SUPPLEMENTARY MATERIAL

Supplementary Discussion

The model has neuron interaction patterns similar to the brain. During the experiment, Simple Visual Cortex Region (SVCR) cells are differentially activated depending on the topographic location of dots in the input stimulus. These neurons will spike (or burst), as a function of voltage summation thresholds in the model (see Methods for corresponding synaptic equations and Figure S1 for an illustration). SVCR cells are only active for 600 ms during presentation of the simulated stimulus (Figure S2A).

The SVCR inputs are combined with the rhythmic inputs from the simulated ascending (ASC1) brainstem next-gen neural mass model (NG-NMM), so that the predominant response of the SVCR thus exhibits a beta-gamma temporal pattern: brief (40 ms) bursts of high-frequency spiking, separated by approximately 70 ms; i.e., gamma bursts carried on beta rhythms. The activity occurs only on 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 (ANTC). These continue to spike due to inputs from ASC1 in a cortico-striatal-thalamo-cortical feedback loop (Figure S2B).

The detailed structure of our model is described in Figure 1 and the Methods. It includes multiple subcircuits whose behavior leads to computational processes. An example of the Striatal Subcircuit and how Tonically Active Neurons (TANs) lead to exploration/explotation is shown in Figure S3.

Calculating and Visualizing Metrics on Multiple Scales. The complex time-variable behavior of our model can be summarized with a few statistical metrics that allow for near-direct comparison to NHP empirical data from (1). Local Spike Summation (LSS) is a simple histogram of the number of spikes across a field of target model neurons (see Methods). The Average Local Membrane Potential is defined as $ALMP(t) = \frac{1}{N}\Sigma_i V^i(t)$, i.e., the average of the membrane potentials (V^i) of each of N neurons. While extremely simple, these can be compared to the Local Field Potential (LFP) measured empirically from NHPs (1). These three measures (simulated LSS, simulated ALMP, and empirical LFP) are shown in Figure S4 for both the pre-frontal cortex and the striatum for a single trial. These are represented directly as well as with spectrograms to show the dominant frequencies of activity.

Inspection of Supplemental Figure S4 shows that simulated model spiking and membrane potentials have many similarities to empirical measurements. In the simulated and empirical anterior cortical neurons, spikes are modulated by the rhythmic model ascending input oscillating at a mean rate of 16Hz. At the stimulus offset (600 msec after onset), the response in anterior neurons greatly lessens, but does not go to zero activity, due to cortico-striatal-thalamo-cortical feedback.

In simulated and empirical anterior cortical neurons, there is an initial strong response (150-200 ms) and then a 100 ms silent period. This behavior was first observed in the model and then subsequently confirmed in the empirical data. (We show that the behavior during this first 200 ms also very strongly correlates with final action selection in both the simulation and empirical data.) Further study of this simulation activity is detailed in Supplemental Figure S5

A more detailed investigation of this behavior in the model (shown in Figure S5) provided a specific biological mechanism for the post-stimulus silence: the strong initial onset response activates a powerful lateral inhibitory response, building up transient inhibition over the first 50 ms; this in turn strongly inhibits further excitation in the model (for 100 ms) by hyperpolarizing the excitatory neurons (see Figure S5 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 ms. Figure S5 demonstrates that this explains the behavior of the simulation. It is hypothesized that these mechanisms may also directly explain the corresponding post-stimulus-onset silence in the empirical data, which could be confirmed with future investigation. This hypothesis is the first example of how our brain model shows properties similar to the brain that are so specific, that it can provide novel explanations and predictions.

Additional Figures illustrating properties of incongruent neurons. Visual stimuli based on two categories can become increasingly categorized by changing synaptic weights as illustrated in Figure S6. A quantitative calculation of these changes is the "projection angles" between dendritic vectors in the 225-dimensional input space. Learning consists of reweighting that leads to rotation of these vectors, shown in Figure S7.

The synaptic plasticity that leads to changing weights also highly aligns individual neurons to specific actions. Figure 4A shows how the spiking activity of individual neurons allows them to be assigned into four categories corresponding to the upcoming choice and outcome: A-correct, A-incorrect, B-correct, and B-incorrect. Figure S8 expands on this categorization by showing how the categorization is related to activity at 200 ms and 1600 ms for the simulated and empirical data in both the cortex and striatum. Figure S9 shows the statistics of spiking frequency averaged over neurons in these categories.

Mechanistic explanation for emergence of Incongruent neurons. We then provide a preliminary analysis regarding the emergence of ICNs, viz., even prior to learning, due to the initial synaptic connectivity of ICNs, their firing responses show relatively low discrimination between the two categories as compared to the congruent neurons. Synaptic connectivity of the neurons determines their receptive fields. This can be demonstrated directly in the simulation where precise synaptic connectivity is known and indirectly from the empirically recorded activity from the real neurons in NHPs. In the simulation, a single excitatory anterior cortical neuron can only "see" 18 out of 225 pixels due to the sparse connections to the simulated visual cortex neurons. It is thus possible for category A and B dot patterns (which are random variations around a template) to appear to overlap. (see Figures S6 and S7). Thus the distance between categories, within the subspace of inputs available to a

particular neuron, can be quite small, leading to incongruent outputs even though the ensemble of cells together has a high accuracy rate. The limited ability of neurons to "see" is shown in detail in Figure S10 A and B. Examples of the receptive fields of the four categories of simulated neurons are shown while it is also shown quantitatively that "congruent" neurons overlap with more dots and with higher A vs. B selectivity than incongruent neurons. Another indirect estimate of receptive fields of neurons is given by quantitative measure of preference of a neuron for selecting A or B, using the discrimination index defined by (1) and discussed in the Methods. Figure S10 C and D show that "correct" neurons have more biased discrimination indices at the beginning of trials in both the simulated and empirical observations. Thus, neurons with a visual receptive field that starts out biased towards a particular category will evolve to prefer that category highly; incongruent neurons are more likely to start as already ambiguous.

Evolution of synaptic weights with time. The physiological and behavioral changes in the model as a result of learning process, is primarily supported by the synaptic weight changes between the cell assemblies. The synaptic plasticity in the form of long term potentiation (and depression) occur in the model in cortico-cortical connections between visual cortex and the anterior cortex, corticostriatal conections between anterior cortex and striatum, and thalamo-cortical connections between thalamus and anterior cortex (see Table 2 in the main text). These changes in the synaptic weights during a single learning session are demonstrated in Figure S11. A spatial pattern emerges in the synaptic weight distributions as a result of successive stimuli of visual patterns (Figure S11 A, C and E). While cortico-cortical and thalamo-cortical weights show a relatively gradual and monotonous evolution (Figure S11 B and D), cortico-striatal synapses show rapid change in increasing as well as decreasing directions (Figure S11F) and shows a more direct correspondence with the behavioural changes (also illustrated in Figure 2B in the main text).

Influence of individual circuit elements on learning: 'lesion' simulations. Since the primary feature of the model is the emergent behaviour from the complex interactions between multiple circuit elements, it is natural to try to characterize the extent of influence that each constituent part of the whole holds in the over all learning behavior demonstrated in this paper. This question however is not very straight forward to answer since the model consists of multiple non-linear interactions and feed-back loops. A systematic deconstruction of this particular model may require a whole different study and thus might be outside the scope of this work which primarily intends to demonstrate the explanatory and predictive power of our modeling approach through an example of category learning paradigm. However, we still demonstrate few lesion simulations where we disable some very central components of cortico-striatal circuit and examine whether and how the learning behaviour breaks down. The elements chosen here for the lesion are those which have subtle and non-trivial functional roles in the circuit viz. SNc, Striosome and TAN neurons.

SNc and Striosome are indispensable circuit elements for learning. Figure S12 shows examples of single learning sessions run with disabling SNc (Figure S12 middle column) and Striosome (Figure S12 right column) respectively, and compares it with the learning session run with the complete circuit (Figure S12 left column). Disabling SNc stops the dopamine that is required for the reward driven LTP/LTD in corticostriatal synapses, which leads corticostriatal synapses to decay gradually (Figure S12 K). This causes TAN activity to dominate (Figure S12 H) and make the model come up with random choices (Figure S12 H). This leads beakdown of learning behavior. Similarly, removing striosomes cause the lack of TAN and SNc suppression as the learning progresses. This makes the response choices largely random (driven by TAN, see Figure S12 I) and makes the corticostriatal synaptic weights highly unstable (see Figure S12 L) due to excessive release of DA. This eventually causes the response to lock into a single choice (Figure S12 F). Eventually the learning behavior breaks. Thus, both SNc and striosome cause the break down of learning behaviour when removed, but through completely distinct mechanisms.

TAN neurons make the model more robust. The influence of TAN neurons is not as straight forward as that of SNc or striosome. Infact the parameter set that was used to generate the main results (Table 2 of Main text) has a very stable learning behavior even without the TAN activity (see Figure S13 A-D). Though in the absence of TAN neurons, there is no clear transition between exploration and exploitation. There might be two reasons for the learning remaining somewhat intact even without TAN. One is that the learning paradigm is relatively simple, with just two categories and clear separation between the two category dot patterns. Second reason is the specific shape of the D1/D2 receptor factor F^{DA} that was used in modeling the corticostriatal synaptic weight change (See Equations 28-29 in the Main Text). This receptor factor determines the direction and magnitude of synaptic weight change as a function of the DA release during a trial. As described in the Methods of the Main Text, the shape of the receptor factor is phenomenologically determined (2) and resembles the BCM curve (3, 4) for synaptic plasticity. Positive shift from the baseline DA release shows positive error prediction ϵ leading to weight increase while negative shift from baseline DA release represents negative prediction error leading to weight reduction (see Figure S3 H). The relative magnitudes of the positive and negative weight changes is currently determined through parameter tuning, and thus there can be multiple choices for that parameters, as shown in Figure S3 H. In our current parameter setting that generated the main results (referred as "punishment biased condition" in Figure S3 and S13), the effect of negative prediction error considerably outweighs the effect of positive prediction error (see punishment biased condition in Figure S3 H). We can alter this by tuning a single parameter α. As described in the Methods in Main Text and also in Table 2 of Main Text, receptor factor function is $F^{DA}(y) = K^{DA} \cdot y \cdot (y + \theta_M) \cdot \sigma'(\alpha(y + \theta_M))$ where $y = DA_{SNc} - \bar{DA} = \epsilon$ (prediction error). Here the sigmoidal function is defined as $\sigma(x) = \frac{1}{1+e^{-x}}$. The original value of scaling factor α used in the model is $\alpha = 2.5$. We altered it to $\alpha = 1.2$ to get qualitatively distinct shape of receptor factor in which the effect of positive error outweighs that

of negative error (see "reward biased condition" in Figure S3 H). For reward biased condition, we also increase the baseline activity of TAN neurons κ^{TAN} (see Equation 14 in Main Text) from 100 to 1200. Both these changes in the reward biased condition do not alter the main results that are in accordance with the data. But they make a crucial change in the learning behavior when the TAN neurons are disabled. Under reward biased conditions, learning behavior is destabilized in the absence of TAN, as shown in Figure S13 E-H, while remains intact under the punishment biased conditions even in the absence of TAN. In the absence of TAN under reward biased condition, even though the model initially learns the categories and responds correctly up until initial few hundred trials, it almost always breaks down due to unstable synaptic weights which tend to strengthen only one striatal block over other (Figure S13 G) leading to response locking to a single choice irrespective of the stimulus (Figure S13 F). This suggests that the TAN and SNc systems may be acting with some redundancy, such that slight disruption TANs may be counterbalanced by SNc, i.e., SNc enables learning even when TAN is perturbed.

Even though this effect needs a more systematic deeper analysis which is outside the scope of this study, in Figure S14 we demonstrate through multiple realizations of reward biased condition that TAN disabled condition leads to unstable learning whereas TAN neurons keep the learning stable even if both the conditions (with and without TAN) are subjected to the same stimulus set. Therefore, without TAN neurons, the model shows stable learning under a limited parameter range of corticostriatal learning (only under punishment based condition), while with the TAN neurons, the model is stable over a much extended parameter range (under both, punishment as well as reward biased conditions). This suggests that the primary influence of TAN neurons in the model is making the learning robust and far less vulnerable to perturbations.

Supplemental Methods

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 S1). These consist of attributes that are highly cited in the literature.

Neuronal Dynamics. Each neuron is modeled using ionic conductance based point neuron models, with the Hodgkin-Huxley formalism (5). An overview of this behavior is shown in the Methods with specific details discussed here. 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. (6) : $\bar{g}_{Na} = 52$, $\bar{g}_K = 20$, $V_{Na} = +55mV$, $V_K = -90mV$, $V_L = -60mV$ and $\phi = 5$.

Different threshold values for used for pyramidal and inhibitory cells. For excitatory pyramidal cells we used the following (6):

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

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

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

Similarly, for inhibitory neurons we used the threshold values implemented in ref. (7):

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

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

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

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

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

where W^{ij} is the synaptic weight from neuron j to i, a quantity which estimates density of dendritic spines and is modified by Long Term Potentiation (LTP) and Long Term Depression (LTD). V_{syn}^{j} is synaptic reversal potential for presynaptic neuron, which for excitatory neuron is $0mV$ and for inhibitory neuron is $-70mV$. The synaptic conductance g_{syn}^{j} is a function of spiking activity of the presynaptic neuron. As suggested in reference (8), synaptic conductance g_{syn} is one of the state variables for each neuron whose dynamics is described by damped oscillator equation:

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

which is driven by impulse force $g_\infty(V)$ defined as follows (9)

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

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 two 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.

Ascending Systems and Intralaminar Thalamic Nuclei. The simulated ascending system (ASC1) is an extreme simplification of ascending systems such as the cholinergic basal forebrain and noradrenergic locus coeruleus; it produces rhythmic activity via a neural mass model. That activity becomes simulated modulatory input to cortex, driving a beta rhythm in the cortex (10). The ASC1 simulation thus functions as a central pattern generator. It has been modelled using a simple 'Next Generation' neural mass model (11, 12) which is a mean field approximation of a network of rhythmic neurons (13). 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

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

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

Similarly, Intralaminar Thalamic Nuclei (ITN) were modeled as a neural mass which targets the TAN neurons (14) 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 (15) either to pyramidal cells or interneurons (10, 16–22). The resulting rhythmicity, in multiple EEG frequency bands, shows evidence of both intrinsic and extrinsic influences (23–26). 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) (27–29). The potential applicability of the present work to Parkinson's and other clinical conditions is of great interest.

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) (11) 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 (17, 18). By contrast, the locus coeruleus (LC) 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 (30). 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 (31). 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.

The basic rhythmic input is thus modeled at present by the next-generation neural mass models (NG-NMM). Notably, all other rhythmic effects in the present work: corticostriatal 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.

Corticostriatal decision making and synaptic plasticity.

The spatio-temporal spiking patterns in the above described cortico-striatal-thalamic-cortico (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 corticostriatal computation at a much coarseer scale of length and time 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 the striatum whereas the latter project only ipsilaterally to the striatum and other structures (32). 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 (33–35).

In principle, all of the following mechanisms can be entirely modelled using only the physiological phenomena at the neuronal length and time scales. 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, Main Text. Following is the detailed description the model components.

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; this 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., (36–41)).
- 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 (42, 43).
- 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 the SNc and in the VTA. The SNc response participates in feedback to striatal MSNs, as part of a reinforcement step that causes changes to corticostriatal 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 (44–48)..
- 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., (49, 50)), 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 (42, 43), 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:

$$\Delta W^{ij} = K^{ion}\Big[S^i \cdot S^j\Big] \cdot (\bar{W} - W^{ij}) \cdot \mathcal{F}\Big(\mathcal{A}\Big) \qquad [17]$$

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

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

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.

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(51).

Our neural circuits are modeled using the Julia ModelingToolkit (MTK) component-based block-modeling system (**?**). MTK generates code using a symbolic preprocessing system that allows for automating transformations such as index-reduction of differential-algebraic equations via Pantelides algorithm (52, 53), alias elimination, tearing of nonlinear systems, symbolic simplification (54, 55), 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 (56) to generate system Jacobians in asymptotically less time than traditional techniques based on symbolic, automatic (57), 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 (58). 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 (59)) 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 are openly available, as stand-alone implementation (github.com/Neuroblox/corticostriatal-circuit-notebook) and also as a model designed in Neuroblox.jl (Neuroblox.org). (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.)

Calculating Phase Locking Synchronization (PLV). 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 (1). First, the wavelet transform of each LSS timeseries was computed by convolution of LSS with Morlet wavelet (60) 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/`. The Wavelet transform of a single LSS gives a phase time series $\psi(t)$ for every frequency f. Frequency specific 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 timeseries would measure a PLV as 1 while completely asynchronous time series would give PLV as 0. As done in the neurophysiological analysis in (1), 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 (1), we subtracted from the observed PLVs the mean PLV obtained from randomised surrogate ensemble. The 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 the discrimination index d' as also used in (1). The 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 $s_p = \sqrt{\frac{s^2_{corr}\cdot n_{corr} + s^2_{incorr}\cdot n_{incorr}}{n_{corr}+n_{incorr}-2}}$. 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'.

Supplementary Tables

Biomimetic Computational Primitives	Computational Operations	Regions	Key Publications
Local-Feedback Lateral Inhibiton Circuit (L-FLIC)	Lateral Inhibition	Anterior, Posterior Neocortex	(61–102)
Matrisomal MSN (MMSN)	Action Selection	Caudate, Putamen	(63, 103–113)
Striosomal MSN (SMSN)	Reward Prediction	Caudate, Putamen	(44–47, 63, 103–113)
Ascending Type 1 (ASC1)	Central Pattern Generation	Pons, Brainstem	(10, 16–22, 114–120)
Superficial Glu−GABA−DA Plasticity (GLTP)	Long-Term Potentiation, Reversal	Cortex, VTA	(36–42, 49, 50, 73, 75, 121–127)
Cholinergic Tonically Active Neurons (CTAN)	Exploration/Exploitation	Caudate, Putamen	(103, 104, 128–140)

Table S1. Current Blocks for Corticostriatal Model.

Supplementary Figures

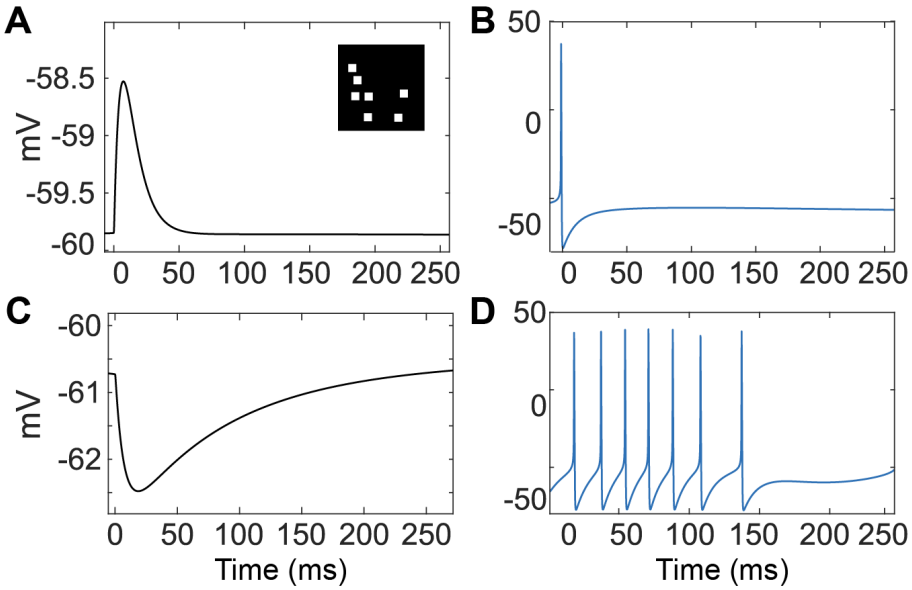

Fig. S1. Examples of synaptic potentials calculated for excitatory and inhibitory neurons. (A) Sample Excitatory Post-Synaptic Potential (EPSP); inset: Sample visual stimulus dot pattern. (B) Sample action potential. (C) Sample Inhibitory Post-Synaptic Potential (IPSP). (D) Sample spiking response to visual stimulus

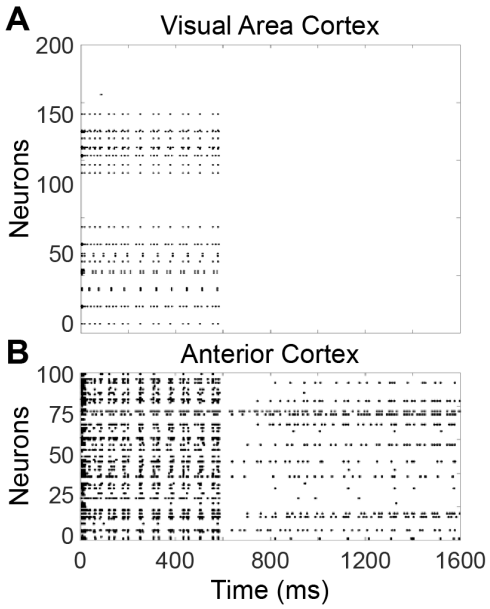

Fig. S2. Raster plot showing spiking activity of each simulated excitatory neuron during a single category learning trial within (A) the simple visual cortex region (SVCR) and (B) Anterior Cortex (ANTC) region. The visual stimulus is removed at 600 milliseconds (ms), but a cortico-striatal-thalamo-cortico feedback loop maintains activity in the PFC.

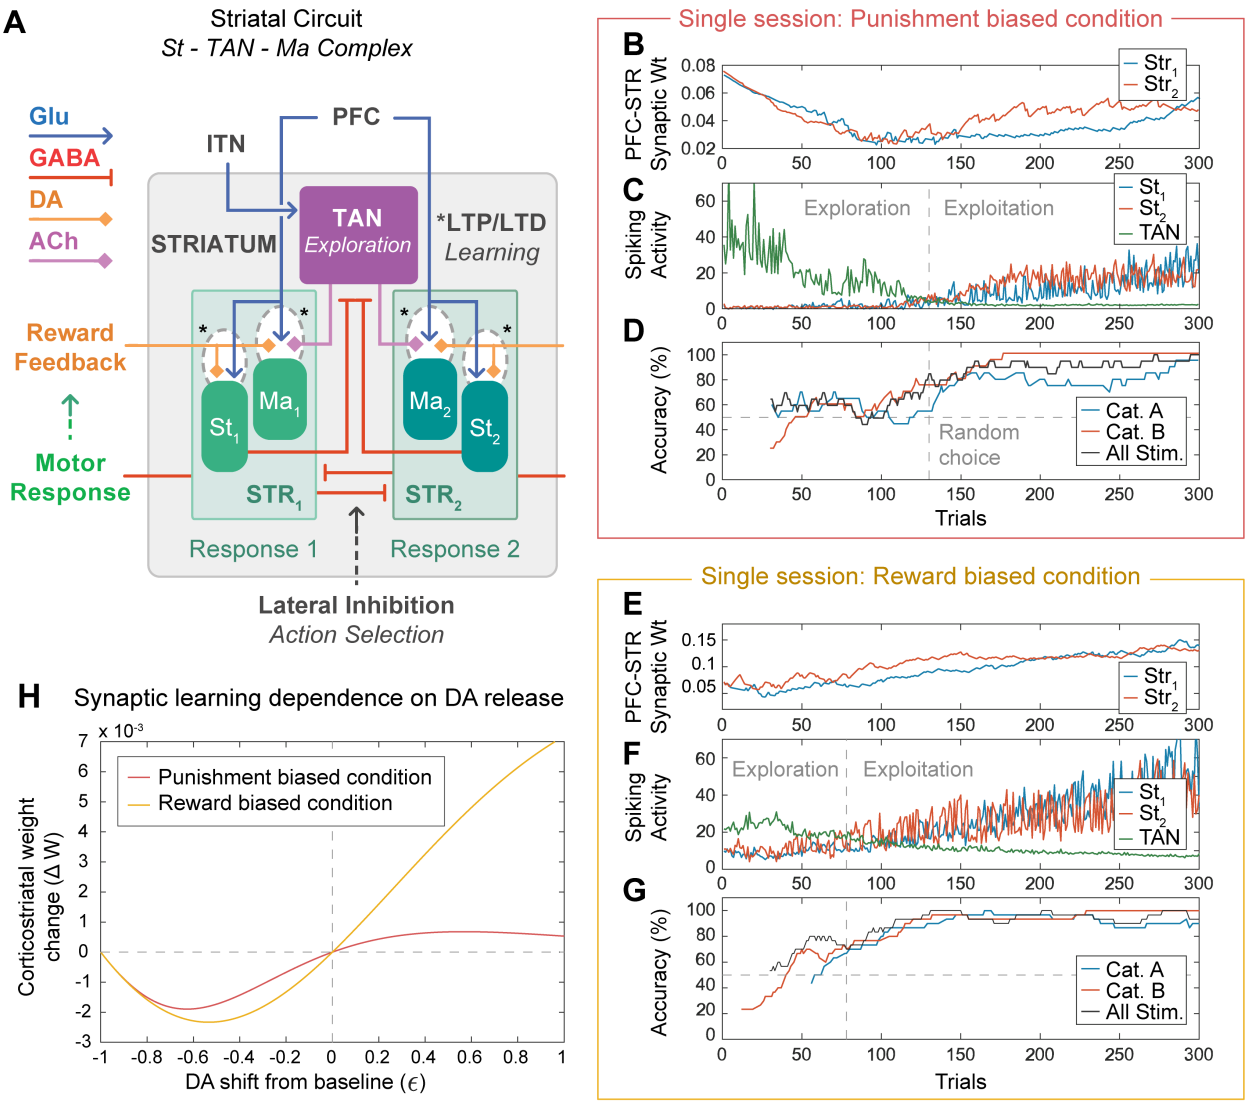

Fig. S3. The role of TAN activity in action selection strategy: exploration to exploitation. (A) shows the circuit diagram of striatal complex consisting of matrisome (Ma), striosome (St) and TAN. Cortical inputs are received by Ma and St. St inhibit TAN while TAN modulate Ma. High TAN activity introduces stochasticity in action selection (exploration). Whereas low TAN activity implies action selection is dominated by corticostriatal synaptic weights, thus the strategy is more deterministic (exploitation). (B) Corticostriatal 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. The transition from exploration to exploitation is thus illustrated. (D) Performance overall (black) and category wise (red and blue.) The parameters used in "punishment biased condition" in (B-D) are listed in Table 2 in the Main Text. We show another learning session with parameters corresponding to "reward biased condition" in panels (E-F). See Section for parameter values. Compare panels (B) & (E), (C) & (F) and (D) & (G). This represents an example of multiple model parameter sets producing qualitatively similar behavior that matches the available data. Key difference is seen in the relative spiking activity in TAN neurons and the striatal blocks, where punishment biased condition shows much stronger TAN activity compared to striatal activity (see panel C) where as reward biased condition shows stronger striatal activity in comparison (see panel F). (H) shows the difference in corticostriatal plasticity under the two parameter sets, that distinguishes session shown in (B-D) from that shown in (E-G). It shows the D1/D2 receptor factor (see Equation 28-29 in Main Text) that describes the dependence of corticostriatal synaptic potentiation/depression on the DA release from SNc on cortico-striatal synapses. For positive shift in DA release (positive reward prediction error due to correct response), the synaptic weight increases whereas negative shift in DA release (negative error due to incorrect response) causes synaptic weight to decrease. Punishment and reward biased conditions are distinct from one another in the relative magnitudes of the effects of positive and negative reward prediction errors, where former has stronger effect of negative error than positive error, while latter has much stronger effect from positive prediction error. This causes the corticostriatal synaptic weights grow at relatively higher rate in reward biased condition (see panel E and B). Even though the behavior shown by the two model settings look qualitatively similar, the key difference emerges when we consider the lesion study for TAN, as we show in Section and Figure S13. Source data are provided as a Source Data file.

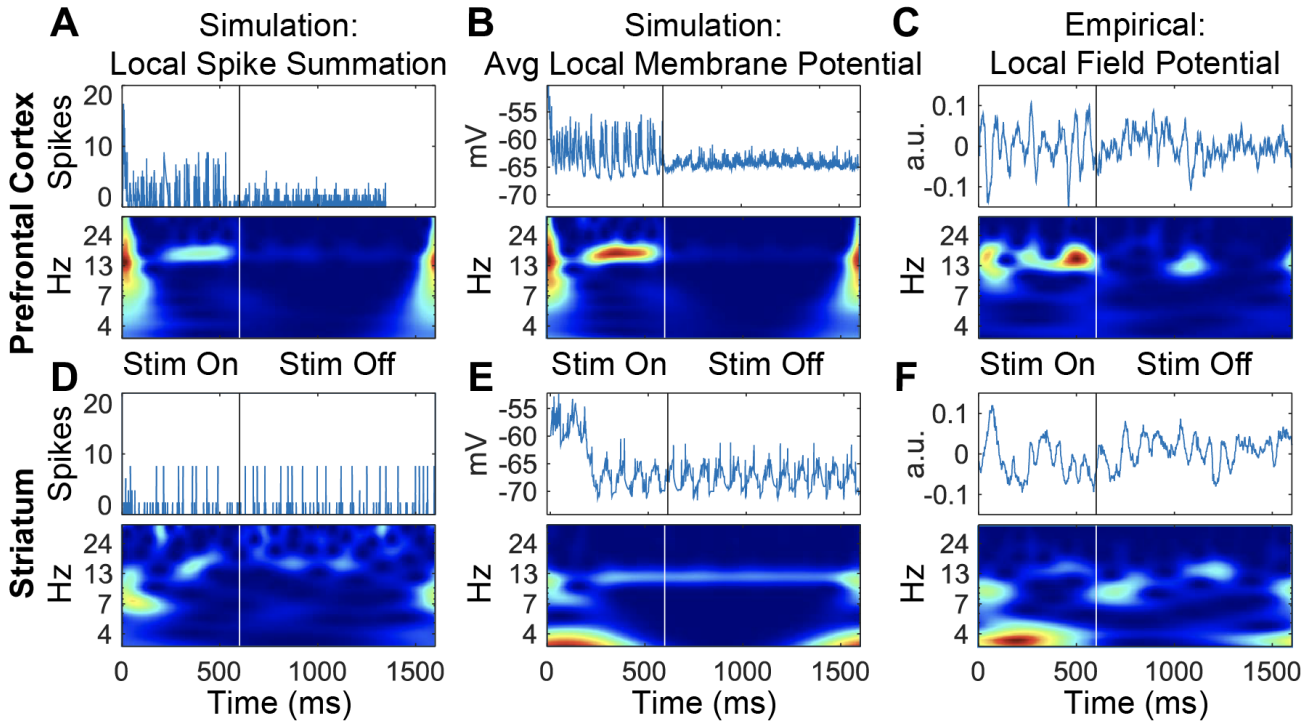

Fig. S4. Summaries of neuronal behavior show that model physiologically richly matches empirical data (which was not included in any model training) even for single trials. Each summary shows behavior from the time of stimulus onset (Time = 0) through stimulus offset (600 ms) until final action (1600 ms). The top shows the summary statistic (spikes, millivolts, and arbitrary units) and the bottom shows a spectrogram (see Methods). (A) Local Spike Summation (LSS) histogram of model anterior-cortex neuron spikes. (B) Average local membrane potential (ALMP) straightforwardly calculated by averaging membrane potentials of all cells. (C) Sample empirical local field potential (LFP) measured from macaque cortical electrode. (D) Model LSS response in simulated striatal MSNs. (E) Striatal model ALMP. (F) LFP measured from striatal electrode in macaque. See text for discussion. Source data are provided as a Source Data file.

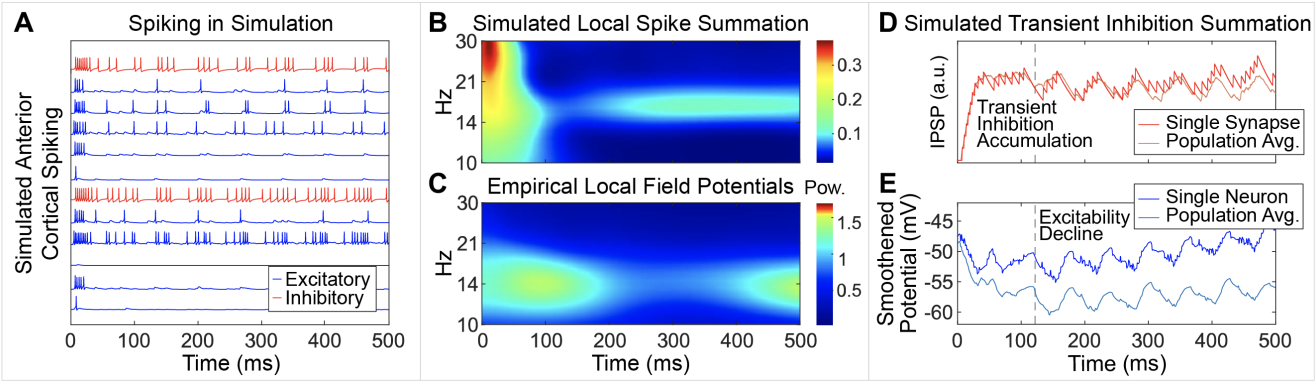

Fig. S5. Cortical spiking shows a post-stimulus-onset silence that was subsequently found in empirical data; investigation of the model provides a specific explanation for this silence that makes a novel prediction for future empirical data. (A) Spiking activity in simulated PFC neurons, colored blue for excitatory neurons and red for inhibitory neurons. 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 the PFC, 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 (corresponding to PFC), averaged across all trials, shows the same qualitative behavior. (D) Inhibitory synaptic conductivity averaged over entire neural population in simulated PFC shows the accumulation of inhibition during 100 ms 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).

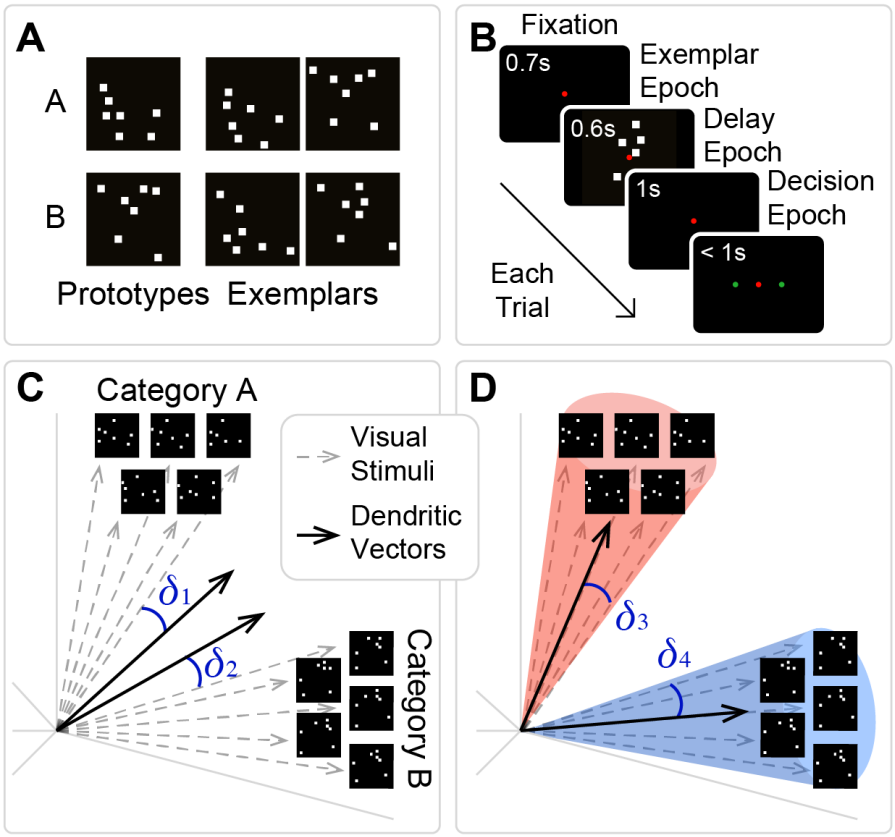

Fig. S6. Synaptic increments achieve natural internal clustering of learned image patterns. (A) Two example categories of visual stimuli. Stimuli are the unique exemplars that were created by distortion of the two prototypes. (B) Schematic illustrating the time course of a single learning trial. See references (1, 141) 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).

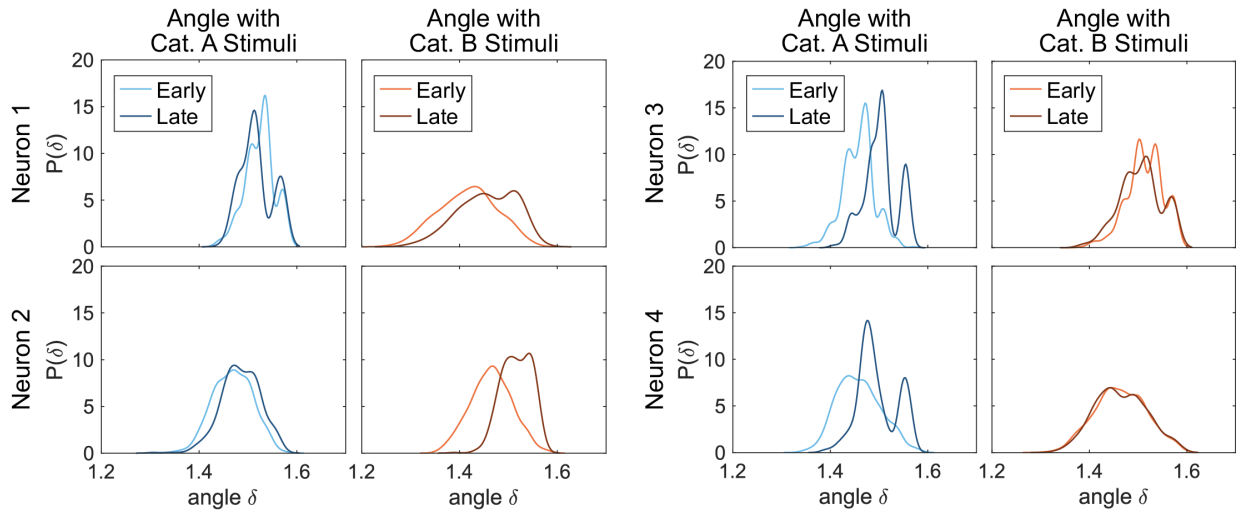

Fig. S7. Learning rotates the dendritic vectors of model anterior cortical (PFC) neurons closer to one category and away from the other category. Each panel shows histogram of projection angles between dendritic vector of single neuron and all visual stimulus vectors of a single category (see main text 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.

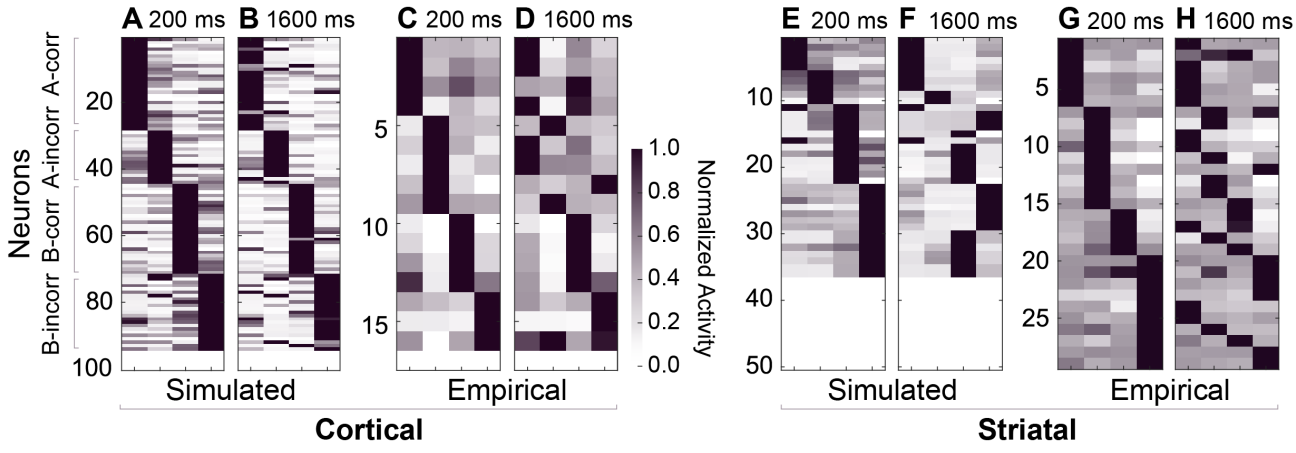

Fig. S8. Cortical and striatal cell spiking preferences during late learning trials, modeled and empirical, in four conditions by category and outcome: category A vs. B; outcome correct vs. incorrect. (See vertical axis descriptors at left.) The grayscale shows the activity normalized for that neuron at that time. (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. White regions indicate units with no spiking activity.

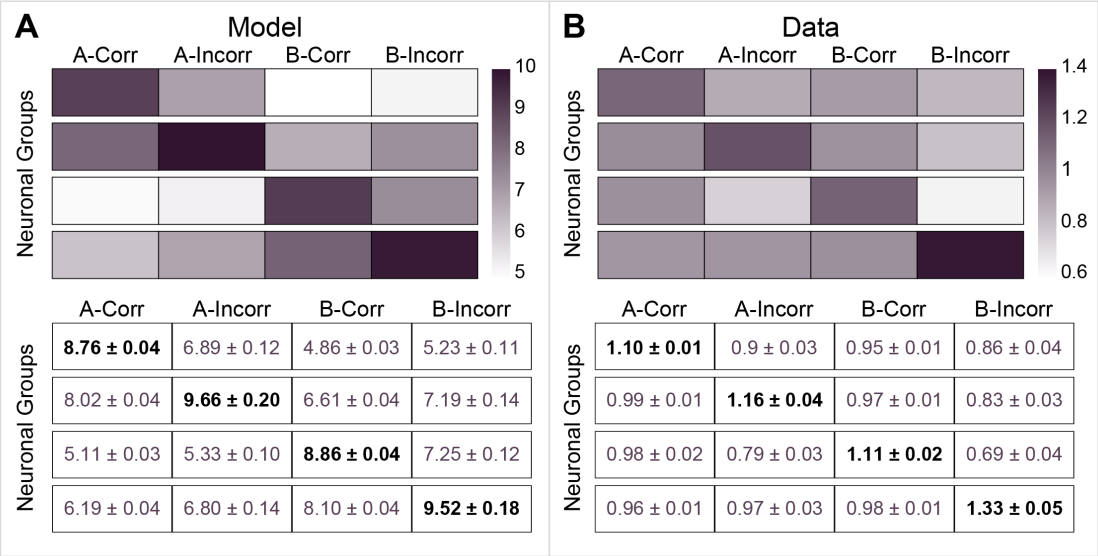

Neuronal Groups	A-Corr	A-Incorr	B-Corr	B-Incorr
A-Corr	**8.76 ± 0.04**	6.89 ± 0.12	4.86 ± 0.03	5.23 ± 0.11
A-Incorr	8.02 ± 0.04	**9.66 ± 0.20**	6.61 ± 0.04	7.19 ± 0.14
B-Corr	5.11 ± 0.03	5.33 ± 0.10	**8.86 ± 0.04**	7.25 ± 0.12
B-Incorr	6.19 ± 0.04	6.80 ± 0.14	8.10 ± 0.04	**9.52 ± 0.18**

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

Fig. S9. 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.

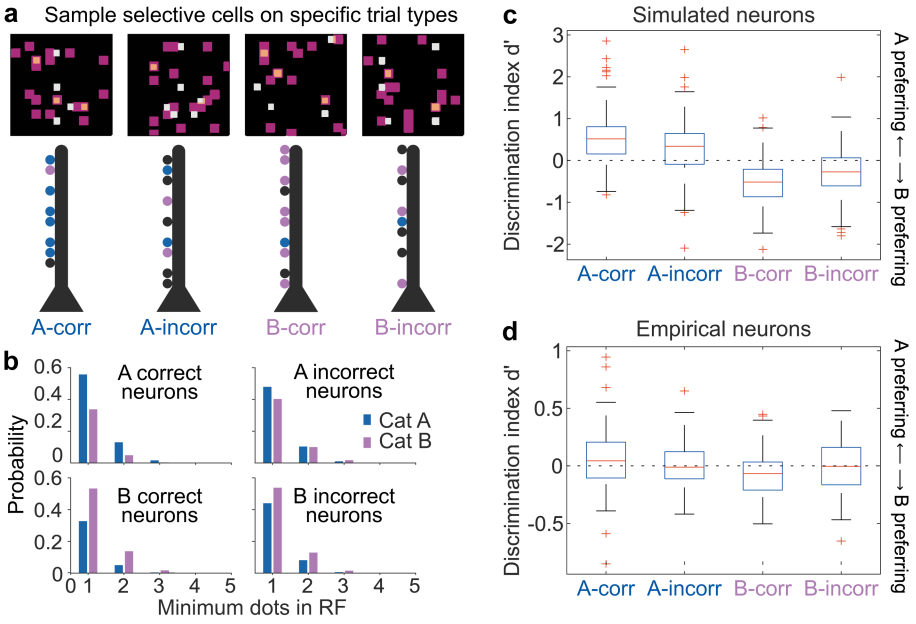

Fig. S10. Neuron selectivity for stimuli and for trial outcomes, in simulated and in empirical data. (A) Sample neuron receptive fields (pink shaded regions) for simulated neurons are shown superimposed on sample visual input patterns (white and orange dots : orange dots overlap with the receptive field and white dots lie outside the receptive field), for each of four types (A-correct, A-incorrect, B-correct, B-incorrect). (B) Selective cell responsiveness of simulated neurons are quantified as the probability (given the somewhat random nature of the stimuli) with which stimuli of both categories overlap with the receptive fields (RF) of each of the cell types, where the degree of overlap is given by minimum number of dots occurring within receptive field. (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). In all box plots, center lines represent medians, box limits represent 1st and 3rd quartiles, whiskers represent 1.5 times inter quartile range and the individual point markers represent the outliers.

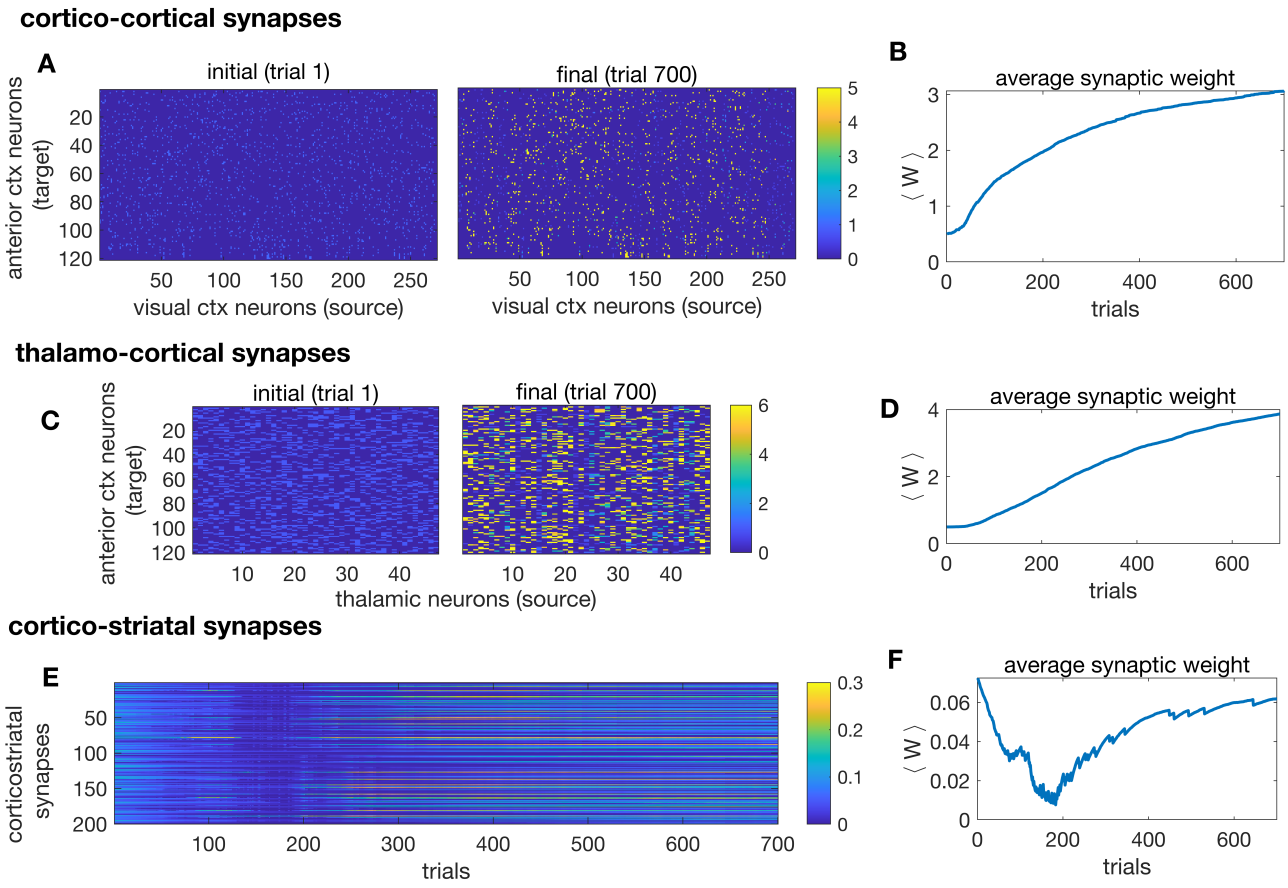

Fig. S11. Time evolution of synaptic weights in the simulation. (A) Shows the initial (left) and final (right) weight matrices of cortico-cortical synapses from simulated visual cortex neurons to simulated anterior cortex neurons. (B) The monotonous increase in the average synaptic weight of cortico-cortical synapses with time (trials) is shown, representing long term potentiation (LTP). (C) Shows the initial (left) and final (right) weight matrices of thalamo-cortical synapses from simulated thalamic neurons to simulated anterior cortex neurons. (D) The monotonous increase in the average synaptic weight of thalamo-cortical synapses with time (trials) is shown, representing long term potentiation (LTP). (E) Shows evolution of all the cortico-striatal synaptic weights with time. (F) Shows rapidly varying average weight of cortico-striatal synapses, representing long-term depression (LTD) as well as potentiation (LTP).

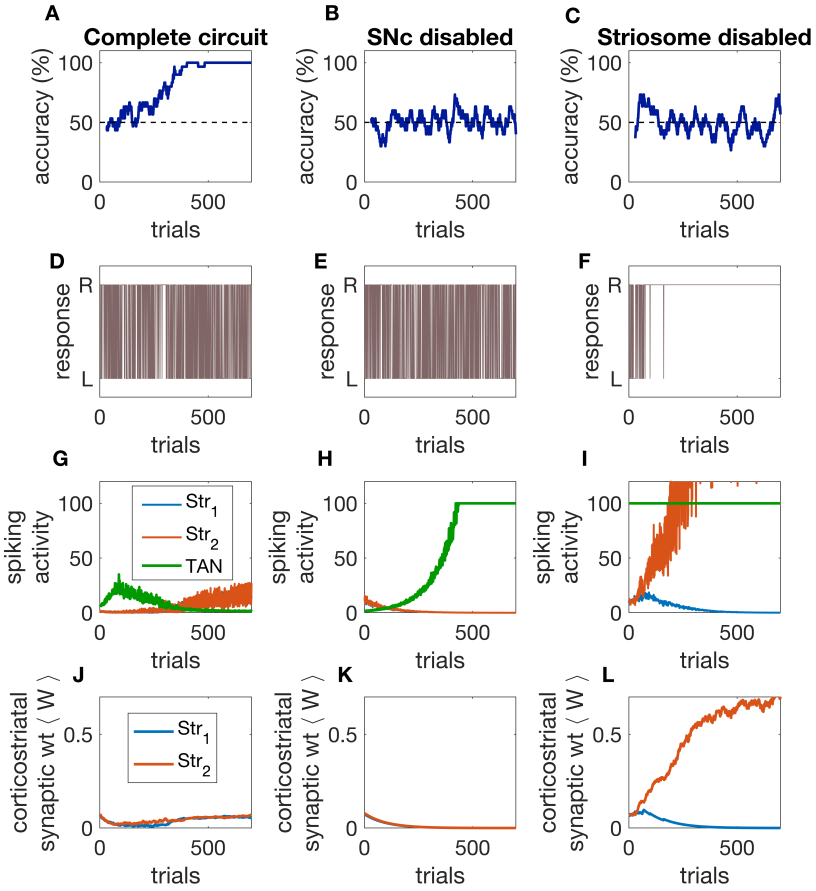

Fig. S12. Influence of SNc and Striosome in the corticostriatal circuit model. (A-C) Behavioural performance shown by the model having (A) complete circuit as shown in main text Figure 1 C, (B) SNc disabled and (C) Striosome disabled. Disabling either of the two elements fails to perform category learning task. (D-F) Trial by trial response (left/right saccade) given by the model having complete ciruit, SNc disabled and striosome disabled respectively. (G-I) Spiking activity (no. of spikes during first 200 ms of stimulus on) in TAN neurons and the two STR blocks, shown by the model having complete ciruit, SNc disabled and striosome disabled respectively.(J-L) Time evolution of average corticostriatal synaptic weights to the two striatal blocks in the model having complete ciruit, SNc disabled and striosome disabled respectively. Note that absence of SNc leads to loss of synaptic potentiation (see panel K) and thus loss of striatal activity and dominance of TAN (see panel H). This leads to almost random action choices (panel E) as reflected in the behavioral performance (panel B). On the otherhand, in the absence of striosome, TAN activity is never suppressed (panel I), causing lack of transition from exploration phase to explitation phase (see Figure S3). Additionally, the SNc is also never suppressed as learning advances, making the synaptic weights unstable. The balance of synaptic weights between the two striatal blocks is broken and the model is locked onto a single response choice (panel F). Thus, both the loss of SNc and of striosome break the learning behavior but in distinct ways.

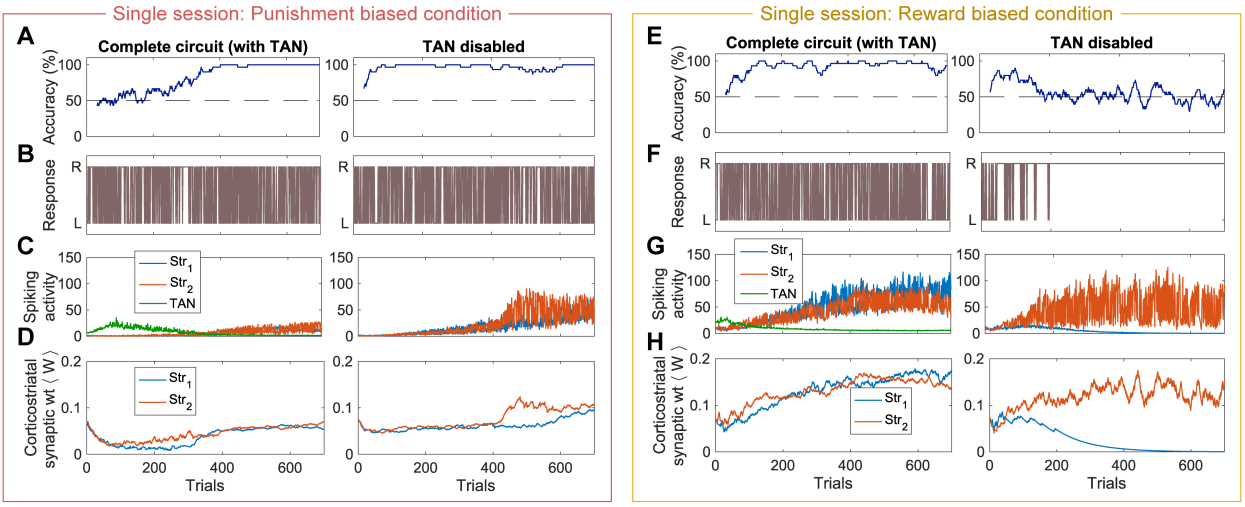

Fig. S13. Influence of TAN neurons in the corticostriatal circuit model under two different parameter settings: punishment biased and reward biased conditions, respectively. Panels (A-D) correspond to punishment biased condition as described in Figure S3, where as panels (E-H) correspond to reward biased condition. (A & E) Behavioural performance shown by the model having complete circuit (left panel) and with TAN neurons disabled (right panel). In reward biased condition, disabling TAN completely destabilizes the category learning behavior. (B & F) Trial by trial response (left/right saccade) given by the model having complete circuit (left panel) and TAN neurons disabled (right panel). For punishment biased condition (A-D), removal of TAN does not cause the breakdown of learning behaviour. However in reward biased condition (E-H), in the absence of TAN, the choice action locks onto a single response irrespective of the stimulus (panel F, right). (C & G) Spiking activity (no. of spikes during first 200 ms of stimulus on) in TAN neurons and the two STR blocks, shown by the model having complete circuit (left panel) and TAN neurons disabled (right panel). In the absence of TAN neurons, an imbalance is developed between the activities of the two STR blocks where activity in one of the blocks becomes constantly dominating over the other block. This reflects in the behavioral locking of choice action onto a single response in panel F, right. (D & H) Time evolution of average corticostriatal synaptic weights to the two striatal blocks in the model having complete circuit (left) and TAN neurons disabled (right). In the absence of TAN neurons, the imbalance of synaptic weights between the two STR blocks mirrors the spiking activity and the response.

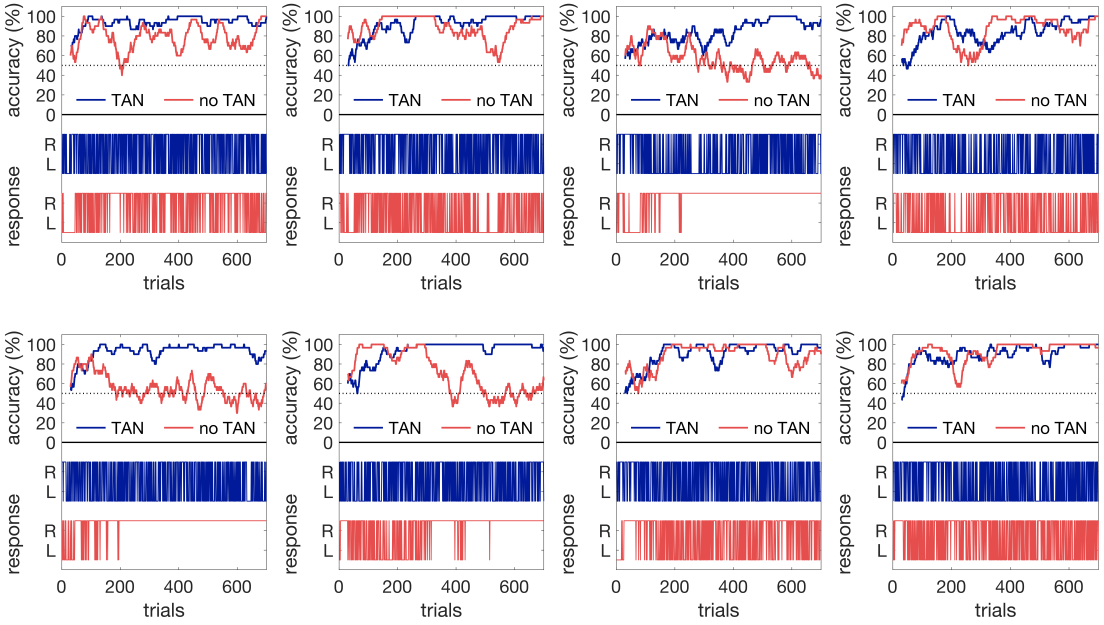

Fig. S14. Influence of TAN neurons demonstrated over multiple realizations of circuit model simulations, under reward biased condition, with and without TAN neurons. Top panels show the behavioral performance for each run and bottom panels how the trial by trial response in each run. Each panel shows a comparison between two learning sessions with all parameters identical and subjected to the same stimulus set, but one session is run with the complete model (blue curve) and the other is run with TAN neurons disabled (red curve). In each example, the runs without TAN neurons have a destabilized learning behaviour with tendency to get locked within a single response for choice action.

References

1. EG Antzoulatos, EK Miller, Increases in Functional Connectivity between Prefrontal Cortex and Striatum during Category Learning. *Neuron* **83**, 216–225 (2014).
2. JNJ Reynolds, JR Wickens, Dopamine-dependent plasticity of corticostriatal synapses. *Neural Networks* **15**, 507–521 (2002).
3. 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).
4. EL Bienenstock, LN Cooper, PW Munro, Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* **2**, 32–48 (1982).
5. AL Hodgkin, AF Huxley, A quantitative description of membrane current and its application to conduction and excitation in nerve. *The J. Physiol.* **117**, 500–544 (1952).
6. XJ Wang, Ionic basis for intrinsic 40 Hz neuronal oscillations:. *NeuroReport* **5**, 221–224 (1993).
7. XJ Wang, Pacemaker Neurons for the Theta Rhythm and Their Synchronization in the Septohippocampal Reciprocal Loop. *J. Neurophysiol.* **87**, 889–900 (2002).
8. C Koch, I Segev, TJ Sejnowski, TA Poggio, eds., *Methods in Neuronal Modeling: From Ions to Networks*, Computational Neuroscience Series. (A Bradford Book, Cambridge, MA, USA), 2 edition, (1998).
9. SR Wicks, CJ Roehrig, CH Rankin, A Dynamic Network Simulation of the Nematode Tap Withdrawal Circuit: Predictions Concerning Synaptic Function Using Behavioral Criteria. *The J. Neurosci.* **16**, 4017–4031 (1996).
10. 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).
11. S Coombes, A 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).
12. A Byrne, RD O'Dea, M Forrester, J Ross, S Coombes, Next-generation neural mass and field modeling. *J. Neurophysiol.* **123**, 726–742 (2020).
13. GB Ermentrout, N Kopell, Parabolic Bursting in an Excitable System Coupled with a Slow Oscillation. *SIAM J. on Appl. Math.* **46**, 233–253 (1986).
14. JB Ding, JN Guzman, JD Peterson, JA Goldberg, DJ Surmeier, Thalamic Gating of Corticostriatal Signaling by Cholinergic Interneurons. *Neuron* **67**, 294–307 (2010).
15. MC Avery, JL Krichmar, Neuromodulatory systems and their interactions: a review of models, theories, and experiments. *Front. Neural Circuits* **11** (2017).
16. BE Jones, Arousal systems. *Front. Biosci.* **8**, 438–451 (2003).
17. RE Brown, JT McKenna, Turning a Negative into a Positive: Ascending GABAergic Control of Cortical Activation and Arousal. *Front. Neurol.* **6**, 135 (2015).
18. 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).
19. BE Kilavik, M Zaepffel, A Brovelli, WA MacKay, A Riehle, The ups and downs of beta oscillations in sensorimotor cortex. *Exp. Neurol.* **245**, 15–26 (2013).
20. 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).
21. Y Kawaguchi, T Shindou, Noradrenergic excitation and inhibition of GABAergic cell types in rat frontal cortex. *J. Neurosci.* **18**, 6963–6976 (1998) Publisher: Society for Neuroscience Section: ARTICLE.
22. H Salgado, et al., Layer-specific noradrenergic modulation of inhibition in cortical layer II/III. *Cereb. Cortex* **21**, 212–221 (2011).
23. H Alitto, Y Dan, Cell-type-specific modulation of neocortical activity by basal forebrain input. *Front. Syst. Neurosci.* **6** (2013).
24. M Jayachandran, et al., Nucleus reuniens transiently synchronizes memory networks at beta frequencies. *Nat. Commun.* **14**, 4326 (2023) Number: 1 Publisher: Nature Publishing Group.
25. A Bibbig, et al., Beta Rhythms (15–20 Hz) Generated by Nonreciprocal Communication in Hippocampus. *J. Neurophysiol.* **97**, 2812–2823 (2007) Publisher: American Physiological Society.
26. LR Silva, Y Amitai, BW Connors, Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Sci. (New York, N.Y.)* **251**, 432–435 (1991).
27. 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) _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1460-9568.2011.07865.x.
28. MM McCarthy, et al., Striatal origin of the pathologic beta oscillations in Parkinson's disease. *Proc. Natl. Acad. Sci.* **108**, 11620–11625 (2011).
29. A Singh, SM Papa, Striatal oscillations in parkinsonian non-human primates. *Neuroscience* **449**, 116–122 (2020).
30. V Breton-Provencher, M Sur, Active control of arousal by a locus coeruleus GABAergic circuit. *Nat. Neurosci.* **22**, 218–228 (2019).
31. JY Hansen, et al., Integrating brainstem and cortical functional architectures (2023).
32. GMG Shepherd, Corticostriatal connectivity and its role in disease. *Nat. Rev. Neurosci.* **14**, 278–291 (2013).
33. 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) Publisher: Society for Neuroscience Section: Articles.
34. 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) Publisher: Society for Neuroscience Section: Brief Communications.

35. CR Gerfen, MN Economo, J Chandrashekar, Long distance projections of cortical pyramidal neurons. *J. Neurosci. Res.* **96**, 1467–1475 (2018) _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jnr.23978.
36. J Larson, P Xiao, G Lynch, Reversal of LTP by theta frequency stimulation. *Brain Res.* **600**, 97–102 (1993).
37. J Larson, E Munkacsy, Theta-burst LTP. *Brain Res.* **1621**, 38–50 (2015).
38. DD Mott, DV Lewis, Facilitation of the Induction of Long-term Potentiation by GABAB Receptors. *Science* **252**, 1718–1720 (1991) Publisher: American Association for the Advancement of Science.
39. DD Mott, DV Lewis, GABAB receptors mediate disinhibition and facilitate long-term potentiation in the dentate gyrus. *Epilepsy research Suppl.* **7**, 119–134 (1992).
40. T Nakashiba, et al., Young Dentate Granule Cells Mediate Pattern Separation, whereas Old Granule Cells Facilitate Pattern Completion. *Cell* **149**, 188–201 (2012).
41. A Terashima, et al., An Essential Role for PICK1 in NMDA Receptor-Dependent Bidirectional Synaptic Plasticity. *Neuron* **57**, 872–882 (2008).
42. EA Kramar, B Lin, CS Rex, CM Gall, G Lynch, Integrin-driven actin polymerization consolidates long-term potentiation. *Proc. Natl. Acad. Sci.* **103**, 5579–5584 (2006).
43. G Lynch, CS Rex, CM Gall, LTP consolidation: Substrates, explanatory power, and functional significance. *Neuropharmacology* **52**, 12–23 (2007).
44. W Schultz, P Dayan, PR Montague, A Neural Substrate of Prediction and Reward. *Science* **275**, 1593–1599 (1997).
45. W Schultz, Predictive reward signal of dopamine neurons. *J. Neurophysiol.* **80**, 1–27 (1998).
46. W Schultz, Updating dopamine reward signals. *Curr. Opin. Neurobiol.* **23**, 229–238 (2013).
47. W Schultz, Dopamine reward prediction-error signalling: a two-component response. *Nat. Rev. Neurosci.* **17**, 183–195 (2016).
48. EO Neftci, BB Averbeck, Reinforcement learning in artificial and biological systems. *Nat. Mach. Intell.* **1**, 133–143 (2019).
49. 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).
50. 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).
51. C Rackauckas, ODE Solver Multi-Language Wrapper Package Work-Precision Benchmarks (MATLAB, SciPy, Julia, deSolve (R)) (2023).
52. C Pantelides, The Consistent Initialization of Differential-Algebraic Systems. *SIAM J. Scl. Stat. Comput.* **9**, 213–231 (1988).
53. 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).
54. 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).
55. H Elmqvist, T Henningsson, M Otter, Innovations for Future Modelica in *Innovations for Future Modelica.* pp. 693–702 (2017).
56. B Guenter, The D* Symbolic Differentiation Algorithm (2007).
57. AG Baydin, BA Pearlmutter, AA Radul, JM Siskind, Automatic Differentiation in Machine Learning: a Survey. *arXiv:1502.05767* (2018).
58. C Rackauckas, N Qing, DifferentialEquations.jl - A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia. *J. Open Res. Softw.* **5** (2017).
59. J Verner, Numerically optimal Runge–Kutta pairs with interpolants. *Numer. Algorithms* **53**, 383–396 (2010).
60. C Torrence, GP Compo, A Practical Guide to Wavelet Analysis. *Bull. Am. Meteorol. Soc.* **79**, 61–78 (1998) Publisher: American Meteorological Society Section: Bulletin of the American Meteorological Society.
61. R Coultrip, R Granger, G Lynch, A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks* **5**, 47–54 (1992).
62. M Abeles, *Corticonics: Neural Circuits of the Cerebral Cortex.* (Cambridge University Press, Cambridge), (1991).
63. 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).
64. SH Hendry, HD Schwark, EG Jones, J Yan, Numbers and proportions of GABA-immunoreactive neurons in different areas of monkey cerebral cortex. *The J. Neurosci. The Off. J. Soc. for Neurosci.* **7**, 1503–1519 (1987).
65. V Braitenberg, A Schuz, *Cortex: statistics and geometry of neuronal connectivity.* (Springer, Berlin, Heidelberg), (1998).
66. J Szentagothai, The 'module-concept' in cerebral cortex architecture. *Brain Res.* **95**, 475–496 (1975).
67. 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).
68. OD Creutzfeldt, HC Nothdurft, Representation of complex visual stimuli in the brain. *Naturwissenschaften* **65**, 307–318 (1978).
69. VB Mountcastle, Brain mechanisms for directed attention. *J. Royal Soc. Medicine* **71**, 14–28 (1978).
70. A Keller, EL White, Triads: a synaptic network component in the cerebral cortex. *Brain Res.* **496**, 105–112 (1989).
71. RAW Galuske, W Schlote, H Bratzke, W Singer, Interhemispheric asymmetries of the modular structure in human temporal cortex. *Science* **289**, 1946–1949 (2000).
72. MS Gazzaniga, Regional differences in cortical organization. *Science* **289**, 1887–1888 (2000).
73. MA Castro-alamancos, BW Connors, Thalamocortical synapses. *Prog. Neurobiol.* **51**, 581–606 (1997).
74. EG Jones, Viewpoint: the core and matrix of thalamic organization. *Neuroscience* **85**, 331–345 (1998).
75. AJ Heynen, MF Bear, Long-term potentiation of thalamocortical transmission in the adult visual cortex in vivo. *J. Neurosci.* **21**,

9801–9813 (2001).
76. G Silberberg, A Gupta, H Markram, Stereotypy in neocortical microcircuits. *Trends Neurosci*. **25**, 227–230 (2002).
77. F Valverde, Structure of the cerebral cortex. Intrinsic organization and comparative analysis of the neocortex. *Revista de neurologia* **34**, 758–780 (2002).
78. EL White, A Peters, Cortical modules in the posteromedial barrel subfield (Sml) of the mouse. *J. Comp. Neurol*. **334**, 86–96 (1993).
79. 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).
80. M Molinari, et al., Auditory thalamocortical pathways defined in monkeys by calcium-binding protein immunoreactivity. *J. Comp. Neurol*. **362**, 171–194 (1995).
81. H Killackey, FF Ebner, Two different types of thalamocortical projections to a single cortical area in mammals. *Brain Behav. Evol*. **6**, 156–169 (1972).
82. 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).
83. 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).
84. 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).
85. 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).
86. 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).
87. H Barbas, N Rempel-Clower, Cortical structure predicts the pattern of corticocortical connections. *Cereb. Cortex* **7**, 635–646 (1997).
88. KS Rockland, Non-uniformity of extrinsic connections and columnar organization. *J. Neurocytol*. **31**, 247–253 (2002).
89. HA Swadlow, AG Gusev, T Bezdudnaya, Activation of a cortical column by a thalamocortical impulse. *J. Neurosci*. **22**, 7766–7773 (2002).
90. DB Bender, Visual activation of neurons in the primate pulvinar depends on cortex but not colliculus. *Brain Res*. **279**, 258–261 (1983).
91. 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).
92. 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).
93. JJ Chrobak, G Buzsaki, Gamma oscillations in the entorhinal cortex of the freely behaving rat. *J. Neurosci*. **18**, 388–398 (1998).
94. M Sarter, JP Bruno, The neglected constituent of the basal forebrain corticopetal projection system: GABAergic projections. *Eur. J. Neurosci*. **15**, 1867–1873 (2002).
95. P Fries, JH Reynolds, AE Rorie, R Desimone, Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**, 1560–1563 (2001).
96. 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).
97. 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).
98. 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).
99. A Rodriguez, J Whitson, R Granger, Derivation and Analysis of Basic Computational Operations of Thalamocortical Circuits. *J. Cogn. Neurosci*. **16**, 856–877 (2004).
100. J Ambros-Ingerson, R Granger, G Lynch, Simulation of paleocortex performs hierarchical clustering. *Science* **247**, 1344–1348 (1990).
101. 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).
102. A Bocincova, TJ Buschman, MG Stokes, SG Manohar, Neural signature of flexible coding in prefrontal cortex. *Proc. Natl. Acad. Sci*. **119**, e2200400119 (2022).
103. 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).
104. R Granger, Engines of the Brain: The Computational Instruction Set of Human Cognition. *AI Mag*. **27**, 15–32 (2006).
105. JJ Canales, et al., Shifts in striatal responsivity evoked by chronic stimulation of dopamine and glutamate systems. *Brain* **125**,

2353–2363 (2002).
106. GE Alexander, MR DeLong, PL Strick, Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu. Rev. Neurosci.* **9**, 357–381 (1986).
107. RL Reep, JV Corwin, Topographic organization of the striatal and thalamic connections of rat medial agranular cortex. *Brain Res.* **841**, 43–52 (1999).
108. 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).
109. M Lundqvist, et al., Gamma and Beta Bursts Underlie Working Memory. *Neuron* **90**, 152–164 (2016).
110. M Lundqvist, et al., Working memory control dynamics follow principles of spatial computing. *Nat. Commun.* **14**, 1429 (2023).
111. A Reiner, N Hart, W Lei, Y Deng, Corticostriatal projection neurons-dichotomous types and dichotomous functions. *Front. Neuroanat.* **4** (2010).
112. AM Graybiel, T Aosaki, AW Flaherty, M Kimura, The basal ganglia and adaptive motor control. *Science* **265**, 1826–1831 (1994).
113. BJ Hunnicutt, et al., A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife* **5**, e19103 (2016).
114. AI Gulyas, N Hajos, TF Freund, Interneurons containing calretinin are specialized to control other interneurons in the rat hippocampus. *J. Neurosci.* **16**, 3397–3411 (1996).
115. JM Blasco-Ibanez, 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).
116. AI Gulyas, L Acsady, TF Freund, Structural basis of the cholinergic and serotonergic modulation of GABAergic neurons in the hippocampus. *Neurochem. Int.* **34**, 359–372 (1999).
117. 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).
118. 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).
119. P Reinagel, D Godwin, SM Sherman, C Koch, Encoding of visual information by lgn bursts. *J. Neurophysiol.* **81**, 2558–2569 (1999).
120. 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).
121. 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).
122. MF Bear, A Kirkwood, Neocortical long-term potentiation. *Curr. Opin. Neurobiol.* **3**, 197–202 (1993).
123. 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).
124. DV Buonomano, MM Merzenich, Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* **21**, 149–186 (1998).
125. 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).
126. G Mongillo, O Barak, M Tsodyks, Synaptic Theory of Working Memory. *Science* **319**, 1543–1546 (2008).
127. L Kozachkov, et al., Robust and brain-like working memory through short-term synaptic plasticity. *PLOS Comput. Biol.* **18**, e1010776 (2022).
128. 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).
129. 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).
130. 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).
131. 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).
132. 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).
133. YF Zhang, SJ Cragg, Pauses in Striatal Cholinergic Interneurons: What is Revealed by Their Common Themes and Variations? *Front. Syst. Neurosci.* **11** (2017).
134. 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).
135. 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).
136. G Mandelbaum, et al., Distinct cortical-thalamic-striatal circuits through the parafascicular nucleus. *Neuron* **102**, 636–652.e7 (2019).
137. BD Bennett, CJ Wilson, Spontaneous activity of neostriatal cholinergic interneurons in vitro. *J. Neurosci.* **19**, 5586–5596 (1999).
138. 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).
139. P Apicella, Leading tonically active neurons of the striatum from reward detection to context recognition. *Trends Neurosci.* **30**,

299–306 (2007).
140. N Abudukeyoumu, T Hernandez-Flores, M Garcia-Munoz, GW Arbuthnott, Cholinergic modulation of striatal microcircuits. *Eur. J. Neurosci.* **49**, 604–622 (2019).
141. EG Antzoulatos, EK Miller, Differences between Neural Activity in Prefrontal Cortex and Striatum during Learning of Novel Abstract Categories. *Neuron* **71**, 243–249 (2011).