonso A, Llinás RR. Subthreshold Na+-dependent theta-like rhythmicity in stellate cells of entorhinal cortex layer II. Nature. 1989;342: 175–177. doi:10.1038/342175a0

16. Dura-Bernal S, Herrera B, Lupascu C, Marsh BM, Gandolfi D, Marasco A, et al. Large-Scale Mechanistic Models of Brain Circuits with Biophysically and Morphologically Detailed Neurons. J Neurosci. 2024;44: e1236242024. doi:10.1523/JNEUROSCI.1236-24.2024

17. Lee K, Pennartz CMA, Mejias JF. Cortical networks with multiple interneuron types generate oscillatory patterns during predictive coding. Neuroscience; 2024. doi:10.1101/2024.10.27.620494

18. Tahvili F, Vinck M, Di Volo M. A cortical microcircuit model reveals distinct inhibitory mechanisms of network oscillations and stability. Neuroscience; 2025. doi:10.1101/2025.02.23.639719

19. Sanchez-Todo R, Bastos AM, Lopez-Sola E, Mercadal B, Santarnecchi E, Miller EK, et al. A physical neural mass model framework for the analysis of oscillatory generators from laminar electrophysiological recordings. NeuroImage. 2023;270: 119938.

doi:10.1016/j.neuroimage.2023.119938

20. Rimehaug AE, Stasik AJ, Hagen E, Billeh YN, Siegle JH, Dai K, et al. Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex. eLife. 2023;12: e87169. doi:10.7554/eLife.87169

21. Del Rosario J, Coletta S, Kim SH, Mobille Z, Peelman K, Williams B, et al. Lateral inhibition in V1 controls neural and perceptual contrast sensitivity. Nat Neurosci. 2025;28: 836–847. doi:10.1038/s41593-025-01888-4

22. Soplata AE, McCarthy MM, Sherfey J, Lee S, Purdon PL, Brown EN, et al. Thalamocortical control of propofol phase-amplitude coupling. Graham LJ, editor. PLOS Comput Biol. 2017;13: e1005879. doi:10.1371/journal.pcbi.1005879

23. Soplata AE, Adam E, Brown EN, Purdon PL, McCarthy MM, Kopell N. Rapid thalamocortical network switching mediated by cortical synchronization underlies propofol-induced EEG signatures: a biophysical model. J Neurophysiol. 2023;130: 86–103. doi:10.1152/jn.00068.2022

24. Carannante I, Johansson Y, Silberberg G, Hellgren Kotaleski J. Data-Driven Model of Postsynaptic Currents Mediated by NMDA or AMPA Receptors in Striatal Neurons. Front Comput Neurosci. 2022;16: 806086. doi:10.3389/fncom.2022.806086

25. Moujahid A, d'Anjou A, Torrealdea FJ, Torrealdea F. Energy and information in Hodgkin-Huxley neurons. Phys Rev E. 2011;83: 031912. doi:10.1103/PhysRevE.83.031912

26. Johnson EC, Jones DL, Ratnam R. A minimum-error, energy-constrained neural code is an instantaneous-rate code. J Comput Neurosci. 2016;40: 193–206. doi:10.1007/s10827-016-0592-x

27. Wang Y, Wang R, Xu X. Neural Energy Supply-Consumption Properties Based on Hodgkin-Huxley Model. Neural Plast. 2017;2017: 1–11. doi:10.1155/2017/6207141

28. Theriault JE, Shaffer C, Dienel GA, Sander CY, Hooker JM, Dickerson BC, et al. A functional account of stimulation-based aerobic glycolysis and its role in interpreting BOLD signal intensity increases in neuroimaging experiments. Neurosci Biobehav Rev. 2023;153: 105373. doi:10.1016/j.neubiorev.2023.105373

29. Lampl I, Yarom Y. Subthreshold oscillations of the membrane potential: a functional synchronizing and timing device. J Neurophysiol. 1993;70: 2181–2186. doi:10.1152/jn.1993.70.5.2181

30. Klaus A, Hjorth J, Hellgren-Kotaleski J. The influence of subthreshold membrane potential oscillations and GABAergic input on firing activity in striatal fast-spiking neurons. BMC Neurosci. 2009;10: P244, 1471-2202-10-S1-P244. doi:10.1186/1471-2202-10-S1-P244

31. Börgers C, Kopell NJ. Gamma Oscillations and Stimulus Selection. Neural Comput. 2008;20: 383–414. doi:10.1162/neco.2007.07-06-289

32. Chauhan K, Neiman AB, Tass PA. Synaptic reorganization of synchronized neuronal networks with synaptic weight and structural plasticity. Gutkin BS, editor. PLOS Comput Biol. 2024;20: e1012261. doi:10.1371/journal.pcbi.1012261

33. Bartos M, Vida I, Jonas P. Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. Nat Rev Neurosci. 2007;8: 45–56. doi:10.1038/nrn2044

34. Frémaux N, Gerstner W. Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules. Front Neural Circuits. 2016;9. doi:10.3389/fncir.2015.00085

35. Mandyam Veerambudi Srinivasan , Simon Barry Laughlin and A. Dubs. Predictive coding: a fresh view of inhibition in the retina. 1982.

36. Rao RPN, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci. 1999;2: 79–87. doi:10.1038/4580

37. Friston K. A theory of cortical responses. Philos Trans R Soc B Biol Sci. 2005;360: 815–836. doi:10.1098/rstb.2005.1622

38. Friston K, Kiebel S. Predictive coding under the free-energy principle. Philos Trans R Soc B Biol Sci. 2009;364: 1211–1221. doi:10.1098/rstb.2008.0300

39. Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ. Canonical Microcircuits for Predictive Coding. Neuron. 2012;76: 695–711. doi:10.1016/j.neuron.2012.10.038

40. Körding KP, Wolpert DM. Bayesian integration in sensorimotor learning. Nature. 2004;427: 244–247. doi:10.1038/nature02169

41. Fletcher PC, Frith CD. Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. Nat Rev Neurosci. 2009;10: 48–58. doi:10.1038/nrn2536

42. Lee JH, Whittington MA, Kopell NJ. Top-Down Beta Rhythms Support Selective Attention via Interlaminar Interaction: A Model. Coombes S, editor. PLoS Comput Biol. 2013;9: e1003164. doi:10.1371/journal.pcbi.1003164

43. Bastos G, Holmes JT, Ross JM, Rader AM, Gallimore CG, Wargo JA, et al. Top-down input modulates visual context processing through an interneuron-specific circuit. Cell Rep. 2023;42: 113133. doi:10.1016/j.celrep.2023.113133

44.   Markov NT, Vezoli J, Chameau P, Falchier A, Quilodran R, Huissoud C, et al. Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. J Comp Neurol. 2014;522: 225–259. doi:10.1002/cne.23458

45.   Richter D, Ekman M, De Lange FP. Suppressed Sensory Response to Predictable Object Stimuli throughout the Ventral Visual Stream. J Neurosci. 2018;38: 7452–7461. doi:10.1523/JNEUROSCI.3421-17.2018

46.   Richter CG, Thompson WH, Bosman CA, Fries P. Top-Down Beta Enhances Bottom-Up Gamma. J Neurosci. 2017;37: 6698–6711. doi:10.1523/JNEUROSCI.3771-16.2017

47.   Bastos AM, Lundqvist M, Waite AS, Kopell N, Miller EK. Layer and rhythm specificity for predictive routing. Proc Natl Acad Sci. 2020;117: 31459–31469. doi:10.1073/pnas.2014868117

48.   Gabhart KM, Xiong Y (Sophy), Bastos AM. Predictive coding: a more cognitive process than we thought? Trends Cogn Sci. 2025; S1364661325000300. doi:10.1016/j.tics.2025.01.012

49.   Miller EK, Brincat SL, Roy JE. Cognition is an emergent property. Curr Opin Behav Sci. 2024;57: 101388. doi:10.1016/j.cobeha.2024.101388

50.   Bastos AM, Vezoli J, Bosman CA, Schoffelen J-M, Oostenveld R, Dowdall JR, et al. Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. Neuron. 2015;85: 390–401. doi:10.1016/j.neuron.2014.12.018

51.   Marr D. Vision. W. H. Freeman; 1982.

52.   Marr D. Vision: a computational investigation into the human representation and processing of visual information. Cambridge, Mass: MIT Press; 2010.

53.   Kopell N, Ermentrout GB, Whittington MA, Traub RD. Gamma rhythms and beta rhythms have different synchronization properties. Proc Natl Acad Sci. 2000;97: 1867–1872. doi:10.1073/pnas.97.4.1867

54.   Sussillo D, Abbott LF. Generating Coherent Patterns of Activity from Chaotic Neural Networks. Neuron. 2009;63: 544–557. doi:10.1016/j.neuron.2009.07.018

55.   Börgers C, Epstein S, Kopell NJ. Gamma oscillations mediate stimulus competition and attentional selection in a cortical network model. Proc Natl Acad Sci. 2008;105: 18023–18028. doi:10.1073/pnas.0809511105

56.   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

57.   Yang H, Han F, Lu W, Wang Q. Computational modeling for gratings stimulated gamma

oscillations in a large-scale cortical neuronal network. Neurocomputing. 2025; 131401. doi:10.1016/j.neucom.2025.131401

58. Lee S, Jones SR. Distinguishing mechanisms of gamma frequency oscillations in human current source signals using a computational model of a laminar neocortical network. Front Hum Neurosci. 2013;7. doi:10.3389/fnhum.2013.00869

59. Sherfey JS, Ardid S, Hass J, Hasselmo ME, Kopell NJ. Flexible resonance in prefrontal networks with strong feedback inhibition. Doiron B, editor. PLOS Comput Biol. 2018;14: e1006357. doi:10.1371/journal.pcbi.1006357

60. Wallace E, Benayoun M, Van Drongelen W, Cowan JD. Emergent Oscillations in Networks of Stochastic Spiking Neurons. Gutkin BS, editor. PLoS ONE. 2011;6: e14804. doi:10.1371/journal.pone.0014804

61. Ardid S, Wang X-J, Gomez-Cabrero D, Compte A. Reconciling Coherent Oscillation with Modulationof Irregular Spiking Activity in Selective Attention:Gamma-Range Synchronization between Sensoryand Executive Cortical Areas. J Neurosci. 2010;30: 2856–2870. doi:10.1523/JNEUROSCI.4222-09.2010

62. Nicola W, Clopath C. Supervised learning in spiking neural networks with FORCE training. Nat Commun. 2017;8: 2208. doi:10.1038/s41467-017-01827-3

63. Druckmann S. A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data. Front Neurosci. 2007;1: 7–18. doi:10.3389/neuro.01.1.1.001.2007

64. Izhikevich EM, Edelman GM. Large-scale model of mammalian thalamocortical systems. Proc Natl Acad Sci. 2008;105: 3593–3598. doi:10.1073/pnas.0712231105

65. Yuille A, Kersten D. Vision as Bayesian inference: analysis by synthesis? Trends Cogn Sci. 2006;10: 301–308. doi:10.1016/j.tics.2006.05.002

66. Aitchison L, Lengyel M. With or without you: predictive coding and Bayesian inference in the brain. Curr Opin Neurobiol. 2017;46: 219–227. doi:10.1016/j.conb.2017.08.010

67. Azabou M, Azar MG, Liu R, Lin C-H, Johnson EC, Bhaskaran-Nair K, et al. Mine Your Own vieW: Self-Supervised Learning Through Across-Sample Prediction. arXiv; 2021. doi:10.48550/ARXIV.2102.10106

68. Nejad KK, Anastasiades P, Hertäg L, Costa RP. Self-supervised predictive learning accounts for cortical layer-specificity. Nat Commun. 2025;16: 6178. doi:10.1038/s41467-025-61399-5

69. Vikhar PA. Evolutionary algorithms: A critical review and its future prospects. 2016 International Conference on Global Trends in Signal Processing, Information

Computing and Communication (ICGTSPICC). Jalgaon, India: IEEE; 2016. pp. 261–265. doi:10.1109/ICGTSPICC.2016.7955308

70. Evolutionary Computation in Bioinformatics. Elsevier; 2003. doi:10.1016/B978-1-55860-797-2.X5000-8

71. Farinati D, Vanneschi L. A survey on dynamic populations in bio-inspired algorithms. Genet Program Evolvable Mach. 2024;25: 19. doi:10.1007/s10710-024-09492-4

72. Hazy TE, Frank MJ, O'Reilly RC. Neural mechanisms of acquired phasic dopamine responses in learning. Neurosci Biobehav Rev. 2010;34: 701–720. doi:10.1016/j.neubiorev.2009.11.019

73. Rescorla, R.A, Wagner, A.R. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. Class Cond II. 1972; 64–99.

74. Frazier-Logue N, Hanson SJ. The Stochastic Delta Rule: Faster and More Accurate Deep Learning Through Adaptive Weight Noise. Neural Comput. 2020;32: 1018–1032. doi:10.1162/neco_a_01276

75. Miller RR, Barnet RC, Grahame NJ. Assessment of the Rescorla-Wagner model. Psychol Bull. 1995;117: 363–386. doi:10.1037/0033-2909.117.3.363

76. Jin Y. Evolutionary Algorithms. Advanced Fuzzy Systems Design and Applications. Heidelberg: Physica-Verlag HD; 2003. pp. 49–71. doi:10.1007/978-3-7908-1771-3_2

77. Hoffmeister F, Bäck T. Genetic Algorithms and evolution strategies: Similarities and differences. In: Schwefel H-P, Männer R, editors. Parallel Problem Solving from Nature. Berlin/Heidelberg: Springer-Verlag; 1991. pp. 455–469. doi:10.1007/BFb0029787

78. Mühlenbein H, Schomisch M, Born J. The parallel genetic algorithm as function optimizer. Parallel Comput. 1991;17: 619–632. doi:10.1016/S0167-8191(05)80052-3

79. Whitley D. An overview of evolutionary algorithms: practical issues and common pitfalls. Inf Softw Technol. 2001;43: 817–831. doi:10.1016/S0950-5849(01)00188-4

80. Slowik A, Kwasnicka H. Evolutionary algorithms and their applications to engineering problems. Neural Comput Appl. 2020;32: 12363–12379. doi:10.1007/s00521-020-04832-8

81. Bastos AM, Loonis R, Kornblith S, Lundqvist M, Miller EK. Laminar recordings in frontal cortex suggest distinct layers for maintenance and control of working memory. Proc Natl Acad Sci. 2018;115: 1117–1122. doi:10.1073/pnas.1710323115

82. Miller EK, Lundqvist M, Bastos AM. Working Memory 2.0. Neuron. 2018;100:

463–475. doi:10.1016/j.neuron.2018.09.023

83. Lundqvist M, Bastos AM, Miller EK. Preservation and Changes in Oscillatory Dynamics across the Cortical Hierarchy. J Cogn Neurosci. 2020;32: 2024–2035. doi:10.1162/jocn_a_01600

84. Jamali S, Bagur S, Bremont E, Van Kerkoerle T, Dehaene S, Bathellier B. Parallel mechanisms signal a hierarchy of sequence structure violations in the auditory cortex. eLife. 2024;13: RP102702. doi:10.7554/eLife.102702

85. Giraud A-L, Arnal LH. Hierarchical Predictive Information Is Channeled by Asymmetric Oscillatory Activity. Neuron. 2018;100: 1022–1024. doi:10.1016/j.neuron.2018.11.020

86. Adesnik H, Bruns W, Taniguchi H, Huang ZJ, Scanziani M. A neural circuit for spatial summation in visual cortex. Nature. 2012;490: 226–231. doi:10.1038/nature11526

87. Golowasch J, Casey M, Abbott LF, Marder E. Network Stability from Activity-Dependent Regulation of Neuronal Conductances. Neural Comput. 1999;11: 1079–1096. doi:10.1162/089976699300016359

88. Prinz AA, Bucher D, Marder E. Similar network activity from disparate circuit parameters. Nat Neurosci. 2004;7: 1345–1352. doi:10.1038/nn1352

89. Uhlenbeck GE, Ornstein LS. On the Theory of the Brownian Motion. Phys Rev. 1930;36: 823–841. doi:10.1103/PhysRev.36.823

90. Wacongne C, Changeux J-P, Dehaene S. A Neuronal Model of Predictive Coding Accounting for the Mismatch Negativity. J Neurosci. 2012;32: 3665–3678. doi:10.1523/JNEUROSCI.5003-11.2012

91. Chien VSC, Maess B, Knösche TR. A generic deviance detection principle for cortical On/Off responses, omission response, and mismatch negativity. Biol Cybern. 2019;113: 475–494. doi:10.1007/s00422-019-00804-x

92. Sherfey JS, Soplata AE, Ardid S, Roberts EA, Stanley DA, Pittman-Polletta BR, et al. DynaSim: A MATLAB Toolbox for Neural Modeling and Simulation. Front Neuroinformatics. 2018;12: 10. doi:10.3389/fninf.2018.00010

93. Billeh YN, Cai B, Gratiy SL, Dai K, Iyer R, Gouwens NW, et al. Systematic Integration of Structural and Functional Data into Multi-scale Models of Mouse Primary Visual Cortex. Neuron. 2020;106: 388-403.e18. doi:10.1016/j.neuron.2020.01.040

94. Roopun AK. concatenation underlies interactions between gamma and beta rhythms in neocortex. Front Cell Neurosci. 2008;2. doi:10.3389/neuro.03.001.2008

95. Sherfey J, Ardid S, Miller EK, Hasselmo ME, Kopell NJ. Prefrontal oscillations modulate

the propagation of neuronal activity required for working memory. Neurobiol Learn Mem. 2020;173: 107228. doi:10.1016/j.nlm.2020.107228

96. Ardid S, Sherfey JS, McCarthy MM, Hass J, Pittman-Polletta BR, Kopell N. Biased competition in the absence of input bias revealed through corticostriatal computation. Proc Natl Acad Sci. 2019;116: 8564–8569. doi:10.1073/pnas.1812535116

97. Xiong Y (Sophy), Donoghue JA, Lundqvist M, Mahnke M, Major AJ, Brown EN, et al. Propofol-mediated loss of consciousness disrupts predictive routing and local field phase modulation of neural activity. Proc Natl Acad Sci. 2024;121: e2315160121. doi:10.1073/pnas.2315160121

98. Friston KJ, Harrison L, Penny W. Dynamic causal modelling. NeuroImage. 2003;19: 1273–1302. doi:10.1016/S1053-8119(03)00202-7

99. David O, Kiebel SJ, Harrison LM, Mattout J, Kilner JM, Friston KJ. Dynamic causal modeling of evoked responses in EEG and MEG. NeuroImage. 2006;30: 1255–1272. doi:10.1016/j.neuroimage.2005.10.045

100. Bastos AM, Litvak V, Moran R, Bosman CA, Fries P, Friston KJ. A DCM study of spectral asymmetries in feedforward and feedback connections between visual areas V1 and V4 in the monkey. NeuroImage. 2015;108: 460–475. doi:10.1016/j.neuroimage.2014.12.081

101. Garrido MI, Friston KJ, Kiebel SJ, Stephan KE, Baldeweg T, Kilner JM. The functional anatomy of the MMN: A DCM study of the roving paradigm. NeuroImage. 2008;42: 936–944. doi:10.1016/j.neuroimage.2008.05.018

102. Katsanevaki C, Bastos AM, Cagnan H, Bosman CA, Friston KJ, Fries P. Attentional effects on local V1 microcircuits explain selective V1-V4 communication. NeuroImage. 2023;281: 120375. doi:10.1016/j.neuroimage.2023.120375

103. Choromanska A, Henaff M, Mathieu M, Arous GB, LeCun Y. The Loss Surfaces of Multilayer Networks.

104. Bastos AM, Donoghue JA, Brincat SL, Mahnke M, Yanar J, Correa J, et al. Neural effects of propofol-induced unconsciousness and its reversal using thalamic stimulation. eLife. 2021;10: e60824. doi:10.7554/eLife.60824

105. Van Derveer AB, Bastos G, Ferrell AD, Gallimore CG, Greene ML, Holmes JT, et al. A Role for Somatostatin-Positive Interneurons in Neuro-Oscillatory and Information Processing Deficits in Schizophrenia. Schizophr Bull. 2021;47: 1385–1398. doi:10.1093/schbul/sbaa184

106. Arnal LH, Giraud A-L. Cortical oscillations and sensory predictions. Trends Cogn Sci. 2012;16: 390–398. doi:10.1016/j.tics.2012.05.003

107. Stoiljkovic M, Kelley C, Nagy D, Hajós M. Modulation of hippocampal neuronal network oscillations by α7 nACh receptors. Biochem Pharmacol. 2015;97: 445–453. doi:10.1016/j.bcp.2015.06.031

108. Mikulasch FA, Rudelt L, Wibral M, Priesemann V. Where is the error? Hierarchical predictive coding through dendritic error computation. Trends Neurosci. 2023;46: 45–59. doi:10.1016/j.tins.2022.09.007

109. O'Toole SM, Oyibo HK, Keller GB. Molecularly targetable cell types in mouse visual cortex have distinguishable prediction error responses. Neuron. 2023;111: 2918-2928.e8. doi:10.1016/j.neuron.2023.08.015

110. Pérez-González D, Lao-Rodríguez AB, Aedo-Sánchez C, Malmierca MS. Acetylcholine modulates the precision of prediction error in the auditory cortex. eLife. 2024;12: RP91475. doi:10.7554/eLife.91475.3

111. Spratling MW. Predictive coding as a model of biased competition in visual attention. Vision Res. 2008;48: 1391–1408. doi:10.1016/j.visres.2008.03.009

112. Buffalo EA, Fries P, Landman R, Buschman TJ, Desimone R. Laminar differences in gamma and alpha coherence in the ventral stream. Proc Natl Acad Sci. 2011;108: 11262–11267. doi:10.1073/pnas.1011284108

113. Adams RA, Pinotsis D, Tsirlis K, Unruh L, Mahajan A, Horas AM, 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. 2022;91: 202–215. doi:10.1016/j.biopsych.2021.07.024

114. Han C, Wang T, Yang Y, Wu Y, Li Y, Dai W, et al. Multiple gamma rhythms carry distinct spatial frequency information in primary visual cortex. Pack C, editor. PLOS Biol. 2021;19: e3001466. doi:10.1371/journal.pbio.3001466

115. Han C, Wang T, Wu Y, Li Y, Yang Y, Li L, et al. The Generation and Modulation of Distinct Gamma Oscillations with Local, Horizontal, and Feedback Connections in the Primary Visual Cortex: A Model Study on Large-Scale Networks. Wang R, editor. Neural Plast. 2021;2021: 1–17. doi:10.1155/2021/8874516

116. Murty DVPS, Shirhatti V, Ravishankar P, Ray S. Large Visual Stimuli Induce Two Distinct Gamma Oscillations in Primate Visual Cortex. J Neurosci. 2018;38: 2730–2744. doi:10.1523/JNEUROSCI.2270-17.2017

117. Samaha J, Bauer P, Cimaroli S, Postle BR. Top-down control of the phase of alpha-band oscillations as a mechanism for temporal prediction. Proc Natl Acad Sci. 2015;112: 8439–8444. doi:10.1073/pnas.1503686112

118. Van Kerkoerle T, Self MW, Dagnino B, Gariel-Mathis M-A, Poort J, Van Der Togt C, et al.

Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. Proc Natl Acad Sci. 2014;111: 14332–14341. doi:10.1073/pnas.1402773111

119. Wang C, Rajagovindan R, Han S-M, Ding M. Top-Down Control of Visual Alpha Oscillations: Sources of Control Signals and Their Mechanisms of Action. Front Hum Neurosci. 2016;10. doi:10.3389/fnhum.2016.00015

120. Schmidt M, Bakker R, Shen K, Bezgin G, Diesmann M, Van Albada SJ. A multi-scale layer-resolved spiking network model of resting-state dynamics in macaque visual cortical areas. Graham LJ, editor. PLOS Comput Biol. 2018;14: e1006359. doi:10.1371/journal.pcbi.1006359

121. Westerberg JA, Xiong YS, Nejat H, Sennesh E, Durand S, Hardcastle B, et al. Adaptation, not prediction, drives neuronal spiking responses in mammalian sensory cortex. Neuroscience; 2024. doi:10.1101/2024.10.02.616378

122. Kiebel SJ, Daunizeau J, Friston KJ. A Hierarchy of Time-Scales and the Brain. Sporns O, editor. PLoS Comput Biol. 2008;4: e1000209. doi:10.1371/journal.pcbi.1000209

123. Wacongne C, Labyt E, Van Wassenhove V, Bekinschtein T, Naccache L, Dehaene S. Evidence for a hierarchy of predictions and prediction errors in human cortex. Proc Natl Acad Sci. 2011;108: 20754–20759. doi:10.1073/pnas.1117807108

124. Chao ZC, Takaura K, Wang L, Fujii N, Dehaene S. Large-Scale Cortical Networks for Hierarchical Prediction and Prediction Error in the Primate Brain. Neuron. 2018;100: 1252-1266.e3. doi:10.1016/j.neuron.2018.10.004

125. Keller GB, Mrsic-Flogel TD. Predictive Processing: A Canonical Cortical Computation. Neuron. 2018;100: 424–435. doi:10.1016/j.neuron.2018.10.003

126. Jiang LP, Rao RPN. Dynamic predictive coding: A model of hierarchical sequence learning and prediction in the neocortex. Rubin J, editor. PLOS Comput Biol. 2024;20: e1011801. doi:10.1371/journal.pcbi.1011801

127. Deistler M, Kadhim KL, Pals M, Beck J, Huang Z, Gloeckler M, et al. Jaxley: differentiable simulation enables large-scale training of detailed biophysical models of neural dynamics. Nat Methods. 2025;22: 2649–2657. doi:10.1038/s41592-025-02895-w

128. Robbins H, Monro S. A Stochastic Approximation Method. Ann Math Stat. 1951;22: 400–407. doi:10.1214/aoms/1177729586

129. Widrow B, Hoff ME. (1960) Bernard Widrow and Marcian E. Hoff, "Adaptive switching circuits," 1960 IRE WESCON Convention Record, New York: IRE, pp. 96-104. In: Anderson JA, Rosenfeld E, editors. Neurocomputing, Volume 1. The MIT Press; 1988. pp. 126–134. doi:10.7551/mitpress/4943.003.0012

130. Bradbury J, Frostig R, Hawkins P, Johnson MJ, Leary C, Maclaurin D, et al. JAX.

131. Markram H, Gerstner W, Sjöström PJ. Spike-Timing-Dependent Plasticity: A Comprehensive Overview. Front Synaptic Neurosci. 2012;4. doi:10.3389/fnsyn.2012.00002

132. Mendoza-Halliday D, Major AJ, Lee N, Lichtenfeld MJ, Carlson B, Mitchell B, et al. A ubiquitous spectrolaminar motif of local field potential power across the primate cortex. Nat Neurosci. 2024;27: 547–560. doi:10.1038/s41593-023-01554-7

133. Pachitariu M, Sridhar S, Pennington J, Stringer C. Spike sorting with Kilosort4. Nat Methods. 2024;21: 914–921. doi:10.1038/s41592-024-02232-7

134. Kass RE, Ventura V, Brown EN. Statistical Issues in the Analysis of Neuronal Data. J Neurophysiol. 2005;94: 8–25. doi:10.1152/jn.00648.2004