New learning principles emerge from biomimetic computational primitives

Anand Pathak[a], Scott L. Brincat[b], Haris Organtzidis[c], Helmut H. Strey[c], Sageanne Senneff[c], Evan G. Antzoulatos[b], Lilianne R. Mujica-Parodi[b,c,d,e], Earl K. Miller[b], and Richard Granger[a1]

[a]Dept. of Psychological and Brain Sciences, Dartmouth College, Hanover, NH; [b]The Picower Institute for Learning and Memory and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA; [c]Dept. of Biomedical Engineering, Laufer Center for Physical and Quantitative Biology, State University of New York at Stony Brook, Stony Brook, NY; [d]Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, MA; [e]Baszucki Endowed Chair for Metabolic Neuroscience, Stony Brook University, NY

This manuscript was compiled on November 7, 2023

To examine the biological building blocks of thought and action, we created biologically realistic local circuits based on detailed well-cited physiological and anatomical characteristics. These biomimetic circuits were then integrated into a large-scale model of cortical-striatal interactions for category learning. The model was not trained on, but nonetheless displayed properties similar to, neurophysiological recordings from non-human primates (NHPs) performing the same task. The model had learning curves similar to the NHPs. It showed how synaptic modifications could induce changes in spiking and synchrony in the brain. The model even made novel predictions that were subsequently found in the brain, including "bad-idea neurons" that signalled impending incorrect choices after learning. This demonstrates how key computational principles can be discovered by modeling local circuitry that mimics the brain.

biological computation | brain rhythms | brain simulation

1. Introduction

What are the elemental components that underlie the brain's computational operations? To begin to answer that, we constructed mechanisms that were greatly simplified from biology, but nonetheless substantially adhered to features of circuit layout and operation directly traceable to the large background literature on these biological circuits (Table 1). The resulting simulation both mimicked complex physiological findings from empirical data, and also uncovered previously unexpected behaviorally-relevant physiological spiking phenomena that were subsequently verified in physiological recordings of NHP subjects.

Many current computational neuroscience efforts aim at reproducing physiological or behavioral experimental results, using various forms of artificial neural networks (ANNs) such as recurrent and convolutional networks (RNNs, CNNs, etc.) (1). But for any given experimental corpus, results can be fitted by numerous methods; thus any particular body of experimental data vastly underspecifies the set of mechanisms that may have generated it. Often, artificial ANNs are proposed as elements underlying brain activity, even though ANNs are constructed from very small sets of simple formal rules, excluding differences between, for instance, excitatory vs. inhibitory cells, different physiological time courses for such cells, different wiring patterns in different brain structures, and much more.

By contrast, there is an enormous corpus of experimental findings on the structural organization of various distinct brain circuits, and on the activity of these circuits over time. Incorporation of these profuse additional biological features can provide many further degrees of constraint on the set of models that suffice as candidates for addressing complex brain data. The vast preponderance of this deep background of biological constraints is typically orthogonal to, or even entirely absent from, network models that are composed of small numbers of straightforward arithmetical rules. Thus, whereas the aims of typical ANN systems (2) align well with the aims of studies we present here, the way these aims are accomplished differs greatly between the two approaches.

In our simulation, each simplified local assembly incorporates well-studied physiological properties of neurons as embedded in anatomically framed local circuits. These "biomimetic primitives" were used to build a large-scale model with the local components connected in accordance with known systems-level wiring. A primary guide was the extensive cortical-striatal loops and associated structures known to play key roles in learning. We included excitatory and inhibitory cells, pyramidal cortical neurons, medium spiny neurons, dopaminergic neurons, cholinergic tonically active neurons. The result was a multi-scale model that incorporated low-level single-neuron details as well as local and large-scale connectivity of real brain circuits (Figure 1). The resulting multi-scale model thus engaged many diverse cortical and subcortical components during behavior.

Multi-scale modeling is increasingly important for understanding the brain. Brain activity occurs on multiple spatial scales from spiking of neurons to local field potentials to coherent electrical fluctuations across large swaths of brain and on multiple timescales from milliseconds to seconds. Questions of causality abound: spikes give rise to waves, yet waves in turn strongly modulate spiking; and both appear affected by disorders and by drugs (3–6). Multi-scale modeling helps us understand the real-time interactions, causal relationships, and emergent properties that cannot be derived from first principles alone or by studying the brain piece by piece.

[1]To whom correspondence should be addressed. E-mail: richard.granger@gmail.com

Our initial test bed was a category learning task. By trial and error, the non-human primates (NHPs) learned to divide a set of dot patterns into two groups based on their configurations (7). Importantly, the model was not trained on any physiological or behavioral data from the NHPs. Rather we constructed the model based on extensively documented biological realism only. The model then was presented with the category learning task, and its activity was evaluated against neurophysiological recordings from non-human primates (NHPs) performing the same task as the model. We examined how the model learned the task, to determine whether it behaved like the actual brains of NHPs.

Again, it is emphasized that this approach is distinct from many standard computational studies in which an artificial neural network is "trained" on some data, and then tested for its ability to extend to new data. Instead, our goal was to determine if the elemental circuits derived here, simultaneously biomimetic and computational, could, on their own, without influence from sample data, act like the actual brain. We found that they did. They even predicted new, previously unknown properties that were then confirmed in the brain data.

2. Results

A. Specific biologically derived simulations embody complex computational learning and performance rules.

A.1. A corticostriatal circuit composed of computational primitives.

The simulation presented here incorporates multiple subcircuits that interact to carry out distinct cortical and striatal learning rules (Figure 1).

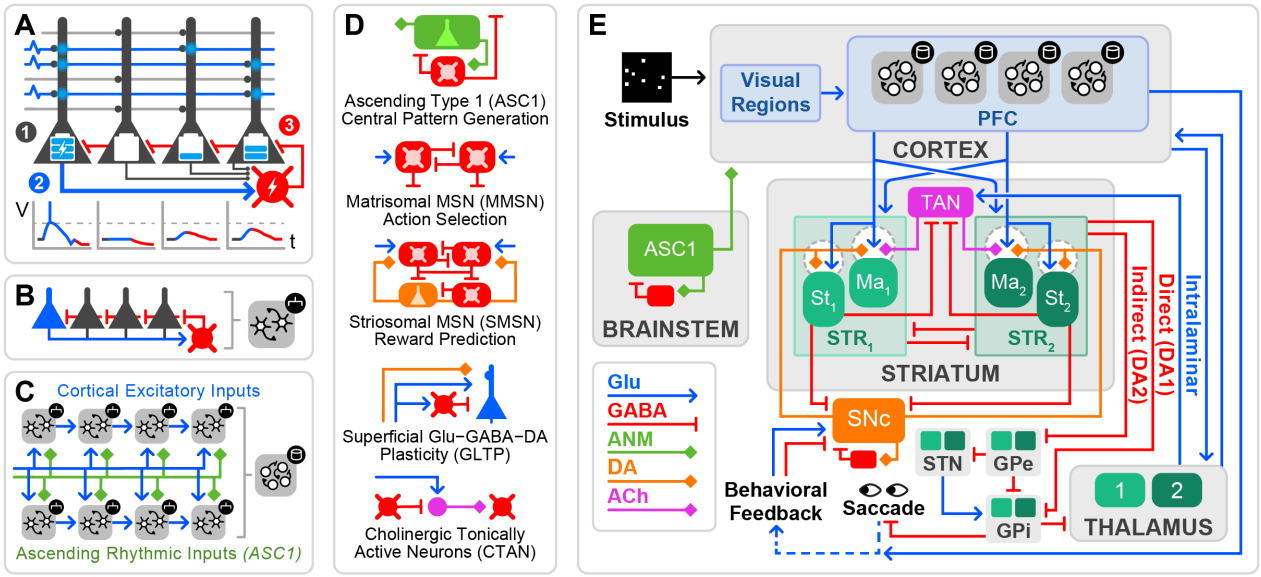

Fig. 1. Multi-scale components are hierarchically incorporated into a large-scale cortical-subcortical model. (A) Lateral inhibition cortical assembly (see text), each composed of several excitatory cells with an individual inhibitory cell. Time-dependent voltage transients within the cells build up over time as a function of the number of active input axons, the number of connected target dendritic synapses, and the weights of those synapses. The overall effect as described in the text, is that the most-activated excitatory cell reaches action potential threshold more rapidly than other cells; once it spikes, it activates its local inhibitory cell which in turn inhibits all cells. Thus only the most-activated excitatory cell can respond, i.e., the 'winner' takes all. (B) A simplified icon corresponding to the lateral-inhibition circuit; these constitute the superficial layer organization of simplified cortical 'columns'. (C) Each such cortical assembly is then sparsely connected with others as described in Methods, and the cortical model receives a diffuse rhythmic input from a simulated ascending system (modeled as a next-gen neural mass model; see Methods). (D) Additional neural assemblies (biomimetic computational primitives) that were derived via similar means and are incorporated into the simulated corticostriatal model; see Table 1. (E) Overall simulation organization. Combined cortical and subcortical structures, interconnected anatomically as described in the literature, form an integrated simulation with extensive feedforward and feedback elements. The model posterior 'visual' regions are presented with visual stimuli ('stimulus'), and the system is presented with simple simulated 'reward' signals ('behavioral feedback') based on the system's output (saccade) behavior (see Results and Methods). Cortical structures (visual posterior; prefrontal anterior) are modeled solely as superficial-layer excitatory-inhibitory assemblies; striatal matrisomes and striosomes are modeled as medium spiny neuron (MSN) assemblies; ascending systems ('BRAINSTEM') consist of excitatory-inhibitory oscillatory interactions modeled as next-gen neural mass models (NG-NMM). (All described in Methods). Whereas cortical elements are modeled as described as spiking individual neurons within connected anatomical structures, several subcortical elements (striatal MSN matrisomes and striosomes, TANs, and SNc) are modeled not as spiking neural elements but as simplified physiologically-derived "rules" whose equations define their input/output behavior (see Methods). Acetylcholine and dopamine units are modeled from the literatures on TANs and SNc respectively. Plasticity both in cortex and in striatum are modeled separately from their respective literature (see Methods). PFC : *Prefrontal Cortex*, ASC1 : *Ascending system; see Methods*, SNc : *Substantia Nigra pars compacta*, STR : *Striatum*, TAN : *Tonically active neurons*, St : *Striosome*, Ma : *Matrisome*, GPe : *Globus Pallidus pars externa*, GPi : *Globus Pallidus pars interna*, STN : *Subthalamic Nucleus*, Glu : *Glutamate*, GABA : *Gamma aminobutyric acid*, ANM : *ascending neuromodulator*, DA : *Dopamine*, ACh : *Acetylcholine*.

Cortex subcircuit. The lateral inhibition assembly shown in Figure 1A-B, is incorporated into an interconnected system by being connected sparsely and randomly with each other (see Methods for exact connection specification). The resulting circuit receives input from a simulated external visual grid input ("Stimulus" in Figure 1E), containing experimental data, as well as from ascending rhythmic inputs ("ASC1" in Figure 1D-E) modeled as next-generation neural mass models (8, 9) as shown in Figure 1C-E. This lateral inhibition circuit is described in further detail in a subsequent section.

Striatal subcircuit. The simulated striatal complex (Figure 1E) consists of matrix (matrisome) assemblies, patch (striosome) assemblies, a substantia nigra pars compacta (SNc), and a cholinergic tonically active neuron (TAN) structure (10, 11). Matrisome and striosome blocks are composed of medium spiny neurons (MSNs) receiving convergent inputs from cortex (glutamate), from TANs (ACh), and from SNc (dopamine) (Figure 1D, BCP: Striosomal MSN Reward Prediction). MSN blocks laterally inhibit each other, such that the one with highest input will respond and others will then be weakened by lateral inhibition (Figure 1D, BCP: Matrisomal MSN Action Selection). MSNs output to pallidum and to thalamus. Components of striatal subcircuit, viz. matrisome, striosome, TAN and SNc are not implemented as spiking neurons in the current model, but rather are equations designed to describe mass input-output transforms (See Methods: Model Description).

TAN subcircuit. Cholinergic tonically active neurons (TANs) in the Figure 1E denote a structure that receives inputs from i) intralaminar thalamic nuclei (described below) and ii) from GABAergic striosomal MSNs (10–12). TAN projection synapses co-localize with corticostriatal synapses, i.e., glutamatergic cortical synapses onto matrix and patch MSNs (see Fig 1E), as well as with dopaminergic SNc synapses. The TAN structure in the simulation is rhythmically active, due to thalamic intralaminar driving input (12). TANs are inhibited by GABAergic input from striosomes, and thus as cortical inputs to striosomal MSNs are strengthened, those striosomes will inhibit TANs more strongly (Figure 1D, BCP: Cholinergic Tonically Active Neurons). In the simulation, this has an effect of initially adding cholinergic "noise" that increases variance in target matrix MSN responses, which we hypothesize adds "exploratory" variability to striatal action selection (see (11)). As learning progresses, corticostriatal synapses to striosomes are strengthened, which increases their ability to inhibit TANs, which lessens this added MSN variance, thus decreasing variability as associations are effectively learned in the corticostriatal system.

Brainstem subcircuit. As described in Methods: Model Description, the current implementation of ascending systems is in the form of next-generation neural mass models (8, 9) with parameter settings which produce rhythmic patterned responses. Proposed work entails elaborating on these, especially to model effects of prominent descending signals (e.g., from cortex to the brainstem system).

A.2. Lateral inhibition in local cortical assemblies produces rich computational operations. Each of the brain circuit assemblies listed above is (i) composed of biological components and (ii) carries out a computational primitive operation. In addition to the brief descriptions above, and the very detailed descriptions of each of these components in the Methods section, we briefly describe here one of those constructs: the lateral inhibition mechanism that is illustrated in simplified form in Figure 1 A and B.

These local assemblies of excitatory and inhibitory neurons in simulated superficial cortical layers are connected such that each input axon makes contact with a fixed number of excitatory target dendrites, and each dendrite receives contacts from a separately fixed number of incoming axons. The resulting hypergeometric connectome (13) ensures the absence of "hot" or "cold" neurons, i.e., cells differently capable of activity due solely to their afferent wiring. All excitatory-excitatory connections are relatively sparse: the probability of any input neuron connecting with any target is 8% in the simulation, contrasting drastically with dense artificial neural network connectivity (of roughly 100%), although still far denser than actual cortical connectivities of roughly 1% (14).

The anatomical organization of lateral inhibition cortical assemblies is shown in Figure 1A. Glutamatergic input axons contact superficial layer (II-III) excitatory cortical target cells sparsely; those cells in turn densely contact local inhibitory cells which are arranged in roughly a 4:1 ratio, and the inhibitory cell reciprocally makes dense contacts with local excitatory cells (15–17). Together, these are superficial layer components within a cortical microcolumn. In Figure 1A item (1), an incoming spatial pattern of spikes arrives, inducing depolarization via those locations at which active input axons contact dendritic synapses. Total voltage over time (voltage-time graph) is summed at a rate that is dependent on the quantity of input activation at a given target cells and the weights of the receiving synapses; the slope of the voltage value over time reflects these summations (bottom voltage-time graphs in Figure 1A). Any cell attaining the spiking threshold emits its own action potentials (item (2) in Figure 1A). These in turn activate local inhibitory cells which then inhibit all local excitatory cells, (including the one that activated the inhibition), driving voltages back down (item (3) in Figure 1A, and bottom voltage-time graphs in Figure 1A).

Thus, a very limited subset of target excitatory cells respond to an input; this effect corresponds to the literature of lateral inhibition, competitive circuitry, and 'winners take all' (18–25). The active synapses of the resulting responding cell will then be increased slightly in weight (see Methods). This makes it even more likely that slight variants of an input will all tend to activate the same target excitatory cell. As a result, similar inputs will increasingly come to elicit identical, not just similar, cortical responses. That response then can be considered a "categorical" response for that set of inputs. This effect is illustrated in Figure 2 C-D: dendritic vectors (vectors containing synaptic weights, with dimensions same as that of the input space) initially pointing in roughly the right direction in the input space will come to point more and more closely to the mean of those inputs that the target cell responds to. Before learning, the synaptic vectors that respond to a given input are those that have the smallest (closest) angle to that input ($\delta 1$, $\delta 2$ in Figure 2 C); the responders strengthen the activated synapses, becoming ever closer to the set of inputs that they successfully respond to ($\delta 3$, $\delta 4$ in Figure 2 D). As the synapses change, those target cells that successfully respond to any member of a closely related set of inputs will come to respond to all of that group of inputs. Any of those similar inputs will elicit the same cortical response; that response corresponds to a 'category name' for that set of inputs. Supplementary Figure S3 illustrates quantified instances of this phenomena observed in the modelled anterior cortical neurons.

A set of information-processing operations is thus carried out by a simple, but biologically derived, small assembly of excitatory and inhibitory cells. Given the few elements composing the circuit presented here, changing nearly any component of the circuit can qualitatively change its information-processing functions; in this sense, the circuit is irre-

ducible for producing these computations; it can be considered to be one functionally 'primitive' element of brain circuitry. It is notable that this arrangement is extraordinarily distinct from artificial neural networks (ANNs) in which (i) cells are not distinguished as excitatory or inhibitory; (ii) there are no physiological time courses, let alone differential durations for excitation (fast) vs. inhibition (slow); and (iii) there are no anatomical connectivity distinctions between, e.g., sparse excitatory-excitatory contacts vs. dense excitatory-inhibitory and inhibitory-excitatory contacts. The computations of the lateral-inhibition circuit described herein, are distinct from those in other (ANN) approaches to lateral inhibition or winner take all (19–25). It is also hypothesized that the circuit introduced here will more closely correspond to physiological measurements, and may differentially respond to potential treatments that, for instance, affect different biochemical components (such as glutamatergic or GABAergic receptors).

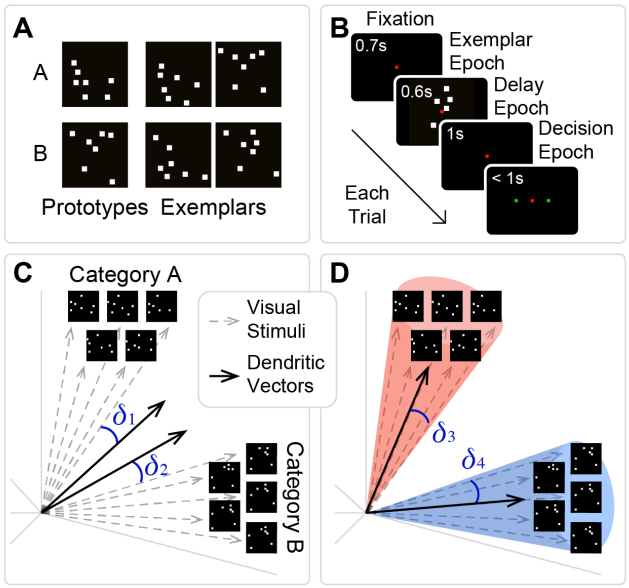

Fig. 2. Synaptic increments achieve natural internal clustering of learned image patterns. (A) Two example categories of visual stimuli. Stimuli are the unique exemplars that were created by distortion of the two prototypes. (B) Schematic illustrating the time course of a single learning trial. See references (7, 26) for details of the experimental set-up. (C) Data are random dot patterns that fall naturally into two distinct groups based on their within-group and between-group similarities. They are shown in an input space defined by the abstract coördinates of the possible positionings of dots within each stimulus image. (Although only three coordinate axes are depicted, the input space in fact is 225-dimensional; see Methods.) Because the number of input axons equals the number of potential target synapses, the synaptic dendrite vectors (solid arrows) can be rendered in the same space. The initial positions of synaptic vectors is random with respect to the locations of possible dot-pattern vectors (gray dashed arrows) in the input space. (D) During learning, synaptic weights are modified (via potentiation), causing the dendritic vectors to move toward the mean of the inputs on which the given dendritic vector "wins," i.e., responds to. After sufficient trials, a given target vector will respond to all and only those input patterns that are closest to the target; this effectively partitions the input space into separate similarity-based categories (red vs. blue shaded areas).

B. The model showed properties that are similar to the brain.

B.1. Simulated spiking and local field potentials (LFPs) were similar to actual spikes and LFPs.

Simulated model spiking, and calculated "fields" in the simulation, exhibited properties that corresponded in many ways with empirical data collected from the brains of experimental NHPs, as seen in side-by-side illustrations of simulated and empirical physiology (Figure 3). For instance, it can be seen that in response to the simulus onset, there is substantial initial spiking in PFC neurons which rapidly falls off; this occurs in both the simulation and in the empirical physiological data. Further correspondences are described here.

Inputs to the simulation are spatial dot patterns (7) (see Figure 2A) that are received (as topographically organized spatial patterns via continuous currents) into a simple "visual" cortical region (SVCR). Simulated neural units use Hodgkin-Huxley models (see Methods).

Each simulated cortical region also receives oscillatory input, which is centered around a 16 Hz oscillation emerging from a next-generation mean-field neural mass model (NG-NMM) (8, 9) (see Methods) representing a highly simplified simulated ascending system (such as locus coeruleus or basal forebrain), termed ASC1 in the simulation ("BRAINSTEM" in Figure 1D); the resulting rhythmic activity, generated by the neural mass model ASC1, becomes the input to a simulated simplified anterior cortical area ("PFC" in Figure 1D). Those simulated anterior responses are then compared with empirical data from prefrontal electrodes in behaving macaques (7).

SVCR cells are differentially activated depending on the topographic location of dots in the input. Cells will spike or burst, as a function of voltage summation thresholds in the model (see Methods for corresponding synaptic equations). Because of this in combination with the rhythmic inputs from the simulated ascending NG-NMM model, the predominant response of the SVCR thus can be described as a beta-gamma temporal pattern: brief (40msec) bursts of high-frequency spiking, separated by approximately 70 msec; i.e., gamma bursts carried on beta rhythms. The activity occurs only on specific activated input cells in a spatial pattern that is determined by the topographic pattern of the dot-pattern visual inputs. The resulting SVCR activity is input to the simulated anterior cortex (model PFC). Supplementary Figure S2 shows raster plots of model SVCR and Anterior cell firing in response to a sample stimulus.

As described, simulated model spiking and LFPs neurophysiology matched brain spiking and LFPs Figure 3 shows them side-by-side (7). Figure 3 A and D shows a simple histogram-summation of spikes across a field of target model PFC neurons and STR neurons respectively, termed local spike summation (LSS, see Methods). It can be seen that in response to the stimulus onset, there is substantial initial spiking in PFC neurons which rapidly falls off to roughly 8 spikes across 100 simulated target PFC cells; this is modulated by the rhythmic model ascending input oscillating at a mean rate of 16Hz. At the stimulus offset (600 msec after onset), the LSS response in anterior cells greatly lessens, but does not go to zero activity, due to cortico-striatal-thalamo-cortical feedback, as will be described.

Similiarities were also seen between simulated and actual LFPs. In Figure 3 B and E the membrane voltages within the model PFC and STR are averaged to yield an average local membrane potential (ALMP) (see Methods). Even though the computed LSS and ALMP values are extremely simple, they are intended to be compared to empirical measures of local field potentials (LFPs) from experimental data. Figure 3 C and F shows LFPs in the PFC and STR of a trained macaque, beginning at the onset of a stim-

ulus; again it is seen that some important aspects of real physiological fields are being captured even in this relatively simplified simulation.

It is of interest to study whether such minimal circuitry may go further, to generate novel circuit responses that can be experimentally tested, and then affirmed or falsified. Examples along these lines will be introduced in several spots in this manuscript. One such example is shown in Figure 4.

In Figure 4 A-C, notably, both simulated and empirical data show strong responses immediately, during the first 150-200 msec after stimulus onset, and then exhibit strong relative silences for more than 100 msec, before any further (late) responses to the stimulus; this initial activity and response can be seen in the simulation in Figure 4 A-B and in the empirical data in Figure 4 C. The figure then shows that the analysis of the simulation revealed the mechanism for the post-stimulus-onset silence: the strong initial onset response activates a powerful lateral inhibitory response, building up transient inhibition over the first 50 msec; this in turn strongly inhibits further excitation in the model (for 100 msec) by hyperpolarizing the excitatory neurons (see Figure 4 D and E). The duration of the silent period is thus directly explained by the duration of GABA-Ra receptor-based inhibitory influences, lasting roughly 100 msec. It is hypothesized that these mechanisms may also directly explain the corresponding post-stimulus-onset silence in the empirical data.

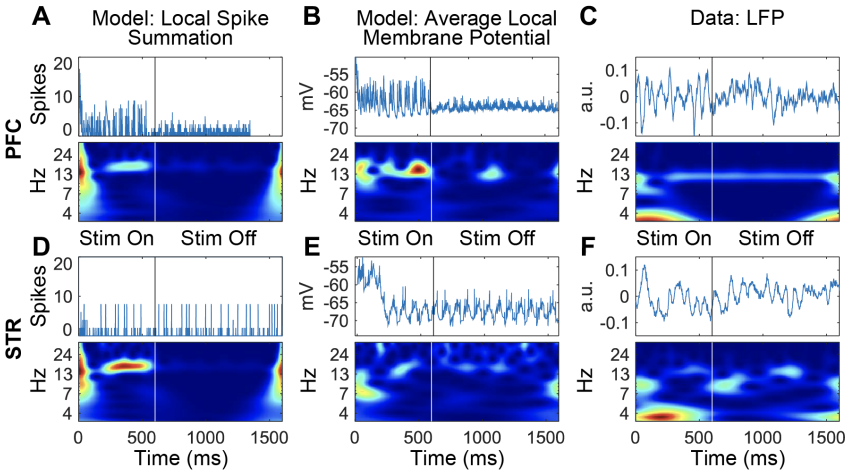

Fig. 3. Model physiology richly matches empirical physiological data. (a) Top: Local Spike Summation (LSS) histogram of model cortical cell spikes in response to stimulus onset (0) and then offset (600msec). Bottom: Spectrogram of LSS response. (b) Average local membrane potential (ALMP) straightforwardly calculated by averaging membrane potentials of all cells, and corresponding spectrogram. (c) Sample empirical local field potential (LFP) measured from macaque cortical electrode, and spectrogram. (d) Model LSS response in sample striatal MSN model neurons, and spectrogram. (e) Striatal model ALMP calculated by averaging membrane potentials of all cells. (f) LFP measured from striatal electrode in macaque, and spectrogram. See text for discussion.

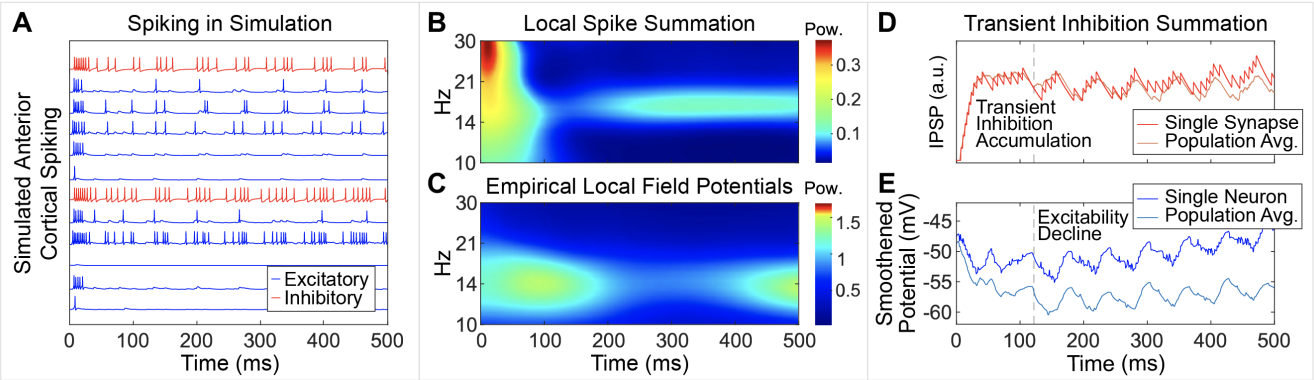

Fig. 4. Model spiking generates field activity that was subsequently found and explained in empirical data. (a) Simulated spiking activity in anterior cortical cell. Note initial burst of spiking activity in a subset of neurons, followed by sparse bursting activity which gradually gets denser. (b) Spectrogram of LSS obtained from simulated spiking activity in anterior cortex, averaged over all trials in a single learning session, shows transient activity after onset and the following brief silence before further activity. (c) Spectrogram of LFP measured from a sample electrode in anterior cortex, averaged across all trials, shows the same behavior. (d) Inhibitory synaptic conductivity averaged over entire neural population in simulated anterior cortex shows the accumulation of inhibition during 100 msec after onset. (e) Subthreshold membrane potential averaged over entire population of simulated anterior cortex shows hyperpolarization as a result of the strong transient inhibition, making the neurons less excitable for a brief window, which results in the intermediate period of relative silence seen in (b).

As will be further discussed, these initial cortical simulated responses within any given trial, unexpectedly turn out to fully identify the spatial pattern of unit firing that will determine the entire remainder of the trial. Those initially-responding cortical units are the ones that selectively activate striatal targets, which in turn determine eventual action selection (saccading left vs. right). Even though the stimulus is present for 600 msec, and there then is a 1000 msec delay period, any physiological activity after the first 200 msec of a trial has little effect on the eventual trial outcome (as discussed later in Figure 7 below).

B.2. Model learning rates were similar to those of non-human primates.

Although the simulation carries out solely physiological operations, in doing so, it successfully learns the visual category task that it is presented with, i.e., it acquires the category characteristics (learning to discriminate category A from B) and it learns (via striatal-based reinforcement) to produce simple appropriate behaviors (saccading left vs. right).

The behavioral responses increase in accuracy across trials (Figure 5 A) in both experimental animals (top) and the simulation (bottom). It should be noted that the training regimens are not quite the same: the animal training incorporated i)

extensive initial shaping and pilot training, absent from the simulation (see Methods); and ii) graded introduction of training instances such that initially only single category deformations were shown, and others added subsequently (see Methods). A remarkable observation is that the somewhat sudden transition from early exploratory behavior to learned behavior, co-occurs with the suppression of TAN activity. These events play an important role in action selection within the simulation [See Supplementary Figure S4].

B.3. Synaptic changes caused spiking and LFP synchrony effects .

Synapses in cortex and in striatum are becoming modified during learning in accordance with specific physiological rules dictating plasticity (see Methods), and these synaptic changes are the cause that underlies both the behavioral and the field phase changes that occur.

Figure 5 B shows changes in averaged synaptic strength across the simulated network for cortico-cortical (top), cortico-striatal (2nd), and thalamocortical (3rd) contacts across trials, as well as behavioral correct trials (bottom), from early training (trials 1-150) to late (trials 500-700), following the behavioral protocol of reference (7). Distinct learning rules obtained for cortico-cortical, thalamo-cortical, and cortico-striatal synapses as detailed in Methods.

In particular, cortico-cortical synapses exhibit long term increases in weight (conductivity), and very recent (within one simulated minute) weight increases can be reversed if they are interfered with by rhythmic beta spiking, in the absence of any gamma bursting (see, e.g., (27)). Details of the reversal effect in cortex are described in Methods under the section "Cortico-cortical and thalamo-cortical and cortico-striatal synaptic plasticity."

By contrast, striatal synapses exhibit increases and decreases as per the combined effect of cortical glutamatergic, substantia nigra dopaminergic, and tonically active neuron cholinergic inputs that converge at each spine of striatal medium spiny neurons (see Methods, and detail in Figure 1). Thus, cortical synapses are seen to slowly increase, and cortico-striatal synapses strengthen and weaken over trials.

In time-window raster plots of activity across the 1600 msec of single sessions across trials (Figure 5 C), during the first 600 msec, the stimulus is shown; during the next 1000 msec, the stimulus is off and the subject is waiting to respond. In each plot, activity for the first 100 trials are shown, followed by the last 200 trials (trials 500-700). The top two plots are the simulations of cortex and striatum, respectively; the bottom two plots are empirical data from (7). In early trials, initial relatively sparse firing builds up; it also can be seen that activity is higher while the stimulus is present. During late trials the spiking activity is significantly denser in simulated cortex and striatum neurons (top two panels), which is in accordance with the empirical data from recorded neurons (bottom two panels). In the working-memory period (600-1600 msec), spiking is maintained in the simulation by cortico-striatal-thalamo-cortical loop activity. This might explain the synchrony behaviour as seen in the following panel.

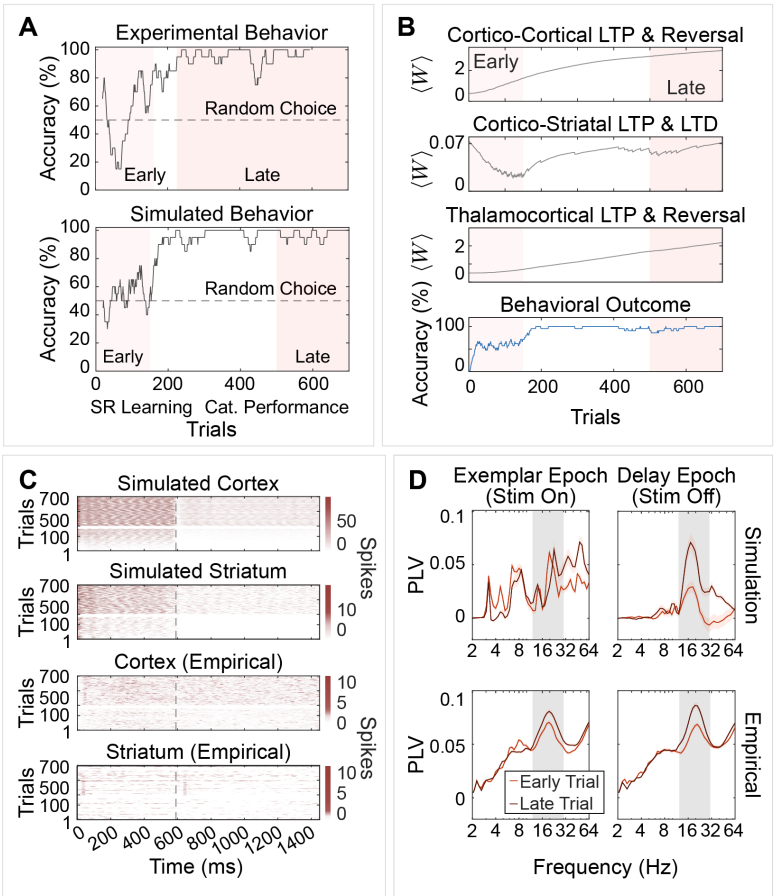

Fig. 5. Cortical and striatal synaptic change underlie behavioral, spiking and synchrony changes. (A) Behavioral learning over trials in macaques (top) and simulation (bottom) exhibit comparable characteristics. Correct behaviors increase slowly over early trials (light pink), and reach plateau performance by late trials (dark pink). (B) Changes to synaptic strength for cortico-cortical, cortico-striatal, and thalamocortical contacts across trials, as well as accuracy of behavioral trials. (C) Simulated (top two) and empirical (bottom two) activity of single sessions across trials. Each row represents Local Spike Summations (see Methods: Field estimation) for the neuronal population (simulation as well as data) over a single trial. Broken lines indicate the time of stimulus offset. (D) Cortico-striatal phase-locking values from simulated (top) and empirical (bottom) data. See text for discussion.

Figure 5 D shows computed phase locking values (PLV) between cortical and striatal activity, from empirical data (bottom) and from the simulation (top), while the stimulus is on (left) and off during the working-memory period (right). It can be seen that the simulation exhibits some characteristics of the data and not others. In particular, during stimulus presentation, the simulation shows a more marked high-theta / low-alpha rhythmic activity (despite there being no oscillatory input whatsoever in the model in that frequency range). Within the low-beta band (14-30 Hz; marked in gray) the model receives rhythmic ascending input at 16 Hz, accounting for the activity in this frequency band that is seen in cortex, but not accounting for synchrony between cortex and striatum. Both the data and the simulation show a modest increase in cortico-striatal synchrony

in late trials (dark red) over early trials (bright red). During the delay, both the empirical and simulated responses show increased cortico-striatal synchrony in this beta frequency band.

B.4. Both the model and the brain showed different synchrony patterns during correct vs incorrect performance.

The synchrony shown in Figure 5 is only obtained for correct trials: i.e., when a subject's saccade (left or right) corresponds to the correct category (visual type A or B). Potential testable predictions are being developed.

Different phase-locking between cortex and striatum arose both in the NHP experimental subjects and in the simulation (Figure 6 A). The (average, z-transformed) difference between error-trial and correct-trial phase locking is shown; in each graph, values above zero denote greater cortico-striatal synchrony on error trials than on correct trials. Slightly higher error-trial synchrony in early trials (light red) is reversed with learning, such that in later trials (dark red), synchrony in the beta frequencies are higher in correct trials than in incorrect (error) trials, both during stimulus viewing (A, left panels) and during working-memory delay period (right panels); and both in the model (bottom panels) and empirical data (top panels).

The simulation provides a candidate explanation for how these changes to phase locking may straightforwardly arise via changes to corticostriatal synaptic weights. Figure 6 B shows the strengths of cortico-striatal synapses for cells that respond on correct trials (blue) vs those that respond on incorrect trials (orange); these weights stay similar to each other throughout training, but then markedly diverge only after extensive learning has occurred (after 450 trials); this is the primary cause of the learning-related relative reduction in error-trial synchrony seen in the simulation (Figure 6 A , bottom).

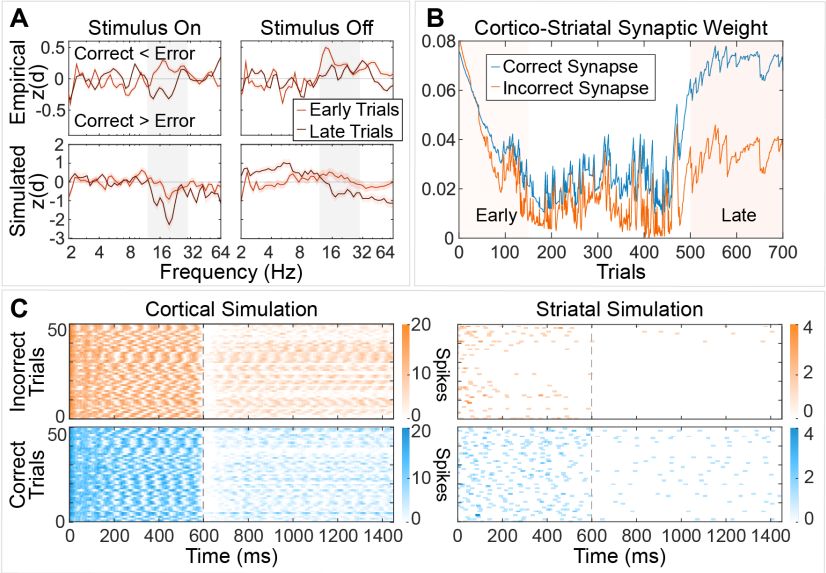

Fig. 6. Striatal synapses drive phase-locking between cortex and striatum (A) Cortico-striatal synchrony on correct vs. incorrect trials. Shown are phase-locking values (PLV) for incorrect minus correct trials; thus above zero are instances where incorrect trials have greater synchrony than correct trials; less than zero means correct trials have greater synchrony. In macaque subjects ("empirical", top), cortico-striatal synchrony in the beta frequency range decreases with learning, from early (light red) to late (dark red) trials, both during stimulus presentation (left) and delay period (right); the simulation exhibits this same effect (bottom left and right). (B) In the simulation, the synchrony decrease is found to be due predominantly to differential synaptic strengthening on those striatal cells responding on correct (blue) vs on incorrect trials (orange). (c) These synaptic effects show up in raster plots of simulation population activity (see Methods: Field estimation) in late trials, after extensive learning, in cortex (left) and striatum (right) on incorrect (top) and correct (bottom) trials. Due to selective weakening of cortico-striatal synapses on incorrectly-responding cells (as per panel b), the striatal responses in incorrect trials become extraordinarily sparse; this is the primary contributor to diminishing cortico-striatal phase-locking values (See text for discussion).

Raster plots of unit firing in the simulation in late trials (after trial 500) shows that both during stimulus presentation (0-600 msec) and delay (600-1600 msec), simulated striatal firing is far lower in error trials (top; orange) than in correct trials (bottom; blue) (Figure 6 C). As proposed, this is due to reduced synaptic strengths (Figure 6 B). This underlies the simulation effect in Figure 6 A in which cortico-striatal synchrony becomes lower for error trials than for correct trials after extensive learning. It is hypothesized that these synaptic changes and resulting cell spiking activity patterns are predictive of causative effects in the macaques; i.e., the simulation is providing a specific candidate explanation of a complex phase-locking phenomenon observed in the experimental subjects.

C. The model predicted new neural properties that were confirmed in the brain .

C.1. "Bad-idea neurons" activated before incorrect behavioral responses in the simulation.

The simulation exhibited a surprising set of physiological responses that emerged as it learned. These were unanticipated even though the experimental physiological records had already been thoroughly studied. The novel simulation responses were subsequently looked for, and verified, in that experimental data where they previously had been missed.

After extended learning trials, each individual simulated cell was found to exhibit a characteristic preference to fire selectively for a stimulus category (A vs. B). Moreover, in addition to that category selectivity, neurons also were found to selectively respond to different outcome-based types of trials: those that ended with a correct action-selection response, i.e., to saccade to the correct side (left or right) in response to the given input stimulus category, versus those trials that ended in incorrect responses.

Furthermore, these two cell activity preferences (category-specific/outcome-specific) were found to strongly interact: individual units exhibited significant selectivity for one of four groups: category A with correct outcome, category A with incorrect outcome, category B with correct outcome, category B with incorrect outcome.

All simulated cells were found to show strongly selective responses to just one of each of those four conditions: figure 7) shows the relative probability of a simulated neuron spiking to each condition. Moreover, the simulated neurons are seen to

display this spike selectivity very early, within the first 200 msec, (Figure 7 A), i.e., the time period during which the stimulus triggers strong initial cortical activity (Figures 3 and 4). In Figure 7 B, these selectively spiking cells are followed across the 1600 msec of a trial that precede behavioral responses (i.e., complete stimulus presentation plus the subsequent no-stimulus delay period); it is seen that they tend to continue their selectivity in the simulation.

C.2. Apparent "bad-idea neurons" also activated before incorrect behavioral responses in the experimental animals.

These findings of highly selective unit responses in the simulation raised the question of whether such activity had occurred in the empirical data of the macaque subjects. Figure 7 C and D exhibit the spiking activity recorded in PFC in the experiments of reference (7), correspondingly sorted for their preferred spiking responses in the four identified (category-specific and outcome-specific) conditions; these empirical patterns compellingly match the simulation findings, especially in the initial 200 msec of each trial (Figure 7 C). (A notable distinction is the smaller number of average spikes produced in the macaques compared with the simulation: 1.33 ± 0.05 spikes in 200 msec empirically vs. 9.52 ± 0.18 spikes per 200 msec in the simulation.)

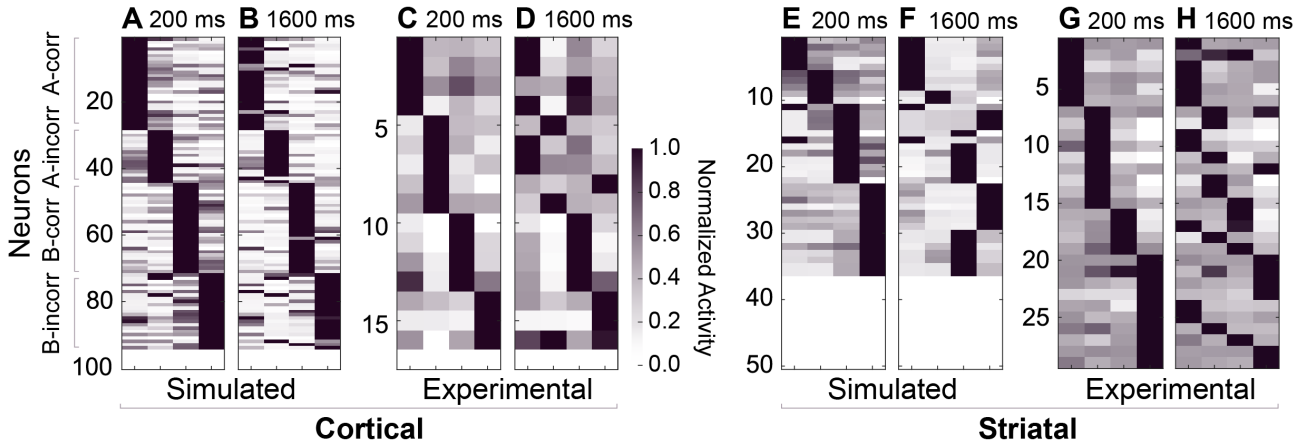

Fig. 7. Cortical and striatal cell spiking preferences during late learning trials, modeled and empirical, in four conditions by category and outcome: category A vs. B; outcome correct vs. incorrect. (See vertial axis descriptors at left.) (A) Simulated cortical spiking preference sorted by four conditions, averaged over first 200 msec. (B) Simulated cortical preference over 1600 msec. (C) Empirical cortical spiking preference in first 200 msec. (D) Empirical cortical preference over 1600 msec. (E) Simulated striatal preference, 200 msec. (F) Simulated striatal preference, 1600 msec. (G) Empirical striatal preference, 200 msec. (H) Empirical striatal preference, 1600 msec. (See text). White regions indicate units with no spiking activity.

Corresponding patterns are seen in striatal MSN activity. Figure 7 E and F show condition-dependent spiking in model MSN cells and Figure 7 G and H in empirically recorded MSN spiking. Notably, for MSNs, the first 200 msec of responses are not reliably continued throughout the 1600 msec pre-response period, in the simulation nor empirically. In the simulation, this may be due to the strong lateral inhibition between response-specific groups of MSNs. This causality is hypothesized to explain the effect in the macaque MSNs; indeed, the empirical responses across 1600 msec correspond even less to the first 200 msec responses than in the simulation. The simulation notably consists of only two separate groups of MSN action-response cells; these are only intended to serve as stand-ins for very large groups of distinct MSN response cells in a real striatum. Only two were constructed for this simulation because these would suffice to produce the only two necessary responses for this particular experiment (saccade left vs. saccade right). In larger simulations under construction, dozens of candidate response groups of MSNs are incorporated, rather than just two. In the empirical recordings, it is possible that the increased variation of responses after the initial 200 msec may be due in part to the far larger numbers of possible MSN groups. This further emphasizes the potentially consequential hypothesis that action selection in striatum is arrived at very rapidly (within 200 msec) after a stimulus, and that subsequent measures of striatal activity may not be indicative of what action will be eventually triggered.

C.3. "Bad-idea neurons" are a good idea.

The identified neuronal spiking preference patterns just described reveal sets of neurons that selectively respond when the simulation (or macaque) is just about to make a particular kind of error. In particular, the "A-incorrect" and "B-incorrect" neurons identified here denote cells whose rapid responses indicate an error that will not actually occur for roughly another 1500 milliseconds.

We show here the underlying mechanisms by which these specialized neural responses arise; we then show that the resulting signals have information processing utility, both in the present experiments and especially in potential extensions to the experiments, which we propose here.

Earlier in this paper it was shown that cortical cells learn via synaptic potentiation in response to the spatial pattern of visual inputs, and that these synaptic changes do not obtain, if subsequent dopamine from VTA arrives at cortex (which causes reversal of recent synaptic increases that have not yet consolidated). That dopamine activation in the simulation is caused specifically when an erroneous response has occurred. Thus synapses should not increase in error trials. How, then, can cells arise that respond selectively to a category in situations where the response will come to be an incorrect saccade?

Each dot-pattern stimulus shown to the simulation (and to the macaques) is a fixed distance away from a pre-set prototype for a given category. Thus, for instance, all "category A" stimuli are a fixed distance from a category-A prototype (which

itself is never shown as a stimulus). The distances among members of category A are very small compared with the distance between category A and category B. Thus it may appear quite unlikely, for instance, that a cell responding to an A stimulus could ever respond to a B stimulus.

However, recall that each target cortical neuron contains only a small subset of potential synaptic connections from the input. In the simulation, of 225 input axons conveying the dot-pattern visual image, each such axon produces just eight synapses onto target cells, and each target dendrite receives just 18 such synapses. Thus each dendrite "sees" only a projection of the 225-dimensional space onto an 18-dimensional subspace. Within those subspaces, it is possible for some category A and category B dot patterns images to overlap (see Supplementary Figure S3 for the distributions of projection angles between dendritic vectors and visual stimuli, showing considerable overlap in angular values between two categories). Thus the distance between images in that subspace may be far smaller than the large distance in the full input space.

During the final 300 trials of a sample 700-trial training session, i.e., once the stimuli are well learned, roughly 20-30 of those trials yield at least some incorrect responses. Within those responding neurons, roughly 15-25% are classified as statistically being "incorrect-A" or "incorrect-B" cells.

Since the dendrite of a simulated cortical cell receives only an 18-dimensional subspace projection of a dot-pattern visual image, the effective 'receptive field' of a given neuron can be visualized, by projecting its subspace inputs back onto the visual field.

Figure 8 A shows four panels, each displaying a sample visual input pattern (white patches) superimposed on the backprojected receptive field (red), thereby highlighting the overlap, i.e., the components of the input that can be 'seen' (orange patches) by this given cell, in light of its synaptic connectivity from the inputs.

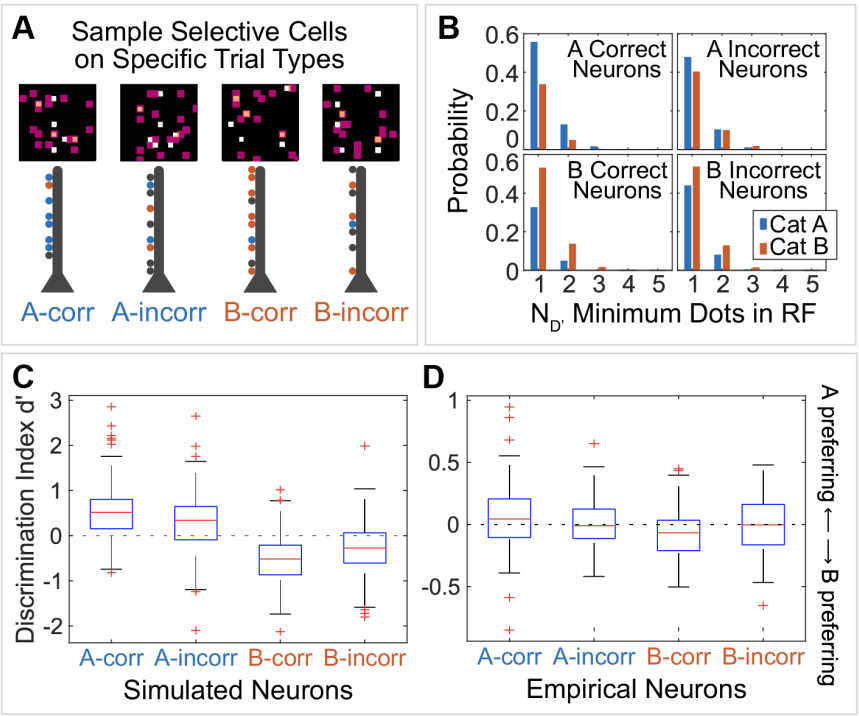

Fig. 8. Neuron selectivity for stimuli and for trial outcomes, in simulated and in empirical data. (A) Sample cell receptive fields are shown superimposed on sample visual input patterns, for each of four cell types and trial types (A correct, A incorrect, B incorrect, B correct). (B) Selective cell responsiveness is quantified as difference in probable response of category-selective (A vs. B) cells to their corresponding visual inputs (category A vs category B) on correct-outcome vs incorrect-outcome trials. (C,D) Calculated discrimination indexes (d') for the four responsiveness types during early trials, showing d' values farther from zero for correct A-preferring and B-preferring cells (1st and 3rd columns), but nearer-zero values for responsiveness to A or B on incorrect trials (2nd and 4th columns). (See text).

For "A-corr," what is shown is the receptive field of a sample "A-correct" cell in response to a visual input pattern that occurred on an A-correct trial in the simulation; the corresponding match-up is shown successively for A-incorrect, B-incorrect, and B-correct inputs and receptive fields. In general, a neuron that has been found to be an A-correct responder will have several overlaps between its receptive field and the visual pattern (A) during this trial (a correct-behavioral-response trial), and correspondingly for the other categories. Thus the A-corr and B-corr neurons have receptive fields with the most overlap between A and B visual patterns respectively, whereas A incorr and B incorr have less overlap with their visual category and correspondingly more with the competing category. The schematic shows sample synaptic responses to A (blue) vs B (orange) inputs on the four trial types (A-corr, A-incorr, B-incorr, B-corr), illustrating this pattern.

Figure 8 B quantifies these selectivities, showing the relative responsiveness of each of the four cell-responsiveness types when presented with a category A or category B visual input. As can be seen on trials that eventually lead to correct behavioral responses, there is a broad distinction between the overlap of the cell's receptive field to the given visual category. On trials with eventual incorrect behavioral responses, there is substantially less cell receptive field match selectivity between visual categories.

Figure 8 C and D show selectivity of cell responses in both the simulation and in macaque data in very early trials, i.e., presumably before substantial learning or plasticity has taken place in the experiment. Shown are the calculated discrimination index d' values for both simulated and empirical cell responses in the four trial types. A-correct and B-correct neurons exhibit preferences (d' values) that are farther from zero for their corresponding visual categories (A vs. B), whereas A-incorrect and B-incorrect neurons exhibit less preference (d' closer to zero).

These results show that cell selectivity in the simulation is in part predetermined, even before learning, via the arrangement

of synaptic connectivities, i.e., the connectome, and its corresponding receptive fields in the input space. (These then are further influenced by synaptic plasticity.) Moreover, the findings suggest that a commensurate pre-learning selectivity is present in the macaques: in these very early trials, before the animals exhibit reliable learning, the categorized cells already behave preferentially toward particular categories (A vs B), if those cells will go on to become A-correct and B-correct responders. Conversely, cells that do not exhibit early preferences (as shown by their near-zero d' values) are those that are more likely to go on to become A-incorrect and B-incorrect cells in later trials.

Although the utility of such cells to the animal may be uncertain, it can be argued that these cells serve a useful function. From the point of view of a given experimenter, the existence of A-incorrect and B-incorrect cells (notionally termed "bad-idea neurons") may simply appear to signal an upcoming behavioral error. However, any subject learning the task has, of course, no prior information that there will only be two types of trials (A vs B), nor only two possible behavioral responses (saccade left or right), let alone that the A visual patterns will all be similar to each other and distinct from B patterns.

It may be surmised that it is most adaptive for a learner to produce multiple categories and subcategories of its learned environmental inputs, maintaining their separability and thus enabling different inputs to be associated with different possible (as yet unknown) outcomes. Imagine just a slightly altered experimental arrangement where, instead, subjects are required to learn several distinct saccade directions for six different categories. We hypothesize that the cells showing up here as bad-idea neurons may form the kernel of further differentiations in responsiveness in that expanded set of experiments. This idea is the subject of ongoing study in further sets of experimental investigations. If these so-called bad-idea cells turn out to have broader efficacy when presented with different experimental conditions, it may be that the tuning of these cells only appears to be a bad idea when viewed with the limited lens of a single set of experiments, missing their potential utility in a broader set of circumstances.

3. Discussion

A biologically derived model exhibited spiking behavior leading to the discovery of "bad-idea neurons," that selectively responded if the simulation would soon respond incorrectly to the stimulus category, i.e., units that preferentially signaled an upcoming behavioral error. Although they had been previously unsuspected, these "bad-idea neurons" were indeed subsequently found in the macaque data. This is a rare instance of a model, never trained on the data, that nonetheless learned a complex categorization task comparably to behaving non-human primates (NHPs), while also exhibiting spiking and field potential synchronies sharing many key characteristics with the NHPs, and then went further to display hidden predictive information encoded in the network, later uncovered in the records of the NHP's brain responses.

The simulation consists of multiple levels of biologically mimetic structures ranging from multiple different synapse types, to single units and local assemblies, to large-scale interacting systems, all constructed from well-studied physiological operations, as embedded in known anatomical circuit layouts (see Table 1). This is in marked contrast to computational systems that are constructed from small sets of highly simplified units, in which, for instance, no distinction exists between excitatory and inhibitory cells, nor the different time duration of their physiological activities, nor their differential circuit wiring in different brain structures. We hypothesize that adding these and other anatomical and physiological constraints assisted in making the model adhere to actual brain circuit activity. Ongoing work is investigating the potential applicability of these simulations to brain dysfunctions, and it is anticipated that the additional biological detail that is represented here will be of direct use.

In addition to these added biological details, the simulation is composed of local circuit units whose computations can be directly specified, i.e., the simulation produces not just cell spiking and local fields and phases, but also specific detailed hypotheses of the computational operations that may be being carried out, within constituent subcircuit structures as well as across large-scale systems that incorporate many such units. Each type of unit exhibits its own computation arising from the physiology of its cell types as they occur within the wiring diagrams of its anatomy; those individual computations then interact to confer larger systems-level effects that arise only from the operation of the model as a whole.

The model never was presented with nor trained on the NHP physiological data on which it was tested; nothing about the task was used in the construction of the simulation. The simulation learned the task over approximately the same learning curve as the NHPs, and more importantly, during the behaviors the model exhibited extensive physiological measures, from individual neuron spiking to fields to cross-region synchrony, that all exhibited strong, and explainable, relations with those of recordings from the NHP's brains.

In addition to the many powerful similarities between the results from the simulation and the NHPs, there also are many notable divergences. For instance, whereas some local-field synchrony effects from the simulation conform with NHP brain measures, some deviate. it is initially hypothesized that the disparities may be due in part to engagement of additional brain structures in the real brains that are absent from, or oversimplified in, the simulations, and it remains unknown whether such additions may integrate well with the existing components or whether unforeseen conflicts may arise. Moreover, many additional factors such as additional transmitters and receptor types, as well as metabolic influences, remain absent from the simulation, and their addition may or may not cause future inconsistencies in the model networks. Ongoing studies are aimed at exploring these expansions of the simulation.

Recurring questions in the literature revolve around how artificial neural networks may be overlain onto real circuitry of the brain; in general, how can modeling tools assist our analyses of what brain units, and assemblies, and circuits, and systems, are actually doing as they process and acquire information. The present work is proffered as one attempted undertaking in which simulated networks are fashioned such that their derivations from biological features, and their intended overlays onto anatomical layouts and physiological patterns of operation, are relatively direct, i.e., more mimetic than artificial.

A primary objective of the present work is an attempt to identify and observe elementary constituent circuit operations that compositionally interact to give rise to more complex physiological phenomena. Data for such phenomena are, in the vast majority of cases, highly underspecified; that is, there are many brain phenomena that might be contributing to empirically measured physiological signals. The work presented here represents an attempt to find candidates for the smallest possible circuits, with fewest added assumptions (biological or mathematical), that can give rise to known physiological characteristics.

In the present studies we found it remarkable that simulated units in these somewhat more biologically-derived networks turn out to perform physiological activity that is actually present in brain circuits, so much so that it predicted a new class of neural spiking behavior that had been previously unsuspected but has been since empirically verified. We hypothesize that combined efforts entailing mathematical tractability on one hand, together with further attempts at biological fidelity on the other hand, may lead to an increasing ability to ascertain the properties and mechanisms that are at play as real brain circuitry directs real behavior in real environmental circumstances.

4. Methods

A. Detailed model description.

Extensive anatomical and physiological features incorporated into the present simulation. A sizeable body of references was used in the conception and construction of the several subcircuits and neural assemblies that constitute the present brain circuit simulation (Table 1, Figure 1D); these consist of attributes that are highly cited in the literature.

Block	Computational Operations		Key Publications
Biomimetic Comp Primitives		*Regions*	
Superficial Cortical (SCORT)	Lateral Inhibition	Anterior, Posterior Neocortex	(14–18, 28–64)
Matrisomal MSN (MMSN)	Action Selection	Caudate, Putamen	(10, 11, 16, 65–73)
Striosomal MSN (SMSN)	Reward Prediction	Caudate, Putamen	(10, 11, 16, 65–77)
Ascending Type 1 (ASC1)	Central Pattern Generation	Pons, Brainstem	(78–92)
Superficial Glu−GABA−DA Plasticity (GLTP)	Long-Term Potentiation, Reversal	Cortex, VTA	(27, 35, 37, 93–107)
Cholinergic Tonically Active Neurons (CTAN)	Exploration/Exploitation	Caudate, Putamen	(10, 11, 108–120)

Table 1. Current Blocks for Corticostriatal Model.

A.1. Spiking neurons and synaptic connections.

Neuronal Dynamics Each neuron is modeled using ionic conductance based point neuron models, with the Hodgkin-Huxley formalism (121). For simplicity, in this model we consider only voltage gated Sodium and Potassium channels. The membrane potential is described by:

$$C_m \frac{dV}{dt} = -g_L(V - V_L) - I_{Na} - I_K + I_{in} - I_{syn} \quad [1]$$

where membrane capacitance $C_m = 1\mu$ F cm^{-2}, leak conductance $g_L = 0.1$ mS cm^{-2}, I_{in} is the input current, which can either be external stimulus or background activity, and I_{syn} is the synaptic input. Sodium current $I_{Na} = \bar{g}_{Na} m^3 h(V - V_{Na})$ and potassium current $I_K = \bar{g}_K n^4 (V - V_K)$ are driven by gating variables $\{m, h, n\}$ whose dynamics are governed by the general form :

$$\frac{d\chi}{dt} = \phi_\chi(\alpha_\chi(V)(1 - \chi) + \beta_\chi(V)\chi) \quad [2]$$

Here, χ is the gating variable and ϕ_χ is the temperature related effect on the time scale.

The following parameters were used as shown in ref. (122) : $\bar{g}_{Na} = 52$, $\bar{g}_K = 20$, $V_{Na} = +55mV$, $V_K = -90mV$, $V_L = -60mV$ and $\phi = 5$.

Different threshold values for used for pyramidal and inhibitory cells. For excitatory pyramidal cells we used the following (122):

$$\alpha_m(V) = \frac{0.1(V + 30)}{1 - exp(-\frac{V+30}{10})}, \qquad \beta_m(V) = 4exp\left(-\frac{V + 55}{10}\right) \quad [3]$$

$$\alpha_h(V) = 0.07exp\left(-\frac{V + 44}{20}\right), \qquad \beta_h(V) = \frac{1}{1 - exp(-\frac{V+14}{10})} \quad [4]$$

$$\alpha_n(V) = \frac{0.01(V + 34)}{1 - exp(-\frac{V+34}{10})}, \qquad \beta_n(V) = 0.125exp\left(-\frac{V + 44}{80}\right) \quad [5]$$

Similarly, for inhibitory neurons we used the threshold values implemented in ref. (123):

$$\alpha_m(V) = \frac{0.1(V+33)}{1-exp(-\frac{V+33}{10})}, \qquad \beta_m(V) = 4exp\left(-\frac{V+58}{10}\right) \tag{6}$$

$$\alpha_h(V) = 0.07exp\left(-\frac{V+51}{20}\right), \qquad \beta_h(V) = \frac{1}{1-exp(-\frac{V+21}{10})} \tag{7}$$

$$\alpha_n(V) = \frac{0.01(V+38)}{1-exp(-\frac{V+38}{10})}, \qquad \beta_n(V) = 0.125exp\left(-\frac{V+48}{80}\right) \tag{8}$$

Synapse

Synaptic current in eq. 1 for a neuron i is the sum of synaptic inputs from all incoming axons $\{j\}$ viz. $I^i_{syn} = \sum_j I^{ij}_{syn}$. Synaptic current from neuron j to neuron i is given by

$$I^{ij}_{syn} = W^{ij} \cdot g^j_{syn}(t)(V^i - V^j_{syn}) \tag{9}$$

where W^{ij} is the synaptic weight from neuron j to i, a quantity which estimates density of dendritic spines and is modified by LTP and LTD (see eq. 34,36). V^j_{syn} is synaptic reversal potential for presynaptic neuron, which for excitatory neuron is $0mV$ and for inhibitory neuron is $-70mV$. The synaptic conductance g^j_{syn} is a function of spiking activity of the presynaptic neuron. As suggested in reference (124), synaptic conductance g_{syn} is one of the state variables for each neuron whose dynamics is described by damped oscillator equation:

$$\begin{aligned} \frac{dz}{dt} &= -\frac{1}{\tau_1}z + g_\infty(V) \\ \frac{dg_{syn}}{dt} &= -\frac{1}{\tau_2}g_{syn} + z \end{aligned} \tag{10}$$

which is driven by impulse force $g_\infty(V)$ defined as follows (125)

$$g_\infty(V) = \frac{\bar{g}_{syn}}{1+e^{-4.394\left(\frac{V-V_{eq}}{V_{range}}\right)}} \tag{11}$$

For excitatory neurons $\bar{g}_{syn} = 3$, $V_{eq} = 10mV$ and $V_{range} = 35mV$ and for inhibitory neurons $\bar{g}_{syn} = 10$, $V_{eq} = 0mV$ and $V_{range} = 35mV$.

Equations 10 provide to different time scales τ_1 and τ_2 for the synaptic conductance waveform which determines the rate of rise and decay respectively. For excitatory neuron we set $\tau_1 = 0.1$ ms and $\tau_2 = 5$ ms while for inhibitory neurons, we used $\tau_1 = 0.1$ ms and $\tau_2 = 70$ ms.

A.2. Brain regions modelled as cell assemblies.

The schematic of the overall circuit is shown in Figure 1. Most of the brain areas shown in the circuit have been modelled as cell assemblies of sparsely connected spiking neurons, where each region has its own specific internal structure and neuronal types comprising them. These structures and neuronal types are derived from the well known anatomical and physiological literature. Below are the specific structures these brain areas modelled here.

Cortical areas

Superficial layers of cortical regions in this model (cortical blocks) comprise large number of local lateral-inhibition circuits (winner-take-all). Each winner-take-all circuit (see Fig. 1 A), comprises 5 excitatory neurons sending synapses to one feedback inhibitory neuron while the inhibitory neuron sends synapses back to the 5 excitatory neurons. The winner-take-all circuits are connected to each other through synapses between the excitatory neurons. We have set the connection density such that each excitatory neuron receives synapse from 1 excitatory neuron from the same cortical block. Within each cortical block, we have 1 *feed-forward* inhibitory neuron which recieves input from excitatory neurons of other cortical regions and sends inhibitory synapses to excitatory neurons of the same cortical block. These feed-forward neurons also receive modulatory inputs from various *Brain Stem* nuclei forming the Ascending Reticular Activating System (78). The rhythmic bursting of feed-forward inhibitory neurons causes the rhythmic bursting of the excitatory neurons and which in turn sets the dominant rhythm in the cortical region.

There two cortical blocks in this simplified model : *Visual Area* (also referred here as Simple Visual Cortical Region, SVCR) and *Prefrontal Cortex*. Visual area consists of 45 winner-take-all circuits comprising ($45X5 =$) 225 excitatory 'target' neurons which receive the stimulus dot patters in the form of input current I_{in}. The visual stimulus consists of ($15X15 =$) 225 pixels (see Figure 2) which are mapped directly onto the 225 visual area neurons, where the I_{in} for each neuron is proportional to gray scale value of the corresponding pixel.

Prefrontal cortex (PFC) comprises of 20 winner-take-all circuits, hence ($20X5 =$) 100 target excitatory cells. The connections between Visual Area and PFC are hypergeometric (13) viz. every neuron have equal number of outgoing and incoming synaptic connections. Every Visual Area neuron makes an outgoing synapse with 8 out of 100 PFC neurons, while each PFC neuron receives incoming synapse from 18 out of 225 Visual Area neuron, thus the connection probability for cortico-cortical connection is 8%.

Striatum

From striatum onwards, we model two blocks ('L' and 'R', say) for each of the downstream brain regions in the loop, viz. *Striatum, GPi, GPe, STN*, and *Thalamus*. Each block in one region connects with the corresponding block in the other region viz. $L(R)$ block of one region connects to $L(R)$ block of another region, thus representing the topographic connections. These two streams of connection correspond to the two choices of response actions (in this case left or right saccade).

Striatum has three components : *Matrisomes* (matrix) and *Striosomes* (patch) MSNs that are GABAergic, and the *Tonically Active Neurons* (TANs) that are cholinergic. These three form a striatal complex (11) which are supposed to play a vital role in reward based action selection (112), we modelled these components along with *SNc* as simplified physiologically-derived "rule" defining their input/output behavior (see following subsection, eq. 17,20,22,24). However we also modelled two blocks of matrisome as spiking neuronal assemblies (*L-STR* and *R-STR*) in order to capture the temporal and spectral properties of the striatal data from the experiments. We also model a single TAN spiking neuron which provides modulatory input to MSNs.

Each matrisomal block consists of single layer of 25 inhibitory neurons, where each neuron gets incoming synapses from 4 PFC neurons. These neurons also get sub-threshold excitation from TAN neuron. Since it was experimentally observed that the TAN neurons are phase locked with the striatal LFPs (126, 127), we propose that the striatal rhythm is modulated by rhythmic bursting of TANs. TAN neurons in turn get the rhythmic input from *Intralaminar Thalamic Nuclei*.

GPe

Two blocks of single layer assemblies of inhibitory neurons, 15 neurons in each block . Each neuron gets connection from 1 striatal neuron of corresponding block in striatum (topographic connections). Each GPe neuron sends synapse to 1 GPi neuron of corresponding block.

STN

Two blocks of single layer assemblies of excitatory neurons, 15 neurons in each block. Each neuron gets connection from 1 GPe neuron of corresponding block in GPe (topographic connections). Each STN neuron sends synapse to 1 GPi neuron of corresponding block.

GPi

Two blocks of single layer assemblies of inhibitory neurons, 25 neurons in each block. Each neuron gets synapse from 1 striatal neuron of corresponding block.

Thalamus

Two blocks of single layer assemblies of excitatory neurons, 25 neurons in each block. Each neuron gets a synapse from 1 GPi neuron of corresponding block. Thalamic neurons project back to PFC target neurons thus completing the loop. Each PFC neuron gets input from 16 thalamic neurons i.e. 8 neurons from each thalamic block.

Ascending Systems and Intralaminar Thalamic Nuclei

The simulated ascending system ASC1, an extreme simplification of ascending systems such as the cholinergic basal forebrain and noradrenergic locus coeruleus, produces rhythmic activity via a neural mass model. That activity becomes simulated modulatory input to cortex, driving a beta rhythm in the cortex (81). The ASC1 simulation thus functions as a central pattern generator. It has been modelled using a simple 'Next Generation' neural mass model (8, 9) which is a mean field approximation of a network of rhythmic neurons (128). A single neural mass representing an ascending system nucleus comprises a synaptically coupled population of excitatory (E) and inhibitory (I) neurons. Their collective activity is defined by the complex valued Kuramoto order parameter Z_e and Z_i respectively, while their synaptic interaction is measured by synaptic conductivities g_{ee}, g_{ei}, g_{ie} and g_{ii} such that g_{ab} is the synaptic conductivity from population b to population a, where $a, b \in \{E, I\}$. The dynamics is described as :

$$\tau_a \frac{dZ_a}{dt} = \mathcal{F}(Z_a; \eta_0^a, \Delta^a) + \sum_b \mathcal{G}(Z_a, g_{ab}; v_{syn}^{ab}) \quad [12]$$

$$\alpha_{ab}^{-1} \frac{dg_{ab}}{dt} = \kappa_{ab} f(Z_b) - g_{ab} \quad [13]$$

where

$$\mathcal{F}(Z; \eta_0, \Delta) = -i\frac{(Z-1)^2}{2} + \frac{(Z+1)^2}{2}[-\Delta + i\eta_0] \quad [14]$$

$$\mathcal{G}(Z, g; v_{syn}) = i\frac{(Z+1)^2}{2} v_{syn} g - \frac{Z^2-1}{2} g \quad [15]$$

and the population firing rate f is given by

$$f(Z) = \frac{1}{\pi\tau} Re\Big(\frac{1-Z^*}{1+Z^*}\Big) \quad [16]$$

The firing rate $f(Z_e)$ of population E directly affects the modulation of the target neurons, in the form of an input current to the target neurons such that $I_{mod} = k_{LC} f(Z_e)$. In this simulation, those targets are simulated somatostatin feed-forward inhibitory cortical neurons.

For ASC1, we used the following parameters : $\tau_e = 52$, $\tau_i = 26$, $\eta_0^e = 10.0$, $\eta_0^i = 0.0$, $\Delta^e = \Delta^i = 0.5$, $v_{syn}^{ee} = v_{syn}^{ie} = 10.0$, $v_{syn}^{ei} = v_{syn}^{ii} = -10.0$, $\alpha_{ee} = \alpha_{ie} = 0.385$, $\alpha_{ei} = \alpha_{ii} = 0.031$, $\kappa_{ee} = \kappa_{ii} = 0$, $\kappa_{ei} = \kappa_{ie} = 15.6$, $k_{LC} = 44$.

Simillarly, Intralaminar Thalamic Nucleus (ITN) is modeled as a neural mass which targets the TAN neurons (12) which then drives the alpha rhythm (12 Hz) in the striatum. We again use the next generation neural mass model with the following parameters : $\tau_e = 72$, $\tau_i = 36$, $\eta_0^e = 10.0$, $\eta_0^i = 0.0$, $\Delta^e = \Delta^i = 0.5$, $v_{syn}^{ee} = v_{syn}^{ie} = 10.0$, $v_{syn}^{ei} = v_{syn}^{ii} = -10.0$, $\alpha_{ee} = \alpha_{ie} = 0.278$, $\alpha_{ei} = \alpha_{ii} = 0.022$, $\kappa_{ee} = \kappa_{ii} = 0$, $\kappa_{ei} = \kappa_{ie} = 21.6$, $k_{ITN} = 100$.

Ascending systems are known to substantially participate in cortical rhythmicity, via modulatory inputs either to pyramidal cells or interneurons (78–85). The resulting rhythmicity, in multiple EEG frequency bands, show evidence of both intrinsic and extrinsic influences (129–132). Differential effects of intrinsic and ascending components may be studied in part by modeling them separately; we here specifically investigate the extent to which simple rhythmic driving from ascending signals can give rise to complex effects of behavioral states quite beyond those of simple rhythmicity: for instance, the phase, amplitude, and synchrony of these rhythms and their dependencies on different behavioral states. In Parkinson's disease, evidence suggests that symptoms are accompanied by changes to the striatal beta rhythm with possible involvement of the subthalamic nucleus and globus pallidus pars interna (GPi) (126, 133, 134). The potential applicability of the present work to Parkinson's and other clinical conditions is being actively pursued. The extraordinary anatomical range and physiological driving power of modulatory ascending systems targeting cortical and subcortical structures suggests their involvement in rhythmic activity; their implication in clinical disruptions is a subject of ongoing study.

We model a simplified ascending system as a next-gen neural mass model (NG-NMM) (8) with parameters set such that calculated local membrane potentials exhibit increased power in the driving frequency range as a Gaussian centered at 16 Hz. The cholinergic basal forebrain, as one ascending example, provides two cell populations with two sets of ascending fibers: cholinergic axons contacting pyramidal cortical cells plus inhibitory axons that contact somatostatin inhibitory feedforward interneurons in cortical and subcortical targets (79, 80). By contrast, the locus coeruleus reportedly projects noradrenergic axons to pyramidal cells; however, there also exists a locus coeruleus-adjacent population of inhibitory cells that reciprocally contact LC neurons, and these inhibitory cells ascend to possible somatostatin targets in cortex (135). Although these LC-adjacent cells are not typically considered a part of LC proper, it is possible that this organization may reflect something analogous to that in basal forebrain. Although brainstem nuclei are typically excluded from human brain imaging, recent work is in accord with extensive prior evidence, indicating extensive alignment of cortical and brainstem activity (136). In the present simulations, an abstracted ascending system is implemented as a stand-in, with the intention of initially exploring the simplest potential effects of such systems. Specialized differential characteristics of cholinergic (79, 80), dopaminergic (137), and noradrenergic (78, 81) systems are being implemented and tested separately and in combination, as part of ongoing research.

The basic rhythmic input is thus modeled at present by the next-gen neural mass models (NG-NMM), and the present simplifications are in the process of being deepened to study these rhythmic phenomena in more depth, including both ascending influences and local intrinsic rhythmic tendencies within cortical fields. Notably, all other rhythmic effects in thie present work: cortico-striatal synchrony; phase-locking values (PLV); learning-related frequency and synchrony changes; rhythmic changes accompanying behavioral states (stimulus, working memory delay, action selection, etc.), all arise solely from modeled unit neuronal activity and synaptic change in the simulation. It is hypothesized that these effects occur as identified in the present simulation, and further elaboration of rhythms will not fundamentally alter these primary findings.

Studies in progress incorporate increasingly detailed simulations of specific ascending system influences as well as local intrinsic oscillatory activity in cortex, in hippocampus, and elsewhere. Initial preliminary findings indicate that these added details continue to conform closely to the current findings presented in the present manuscript, as well as introducing further findings of interest that are the topic of ongoing investigation.

A.3. Cortico-Striatal decision making and synaptic plasticity.

The spatio-temporal spiking patterns in the above described CSTC circuit have been simulated by modelling the interconnected cortical and subcortical regions as synaptic networks of neuronal assemblies, each comprising spiking neurons of Hodgkin-Huxley type (as described above). Now we describe the set of mechanisms through which these spatio-temporal firing patterns generate the learning driven behaviour (response choice) and feedback driven synaptic plasticity (and thus category learning). Although these equations describe our proposed mechanisms for cortico-striatal computation at a much coarse-grained scale of length and time as compared to the physiological scales of spiking neurons and neuroreceptors, they are nevertheless based on well known physiological and anatomical principles found across the literature.

A notable feature of corticostriatal projections is the distinction between cortical intratelencephalic (IT) and pyramidal tract (PT) neurons; the former project ipsi- and contralaterally to striatum whereas the latter project only ipsilaterally to striatum and other structures (138). Moreover, IT and PT neurons exhibit distinct local cortical circuit arrangements and distinct physiological characteristics. Possible differential effects of distinct dopamine receptors has been reported but remains uncertain (139–141). We plan to incorporate these features in future developments of the model.

In principle, all of the following mechanisms can be entirely modelled using only the physiological phenomena at the neuronal length and time scales, as we intend to show in the future editions of this model. However, in this initial rendition of the model we have implemented reasonably simplified population level abstractions of these well defined physiological mechanisms and anatomical constraints. The model is entirely summarized in Table 2. Following is the detailed description the model components.

Population activity in STR and SNc

There are four key sub-populations of subcortical regions: *matrisome* (matrix), *striosome* (patch) and *tonically active neurons* (TANs) in the *striatum* and dopaminergic neurons of *substantia nigra pars compacta* (SNc). We consider one sub-population each for TANs and SNc and two competing sub-populations each for matrisome and striosome, respectively part of the two cortico-striatal loops described in the circuit model. Since matrix and patch neurons are GABAergic neurons with relatively longer time scale ($80ms$), their firing have a cumulative effect over a given time interval. Hence for striosome and matrisome we consider their cumulative spiking activities ρ^{str} and ρ^{mat} respectively. For dopaminergic SNc neurons and cholinergic TANs on the other hand, we consider their instantaneous firing rates R^{TAN} and R^{DA} respectively.

Striosome activity : Striosome ("patch") sub-population is driven by spatio-temporal spiking input from cortex, PFC in this case, and is laterally inhibited by matrisome ("matrix") of the competing sub-population. The cumulative spiking activity ρ_x^{pat} in striosome sub-population x at time t over interval Δt is given by

$$\rho_x^{pat}(t,\Delta t) = PAT\Big(\sum_{i\in Ctx} W_x^i(t)\cdot\big[S^i\big]_{t-\Delta t}^{t}\Big) - LI\Big(\sum_{y\neq x}\rho_y^{mat}\Big) \qquad [17]$$

where, PAT is a generic function with cortical neuronal activity as input and striosomal ("patch") activity as output. Here cortical input is represented by $\big[S^i\big]_{t-\Delta t}^{t}$ that is the number of spikes occurring in PFC neuron i during interval $[t-\Delta t, t]$ and $W_x^i(t)$ is synaptic weight for the synapse(s) from PFC neuron i to striatal sub-population x (including patch and matrix). Lateral inhibition function LI represents the inhibitory inputs coming from all other areas of the striatum.

For the sake of simplicity, in this paper we used discrete time maps for population activity instead of above described continuous dynamics : $t \implies t_n$ where $t_n = t_0 + n\cdot\Delta t_{inh}$. Here t_0 is the time of stimulus onset and interval Δt_{inh} is the time duration of GABAergic inhibition. The discretized form for cumulative striosomal activity (eq. 17) is given by

$$\rho_x^{pat}(t_n) = H_x(t_n)\sum_{i\in PFC} A_x^i(t_n)\cdot\big[S^i\big]_{t_{n-1}}^{t_n} \quad |x\in\{1,2\}, n\in\{1,2,3...\} \qquad [18]$$

and the simplified form for used for lateral inhibition between two striatal regions modeled here is given by

$$H_x(t_n) = \begin{cases} \Theta(\rho_x^{mat}(t_{n-1}) - \rho_y^{mat}(t_{n-1}))|_{(x,y)\in\{(1,2),(2,1)\}} & \text{n>2} \quad [19a] \\ 1 & n\leq 2 \quad [19b] \end{cases}$$

where Θ is the Heavyside function and ρ_x^{mat} is the cumulative spiking activity of matriosome ("matrix") subpopulation x (see below).

TAN activity

TANs are driven by intralaminar regions of *Thalamus* and are inhibited by all striosomal sub-populations. Firing rate R^{TAN} can be generally described by:

$$R^{TAN}(t) = Inhibition\Big(\sum_x \rho_x^{pat}(t,\Delta t)\Big|\kappa^{TAN}\Big) \qquad [20]$$

where, κ^{TAN} is the baseline activity of TAN neurons and $Inhibiton(\rho|\kappa)$ is a monotonically decreasing function of the cumulative spiking activity ρ of the source population, given the baseline spiking activity of the target population is κ. The specific discrete form for TAN activity (eq. 20) we used here is given by

$$R^{TAN}(t_n) = \begin{cases} \kappa^{TAN}, & \text{n=1} \quad [21a] \\ \dfrac{\kappa^{TAN}}{\Big[\sum_x \rho_x^{pat}(t_{n-1})\Big]}, & \text{n>1} \quad [21b] \end{cases}$$

Matrisome activity

Matrisome sub-populations receive simillar inputs as corresponding striosome sub-populations with an additional modulatory input from TANs. Hence the cumulative spiking activity in matrisome (matrix) sub-population x over interval Δt at time t is given by

$$\rho_x^{mat}(t,\Delta t) = MAT\Big(\sum_{i\in PFC} A_x^i(t)\cdot\big[S^i\big]_{t-\Delta t}^{t}\Big) + TAN(R^{TAN}) - LI\Big(\sum_{y\in N}\rho_y^{mat}(t,\Delta t)\Big) \quad \Big|y\neq x \qquad [22]$$

where, MAT is a generic function like PAT but for matrisome, with cortical neuronal activity as input and matrisomal (patch) activity as output and TAN is the contribution to matrisomal activity from ACh modulatory inputs from the TAN neurons Discrete form of eq. 22 used in this paper is as follows

$$\rho_x^{mat}(t_n) = H_x(t_n)\Big[\sum_{i\in PFC} A_x^i(t_n)\cdot\big[S^i\big]_{t_{n-1}}^{t_n} + Pois(R^{TAN}(t_n))\Big] \quad |x\in\{1,2\} \qquad [23]$$

where $Pois(\lambda)$ generates random value from Poisson distribution with mean λ. **Dopaminergic SNc activity** : DA neurons in SNc, like TANs, are inhibited by GABAergic input from striosomal sub-populations. Their activity can be generally discribed by

$$R^{DA}(t) = Inhibition\Big(\sum_x \rho_x^{pat}(t, \Delta t)\Big|\kappa^{DA}\Big) \tag{24}$$

where, κ^{DA} is the baseline activity of DA neurons in SNc.
The specific discretized form for eq. 24 used in the paper is given by

$$R^{DA}(t_n) = \begin{cases} \kappa^{DA}, & \text{n=1} \quad [25a] \\ \dfrac{\kappa^{DA}}{\Big[\sum_x \rho_x^{pat}(t_{n-1})\Big]}, & \text{n>1} \quad [25b] \end{cases}$$

Action selection and Reward feedback
Due to lateral inhibition between striatal sub-populations, one of them continues to dominate while others are suppressed. Thus, during a single trial with a given stimulus, only one of the striatal sub-populations consistently remains active, which then activates specific downstream circuits corresponding to a specific response. Thus the response action $\mathcal{A}$ is chosen at the response time t_r by

$$\mathcal{A} = action_x \qquad \Big|\rho_x^{mat}(t_r) = max\{\rho_1^{mat}(t_r), \rho_2^{mat}(t_r)...\} \tag{26}$$

The specific form of selected action is

$$\mathcal{A} = \begin{cases} left\ saccade & | \ \rho_1^{mat}(t_r) > 0 \quad [27a] \\ right\ saccade & | \ \rho_2^{mat}(t_r) > 0 \quad [27b] \end{cases}$$

After the response, a feedback $\mathcal{F}$ is provided in the form of reward or punishment depending on the correctness of choice

$$\mathcal{F}(\mathcal{A}) = \begin{cases} reward & | \ \mathcal{A}\text{:correct} \quad [28a] \\ punishment/no\ reward & | \ \mathcal{A}\text{:incorrect} \quad [28b] \end{cases}$$

and for this specific case, reward is determined by

$$\mathcal{F}(\mathcal{A}) = \begin{cases} 1 & | \ \mathcal{A}\text{:correct} \quad [29a] \\ 0 & | \ \mathcal{A}\text{:incorrect} \quad [29b] \end{cases}$$

DA release at cortico-striatal synapses
Cumulative amount of dopamine released at the cortico-striatal synapses by SNc neurons over the course of a single trial, i.e. from stimulus onset until feedback reward, is given by

$$DA_{SNc} = \int_{t_0}^{t_f} R^{DA}(t)\, dt + DA^{fb}\Big(\mathcal{F}(\mathcal{A})\Big) \tag{30}$$

where, t_0 is stimulus onset time, t_f is the time of reward feedback and $DA^{fb}\Big(\mathcal{F}\Big)$ is the amount of DA released due to reward feedback $\mathcal{F}$ which is a function of response action $\mathcal{A}$

Specific discretized form of eq. 30 used for this work is as follows

$$\begin{aligned} DA_{SNc} &= \sum_n \big[R^{DA}(t_n)\big] + DA^{fb}\Big(\mathcal{F}\Big) \\ &\approx \Big[\bar{DA} + \big[R^{DA}(t_2) - \kappa^{DA}\big]\Big] + [DA]^{fb}\Big(\mathcal{F}\Big) \end{aligned} \tag{31}$$

where, $[\bar{DA}]$ is the baseline DA release in the absence of any striosomal inhibition of SNc.
The strongest suppression of DA neurons from baseline activity as a result of striosomal inhibition occurs at time t_2, since for $t > t_2$ the lateral inhibition within STR sub-populations suppresses most of the striosomal activity bringing SNc activity back to baseline. Thus the approximation in eq. 31 is justified. In the specific case, DA released as a result of feedback is given by:

$$DA^{fb}\Big(\mathcal{F}\Big) = \begin{cases} \theta^{rew} & | \ \mathcal{F}(\mathcal{A}) = 1 \quad [32a] \\ 0 & | \ \mathcal{F}(\mathcal{A}) = 0 \quad [32b] \end{cases}$$

Since the total amount of DA released during a trial directly affects the cortico-striatal synaptic plasticity and thus affect learning, it can also be expressed in a traditional reinforcement learning framework in terms of value prediction error:

$$DA_{SNc} = \bar{DA} + \epsilon \tag{33}$$

where, error prediction $\epsilon = \mathcal{R} - \mathcal{Q}$ such that reward $\mathcal{R} = DA^{fb}\Big(\mathcal{F}\Big)$ and expected reward $\mathcal{Q} = \kappa^{DA} - R^{DA}(t_2)$

Cortico-Striatal synaptic plasticity

Synaptic plasticity on striatal medium spiny neurons (MSNs) follows a set of rules corresponding to the literature on plasticity in basal ganglia (142). A large body of experimental literature indicate three factors involved in the LTP/LTD of cortico-striatal synapses (143) : phasic increase in DA release; presynaptic activity; and postsynaptic depolarization. Here in this model we also consider the synaptic decay of cortico-striatal synapses. The general form of cortico-striatal change is described as :

$$\frac{dW^{ij}}{dt} = F^{ion}\Big(S^i, S^j\Big) \cdot F^{DA}\Big(DA_{SNc}\Big) - \Gamma_{dec}(W^{ij}) \qquad [34]$$

where, F^{ion} is the synaptic weight change as a function of pre- and post-synaptic activity of cortical and striatal neurons S^i and S^j respectively, while F^{DA} is the D1/D2 receptor factor as a function of the DA from SNc and Γ_{dec} is the synaptic decay.

The simplified specific form of eq. 34 implemented here is given by

$$\Delta W^i_x = H_x \cdot S^i \cdot F^{DA}\Big(DA_{SNc} - \bar{DA}\Big) - \gamma W^i_x \quad \Big| x \in \{striatal\ blocks\} \qquad [35]$$

where H_x determines the state of activity in striatal block x (see eq. 19b), S^i determines activity in cortical neuron i.

Specific form of F^{DA} is derived from the experimental literature whose qualitative behaviour is summarized in (143). Due to its qualitative similarity with the BCM curve for synaptic plasticity, even though it explains a completely different phenomena, we defined our F^{DA} with the same mathematical form (144) as $F^{DA}(y) = K^{DA} \cdot y \cdot (y + \theta_M) \cdot \sigma'(\alpha(y + \theta_M))$ where $y = DA_{SNc} - \bar{DA} = \epsilon$ (error prediction). Here sigmoidal function $\sigma(x) = \frac{1}{1+e^{-x}}$.

Due to the specific form BCM function, it allows both LTP and LTD depending on depending on the value prediction error as quantified by the DA release (see eq. 33)

Cortico-cortical and thalamo-cortical synaptic plasticity

In the simulation, specific rules of synaptic potentiation and reversal of that potentiation are adopted for synapses in superficial-layer anterior cortical cells, corresponding to consistently replicated studies from the plasticity literature (63, 93, 94). Changes in synaptic weights W from cortical cell i to j are described first in general form, illustrating the proposed set of mechanisms to be studied, and then in highly specific form showing a very limited special case of the general form, currently implemented in simulations.

The general form of synaptic change modeled here in is:

$$\frac{dW^{ij}}{dt} = F^{ion}\Big(S^i, S^j\Big) - F^{rev}\Big(DA_{VTA}\Big) \qquad [36]$$

where F^{ion} and F^{rev} are increments and decrements to synaptic weights; S^i and S^j are number of spikes from cortical source cells i and targets j, D_{VTA} is the quantity of DA released from VTA onto the cortical targets. F^{rev} is a function erasing the increment of all recently incremented synapses.

This general form partially captures a set of physiological phenomena from the literature, directly relating to experimentally validated synaptic change:

- Brief bursts of high-gamma (80-100 Hz) spiking, recurring at intervals of 5 Hz (thus, "theta-burst" or "theta-gamma" patterns) induce very long-lasting synaptic potentiation; the phenomenon has been demonstrated to various extents in hippocampal field CA1, and in posterior and anterior neocortical regions. The mechanisms for this are surprisingly well worked out and extensively replicated (see, e.g., (93–98)).
- Reciprocally, individual spikes that recur at 5 Hz intervals ("theta-spike" patterns) cause very recently-induced LTP to be reversed; this adenosine-dependent phenomenon is linked to interference with consolidation of recent LTP, which otherwise remains fragile for several minutes after induction (27, 145).
- Previous work studied hypothetical expansions of these phenomena to carrier waves at high alpha / low beta rhythms (14-16 Hz rather than 5 Hz); the characteristics of these hypothetical expansions will be examined in further detail in upcoming model systems.
- In a cortico-striatal-thalamic-cortical simulation, ongoing bursting, at 14-16 Hz repetitions (beta-burst), is induced by correct behavioral responses to learned stimuli. As described, this low-beta-burst stimulation gives rise to synaptic potentiation in cortex, strengthening responses to learned cues.
- During early training, when an incorrect behavioral response to a given stimulus was made, the mismatch triggered dopamine responses, both in SNc and in VTA. The SNc response participates in feedback to striatal MSNs, as part of a reinforcement step that caused changes to cortico-striatal synapses. This SNc dopamine response is in direct concordance with the widely-agreed upon and highly cited findings of many studies on reinforcement learning in the striatal complex (74–77).
- Meanwhile, the VTA dopamine response is far less well defined in the literature than the SNc response; the simulation implements a well-cited effect of dopamine in prefrontal cortex, in which dopamine activity suppresses cortical excitatory cell gamma bursting and instead allows cortical single spikes (e.g., (99, 100)), i.e., in this case, spiking recurring at beta.

- Thus, incorrect behavioral responses generate beta-spiking, rather than beta-bursting. These incorrect behaviors thereby cause reversal of the potentiation that otherwise would have taken hold in the cortical responses (27, 145), thus preventing the incorrect encoding from being permanently learned. It is notable that this confluence of ascending systems together with rhythmic bursts and spikes, yields a result that can (surprisingly) be interpreted at the level of specific receptor activity on individual synapses.

Having identified physiological-level systems mechanisms for potentiation and reversal, the effects of the general form can be substantially simplified for initial incorporation into a simulation. In the currently implemented model, these general forms are implemented as special cases, such that the change to a cortical weight from cell i to cell j is:

$$\Delta W^{ij} = K^{ion}\Big[S^i \cdot S^j\Big] \cdot (\bar{W} - W^{ij}) \cdot \mathcal{F}\Big(\mathcal{A}\Big) \qquad [37]$$

where K^{ion} is constant and $\bar{W}$ is the maximum value that a weight can attain, while $\mathcal{F}(\mathcal{A})$ is a feedback function of the response such that

$$\mathcal{F}(\mathcal{A}) = \begin{cases} 1 & \mid \mathcal{A}\text{:correct} \\ 0 & \mid \mathcal{A}\text{:incorrect} \end{cases}$$

thus capturing the effects of potentiation on correct trials and lack of potentiation on incorrect trials. We apply the same plasticity model for thalamo-cortical synapses as well.

Short-term augmentation of thalamo-cortical synapses: mechanism for sustained activity

In addition of long term plasticity in thalamo-cortical synapses explained above, there is an additional short term synaptic augmentation that has been studied in the literature (146, 147). This short-term augmentation of time scales ranging from few milliseconds to few seconds have been known to caused due to several factors e.g. the stimulus frequencies and behavioral states of cortical neurons. These short-term synaptic facilitation/augmentation have also been suggested as a mechanism for working memory through persistent reverberation of spiking activity (106). The various molecular mechanisms underlying various types of short-term synaptic augmentation include presynaptic residual calcium (148) and perturbation in postsynaptic firing (149). In our model this short term thalamo-cortical augmentation serves two purposes : (1) sustain the spiking activity in the circuit even after offset of stimulus (2) hold the pre-offset spatial firing pattern in the anterior cortex ("PFC") intact even during the delay epoch after the offset : the "holding" memory. Thus, post stimulus offset, not only the spiking activity within the ciruit is sustained but also the "memory" of the stimulus encoded in the spatial firing pattern in the PFC is sustained. This requires a mechanism which augments the thalamo-cortical synapses based on the recent firing of post-synaptic cortical neurons such that the neurons firing during stimulus would be more susceptible to the thalamic excitation even after stimulus offset. We implement this by introducing an extra factor in thalamo-cortical synaptic conductance : g^i_{stp}, for a synapse from thalamic neuron j onto the cortical neuron i. The synaptic equation 9 gets modified for thalamo-cortical connections as:

$$I^{ij}_{syn} = \Big[g^i_{stp}(t) \cdot W^{ij}\Big] \cdot g^j_{syn}(t)(V^i - V^j_{syn}) \qquad [39]$$

where g^i_{stp}, as in the case of synaptic conductance g_{syn} (see equations 10), increases due to spiking of neuron i:

$$\begin{aligned} \frac{dz}{dt} &= -\frac{1}{\tau_1} z + g_\infty(V) \\ \frac{dg_{stp}}{dt} &= -\frac{1}{\tau_3} g_{stp} + z \cdot k_{stp}(\bar{G}_{stp} - g_{stp}) \end{aligned} \qquad [40]$$

where τ_3 is the decay constant and $\bar{G}_{stp}$ is the asymptotic value (upper bound) of g_{stp}. The intermediate variable z is the same in equations 10 and 40.

Table 2. Model Equations

General form	Specific form
Striosome (patch) activity	
$\rho^{pat}_x(t, \Delta t) = PAT\Big(\sum_{i \in Ctx} W^i_x(t) \cdot \big[S^i\big]^t_{t-\Delta t}\Big) - LI\Big(\sum_{y \neq x} \rho^{mat}_y\Big)$	$\rho^{pat}_x(t_n) = H_x(t_n) \sum_{i \in Ctx} W^i_x(t_n) \cdot \big[S^i\big]^{t_n}_{t_{n-1}} \quad \lvert x \in \{1, 2\}$
$\rho^{pat}_x(t, \Delta t)$: striosome spiking activity in striatal block x	$H_x(t_n) = \begin{cases} \Theta(\rho^{mat}_x(t_{n-1}) - \rho^{mat}_y(t_{n-1})) & \text{n>2} \\ 1 & n \leq 2 \end{cases}$

LI : Lateral inhibition by other striatal areas	$(x,y) \in \{(1,2),(2,1)\}$
TAN activity $R^{TAN}(t) = Inhibition\Big(\sum_x \rho_x^{pat}(t,\Delta t)\Big\vert \kappa^{TAN}\Big)$	$R^{TAN}(t_n) = \begin{cases} \kappa^{TAN}, & \text{n=1} \\ min\Big(\frac{\kappa^{TAN}}{\lambda^{TAN}\sum_x \rho_x^{pat}(t_{n-1})}, \kappa^{TAN}\Big), & \text{n>1} \end{cases}$ $\kappa^{TAN} = 100,\ \lambda^{TAN} = 1$
Matrisome (matrix) activity $\rho_x^{mat}(t,\Delta t) = MAT\Big(\sum_{i\in Ctx} W_x^i(t)\cdot\big[S^i\big]_{t-\Delta t}^{t}\Big) +$ $TAN(R^{TAN}) - LI\Big(\sum_{y\in N}\rho_y^{mat}(t,\Delta t)\Big) \quad \Big\vert y \neq x$ $\rho_y^{mat}(t,\Delta t)$: matrisome spiking activity in striatal block y	$\rho_x^{mat}(t_n) = H_x(t_n)\Big[\sum_{i\in Ctx} W_x^i(t_n)\cdot\big[S^i\big]_{t_{n-1}}^{t_n} + Pois(R^{TAN}(t_n))\Big]$ $Pois(\lambda)$: Poisson random variable of mean λ
SNc activity $R^{DA}(t) = Inhibition\Big(\sum_x \rho_x^{pat}(t,\Delta t)\Big\vert \kappa^{DA}\Big)$	$R^{DA}(t_n) = \begin{cases} \kappa^{DA}, & \text{n=1} \\ min\Big(\frac{\kappa^{DA}}{\lambda^{DA}\sum_x \rho_x^{pat}(t_{n-1})}, \kappa^{DA}\Big), & \text{n>1} \end{cases}$ $\kappa^{DA} = 1,\ \lambda^{DA} = 0.33$
Action selection and reward feedback $\mathcal{A} = action_i \quad \Big\vert \rho_i^{mat}(t_r) = max\{\rho_1^{mat}(t_r), \rho_2^{mat}(t_r)...\}$ $\mathcal{F}(\mathcal{A}) = \begin{cases} reward & \vert\ \mathcal{A}\text{:correct} \\ punishment/no\ reward & \vert\ \mathcal{A}\text{:incorrect} \end{cases}$	$\mathcal{A} = \begin{cases} left\ saccade & \vert\ \rho_1^{mat}(t_r) > 0 \\ right\ saccade & \vert\ \rho_2^{mat}(t_r) > 0 \end{cases}$ $\mathcal{F}(\mathcal{A}) = \begin{cases} 1 & \vert\ \mathcal{A}\text{:correct} \\ 0 & \vert\ \mathcal{A}\text{:incorrect} \end{cases}$
DA release at cortico-striatal synapses $DA_{SNc} = \int_{t_0}^{t_f} R^{DA}(t)\,dt + DA^{fb}$ $DA^{fb} : DA^{fb}\Big(\mathcal{F}(\mathcal{A})\Big)$	$DA_{SNc} = \sum_n \big[R^{DA}(t_n)\big] + DA^{fb}$ $\approx \Big[\bar{DA} + \big[R^{DA}(t_2) - \kappa^{DA}\big]\Big] + DA^{fb}$ $DA^{fb} = \begin{cases} \theta^{rew} & \vert\ \mathcal{F}(\mathcal{A}) = 1 \\ 0 & \vert\ \mathcal{F}(\mathcal{A}) = 0 \end{cases}$ DA_{SNc} **as a measure of 'Prediction Error'** $DA_{SNc} = \bar{DA} + \epsilon \qquad \vert \epsilon : prediction\ error$ $\epsilon = \mathcal{R} - \mathcal{Q} \qquad \vert \mathcal{R} : reward,\ \mathcal{Q} : expectation$ $\mathcal{R} = [DA]^{fb}, \quad \mathcal{Q} = \kappa^{DA} - R^{DA}(t_2)$
Cortico-Striatal synaptic plasticity $\frac{dW^{ij}}{dt} = F^{ion}\Big(S^i, S^j\Big)\cdot F^{DA}\Big(DA_{SNc}\Big) - \Gamma_{dec}(W^{ij})$ F^{ion}: ionotropic receptor factor, S^i: activity of neuron i F^{DA}: DA receptor factor Γ^{dec}: synaptic decay	$\Delta W_x^i = H_x\cdot S^i\cdot F^{DA}\Big(DA_{SNc} - \bar{DA}\Big) - \gamma W_x^i \quad \Big\vert x \in \{str\ blocks\}$ $F^{DA}(y) = K^{DA}\cdot y\cdot(y+\theta_M)\cdot\sigma'(\alpha(y+\theta_M))$ $y = DA_{SNc} - \bar{DA} = \epsilon, \quad \sigma(x) = \frac{1}{1+e^{-x}}$ $K^{DA} = 0.008,\ \alpha = 2.5,\ \gamma = 0.01,\ \theta_M = 1.0$
Cortico-cortical & thalamo-cortical plasticity	

$$\frac{dW^{ij}}{dt} = F^{ion}\Big(S^i, S^j\Big) - F^{rev}\Big(DA_{VTA}\Big)$$ F^{rev}: LTP reversal as function of DA from VTA	$$\Delta W^{ij} = K^{ion}\Big[S^i \cdot S^j\Big] \cdot (\bar{W} - W^{ij}) \cdot \mathcal{F}\Big(\mathcal{A}\Big)$$ cortico-cortical : $K^{ion} = 5 \times 10^{-5}$, $\bar{W} = 5$ thalamo-cortical : $K^{ion} = 10^{-5}$, $\bar{W} = 6$
Thalamo-cortical short-term synaptic augmentation Synaptic Current $$I^{ij}_{syn} = \left[g^i_{stp}(t) \cdot W^{ij}\right] \cdot g^j_{syn}(t)(V^i - V^j_{syn})$$	Conductance Dynamics $$\frac{dg_{stp}}{dt} = -\frac{1}{\tau_3} g_{stp} + z \cdot k_{stp}(\bar{G}_{stp} - g_{stp})$$ $$\frac{dz}{dt} = -\frac{1}{\tau_1} z + g_\infty(V)$$ $\tau_3 = 2000ms$, $\tau_1 = 0.1ms$, $k_{stp} = 0.5$, $\bar{G}_{stp} = 0.5$

B. Computational Methods. One of the practical difficulties in multi-scale computational neuroscience has been the limitation due to the tension between model flexibility and simulation speed. For this reason, we chose to code our mathematical models in Julia, which shows *orders of magnitude speed improvements in ODE systems compared to Python and MATLAB*(150). Our neural circuit are modeled using the Julia ModelingToolkit (MTK) component-based block-modeling system (151). MTK generates code using a symbolic preprocessing system that allows for automating transformations such as index-reduction of differential-algebraic equations via Pantelides algorithm (152, 153), alias elimination, tearing of nonlinear systems, symbolic simplification (154, 155), and more to produce differential equation and nonlinear systems which have improved numerical stability over naive approaches. In addition, MTK uses sophisticated algorithms such as the D* algorithm (156) to generate system Jacobians in asymptotically less time than traditional techniques based on symbolic, automatic (157), and numerical differentiation. Code generation improvements are coupled with advancements in the numerical solving of differential equations provided by the Julia-based DifferentialEquations.jl methods (158). In particular, we created a MTK library of modular computational dynamical building blocks that consists of hierarchical composites of indiviual Hodgkin-Huxley neurons (cortical blocks), next-generation neural mass models, and various reinforcement computational primitives that can be assembled using a directed graph structure. Once all the connections are defined, we transform the directed graph into a set of symbolic Ordinary Differential Equations (several thousand ODEs) that are then symbolically simplified and compiled into numerically optimized functions to interface with the Julia DifferentialEquations.jl ODE solver packages. We chose to use the Vern7 solver (Verner's "Most efficient" 7/6 Runge-Kutta method /citeverner2010) to optimize both speed and accuracy of the solution. For our model system with 600 Hodgkin-Huxley neurons, simulation of spiking activity over a single trial duration of 1600 ms, takes only 50-60 seconds.

Our computational models will be made available as stand-alone implementation (github.com/Neuroblox/corticostriatal-circuit-notebook) and also as a model designed in Neuroblox.jl (Neuroblox.org), a comprehensive computational neuroscience tool to tackle large multi-scale brain simulations using Julia. (Within this system, ModelingToolkit.jl automatically transforms equations into simpler forms for numerical solving, and the system can be used interactively via a graphical user interface.)

C. Experimental Task Design.
The category learning task and the corresponding neurophysiological recording with the macaques, on which we specifically model our simulations, is described in detail in references (26) and (7). The brief overview of the task is as follows. A single category learning session is consists of 600-800 trials where in each trial, initially the animal has to fixate on a central target onscreen. While fixating, a randomly chosen category exemplar appears on the screen for 600 ms. Each exemplar is a constellation of 7 dots derived from either of the two category prototypes (see SI, (7)). The exemplar epoch (stimulus on) is followed by 1000 ms of delay epoch (stimulus off) in which the animal again fixates on the center and waits for the response cue, after which two saccade targets appear on the left and right of center, each corresponding to one of the categories. The animal responds with a direct saccade to one of the targets indicating the correct category of the exemplar, thus receives a reward (juice drops) on correct response and no reward for incorrect response. The exemplar set for each category is gradually increased starting from 1 for each category, and is doubled in size after each bock of trials. Each block of trials is over once the accuracy of the animal reaches atleast 80% over last 20 trials of that block, after which th next block begins with doubled set of exemplars. Based on this demarcation, first two blocks (100-150 trials) are considered early learning trials (SR phase in (7)) while trials from block 6 onward are considered late learning stage (Category Performance phase in (7)). For every session, a new set of category prototype patterns are used.

D. Simulation Protocol.
We simulated the category learning trial sessions as performed in the NHPs on our cortico-striatal model. The set of random dot pattern stimuli as used in the experiments were also was used in the simulation as the input. Since the visual area in this model does not represent the primary visual area, but some higher area in the hierarchy of visual processing, we added a spatial filter to the dot patterns that approximate the gaussian smoothening as follows. The images are generated in gray-scale with total (15 X 15 =) 225 pixels, where each pixel has brightness value within $[0, 1]$ (0 :black, 1 : white). For each pixel (i, j) corresponding to the dot location, the brightness value is 1. For the four adjacent pixels $(i-1, j), (i+1, j), (i, j-1), (i, j+1)$,

brightness value is 0.6. For all other pixels, brightness is 0. As mentioned in the model description, the pixels map one-to-one on the simulated visual area cortical units, where each unit (neuron) receives a steady input current proportional to the brightness value of corresponding pixel.

In accordance with the experiments, each trial was simulated for 1600 msec, in which the stimulus onset is between 0-600 msec, while during 600-1600 msec there is no visual input. At the end of 1600 msec the circuit model output determines the action choice, and based on the outcome a reward feedback is instantaneously fed into the model (see Model Description above). The synaptic weights are discretely updated after the reward feedback depending on the spatio-temporal spiking patterns within the circuit and the level of dopamine release from SNc and VTA respectively during the trial interval. Succeedingly the next trial cycle begins with the onset of the next visual stimulus. Right from the first trial onwards the stimuli are chosen from either of the two categories randomly out of the entire set, as opposed to gradual doubling of the stimulus set in the experiments (see (7)).

E. Analysis of Model Physiological Response.

Field Estimate.

We obtain a very simple neural mass activity from the spiking activities of various blocks in the model simulating specific brain regions, viz. Local Spike Summation (LSS), which is simply the spike histogram.

$$LSS(t, \Delta t) = \sum_{j=1}^{N} \sum_{m \in \mathcal{Z}} \int_{t-\frac{\Delta t}{2}}^{t+\frac{\Delta t}{2}} \delta(t - T_j^m) dt \qquad [48]$$

where, T_j^m is the mth spike-time for j^{th} neuron while neuron j belongs to a specific block (brain region such as PFC, STR etc.). For all PLV analyses of LSS, we used time bin $\Delta t = 1ms$. For rastor plot representations in Figures 5 and 6 we used $\Delta t = 20ms$. We also calculate Average Local Membrane Potential (ALMP) as an alternate surrogate for local field: $ALMP(t) = \frac{1}{N} \sum_i V^i(t)$ where N is the neuronal population size and $V^i(t)$ is membrane potential of neuron i.

Phase Locking Synchronisation.

Wavelet analysis of the Local Spike Summation (LSS) and subsequent calculation of synchrony measure Phase-Locking Value (PLV) was performed using the same methods as was followed in the measured LFPs in reference (7). First wavelet transform of each LSS timeseries was carried by convolution of LSS with Morlet wavelet (159) at 5 octaves between 2-64 Hz at frequency resolution of 0.1 octave. For this analysis, we used MATLAB based software available at `http://atoc.colorado.edu/research/wavelets/`. Wavelet transform of a single LSS gives phase time seriese $\psi(t)$ for every frequency f. Frequency spefific PLV between two LSS timeseries is given by $PLV = \left| \frac{1}{n} \sum_{t1}^{tn} e^{i(\psi_1(t) - \psi_2(t))} \right|$. Two perfectly phase synchronised time series would measure a PLV as 1 while completely asynchronous time series would give PLV as 0. As done in the neurophysiological analysis in (7), we calculate the PLV over the last 500 ms of exemplar (stimulus on) and delay (stimulus off) epochs, viz. between 100 to 600 ms and between 1100 to 1600 ms respectively. In order to compare the PLV values of simulated LSS with that of the empirical LFPs in reference (7), we subtracted from the observed PLVs the mean PLV obtained from randomised surrogate ensemble. Surrogate ensemble was obtained by randomising the trial orders of the LSS time series from PFC ans STR and we took 1000 realisations of these randomisations.

Differential synchrony behaviour between two sets of trials (correct and incorrect) is measured by discrimination index d' as also used in (7). Discrimination index is the difference between mean PLVs of two sets of trials normalised by their combined standard deviation : $d' = \frac{\langle PLV_{incorr} \rangle - \langle PLV_{corr} \rangle}{s_p}$, where $s_p = \sqrt{\frac{s_{corr}^2 \cdot n_{corr} + s_{incorr}^2 \cdot n_{incorr}}{n_{corr} + n_{incorr} - 2}}$. We also remove the sampling bias by converting the observed discrimination d' into Z-scores with respect to another randomised surrogate ensemble. This ensemble is obtained by randomly shuffling the trials between two sets and then calculating d'. We used 400 realisations of randomised surrogates for d'.

Normalised single unit activity

In order to determine the condition specific preference of single units (neurons) based on their average spiking behavior (as shown in Figure 7), we first sort the late learning trials i.e. trials 400-700 into four groups or conditions (C_1, C_2, C_3, C_4, say) based on stimulus category (A vs B) and outcome (correct vs incorrect) viz. A-correct, A-incorrect, B-correct and B-incorrect. For each unit (represented by a single row in Figure 7), we calculate the average number of spikes occurring in that unit during a specific time interval (0-200 msec or 0-1600 msec) for each of the four trial conditions. The unit is considered to prefer the trial condition that shows maximum value for average spikes. For sorting the units based on their prefered trial conditions, we only need the relative spiking activity for each trial condition with respect to the maximum spiking activity. Therefore we normalise the condition specific spiking activity by dividing the average spiking activity in each condition by the maximum average spiking activity over all conditions. Thus, normalised spiking activity N_c^i for neuron i for trial condition $C \in C_1, C_2, C_3, C_4$ is given by:

$$N_C^i = \frac{S_C^i}{max(S_{C_1}^i, S_{C_2}^i, S_{C_3}^i, S_{C_4}^i)} \qquad [49]$$

where S_C^i is the average number of spikes in neuron i over all trials of condition C.

Competing Interests.

The authors declare the following competing interests: authors RG, EKM, LRMP, and HHS are co-Founders of Neuroblox Inc., a company spun out of SUNYSB, MIT, and Dartmouth to develop a commercial-grade software platform for multi-scale computational neuroscience with applications to diagnosis and treatment of brain-based disorders.

Accessing Code/Tools.

All methods shown here, which include an easy-to-use GUI as well as documentation and tutorials, are made freely available through the Neuroblox computational neuroscience platform (neuroblox.org). Neuroblox can be accessed through the JuliaHub repository (juliahub.com).

Author Contributions.

Conceptualization: AP, HHS, LRMP, EKM and RG; Data curation: SLB, EGA and EKM; Formal analysis: AP and RG; Funding acquisition: HHS, LRMP, EKM and RG; Investigation: AP and RG; Methodology: AP and RG; Project administration: LRMP; Software: HO, HHS and SS; Supervision: HHS, LRMP, EKM and RG; Validation: AP; Visualization: AP, LRMP and RG; Writing - original draft: AP and RG; Writing - review & editing: AP, SLB, HHS, LRMP, EKM and RG.

ACKNOWLEDGMENTS. We thank Annabel Driussi for her valuable support with graphics. We thank Anthony Chesebro for useful discussions. The research presented here was funded by the Baszucki Brain Research Fund, United States (LRMP). This work was supported in part by the Office of Naval Research, United States (RG).

References

1. Y LeCun, Y Bengio, G Hinton, Deep learning. *Nature* **521**, 436–444 (2015).
2. A Doerig, et al., The neuroconnectionist research programme. *Nat. Rev. Neurosci.* **24**, 431–450 (2023).
3. S Bhattacharya, SL Brincat, M Lundqvist, EK Miller, Traveling waves in the prefrontal cortex during working memory. *PLOS Comput. Biol.* **18**, e1009827 (2022).
4. ZW Davis, et al., Spontaneous traveling waves naturally emerge from horizontal fiber time delays and travel through locally asynchronous-irregular states. *Nat. Commun.* **12**, 6057 (2021).
5. EP Stephen, et al., Broadband slow-wave modulation in posterior and anterior cortex tracks distinct states of propofol-induced unconsciousness. *Sci. Reports* **10**, 13701 (2020).
6. EV Lubenov, AG Siapas, Hippocampal theta oscillations are travelling waves. *Nature* **459**, 534–539 (2009).
7. EG Antzoulatos, EK Miller, Increases in functional connectivity between prefrontal Ccrtex and striatum during category learning. *Neuron* **83**, 216–225 (2014).
8. S Coombes, Á Byrne, Next generation neural mass models in *Nonlinear dynamics in computational neuroscience*, eds. F Corinto, A Torcini. (Springer International Publishing, Cham), pp. 1–16 (2019).
9. Á Byrne, RD O'Dea, M Forrester, J Ross, S Coombes, Next-generation neural mass and field modeling. *J. Neurophysiol.* **123**, 726–742 (2020).
10. R Granger, Brain circuit implementation: High-precision computation from low-precision components in *Toward replacement parts for the brain*, eds. T Berger, DL Glanzman. (The MIT Press), pp. 277–294 (2005).
11. R Granger, Engines of the brain: The computational instruction set of human cognition. *AI Mag.* **27**, 15–32 (2006).
12. JB Ding, JN Guzman, JD Peterson, JA Goldberg, DJ Surmeier, Thalamic gating of corticostriatal signaling by cholinergic interneurons. *Neuron* **67**, 294–307 (2010).
13. AC Felch, RH Granger, The hypergeometric connectivity hypothesis: Divergent performance of brain circuits with different synaptic connectivity distributions. *Brain Res.* **1202**, 3–13 (2008).
14. V Braitenberg, A Schüz, *Cortex: statistics and geometry of neuronal connectivity.* (Springer, Berlin, Heidelberg), (1998).
15. M Abeles, *Corticonics: Neural circuits of the cerebral cortex.* (Cambridge University Press, Cambridge), (1991).
16. D Fitzpatrick, JS Lund, DE Schmechel, AC Towles, Distribution of GABAergic neurons and axon terminals in the macaque striate cortex. *The J. Comp. Neurol.* **264**, 73–91 (1987).
17. SH Hendry, HD Schwark, EG Jones, J Yan, Numbers and proportions of GABA-immunoreactive neurons in different areas of monkey cerebral cortex. *J. Neurosci.* **7**, 1503–1519 (1987).
18. R Coultrip, R Granger, G Lynch, A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks* **5**, 47–54 (1992).
19. ZH Mao, SG Massaquoi, Dynamics of winner-take-all competition in recurrent neural networks With lateral inhibition. *IEEE Transactions on Neural Networks* **18**, 55–69 (2007).
20. S Kaski, T Kohonen, Winner-take-all networks for physiological models of competitive learning. *Neural Networks* **7**, 973–984 (1994).
21. SI Amari, A mathematical approach to neural systems in *Systems Neuroscience*, ed. J Metzler. (Academic Press), pp. 67–117 (1977).
22. DJ Pinto, GB Ermentrout, Spatially structured activity in synaptically coupled neuronal networks: II. Lateral inhibition and standing pulses. *SIAM J. on Appl. Math.* **62**, 226–243 (2001).
23. E Oja, Neural networks, principal components, and subspaces. *Int. J. Neural Syst.* **01**, 61–68 (1989).
24. S Grossberg, Competitive learning: From interactive activation to adaptive resonance. *Cogn. Sci.* **11**, 23–63 (1987).
25. C von der Malsburg, Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* **14**, 85–100 (1973).
26. EG Antzoulatos, EK Miller, Differences between neural activity in prefrontal cortex and striatum during learning of novel abstract categories. *Neuron* **71**, 243–249 (2011).
27. EA Kramár, B Lin, CS Rex, CM Gall, G Lynch, Integrin-driven actin polymerization consolidates long-term potentiation. *Proc. Natl. Acad. Sci.* **103**, 5579–5584 (2006).
28. J Szentágothai, The 'module-concept' in cerebral cortex architecture. *Brain Res.* **95**, 475–496 (1975).
29. DH Hubel, TN Wiesel, S LeVay, HB Barlow, RM Gaze, Plasticity of ocular dominance columns in monkey striate cortex. *Philos. Transactions Royal Soc. London. B, Biol. Sci.* **278**, 377–409 (1997).
30. OD Creutzfeldt, HC Nothdurft, Representation of complex visual stimuli in the brain. *Naturwissenschaften* **65**, 307–318 (1978).
31. VB Mountcastle, Brain mechanisms for directed attention. *J. Royal Soc. Medicine* **71**, 14–28 (1978).
32. A Keller, EL White, Triads: a synaptic network component in the cerebral cortex. *Brain Res.* **496**, 105–112 (1989).
33. RAW Galuske, W Schlote, H Bratzke, W Singer, Interhemispheric asymmetries of the modular structure in human temporal cortex. *Science* **289**, 1946–1949 (2000).
34. MS Gazzaniga, Regional differences in cortical organization. *Science* **289**, 1887–1888 (2000).
35. MA Castro-alamancos, BW Connors, Thalamocortical synapses. *Prog. Neurobiol.* **51**, 581–606 (1997).
36. EG Jones, Viewpoint: the core and matrix of thalamic organization. *Neuroscience* **85**, 331–345 (1998).
37. AJ Heynen, MF Bear, Long-term potentiation of thalamocortical transmission in the adult visual cortex in vivo. *J. Neurosci.* **21**, 9801–9813 (2001).
38. G Silberberg, A Gupta, H Markram, Stereotypy in neocortical microcircuits. *Trends Neurosci.* **25**, 227–230 (2002).
39. F Valverde, Structure of the cerebral cortex. Intrinsic organization and comparative analysis of the neocortex. *Revista de neurologia* **34**, 758–780 (2002).
40. EL White, A Peters, Cortical modules in the posteromedial barrel subfield (Sml) of the mouse. *J. Comp. Neurol.* **334**, 86–96 (1993).
41. A Peters, BR Payne, J Budd, A numerical analysis of the geniculocortical input to striate cortex in the monkey. *Cereb. Cortex* **4**, 215–229 (1994).
42. M Molinari, et al., Auditory thalamocortical pathways defined in monkeys by calcium-binding protein immunoreactivity. *J. Comp. Neurol.* **362**, 171–194 (1995).
43. H Killackey, FF Ebner, Two different types of thalamocortical projections to a single cortical area in mammals. *Brain Behav. Evol.* **6**, 156–169 (1972).
44. M Herkenham, New perspectives on the organization and evolution of nonspecific thalamocortical projections in *Sensory-Motor Areas and Aspects of Cortical Connectivity*, Cerebral Cortex, eds. EG Jones, A Peters. (Springer US, Boston, MA), pp. 403–445 (1986).
45. HJ Groenewegen, HW Berendse, JG Wolters, AHM Lohman, The anatomical relationship of the prefrontal cortex with the striatopallidal system, the thalamus and the amygdala: evidence for a parallel organization in *Progress in Brain Research*, The Prefrontal Its Structure, Function and Cortex Pathology, eds. HBM Uylings, CG Van Eden, JPC De Bruin, MA Corner, MGP Feenstra. (Elsevier) Vol. 85, pp. 95–118 (1991).
46. J DeFelipe, EG Jones, Parvalbumin immunoreactivity reveals layer IV of monkey cerebral cortex as a mosaic of microzones of thalamic afferent terminations. *Brain Res.* **562**, 39–47 (1991).
47. T Van Groen, JM Wyss, Projections from the anterodorsal and anteroveniral nucleus of the thalamus to the limbic cortex in the rat. *J. Comp. Neurol.* **358**, 584–604 (1995).
48. TF Freund, KaC Martin, D Whitteridge, Innervation of cat visual areas 17 and 18 by physiologically identified X- and Y- type thalamic afferents. I. Arborization patterns and quantitative distribution of postsynaptic elements. *J. Comp. Neurol.* **242**, 263–274 (1985).
49. H Barbas, N Rempel-Clower, Cortical structure predicts the pattern of corticocortical connections. *Cereb. Cortex* **7**, 635–646 (1997).
50. KS Rockland, Non-uniformity of extrinsic connections and columnar organization. *J. Neurocytol.* **31**, 247–253 (2002).

51. HA Swadlow, AG Gusev, T Bezdudnaya, Activation of a cortical column by a thalamocortical impulse. *J. Neurosci.* **22**, 7766–7773 (2002).
52. DB Bender, Visual activation of neurons in the primate pulvinar depends on cortex but not colliculus. *Brain Res.* **279**, 258–261 (1983).
53. ME Diamond, M Armstrong-James, MJ Budway, FF Ebner, Somatic sensory responses in the rostral sector of the posterior group (POm) and in the ventral posterior medial nucleus (VPM) of the rat thalamus: Dependence on the barrel field cortex. *J. Comp. Neurol.* **319**, 66–84 (1992).
54. ME Diamond, M Armstrong-James, FF Ebner, Somatic sensory responses in the rostral sector of the posterior group (POm) and in the ventral posterior medial nucleus (VPM) of the rat thalamus. *J. Comp. Neurol.* **318**, 462–476 (1992).
55. JJ Chrobak, G Buzsáki, Gamma oscillations in the entorhinal cortex of the freely behaving rat. *J. Neurosci.* **18**, 388–398 (1998).
56. M Sarter, JP Bruno, The neglected constituent of the basal forebrain corticopetal projection system: GABAergic projections. *Eur. J. Neurosci.* **15**, 1867–1873 (2002).
57. P Fries, JH Reynolds, AE Rorie, R Desimone, Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**, 1560–1563 (2001).
58. A Rozov, N Burnashev, B Sakmann, E Neher, Transmitter release modulation by intracellular Ca2+ buffers in facilitating and depressing nerve terminals of pyramidal cells in layer 2/3 of the rat neocortex indicates a target cell-specific difference in presynaptic calcium dynamics. *The J. Physiol.* **531**, 807–826 (2001).
59. A Knoblauch, G Palm, Scene segmentation by spike synchronization in reciprocally connected visual areas. I. Local effects of cortical feedback. *Biol. Cybern.* **87**, 151–167 (2002).
60. B Pesaran, JS Pezaris, M Sahani, PP Mitra, RA Andersen, Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat. Neurosci.* **5**, 805–811 (2002).
61. A Rodriguez, J Whitson, R Granger, Derivation and Analysis of Basic Computational Operations of Thalamocortical Circuits. *J. Cogn. Neurosci.* **16**, 856–877 (2004).
62. J Ambros-Ingerson, R Granger, G Lynch, Simulation of paleocortex performs hierarchical clustering. *Science* **247**, 1344–1348 (1990).
63. R Granger, J Whitson, J Larson, G Lynch, Non-Hebbian properties of long-term potentiation enable high-capacity encoding of temporal sequences. *Proc. Natl. Acad. Sci.* **91**, 10104–10108 (1994).
64. A Bocincova, TJ Buschman, MG Stokes, SG Manohar, Neural signature of flexible coding in prefrontal cortex. *Proc. Natl. Acad. Sci.* **119**, e2200400119 (2022).
65. JJ Canales, et al., Shifts in striatal responsivity evoked by chronic stimulation of dopamine and glutamate systems. *Brain* **125**, 2353–2363 (2002).
66. GE Alexander, MR DeLong, PL Strick, Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu. Rev. Neurosci.* **9**, 357–381 (1986).
67. RL Reep, JV Corwin, Topographic organization of the striatal and thalamic connections of rat medial agranular cortex. *Brain Res.* **841**, 43–52 (1999).
68. J Feingold, DJ Gibson, B DePasquale, AM Graybiel, Bursts of beta oscillation differentiate postperformance activity in the striatum and motor cortex of monkeys performing movement tasks. *Proc. Natl. Acad. Sci.* **112**, 13687–13692 (2015).
69. M Lundqvist, et al., Gamma and Beta Bursts Underlie Working Memory. *Neuron* **90**, 152–164 (2016).
70. M Lundqvist, et al., Working memory control dynamics follow principles of spatial computing. *Nat. Commun.* **14**, 1429 (2023).
71. A Reiner, N Hart, W Lei, Y Deng, Corticostriatal projection neurons âĂŞ dichotomous types and dichotomous functions. *Front. Neuroanat.* **4** (2010).
72. AM Graybiel, T Aosaki, AW Flaherty, M Kimura, The basal ganglia and adaptive motor control. *Science* **265**, 1826–1831 (1994).
73. BJ Hunnicutt, et al., A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife* **5**, e19103 (2016).
74. W Schultz, P Dayan, PR Montague, A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
75. W Schultz, Predictive reward signal of dopamine neurons. *J. Neurophysiol.* **80**, 1–27 (1998).
76. W Schultz, Updating dopamine reward signals. *Curr. Opin. Neurobiol.* **23**, 229–238 (2013).
77. W Schultz, Dopamine reward prediction-error signalling: a two-component response. *Nat. Rev. Neurosci.* **17**, 183–195 (2016).
78. BE Jones, Arousal systems. *Front. Biosci.* **8**, 438–451 (2003).
79. RE Brown, JT McKenna, Turning a negative into a positive: Ascending gabaergic control of cortical activation and arousal. *Front. Neurol.* **6**, 135 (2015).
80. TF Freund, V Meskenaite, Gamma-Aminobutyric acid-containing basal forebrain neurons innervate inhibitory interneurons in the neocortex. *Proc. Natl. Acad. Sci.* **89**, 738–742 (1992).
81. EM Vazey, DE Moorman, G Aston-Jones, Phasic locus coeruleus activity regulates cortical encoding of salience information. *Proc. Natl. Acad. Sci.* **115**, E9439–E9448 (2018).
82. BE Kilavik, M Zaepffel, A Brovelli, WA MacKay, A Riehle, The ups and downs of β oscillations in sensorimotor cortex. *Exp. Neurol.* **245**, 15–26 (2013).
83. MA Sherman, et al., Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice. *Proc. Natl. Acad. Sci.* **113**, E4885–E4894 (2016).
84. Y Kawaguchi, T Shindou, Noradrenergic excitation and inhibition of GABAergic cell types in rat frontal cortex. *J. Neurosci.* **18**, 6963–6976 (1998).
85. H Salgado, et al., Layer-specific noradrenergic modulation of inhibition in cortical layer II/III. *Cereb. Cortex* **21**, 212–221 (2011).
86. AI Gulyás, N Hájos, TF Freund, Interneurons containing calretinin are specialized to control other interneurons in the rat hippocampus. *J. Neurosci.* **16**, 3397–3411 (1996).
87. JM Blasco-Ibáñez, FJ Martinez-Guijarro, TF Freund, Enkephalin-containing interneurons are specialized to innervate other interneurons in the hippocampal CA1 region of the rat and guinea-pig. *Eur. J. Neurosci.* **10**, 1784–1795 (1998).
88. AI Gulyás, L Acsády, TF Freund, Structural basis of the cholinergic and serotonergic modulation of GABAergic neurons in the hippocampus. *Neurochem. Int.* **34**, 359–372 (1999).
89. TW Kim, et al., Selective localization of amyloid precursor-like protein 1 in the cerebral cortex postsynaptic density. *Mol. Brain Res.* **32**, 36–44 (1995).
90. JJ Zhu, BW Connors, Intrinsic firing patterns and whisker-evoked synaptic responses of neurons in the rat barrel cortex. *J. Neurophysiol.* **81**, 1171–1183 (1999).
91. P Reinagel, D Godwin, SM Sherman, C Koch, Encoding of visual information by lgn bursts. *J. Neurophysiol.* **81**, 2558–2569 (1999).
92. NN Odean, M Sanayei, MN Shadlen, Transient oscillations of neural firing rate associated with routing of evidence in a perceptual decision. *J. Neurosci.* **43**, 6369–6383 (2023).
93. J Larson, P Xiao, G Lynch, Reversal of LTP by theta frequency stimulation. *Brain Res.* **600**, 97–102 (1993).
94. J Larson, E Munkácsy, Theta-burst LTP. *Brain Res.* **1621**, 38–50 (2015).
95. DD Mott, DV Lewis, Facilitation of the induction of long-term potentiation by GABAB receptors. *Science* **252**, 1718–1720 (1991).
96. DD Mott, DV Lewis, GABAB receptors mediate disinhibition and facilitate long-term potentiation in the dentate gyrus. *Epilepsy research Suppl.* **7**, 119–134 (1992).
97. T Nakashiba, et al., Young dentate granule cells mediate pattern separation, whereas old granule cells facilitate pattern completion. *Cell* **149**, 188–201 (2012).
98. A Terashima, et al., An essential role for PICK1 in NMDA receptor-dependent bidirectional synaptic plasticity. *Neuron* **57**, 872–882 (2008).
99. N Gorelova, JK Seamans, CR Yang, Mechanisms of dopamine activation of fast-spiking interneurons that exert inhibition in rat prefrontal cortex. *J. Neurophysiol.* **88**, 3150–3166 (2002).
100. P Zhong, Z Yan, Distinct physiological effects of dopamine D4 receptors on prefrontal cortical pyramidal neurons and fast-spiking interneurons. *Cereb. Cortex* **26**, 180–191 (2016).
101. A Kirkwood, SM Dudek, JT Gold, CD Aizenman, MF Bear, Common forms of synaptic plasticity in the hippocampus and neocortex in vitro. *Science* **260**, 1518–1521 (1993).
102. MF Bear, A Kirkwood, Neocortical long-term potentiation. *Curr. Opin. Neurobiol.* **3**, 197–202 (1993).
103. MA Castro-Alamancos, JP Donoghue, BW Connors, Different forms of synaptic plasticity in somatosensory and motor areas of the neocortex. *J. Neurosci.* **15**, 5324–5333 (1995).
104. DV Buonomano, MM Merzenich, Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* **21**, 149–186 (1998).
105. K Seki, M Kudoh, K Shibuki, Sequence dependence of post-tetanic potentiation after sequential heterosynaptic stimulation in the rat auditory cortex. *The J. Physiol.* **533**, 503–518 (2001).
106. G Mongillo, O Barak, M Tsodyks, Synaptic theory of working memory. *Science* **319**, 1543–1546 (2008).
107. L Kozachkov, et al., Robust and brain-like working memory through short-term synaptic plasticity. *PLOS Comput. Biol.* **18**, e1010776 (2022).
108. T Aosaki, et al., Responses of tonically active neurons in the primate's striatum undergo systematic changes during behavioral sensorimotor conditioning. *J. Neurosci.* **14**, 3969–3984 (1994).
109. T Aosaki, M Kimura, AM Graybiel, Temporal and spatial characteristics of tonically active neurons of the primate's striatum. *J. Neurophysiol.* **73**, 1234–1252 (1995).
110. A Raz, A Feingold, V Zelanskaya, E Vaadia, H Bergman, Neuronal synchronization of tonically active neurons in the striatum of normal and parkinsonian primates. *J. Neurophysiol.* **76**, 2083–2088 (1996).
111. H Inokawa, N Matsumoto, M Kimura, H Yamada, Tonically active neurons in the monkey dorsal striatum signal outcome feedback during trial-and-error search behavior. *Neuroscience* **446**, 271–284 (2020).
112. Y Shimo, O Hikosaka, Role of tonically active neurons in primate caudate in reward-oriented saccadic eye movement. *The J. Neurosci.* **21**, 7804–7814 (2001).
113. YF Zhang, SJ Cragg, Pauses in Striatal Cholinergic Interneurons: What is Revealed by Their Common Themes and Variations? *Front. Syst. Neurosci.* **11** (2017).
114. AC Martel, et al., Targeted transgene expression in cholinergic interneurons in the monkey striatum using canine adenovirus serotype 2 vectors. *Front. Mol. Neurosci.* **13** (2020).
115. YD Van der Werf, MP Witter, HJ Groenewegen, The intralaminar and midline nuclei of the thalamus. Anatomical and functional evidence for participation in processes of arousal and awareness. *Brain Res. Rev.* **39**, 107–140 (2002).
116. G Mandelbaum, et al., Distinct cortical-thalamic-striatal circuits through the parafascicular nucleus. *Neuron* **102**, 636–652.e7 (2019).
117. BD Bennett, CJ Wilson, Spontaneous activity of neostriatal cholinergic interneurons in vitro. *J. Neurosci.* **19**, 5586–5596 (1999).
118. O Hikosaka, M Sakamoto, S Usui, Functional properties of monkey caudate neurons. I. Activities related to saccadic eye movements. *J. Neurophysiol.* **61**, 780–798 (1989).
119. P Apicella, Leading tonically active neurons of the striatum from reward detection to context recognition. *Trends Neurosci.* **30**, 299–306 (2007).
120. N Abudukeyoumu, T Hernandez-Flores, M Garcia-Munoz, GW Arbuthnott, Cholinergic modulation of striatal microcircuits. *Eur. J. Neurosci.* **49**, 604–622 (2019).
121. AL Hodgkin, AF Huxley, A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* **117**, 500–544 (1952).
122. XJ Wang, Ionic basis for intrinsic 40 Hz neuronal oscillations:. *NeuroReport* **5**, 221–224 (1993).
123. XJ Wang, Pacemaker neurons for the theta rhythm and their synchronization in the septohippocampal reciprocal loop. *J. Neurophysiol.* **87**, 889–900 (2002).
124. C Koch, I Segev, TJ Sejnowski, TA Poggio, *Methods in neuronal modeling: From ions to networks*, Computational Neuroscience Series. (Bradford), 2 edition, (1998).
125. SR Wicks, CJ Roehrig, CH Rankin, A dynamic network simulation of the nematode tap withdrawal circuit: Predictions concerning synaptic function using behavioral criteria. *J. Neurosci.* **16**,

4017–4031 (1996).
126. A Singh, SM Papa, Striatal oscillations in parkinsonian non-human primates. *Neuroscience* **449**, 116–122 (2020).
127. M Deffains, et al., Subthalamic, not striatal, activity correlates with basal ganglia downstream activity in normal and parkinsonian monkeys. *eLife* **5**, e16443 (2016).
128. GB Ermentrout, N Kopell, Parabolic bursting in an excitable system coupled with a slow oscillation. *SIAM J. on Appl. Math.* **46**, 233–253 (1986).
129. H Alitto, Y Dan, Cell-type-specific modulation of neocortical activity by basal forebrain input. *Front. Syst. Neurosci.* **6** (2013).
130. M Jayachandran, et al., Nucleus reuniens transiently synchronizes memory networks at beta frequencies. *Nat. Commun.* **14**, 4326 (2023).
131. A Bibbig, et al., Beta rhythms ($15 - 20$ Hz) generated by nonreciprocal communication in hippocampus. *J. Neurophysiol.* **97**, 2812–2823 (2007).
132. LR Silva, Y Amitai, BW Connors, Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Science* **251**, 432–435 (1991).
133. Y Tachibana, H Iwamuro, H Kita, M Takada, A Nambu, Subthalamo-pallidal interactions underlying parkinsonian neuronal oscillations in the primate basal ganglia. *Eur. J. Neurosci.* **34**, 1470–1484 (2011).
134. MM McCarthy, et al., Striatal origin of the pathologic beta oscillations in Parkinson's disease. *Proc. Natl. Acad. Sci.* **108**, 11620–11625 (2011).
135. V Breton-Provencher, M Sur, Active control of arousal by a locus coeruleus GABAergic circuit. *Nat. Neurosci.* **22**, 218–228 (2019).
136. JY Hansen, et al., Integrating brainstem and cortical functional architectures. *bioRxiv* (2023).
137. C Liu, et al., An action potential initiation mechanism in distal axons for the control of dopamine release. *Science* **375**, 1378–1385 (2022).
138. GMG Shepherd, Corticostriatal connectivity and its role in disease. *Nat. Rev. Neurosci.* **14**, 278–291 (2013).
139. S Gee, et al., Synaptic activity unmasks dopamine D2 receptor modulation of a specific class of layer V pyramidal neurons in prefrontal cortex. *J. Neurosci.* **32**, 4959–4971 (2012).
140. HJ Seong, AG Carter, D1 receptor modulation of action potential firing in a subpopulation of layer 5 pyramidal neurons in the prefrontal cortex. *J. Neurosci.* **32**, 10516–10521 (2012).
141. CR Gerfen, MN Economo, J Chandrashekar, Long distance projections of cortical pyramidal neurons. *J. Neurosci. Res.* **96**, 1467–1475 (2018).
142. AC Kreitzer, RC Malenka, Striatal plasticity and basal ganglia circuit function. *Neuron* **60**, 543–554 (2008).
143. JNJ Reynolds, JR Wickens, Dopamine-dependent plasticity of corticostriatal synapses. *Neural Networks* **15**, 507–521 (2002).
144. N Intrator, LN Cooper, Objective function formulation of the BCM theory of visual cortical plasticity: Statistical connections, stability conditions. *Neural Networks* **5**, 3–17 (1992).
145. G Lynch, CS Rex, CM Gall, LTP consolidation: Substrates, explanatory power, and functional significance. *Neuropharmacology* **52**, 12–23 (2007).
146. MA Castro-Alamancos, BW Connors, Cellular mechanisms of the augmenting response: short-term plasticity in a thalamocortical pathway. *J. Neurosci.* **16**, 7742–7756 (1996).
147. MA Castro-Alamancos, BW Connors, Short-term plasticity of a thalamocortical pathway dynamically modulated by behavioral state. *Science* **272**, 274–277 (1996).
148. H Kamiya, RS Zucker, Residual Ca2 + and short-term synaptic plasticity. *Nature* **371**, 603–606 (1994).
149. K Ibata, Q Sun, GG Turrigiano, Rapid synaptic scaling induced by changes in postsynaptic firing. *Neuron* **57**, 819–826 (2008).
150. C Rackauckas, ODE Solver Multi-Language Wrapper Package Work-Precision Benchmarks (MATLAB, SciPy, Julia, deSolve (R)) (2023).
151. Y Ma, et al., ModelingToolkit: A Composable Graph Transformation System for Equation-Based Modeling. *arXiv* **2103.05244** (2021).
152. C Pantelides, The Consistent Initialization of Differential-Algebraic Systems. *SIAM J. Scl. Stat. Comput.* **9**, 213–231 (1988).
153. K Shimako, M Koga, An Extension of Pantelides Algorithm for Consistent Initialization of Differential-Algebraic Equations Using Minimally Singularity. *2020 59th Annu. Conf. Soc. Instrum. Control. Eng. Jpn. (SICE)* pp. 1676–1681 (2020).
154. M Otter, H Elmqvist, Transformation of Differential Algebraic Array Equations to Index One Form in *Transformation of Differential Algebraic Array Equations to Index One Form.* pp. 565–579 (2017).
155. H Elmqvist, T Henningsson, M Otter, Innovations for Future Modelica in *Innovations for Future Modelica.* pp. 693–702 (2017).
156. B Guenter, The D* Symbolic Differentiation Algorithm (2007).
157. AG Baydin, BA Pearlmutter, AA Radul, JM Siskind, Automatic Differentiation in Machine Learning: a Survey. *arXiv:1502.05767* (2018).
158. C Rackauckas, N Qing, DifferentialEquations.jl - A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia. *J. Open Res. Softw.* **5** (2017).
159. C Torrence, GP Compo, A practical guide to wavelet analysis. *Bull. Am. Meteorol. Soc.* **79**, 61–78 (1998).

SUPPLEMENTARY MATERIAL

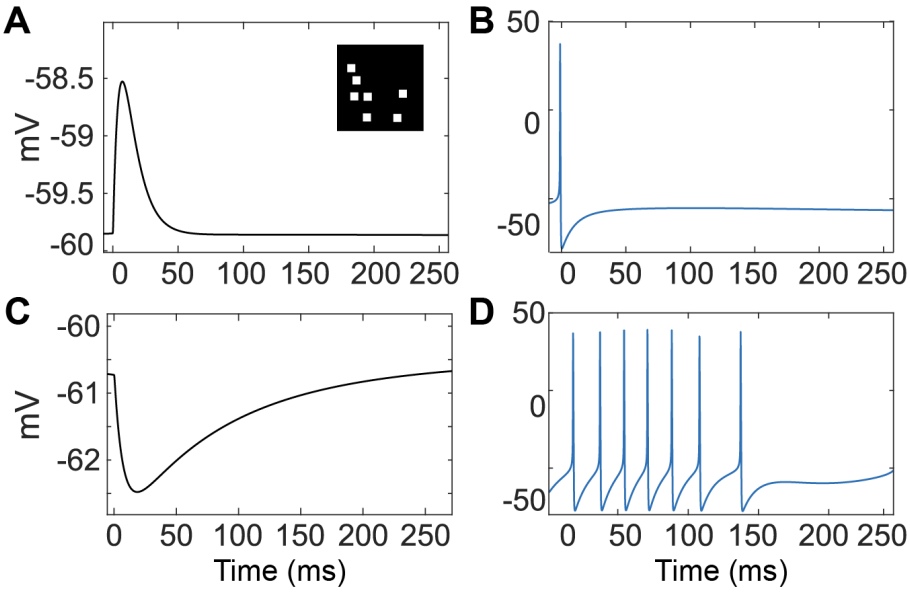

Fig. S1. a) Sample EPSP (inset: Sample visual stimulus dot pattern. (b) Sample action potential. (c) Sample IPSP. (d) Sample spiking response to visual stimulus

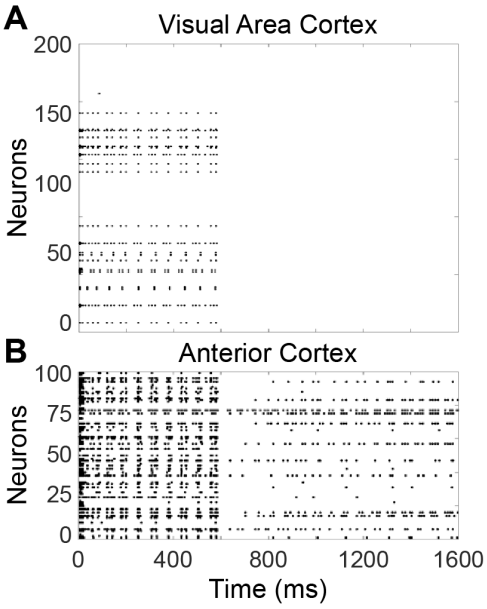

Fig. S2. Raster plot showing spiking activity during a single category learning trial within a simulated (A) visual area cortex and (B) Anterior Cortex ("PFC"). Stimulus is off-set at 600 ms, therefore no activity is seen in the model visual area cortex after t=600ms.

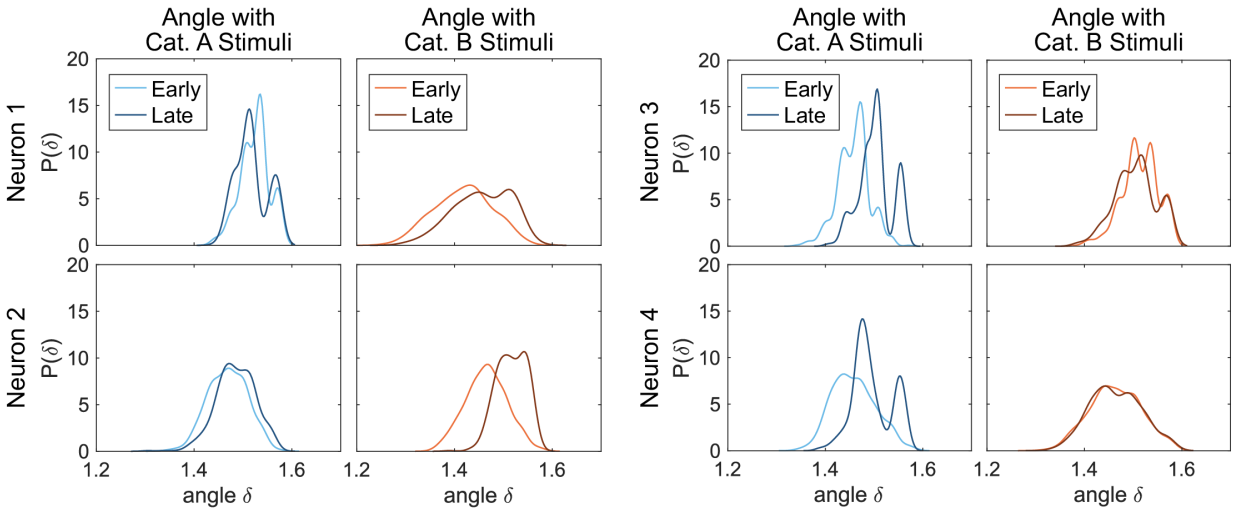

Fig. S3. Learning rotates the dendritic vectors of model anterior cortical (PFC) neurons closer to one category and away from the other category. Each panel shows histogram of projection angles between dendritic vector of single neuron and all visual stimulus vectors of a single category (see main text Figure 2). (A) shows two representative neurons whose dendritic vectors rotate towards category A stimuli and away from category B stimuli after the learning session (blue curves in left two panels shift to smaller values while red curves in the right two panels shift to higher values). (B) shows two representative neurons whose dendritic vectors behave the other way i.e. rotate away from category A and towards category B.

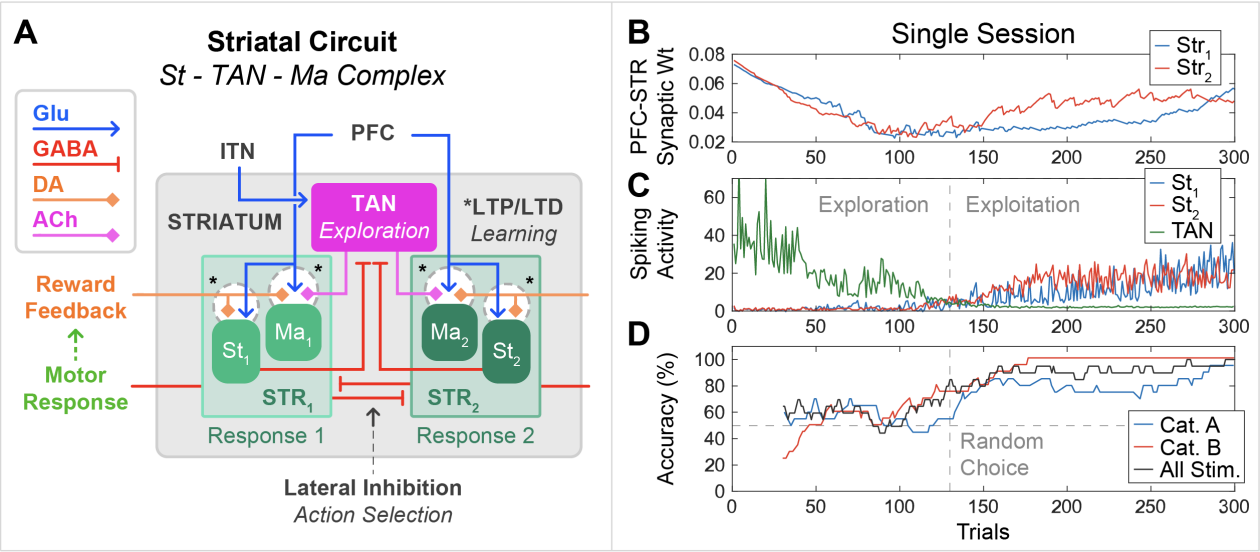

Fig. S4. Role of TAN activity in action selection strategy: exploration to exploitation. (A) shows the circuit diagram of striatal complex consisting of matrisome, striosome and TAN. Cortical inputs received by Ma and St. St inhibit TAN while TAN modulate Ma. High TAN activity introduces stochasticity in action selection (exploration). Whereas low TAN activity implies action selection is dominated by cortico-striatal synaptic weights, thus the strategy is more deterministic (exploitation) (B) Cortico-striatal synaptic weight to the two striatal blocks (corresponding to two response choice). They first weaken during early low-performance stage, while strengthens during the late learning stages. The weights do not increase symmetrically for both striatal blocks. (C) Spiking activity (no. of spikes during first 200 ms of stimulus on) in TAN and two STR blocks. When learning stage transitions, TAN activity goes down and STR activities dominate. Exploratio to exploitation clearly visible in these activities. (D) Performance overall (black) and category wise (red and blue.) Performance is also asymmetric in categories, just like synaptic weights in (B)

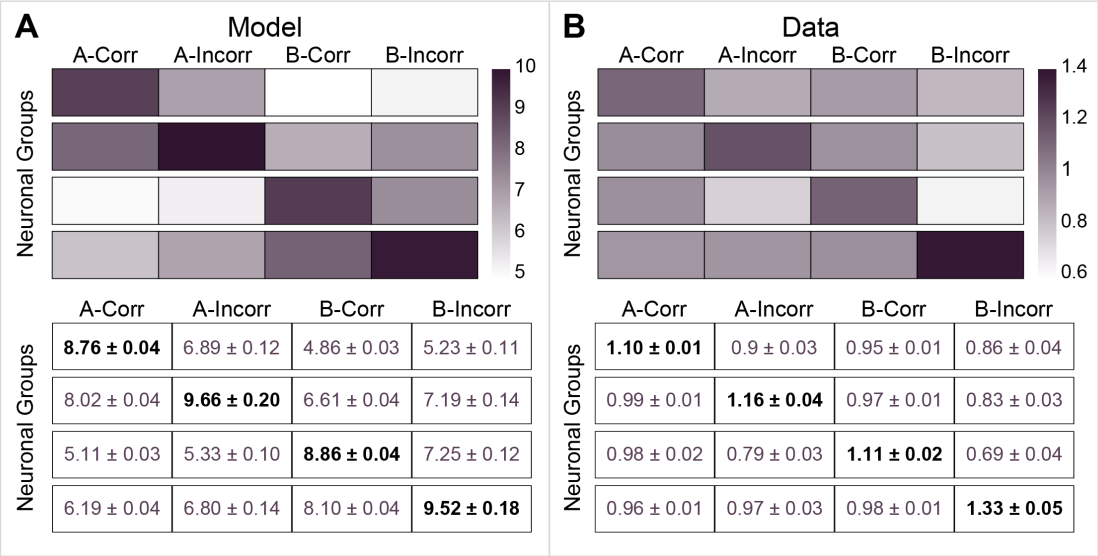

Neuronal Groups	A-Corr	A-Incorr	B-Corr	B-Incorr
	8.76 ± 0.04	6.89 ± 0.12	4.86 ± 0.03	5.23 ± 0.11
	8.02 ± 0.04	**9.66 ± 0.20**	6.61 ± 0.04	7.19 ± 0.14
	5.11 ± 0.03	5.33 ± 0.10	**8.86 ± 0.04**	7.25 ± 0.12
	6.19 ± 0.04	6.80 ± 0.14	8.10 ± 0.04	**9.52 ± 0.18**

Neuronal Groups	A-Corr	A-Incorr	B-Corr	B-Incorr
	1.10 ± 0.01	0.9 ± 0.03	0.95 ± 0.01	0.86 ± 0.04
	0.99 ± 0.01	**1.16 ± 0.04**	0.97 ± 0.01	0.83 ± 0.03
	0.98 ± 0.02	0.79 ± 0.03	**1.11 ± 0.02**	0.69 ± 0.04
	0.96 ± 0.01	0.97 ± 0.03	0.98 ± 0.01	**1.33 ± 0.05**

Fig. S5. Mean spikes in PFC neurons between 0-200 ms : all sessions Average spiking activity of four cortical neuron groups as defined by their preferential responsiveness to four types of trials based on category and outcomes for all simulated sessions (A) and for all recorded experimental sessions (B). Row represents a given type of neurons and columns represent their response to a given trial type : A correct, A incorrect, B correct and B incorrect. The tables below the heat maps show the actual numerical values of the average spiking activity and their SEM, in response to the trial during first 200 msec of stimulus onset. Notice the low SEM values indicating the robustness of reliability of the preferential responses.