

Local, calcium- and reward-based synaptic learning rule that enhances dendritic nonlinearities can solve the nonlinear feature binding problem

Zahra Khodadadi^{1*}, Daniel Trpevski¹, Robert Lindroos^{1†}, Jeanette Hellgren Kotaleski^{1,2†}

¹ Department of Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden

² Department of Neuroscience, Karolinska Institutet, Stockholm, Sweden

***For correspondence:**

zahra.khodadadi@scilifelab.se (ZK)

[†] These authors contributed equally to this work

Abstract

This study explores the computational potential of single striatal projection neurons (SPN), emphasizing dendritic nonlinearities and their crucial role in solving complex integration problems. Utilizing a biophysically detailed multicompartmental model of an SPN, we introduce a calcium-based, local synaptic learning rule that leverages dendritic plateau potentials. According to what is known about excitatory corticostriatal synapses, the learning rule is governed by local calcium dynamics from NMDA and L-type calcium channels and dopaminergic reward signals. In addition, we incorporated metaplasticity in order to devise a self-adjusting learning rule which ensures stability for individual synaptic weights. We demonstrate that this rule allows single neurons to solve the nonlinear feature binding problem (NFBP), a task traditionally attributed to neuronal networks. We also detail an inhibitory plasticity mechanism, critical for dendritic compartmentalization, further enhancing computational efficiency in dendrites. This *in silico* study underscores the computational capacity of individual neurons, extending our understanding of neuronal processing and the brain's ability to perform complex computations.

Introduction

Classically, single neurons in the nervous system have been thought to operate as simple linear integrators where the nonlinearity of dendrites can be neglected (McCulloch and Pitts, 1943). Based on this simplification, powerful artificial neural systems have been created, outperforming humans on multiple tasks (Silver *et al.*, 2018). However, in recent decades it has been shown that active dendritic properties participate in shaping neuronal output. Not only the axon spikes, but also the dendrites can display nonlinear integration of input signals (Antic *et al.*, 2010). Dendritic nonlinearities endow a neuron with the ability to perform sophisticated dendritic computations, expanding its computational power beyond what is available with the somatic voltage threshold and making it similar to a multilayer artificial neural network (Poirazi, Brannon and Mel, 2003).

A dendritic nonlinearity common among projection neurons in several brain areas is the NMDA-dependent plateau potential (Oikonomou *et al.*, 2014). Plateau potentials are regenerative, all-or-none, supralinear voltage elevations triggered by spatio-temporally clustered glutamatergic input (Schiller *et al.*, 2000; Polsky, Mel and Schiller, 2004; Losonczy and Magee, 2006; Major *et al.*, 2008; Larkum *et al.*, 2009; Lavzin *et al.*, 2012; Xu *et al.*, 2012). Such plateaus require that nearby spines are coactivated, but the requirement is perhaps somewhat loose as even single branches have been proposed to act as computational units (Losonczy and Magee, 2006; Branco and Häusser, 2010). With that said, multiple dendritic regions, preferentially responsive to different input values or features, are known to form with close dendritic proximity (Jia *et al.*, 2010; Chen *et al.*, 2011; Varga *et al.*, 2011). Such functional synaptic clusters are present in multiple species, developmental stages and brain regions (Kleindienst *et al.*, 2011; Takahashi *et al.*, 2012; Winnubst *et al.*, 2015; Wilson *et al.*, 2016; Iacarus, Gasler and Hofer, 2017; Scholl, Wilson and Fitzpatrick, 2017; Niculescu *et al.*, 2018; Kerlin *et al.*, 2019; Ju *et al.*, 2020). Hence, multiple features are commonly clustered in a single dendritic branch, indicating that this could be the neural substrate where combinations of simple features into more complex items occur.

Combinations of features in dendritic branches further provide single neurons with the possibility to solve linearly non-separable tasks, such as the nonlinear feature binding problem (NFBP) (Tran-Van-Minh *et al.*, 2015; Gidon *et al.*, 2020). In the most elementary form, the NFBP consists of discriminating between two groups of feature combinations. The problem is nonlinear since the neuron should learn to respond only to specific feature combinations, even though all features are represented by the same amount of synaptic input. Common example features used are two different shapes combined with two different colors, giving in total four combinations of which the neuron should respond to only two feature combinations (exemplified in Fig. 1A and 1B).

As a task, the NFBP is relevant to brain regions which perform integration of multimodal input signals, or signals representing different features of the same modality (Roskies, 1999). It is usually illustrated with examples from the visual system, as in Fig. 1A (Roskies, 1999; von der Malsburg, 1999; Tran-Van-Minh *et al.*, 2015). A region that integrates multimodal inputs, such as sensory information and motor-related signals, is the input nucleus of the basal ganglia, the striatum (Reig and Silberberg, 2014; Johansson and Silberberg, 2020), and this system will be used in the present modeling study. We will here, however, continue to illustrate the NFBP with the more intuitive features borrowed from the visual field, although for dorsal striatum these features would rather map onto different sensory- and motor-related features. Plateau potentials and some clustering of input have been demonstrated in SPNs (Plotkin, Day and Surmeier, 2011; Oikonomou *et al.*, 2014; Du *et al.*, 2017; Hwang *et al.*, 2022; Day *et al.*, 2024; Sanabria *et al.*, 2024).

In addition to integrating convergent input from the cortex and the thalamus, the striatum is densely innervated by midbrain dopaminergic neurons which carry information about rewarding stimuli (Schultz, 2007; Matsuda *et al.*, 2009; Surmeier *et al.*, 2010). As such, the striatum is thought to be an important site of reward learning, associating actions with outcomes based on neuromodulatory cues. In this classical framework, peaks in dopamine (Da) signify rewarding outcomes and pauses in dopamine represent omission of expected rewards (Schultz, Dayan and Montague, 1997). Dopamine signals further control the synaptic plasticity of corticostriatal synapses on the SPNs (Fig. 1C). In direct pathway SPNs (dSPN) expressing the D₁ receptor, a dopamine peak stimulates the D₁ receptor, which together with significant calcium influx through NMDA receptors triggers synaptic strengthening (long-term potentiation - LTP). Conversely, when little/no dopamine is bound to the D₁ receptors, as during a dopamine pause, and there is significant calcium influx through L-type calcium channels, synaptic weakening occurs (long-term depression - LTD) (Shen *et al.*, 2008; Fino *et al.*, 2010; Plotkin *et al.*, 2013) (see Fig. 1D).

If dopamine peaks are associated with the relevant feature combinations in the NFBP and dopamine pauses with the irrelevant ones, it can trigger LTP in synapses representing the relevant feature combinations and LTD in those representing irrelevant combinations. If, after learning, the relevant feature combinations have strong enough synapses so they can evoke plateau potentials while the irrelevant feature combinations have weak enough synapses so they don't evoke plateaus, the outcome of this learning process should be a synaptic arrangement that could solve the NFBP (Fig. 1D) (Tran-Van-Minh *et al.*, 2015). In line with this, it has been demonstrated that the NFBP can be solved in abstract neuron models where the soma and dendrites are represented by single electrical compartments and where neuronal firing and plateau potentials are phenomenologically represented by

instantaneous firing rate functions (Legenstein and Maass, 2011; Schiess, Urbanczik and Senn, 2016). Good performance on the NFBP has also been demonstrated with biologically detailed models (Bicknell and Häusser, 2021). This solution used a multicompartmental model of a single pyramidal neuron, including both excitatory and inhibitory synapses and supralinear NMDA depolarizations. Synapses representing different features were randomly dispersed throughout the dendrites and a phenomenological learning rule - dependent on somatic spike timing and high local dendritic voltage - were used to optimize the strength of the synapses. The solution did, however, depend on a form of supervised learning as somatic current injections were used to raise the spiking probability of the relevant feature combinations.

In this article we develop a local, calcium-dependent synaptic learning rule based on the known synaptic machinery of the corticostriatal synapse onto dSPNs (Shen *et al.*, 2008), which, guided by the reward signals from dopaminergic neurons, can solve the NFBP in a multicompartment model of a dSPN. The learning rule is initially tested assuming pre-existing clustered synapses for each individual feature, potentiating synaptic clusters to the point where they can evoke robust plateau potentials for the relevant feature combinations of the NFBP, and weakening synapses representing irrelevant features. We also apply the learning rule on more randomly distributed synapses, and results suggest that branch-specific plasticity might be important for the single neuron to solve nonlinear problems. Although brain systems integrating multimodal inputs, such as the striatum, somehow solve nonlinear problems at the network/systems level, it is not known whether any individual neuron in the brain actually solves the NFBP on a regular basis. Our investigation suggests, however, that single SPNs have the computational capacity to solve linearly non-separable tasks by utilizing information of the organism's success (here represented by a dopamine peak) and failures (here a dopamine pause). This might also generalize to other projection neurons that can display dendritic plateaus, such as pyramidal neurons. Depending on neuron and synapse type, however, the specific feedback mechanisms (which in this case is dopamine) will have to be mapped to another neuromodulatory signal.




Figure 1: Learning Mechanisms in direct pathway Striatal Projection Neurons (dSPNs) for the Nonlinear Feature Binding Problem (NFBP)

A: Inputs and assumed supralinearity that could solve the NFBP: The NFBP is represented with an example from visual feature binding. In the simplest form of the NFBP, a stimulus has two features, here shape and form, each with two possible values, strawberry and banana, and red and yellow, respectively. The NFBP consists of responding with neuronal spiking to two of the feature combinations, corresponding to the relevant stimuli (red strawberry and yellow banana), and remaining silent for the other two feature combinations which represent the irrelevant stimuli (yellow strawberry and red banana). Assuming that each feature is represented with locally clustered synapses, a solution of the NFBP can be achieved when the co-active clusters on a single dendrite, representing the features of a relevant stimulus, evoke a plateau potential, thus supralinearly exciting the soma. Conversely, co-activation of synaptic clusters for the irrelevant combinations should not evoke plateau potentials.

B: Dendritic Learning: Illustration of how synaptic plasticity in SPNs may contribute to solving the NFBP for a pre-existing arrangement of synaptic clusters on two dendrites. A plasticity rule which strengthens only synaptic clusters representing relevant feature combinations, so that they produce robust supralinear responses, while weakening synapses activated by irrelevant feature combinations, could solve the NFBP.

C: Dopamine (Da) Feedback in Learning: dopaminergic feedback from the midbrain to the striatum (Str) guides the learning process, differentiating between positive feedback for relevant stimuli and negative feedback for irrelevant stimuli. Positive feedback represented by dopamine peaks is necessary for LTP, and negative feedback represented by a dopaminergic pause is necessary for LTD.

D: Signaling pathways underlying synaptic plasticity in dSPNs: Illustrations of molecular components at the corticostriatal synapse that modify synaptic strength (redrawn from Shen et al., 2009). NMDA calcium influx, followed by stimulation of D₁ dopamine receptors (D₁Rs), triggers LTP (while inhibiting the LTD cascade). L-type

calcium influx and activation of metabotropic glutamate receptors (mGluRs) when D₁Rs are free of D₂ triggers LTD (while counteracting the LTP cascade).

Results

Characterization of the dendritic NMDA-dependent nonlinearities in the model

The nonlinear, sigmoidal voltage sensitivity of the NMDA receptors is a crucial element for formation of dendritic plateau potentials. We use a model for generating plateau potentials first presented in Gao *et al.*, (2021) and adjusted to SPNs (Trpevski *et al.*, 2023). To produce robust, all-or-none plateau potentials, glutamate spillover from the synaptic cleft that activates extrasynaptic NMDA receptors is included in the model. Glutamate spillover occurs when the total synaptic weight (normalized value) of the activated synapses reaches a threshold value (see Methods). The threshold value here is set to be equivalent to the total weight of 10 clustered synapses with weights of 0.2 each (weight 0.2 corresponds to 0.5 nS). Figure 2A shows the somatic membrane potential following synaptic activation of a cluster of synapses of increasing size and the corresponding local spine membrane potential (averaged over all spines in the cluster) as well as the NMDA and L-type calcium accumulated in a single spine (also averaged over all spines in the cluster). A plateau potential is generated when a critical level of total NMDA conductance in a dendritic segment is reached (accomplished here by the addition of more synapses in a cluster and by glutamate spillover). Reaching the spillover threshold produces a sudden and robust increase in NMDA conductance, caused by the activation of extrasynaptic NMDA receptors (where the clearance of glutamate is assumed to be slower), thus triggering an all-or-none plateau potential as investigated in Trpevski *et al.*, (2023).

The synaptic input to the neuron is provided through the activation of a cluster of synapses at the indicated location in Fig. 2B, and gives the voltage and calcium responses in Fig. 2A. Figure 2C shows the maximal amplitudes of the somatic voltage and of the NMDA and L-type calcium signals in the dendritic spines, averaged over 10 trials and over 10 dendrites, and shown for three different synaptic weights in the clusters. The “baseline” results in Fig. 2C are within the range of the initial synaptic weights of excitatory synapses in all remaining figures in the article, and thus illustrate a possible initial situation before learning. Stronger and weaker synapses require a smaller and a larger cluster to trigger a plateau potential, respectively. In the simulation using “strengthened” synapses the synaptic weights in the cluster are 40% greater than in the “baseline” case, and hence need fewer synapses to trigger plateau potentials. Conversely, with weaker synapses where weights are 25% smaller than the “baseline” case, more synapses are needed to evoke a plateau.

To summarize, the dSPN model exhibits the dendritic nonlinearities required for solving the NFBP. If a cluster of strengthened synapses can reliably generate robust plateau potentials, this increases the likelihood for somatic spiking compared to when a more gradual NMDA dependent nonlinearity occurs in the dendrite. This also means that the neuron can reliably spike following the activation of a set of synapses strengthened (see illustration in Fig. 1). Conversely, a cluster of weakened synapses will most likely not generate plateau potentials,

and thus the neuron will spike with much lower probability following activation of such a cluster.




Figure 2: Characterization of dendritic plateau behavior in the model. A: Somatic voltage, spine voltage, NMDA calcium ($[Ca]_{NMDA}$), and L-type calcium ($[Ca]_{L-type}$) evoked by a cluster varying in size from 1 to 20 synapses. A plateau potential is evoked after a threshold level of NMDA conductance is exceeded, here set at 10 synapses with a weight of 0.2 each (corresponding to the "baseline" weights in C). The traces for spine voltage, $[Ca]_{NMDA}$, and $[Ca]_{L-type}$ are averages over all activated spines in the cluster.

B: Schematic of the neuron morphology with an arrow indicating the stimulated dendritic branch in A.

C: Maximal amplitude of the measures shown in A averaged over 10 dendritic locations. The curves represent clusters with different synaptic weights: baseline (0.2), strengthened (0.28) and weakened weights (0.15). Somatic voltages higher than the action potential threshold were set to -50 mV. A synaptic background noise is used in all simulations to elevate the membrane potential to ranges seen in vivo (Reig and Silberberg, 2014).

Characterization of the plasticity rule

To characterize the learning rule we start with a simple setup where three features (each illustrating either a color or a shape) are distributed onto two dendritic branches, so that one relevant and one irrelevant feature combination can be represented per dendrite (see Fig. 3A). Each feature is represented with 5 synapses, and we start with assuming that those synapses are already organized in pre-existing clusters. In addition to background synaptic noise inputs (compare Fig. 1), we also added 108 distributed glutamatergic synapses that were activated together with all four stimuli, i.e. they were feature-unspecific. SPNs have a very hyperpolarized resting potential, and the additional feature-unspecific synapses allow the neurons to spike often enough in the beginning of the learning process so that a dopamine feedback signal would be elicited and trigger learning in the activated synapses. The neuron model was then activated with a sequence of feature combinations (960 here) including equal amounts of relevant (i.e. 'red strawberry' and 'yellow banana') and irrelevant feature combinations (i.e. 'yellow strawberry' and 'red banana'). Thus, since the neuron always spikes initially due to the feature unspecific synapses, dopamine peaks and dopamine pauses arrive equally much. Fig. 3A shows an illustration of the setup. When the neuron spikes for the relevant feature combinations, dopamine rewards are delivered, triggering LTP in those active synapses with NMDA Ca levels within the LTP kernel, while spiking for the irrelevant feature combinations elicits a dopamine pause as feedback, instead triggering LTD as a function of L-type Ca. As learning progresses, the cell learns to respond only to the relevant feature combinations (see example in Fig. 3B). Initially all four stimuli, 'yellow banana' and 'red banana' in dendrite 1, and 'red strawberry' and 'yellow strawberry' in dendrite 2, elicited robust supralinear responses (as they together reached the threshold for glutamate spillover in our model). After learning the neuron could differentiate between the two sets of stimuli. The relevant feature combinations associated with a reward continued to provoke a plateau potential, eliciting somatic spiking. In contrast, the neuron's response to irrelevant feature combinations was notably decreased following LTD in the synapses on the dendrite where this feature is irrelevant.

Figure 3C shows the evolution of synaptic conductances during the learning process for dendrite 1 (dendrite 2 is not shown, but has the same behavior). The synapses representing the relevant feature combination in this dendrite ('yellow' and 'banana') are typically strengthened, eventually encoding this stimulus robustly. Conversely, the synapses for the feature 'red', activated during the irrelevant feature combination ('red banana') are all weakened, making the dendrite only weakly responsive to this stimulus following learning. Note that LTD also occurs in some of the synapses representing the features yellow as well as 'banana' as these features are also parts of the irrelevant stimuli (yellow strawberry and 'red banana', respectively). Nevertheless, the reward process still causes the remaining yellow synapses to strengthen, albeit to a lower level than the synapses for the feature 'banana' due to the LTD that has occurred (when combined with strawberry). This is because of the learning rule that successively slides the LTP NMDA Ca dependent plasticity kernel over training (see Methods). This means that our learning rule tends to stabilize the number of synapses that are needed to perform the task, but not necessarily all the synapses carrying the relevant features. Depending on the initial local Ca response in the synapse to relevant and irrelevant stimuli, individual synapses might either be preferentially recruited into the LTP or LTD process.

The NMDA calcium levels shown in Fig. 3D, left panel, shows examples on how the different synapses stabilize at their particular conductances/weights (for three example synapses marked with arrows in 3C). With repeated dopamine rewards, the metaplasticity kernel moves towards lower calcium levels, while a dopamine pause moves the kernel to higher calcium levels (i.e. it becomes easier for a synapse to strengthen for a slightly lower local calcium level if a reward has been seen often and also if the synapse is not being activated too often during reward omissions). Because the rate of metaplasticity kernel adaptation towards higher calcium levels is faster when the neuron makes errors (omitted rewards) on the NFBP, the kernels for all features move higher. This results in LTD in the 'red' synapses on the dendrite, while LTP develops in most of the 'yellow' and 'banana' synapses in dendrite 1. The upward shift in the kernel serves to prevent excessive strengthening in the weakened 'red' synapses on dendrite 1 when rewards for 'red strawberry' arrive (due to learning in dendrite 2), as they no longer fall within the optimal calcium range for LTP. In contrast to this, most synapses for 'yellow' and 'banana' are strengthened initially, but as their LTP kernels move downwards due to rewards and a sufficiently high NMDA calcium locally, no further LTP occurs. This means that in our learning rule, synapses tend to stop changing their weights when the neuron can perform the NFBP.

The feature-unspecific synapses exhibit different types of behaviors, which are primarily determined by the initial positioning of their calcium levels with respect to the optimal learning zone of the LTP plasticity kernel (Fig. 3C and 3D, right panels). The common trend of gradual weakening of these synapses, triggered by co-activation with irrelevant stimuli and subsequent dopaminergic pauses, is due to the metaplasticity kernel shifting away from the calcium level necessary for LTP. This shift increases the likelihood of LTD (initially during the training this occurs half of the time) This general weakening is exemplified by the purple trace in the right panel of Fig. 3D.

However, not all feature-unspecific synapses follow this trend. Some synapses, like the one depicted in blue, are situated close enough to the input clusters to experience increased NMDA Ca levels, enabling them to strengthen their synaptic connections. Finally, there are also a few feature-unspecific synapses that randomly happen to trigger calcium levels within the optimal learning region of the LTP plasticity kernel, causing enhanced LTP at the beginning (Fig. 3D, right panel, green trace). However, as the LTD amplitude is proportional to the L-type calcium level, the LTD-drive eventually becomes too high, pushing the calcium level below the LTP-kernel from which a period of LTD-only follows until synapses stabilize. To summarize, initially all synapses, clustered and feature-unspecific, may experience some LTP as dopamine rewards are delivered initially half of the time. However, as omitted rewards occurring for irrelevant stimuli cause the LTP metaplasticity kernel to slide higher, only clustered synapses signaling relevant feature combinations (and synapses close to them), where supralinear voltage effects cause high NMDA calcium, remain in the optimal learning region of the LTP kernel and a few undergo LTP before stabilizing. Note that learning is always "on", i.e. weight updates happen for each training example, meaning that there are no separate training and testing phases. The reason why synapse weights stop changing and stabilize is because the neuron performance improves so fewer and fewer omitted rewards occur (and this halts the LTD process) while at the same time the LTP kernel slides downwards when the neuron responds correctly most of the time and thus receives a reward (and this halts the LTP process).




Figure 3: Example of setup and learning-induced synaptic plasticity A: Illustration of input configuration. Upper panel shows the arrangement of the features in two dendrites. Each dendrite has synaptic clusters for three features allowing the representation of each of the four feature combinations if seen from the whole neuron's perspective. Also, the feature combinations allow for only one relevant feature combination per dendrite. The middle panel illustrates the stimulation protocol used in the simulation where stimuli presentation is followed by a dopaminergic feedback signal only if the neuron spikes. The two features representing a stimulus are active within 20 ms, and dopamine feedback, lasting for 50 ms, is delivered 300 ms after the beginning of the stimulus. Two stimuli are spaced 800 ms apart, to allow for the calcium dynamics to reach baseline levels. The bottom panel illustrates the stimulus sequence over the full learning task: all stimuli are equally present in a sequence of 12 stimuli.

B: Example voltage in the soma and the two dendrites before and after learning. Each dendrite stops responding to the irrelevant feature represented by its synaptic clusters.

C: Evolution of synaptic conductances throughout learning. The left panel shows the conductances of the clustered synapses in one of the dendrites (dendrite 1, d_1). B, Y and R stand for 'banana', 'yellow' and 'red', respectively. The right panel shows the distributed feature-unspecific synapses. The initial synaptic conductances are set to 0.25 ± 0.05 (around 0.625 nS). The purple trace exemplifies a synapse that is weakened, while the green and blue traces exemplify synapses that are close to the clusters (blue) or by chance (green) have a sufficiently high local NMDA calcium level for LTP to dominate.

D: Peak calcium (dots) and plasticity kernel dynamics (solid lines) during learning. The left panel shows the NMDA calcium of a single 'banana' and a single 'yellow' synapse that undergo LTP as well as for a single red feature synapse undergoing LTD in dendrite 1 (marked with arrows in the left panel in C). The right panel shows the NMDA calcium for the feature-unspecific synapses identified in the right panel of C.

Characterization of input combinations that can be learned

After showing that the SPN can learn to separate relevant from irrelevant stimuli in the single example in Fig. 3, we next generalized the setup by giving different innervations of the four features to the two dendrites and recorded the performance of the SPN on the NFBP as learning progressed. Out of all possible feature innervations to two dendrites, we used only the innervations containing both relevant feature combinations arranged at least one per dendrite, as these are the cases with enough feature innervation to possibly solve the NFBP (for an illustration see Fig. 4A). We then estimated how the performance developed over

training, i.e. whether the SPN spiked for the right feature combinations and was silent for the irrelevant ones. Performance of 100% indicates that the neuron spikes only for the relevant stimuli and is silent for the irrelevant stimuli. Performance of 50% can indicate two situations: i) either the neuron spikes for all four stimuli, or ii) is silent for all four stimuli. Performance of 75% indicates spiking behavior between the following two cases: i) the SPN spikes for only one of the relevant feature combinations, remaining silent for the other three, or ii) it is silent for one of the irrelevant feature combinations while spiking for the other three. In this study we consider the NFBP solved when the performance is greater than 87.5%, exemplified by the situation where the SPN always spikes for one relevant stimulus, and at least half of the time for the other relevant stimulus, while remaining silent for the irrelevant stimuli.

In the cases with two or three features innervating one dendrite (where one relevant feature combination is present per dendrite) the mean performance is above 90%, indicating that both stimuli are learned (Fig. 4B, purple traces). In the cases when all four features can innervate at least one of the dendrites, typically only one relevant feature combination (i.e. 'red strawberry' or 'yellow banana') could be correctly learned (Fig. 4B, light blue traces). In this case the same feature combination is usually encoded in both dendrites. In Fig. 4C the performance over the last 160 training examples of the simulation is shown, for each of the four stimuli separately. Fig. 4C₁ shows the performance for the purple traces in Fig. 4B, while Fig. 4C₂ shows the same for the light blue traces in Fig. 4B. This illustrates that for the purple traces in Fig. 4B, the NFBP is solved by the neuron, and the mistakes are usually made by firing for some of the irrelevant feature combinations. For the light blue traces in Fig. 4B, the NFBP is usually not solved, with somatic spiking for the relevant stimuli occurring only around 75% of the time. Additionally, spiking for the irrelevant stimuli is increased.

Optimal learning is achieved at intermediate distances from soma through excitatory plasticity

We also investigated the impact of synaptic positioning on learning (Fig. 4 D) when using the same settings as in Fig. 3, but varying dendritic locations. Our results predict that the best performance on the NFBP is obtained with synaptic clusters positioned at intermediate somatic distances from the soma (Fig. 4E). From Fig. 4E one can infer that after learning the proximal synapses have actually decreased when trained on the NFBP as the performance is around 50%. The neuron stays silent for the irrelevant feature combinations (blue dots), which is the correct response, but also the neuron can't respond to the relevant feature combinations (red dots). For successively more distal synapses the performance increases and then slightly decreases for the most distal clusters that sometimes fail to evoke somatic spiking for the correct feature combinations (red dots). This result can be conceptually explained in the following way. The electrotonic properties of dendrites dictate that synapses near the soma, in the most proximal regions, are less capable of inducing supralinear potentiation underlying plateau potentials (Du *et al.*, 2017). This is due to the soma acting as a current sink, resulting in smaller localized voltage changes (and hence a lower input resistance in accordance with Ohm's law). Consequently, these synapses cannot easily evoke dendritic nonlinearities necessary for solving the NFBP, and hence the performance with proximal clusters is low for NFBP. Note that in our simulations we allow glutamatergic synapses on spines quite close to the soma, although SPN dendritic spines are relatively rare at more proximal distances than 40-50 μm from the soma (Wilson *et al.*, 1983).

In contrast, the most distal dendritic regions are electrically more isolated and have a higher local input resistance, enabling larger voltage changes locally and thus also higher local calcium concentrations when synapses are activated in our model. This allows even a small number of active synapses to generate local supralinear NMDA-dependent responses. Such ease of elevating the local calcium, seems advantageous, but in fact results in decreased performance on the NFBP for the following reasons. In the distal synaptic clusters related to irrelevant stimuli, local calcium levels might stay in the optimal learning region of the LTP kernel despite omitted rewards occurring regularly and thus the weakening of the synapses takes longer. This causes increased somatic stimulation for irrelevant stimuli, sometimes leading to spiking over a longer period during the training. In addition, distally evoked dendritic signals, including plateau potentials from strengthened synaptic clusters, naturally attenuate more before they reach the soma, and sometimes fail to elicit somatic spiking. As a consequence, compared to more optimally placed clusters, the proportion of negative feedback signals is increased and that of positive feedback is decreased. The altered feedback results in a decreased capacity to selectively reinforce the active synapses to only relevant stimuli and weaken synapses activated with irrelevant stimuli.

The ideal learning zone for NFBP in our simulations thus lies at an intermediate somatic distance, roughly 100-150 μm from the soma, where synapses can effectively contribute to learning the NFBP (Fig. 4 E). In this zone synaptic changes are more likely to impact the neuron's firing probability as the dendritic plateau potential at this location causes a larger elevation of the somatic potential, and thus synapses at this distance benefit more from the dopamine feedback loop. Note that the prediction that proximal and very distal synapses are less likely to contribute to the solving of the NFBP doesn't imply they are not important for more 'linear' learning contexts. For instance, if we had trained the neuron to only respond to one single stimulus, such as 'red strawberry', both proximal and very distal synapses representing that stimulus would of course be able to both strengthen or weaken as well as contribute to spiking of the neuron following learning.




Figure 4: Impact of feature combinations and synaptic cluster locations on NFBP learning performance
A-C: NFBP configurational analysis. Configurations are categorized based on the number of features per dendrite, those with two or three features (purple traces) and those with all four features on at least one dendrite (light blue traces).
A: Illustration of the setup of the task where two, three or four features are given in two dendritic locations.
B: Performance trajectories for the combinatorial task illustrated in A. Traces of the performance of the model over time for 31 unique feature configurations. Light blue lines show combinations with four features in at least one dendrite, while purple lines show input with maximally 3 input in the local dendritic branches. Right-side histograms display the distribution of the end performance.
C: Outcome results split on stimuli for configurations where maximally three features in one dendritic location (C_1) or at least one dendrite has all four features present in a single dendrite (C_2).
D and E: End performance in a three-feature configuration as a function of cluster location. D illustrates the location of the inputs, and E shows end performance in the three-pattern configuration as a function of somatic distance of the synapse clusters. Initial synaptic weight is 0.25 ± 0.05 in all the simulation experiments.

The possible role of inhibitory plasticity in learning

In our initial simulations we assumed that only the excitatory synapses could undergo plasticity during learning, and we identified two critical observations that highlight possible areas for improvement. The first is a vulnerability to noise, which resulted in a performance around 90 percent, as illustrated in Fig. 4B. The second is decrease in performance observed across very distal synapses, as detailed in Fig. 4E. These findings prompted us to explore the potential role of inhibitory synapses in addressing these challenges.

We therefore developed a phenomenological inhibitory plasticity rule to enhance dendritic nonlinearities (detailed in the Methods section). The rule is designed to compartmentalize a dendrite so that it mostly responds to excitatory inputs that cause the strongest dendritic activation. It achieves this by weakening inhibitory connections at highly active excitatory synapses based on the local increase in voltage-dependent calcium, reinforcing their dominance, and strengthening them where excitatory activity is lower (and thus as well local calcium concentration, see Methods for details).

To demonstrate our inhibitory plasticity rule, we use the same excitatory synapse setup as in Fig. 4 to which we add four inhibitory synapses near each cluster in the middle of the dendritic branch, representing each of the four features (Fig. 5A). Thus, a single feature activates both the excitatory and inhibitory synapses. To achieve a level of depolarization and spike probability comparable to that in our excitatory-only setup, we increased the distributed synaptic input from 108 to 144. Alongside this, we began with low inhibitory synaptic weights. This was key to maintaining higher baseline activity in our model as starting with strong inhibitory weights could excessively suppress excitatory activity as e.g. inhibitory inputs close to clustered synapses effectively can counteract the NMDA-dependent nonlinearities (Doron *et al.*, 2017; Du *et al.*, 2017; Dorman, Jędrzejewska-Szmek and Blackwell, 2018). This setup ensured that we could clearly observe how learning modified synaptic connections without initial excessive inhibition preventing the dendritic nonlinearities initially.

As in the example in Fig. 3, the conductances of the excitatory inputs representing 'yellow' and 'banana' increase to cluster the 'yellow banana' pairing, while the weights of synapses representing the feature 'red' decrease (Fig. 5D, left panel). Conversely, the inhibitory synapses associated with the 'yellow' and 'banana' features, linked to excitatory synapses, are instead further weakened, while those linked to 'red' and 'strawberry' features in the same compartment are strengthened (Fig. 5D, right panel). This outcome of the inhibitory plasticity rule effectively prevents the 'red banana' and 'yellow strawberry' stimuli from triggering spikes due to the excitatory inputs to this particular dendrite, thereby compartmentalizing dendrite 1 to be responsive to mainly 'yellow banana'.

We also show the dynamics of peak calcium levels associated with both excitatory and inhibitory synapses for each task (Fig. 5E). For excitatory synapses, for which strengthening NMDA calcium influx is a key factor, the patterns in peak calcium levels behave as in the example in Fig. 3D (Fig. 5E, right panel). On the inhibitory side, the peak calcium levels, influenced by voltage-gated calcium channels as shown in the right panel of Fig. 5E, displayed different behavior, resulting from the inhibitory plasticity rule. The threshold levels of calcium for inhibitory synapses followed the highest excitatory synaptic activity, which in our task corresponded to the 'yellow banana' input. Conversely, the minimum threshold level was able to surpass all other synaptic activities, effectively designating it for 'yellow banana' only.

We also compared the performance for different feature configurations with added inhibitory synapses to the results for the excitatory-only setup from Fig. 4B. The results show not only that learning with inhibitory plasticity is faster, i.e. requires less training examples, but also achieves high performance, nearing 100%. This demonstrates the potential impact of

inhibitory synapses in enhancing learning capabilities, as well as in performing the NFBP computation (Fig. 5B).

Inhibitory plasticity additionally produced a marked improvement in performance when the location of the synaptic clusters was varied, especially prominent at very distal locations (Fig. 5C). Incorporating inhibitory plasticity allows for more rapid, robust, and accurate learning. Unlike excitatory plasticity, which relies on dopaminergic feedback signals, our inhibitory plasticity model follows a passive rule that just aligns with the local excitatory activity. It reinforces the most active excitatory synapses within a dendritic branch by decreasing the inhibitory synapses corresponding to the same features there, and conversely, it strengthens inhibitory synapses for features for which excitatory activity is less prominent. This 'winner-takes-all' strategy not only accelerated the learning process by requiring fewer training examples but also enhanced performance consistency across different dendritic locations, including the distal regions where excitatory-only models faltered to some extent when challenged with the NFBP. The inhibitory rule's ability to function without dopaminergic feedback streamlines the learning process, leading to higher accuracy and a more stable and robust set of synaptic modifications, regardless of the synaptic cluster's position along the dendrite. This demonstrates the influence inhibitory synapses could potentially exert in fine-tuning dendritic responsiveness and refining the neuronal circuitry critical for learning.




Figure 5: Effects of inhibitory inputs on performance

A: Dendritic input configuration with inhibitory synapses added. Illustrations depict two dendritic branches, each with synaptic connections from three excitatory and four inhibitory features, but the general setup is the same as in Fig. 4.

B: Displays average performance for configurations with varying pattern configurations as a comparison between the setup with (orange) and without inhibitory plasticity (blue).

C: Shows task-specific performances for dendritic locations of the clustered synapses, with individual dots and curves representing fitted performance curves with (solid) and without plastic inhibition (dashed).

D: Synaptic conductance changes during learning. Left panel shows excitatory synaptic conductances in dendrite 1 during learning. The right panel shows the inhibitory synaptic conductances. B, Y, R and S stands for 'banana', 'yellow', 'red' and 'strawberry'.

E: Peak calcium (dots) and plasticity thresholds dynamics (lines) over the learning. The left panel shows examples of the postsynaptic calcium amplitude of the clustered spines in dendrite 1 (marked with arrows on the left panel in D). The right panel tracks calcium levels at inhibitory synapses. The Upper threshold (T_{max}) captures peak calcium levels while the lower threshold (T_{min}) identifies the next highest. Stability achieved at both thresholds enhances contrast between the activity levels and makes the dendrites respond preferentially to one relevant feature combination at that dendritic site.

Dendritic Branch Plasticity and Synaptic Learning – Investigating the Role of Randomly Distributing Excitatory and Inhibitory Synapses on the Dendrites

We finally challenged our plasticity rule by relaxing the assumption that single features are represented by pre-clustered synapses on specific dendritic branches. The synaptic plasticity rule was therefore investigated by randomly distributing 200 excitatory synapses signaling the different features across 30 dendrites. Each feature was represented by 40 excitatory synapses and an additional 40 feature-unspecific excitatory synapses were used. Figure 6A illustrates the setup and exemplifies the pre- and post-learning synaptic weights for both excitatory and inhibitory synapses. Our objective was to examine the learning dynamics in the absence of assumed synaptic clustering and to determine the capability of the single neuron to learn the NFBP. We initiated the experiment with synaptic conductance set to 0.3 ± 0.1 (around 0.75 nS), focusing first on purely excitatory synapses. This higher initial weight and increased variance, compared to the initial weights of 0.25 ± 0.05 (around 0.75 nS) with clustered synapses, were chosen to compensate for the reduced efficacy of

non-clustered synaptic inputs in producing sufficient depolarization and calcium influx, essential for effective learning and synaptic plasticity.

Figure 6B (right panel) shows limited learning capacity on the NFBP, which corresponded with our hypothesis: that without clustering-induced supralinearities due to spillover in the dendrites, distributed synapses would be too weak to trigger high enough differences in local calcium levels for separating the correct and incorrect feature combinations. To confirm that the lack of dendritic nonlinearities is the cause for not learning the NFBP, we extended this concept by introducing a hypothetical branch-specific spillover mechanism assuming that spillover occurs if synapses are on the same dendritic branch. This hypothetical construct assumes that concurrent activation of a critical mass of co-activated synapses in a single dendritic branch would trigger both plateau potentials and a significant rise in calcium levels despite that the synapses are not spatially pre-clustered, endowing the dendrites with supralinear responses even without closely clustered synapses. We also tested a more gradual spillover model proposing that each synapse could contribute incrementally to a cumulative spillover effect, instead of using a sharp glutamate threshold (see Trpevski et al, 2023). This modification provided a reduced level of nonlinearity compared to the branch-specific spillover model but still led to an increase in learning effectiveness. As shown in Fig. 6B, adding the hypothetical nonlinearities to the model increases the performance towards solving part of the NFBP, i.e. learning to respond to one relevant feature combination only. The performance increases with the amount of nonlinearity.

We next extended our investigation by including 60 inhibitory synapses, 15 for each feature, dispersed randomly over the 30 dendrites. Performance improved with the introduction of these additional nonlinearities, especially in combination with the branch-specific spillover mechanisms. This finding emphasizes the importance of dendritic branch plasticity and calcium nonlinearity in coordinating synaptic modifications for both excitatory and inhibitory inputs. In Fig. 6C, we report that a subset of neurons (5 out of 31) successfully solved the NFBP (reaching a performance of at least 87.5%).

For a more granular analysis of the successful cases, we selected an example that learned the NFBP. In Fig. 6D, we show the somatic and dendritic voltages for four dendritic branches where we found successful encodings of the relevant stimuli. The voltage traces show the somatic and dendritic responses to all four stimuli after learning. For each dendrite we also estimated the cumulative synaptic conductances for both excitatory and inhibitory synapses at the dendritic branch midpoint. Notably, in dendrites numbered 1 and 4, we observed an enhancement of excitatory inputs for the 'yellow banana' pattern and inhibitory inputs for the 'red strawberry' pattern. Conversely, dendrites 2 and 3 displayed the opposite arrangement.

These results show that for an adequate random innervation of distributed synapses, where the necessary features for a relevant stimulus innervate the same dendritic branch with enough synapses, that stimulus can be stored on that dendrite. In this way, the NFBP can also be learned if the two relevant stimuli are encoded on different dendrites and each of them can trigger a supralinear dendritic response. Since on average the two features representing a relevant stimulus do not innervate a single dendrite with enough synapses, the stimulus is not stored in a single dendrite, but is distributed across the dendritic tree. And since distributed synapses summate more linearly at the soma, only one of the relevant stimuli can be typically encoded by the dendritic tree.




Figure 6: Performance analysis of learning using distributed synaptic inputs

A: Example illustration of synaptic distribution before (top) and after learning (bottom) of the 200 excitatory and 60 inhibitory inputs.

B: Learning performance with two spillover models, branch-specific thresholded and accumulative, without (left) and with plasticity of inhibitory synapses (right).

C: Performance trends of 31 distributions with branch-specific thresholded spillover and inhibitory plasticity.

D: Example of summed synaptic conductances (left) and voltage (right) in the soma, and four example dendrites (d1-d4) of one model following successful learning of the NFBP. The sums of both excitatory (Ex) and inhibitory (Inh) inputs are shown.

Discussion

In this article we studied whether single neurons can solve linearly non-separable computational tasks, represented by the NFBP, by using a biophysically detailed multicompartiment dSPN model. Based on the synaptic machinery of corticostriatal synapses onto dSPNs, we propose a learning rule that uses local synaptic calcium concentration and dopamine feedback signals: rewards for relevant stimuli and omitted rewards for irrelevant stimuli. Assuming first that single features in the NFBP are represented by clustered synapses, we show that the learning rule can solve the NFBP by strengthening (or stabilizing) synaptic clusters for relevant stimuli and weakening clusters for the irrelevant stimuli. The feature combinations for the relevant stimuli, stored in strengthened synaptic clusters, trigger supralinear dendritic responses in the form of plateau potentials, which is an important ingredient for the solution of the NFBP, as plateaus significantly increase the likelihood of neuronal spiking in SPNs in a robust way (Du *et al.*, 2017).

The location of the synaptic clusters along the dendrites influenced the performance on the NFBP. In our model the region for optimal performance was 100-150 μm away from the soma, at about the same distance as where the somatic depolarization and induced spike probability, following activation of clustered input, is largest in an older version of the model (Lindroos and Hellgren Kotaleski, 2021). Clusters placed further away produce smaller somatic depolarizations, due to dendritic filtering (Major *et al.*, 2008), and as a consequence does not control the likelihood of somatic spiking as decisively. As the supralinear response is necessary for discriminating the relevant from the irrelevant stimuli, the performance with distally placed clusters decreases somewhat.

We further verified that supralinear dendritic responses were necessary to solve the NFBP by using randomly distributed synapses instead of clustered for each feature. In this scenario, only one relevant stimulus is sometimes learned in a dendritic branch by the randomly distributed synapses. The random setup was tested with two glutamate spillover models: The first model assumed that thresholded glutamate spillover is branch-specific, building on the notion that the single branch act as a single computational unit (Losonczy and Magee, 2006; Branco and Häusser, 2010). The second assumed that spillover is accumulative (Trpevski *et al.*, 2023). Both spillover versions increase the performance of our neuron model, further certifying the necessity of supralinear responses for solving the NFBP. In the real dendritic branches, however, also diffusion of signaling molecules within the branch likely contributes significantly both to the plasticity in already existing synapses as well as to local structural plasticity, both of which could increase branch specific supralinearities.

By using a phenomenological inhibitory plasticity rule based on the BCM formalism (Bienenstock, Cooper and Munro, 1982), we also show that inhibitory synapses can significantly improve performance on the NFBP. This is in line with earlier theoretical studies where negative synaptic weights were required to solve the NFBP (Schiess, Urbanczik and Senn, 2016). In our setup with pre-existing synaptic clusters, inhibitory synapses made learning faster and increased performance by inhibiting supralinear NMDA responses for the irrelevant stimuli. This was specifically true in distal dendrites where the input impedance is higher (Branco, Clark and Häusser, 2010). The threshold for plateau initiation is also lower in distal dendrites compared to proximal (Losonczy and Magee, 2006) which likely will further extend the influence of inhibition in this region (Doron *et al.*, 2017; Du *et al.*, 2017). Similarly, in the scenario with distributed synapses, inhibition enables one of the relevant stimuli to be reliably encoded in the dendritic branch by strengthening the inhibitory synapses for features different from those of the encoded stimulus. Together, it therefore seems like inhibition not only has a role in learning (Chen *et al.*, 2015; Cichon and Gan, 2015), but also improves the ability of the neuron to discriminate between stimuli with shared features.

Although our learning rule only occasionally solves the NFBP when used with randomly distributed synapses, it can always learn to perform a linearly separable task, such as learning to respond to only one relevant stimulus (such as red strawberry). Moreover, the learning rule is general enough so that in addition to the feature-specific inputs related to the task, it can handle feature-unspecific inputs that might or might not be related to the NFBP. Finally, the learning rule is always “on”, continuously updating synapses with each stimulus presentation, which is a more realistic mechanism compared to using separate training and

testing phases as in the field of machine learning. The synaptic weights automatically stabilize in the model when the performance improves. That is, rewards are then seen very regularly as the neuron has learned to spike for the relevant stimuli, while omitted rewards seldomly occur as the neuron stays silent when the irrelevant stimuli are provided.

When formulating the learning rule in this article, our goal was to base it on what is known regarding the synaptic machinery in corticostriatal synapses. This implied that the learning rule is based on the local calcium activity and on dopamine signals. Feedback from the dopamine system can also be viewed as an innate, evolutionarily encoded “supervisor”, which instructs neurons which feature combinations are beneficial and which ones should be avoided. However, in our case we do not use additional excitation to promote somatic spiking for only the relevant feature combinations, and in that sense the learning rule does not require a supervised learning paradigm. Since the SPNs rest at very hyperpolarized membrane potential, our setup includes distributed excitatory inputs which are feature-unspecific in order to make sure that the neuron spikes for all stimuli, especially at the beginning of training. These additional inputs are on average weakened as learning progresses (as they are activated for all stimuli and thus often receive negative feedback). That general or noisy inputs are reduced during learning is in line with the observed reduction of execution-variability during motor-learning as a novice becomes an expert (Kawai *et al.*, 2015). Also, the underlying neuronal representation of corticostriatal synapses undergo a similar change during learning (Santos *et al.*, 2015).

Since each synapse has its own calcium response, it is important for the learning rule to be able to follow individual synaptic activities. The LTP plasticity kernel in our model is for this reason itself plastic, meaning that it changes its calcium dependence as a function of the history of reward and punishment. This setup helps the model separate clustered synaptic input from the feature-unspecific input at the beginning of training, as only clustered synapses will see enough calcium to fall within the plastic range. We further use an asymmetric metaplasticity rule, where negative feedback causes a larger shift of the LTP kernel than a positive. This was necessary in order to stop LTP in synapses participating in LTD. Similarly to the classical loss-aversion tendency described in economic decision theory (Kahneman and Tversky, 1979), the model hence predicts that negative feedback will have a bigger impact in changing a well learned behavior on the single cell level than a positive feedback. Dopamine signaling has also been linked as a neural substrate to the decision making theory mentioned above (Stauffer *et al.*, 2016).

It is not known whether single neurons solve the NFBP or other linearly non-separable tasks. However, many brain nuclei receive convergent inputs from numerous other brain nuclei, acting as integratory hubs (van den Heuvel and Sporns, 2013). Since feature binding evidently occurs in the brain, and functional clusters for single features such as visual stimulus orientation, receptive fields, color, or sound intensity exist on single neurons (Chen *et al.*, 2011; Wilson *et al.*, 2016; Iacuruso, Gasler and Hofer, 2017; Scholl, Wilson and Fitzpatrick, 2017; Ju *et al.*, 2020), it is possible that the NFBP is a relevant task for neurons to solve. How brain regions with different synaptic machinery than the striatal dSPN might solve the NFBP remains a question, and reliance on other neuromodulatory signals may be part of the answer. For example, in the striatum, the indirect pathway SPNs (iSPN) have analogous synaptic machinery to the one in dSPNs, requiring calcium influx from the same sources for LTP and LTD, but are differently responsive to dopamine (Shen *et al.*, 2008). In

iSPNs a dopaminergic pause, together with a peak in adenosine, is required to trigger LTP, whereas a dopamine peak without peaks in adenosine rather promotes LTD (Shen *et al.*, 2008; Nair *et al.*, 2015). Therefore, we expect that an analogously formulated learning rule will also solve the NFBP in iSPNs, activating them for irrelevant feature combinations to e.g. suppress movement, and suppressing their activity for relevant feature combinations to facilitate movement. In addition, LTP in dSPNs might require co-activation of other neuromodulatory systems, such as a coincident acetylcholine pause with the dopamine peak, which we have not explicitly included in the model (Nair *et al.*, 2015; Bruce *et al.*, 2019; Reynolds *et al.*, 2022).

Method

In this paper we introduce a local, calcium- and reward-based synaptic learning rule, constrained by experimental findings, to investigate learning in SPNs. The learning rule operates based on the local calcium concentration in conjunction with plateau potentials and other dendritic nonlinearities and shows that such events enable learning of the nonlinear feature binding problem (NFBP). The learning rule is embedded in a biophysically detailed model of a dSPN built and simulated in the NEURON software, v8.2 (Carnevale *et al.*, 2006). Here we will focus on the setup of the learning rule and only give a short summary of the neuron and synapse models, focusing on changes compared to previously published versions. For a detailed description of the neuron model setup, see Lindroos *et al.*, (2018), Lindroos & Hellgren Kotaleski, (2021), and Trpevski *et al.*, (2023).

Neuron model

In short, the dSPN model used here was sourced from a collection of biophysically detailed models, including a reconstructed morphology and all of the most influential ion channels, including six calcium channels, each with its own voltage dependence and dendritic distribution (Lindroos and Hellgren Kotaleski, 2021). In accordance with Trpevski *et al.*, (2023), the model was further extended with synaptic spines on selected dendrites. Each spine was modeled as two additional compartments consisting of a neck and a head region, and contains voltage-gated calcium channels of types R ($\text{Ca}_v2.3$), T ($\text{Ca}_v3.2$ and $\text{Ca}_v3.3$), and L ($\text{Ca}_v1.2$ and $\text{Ca}_v1.3$); the addition of explicit spines did not change the basic behavior of the model, such as the response to current injections, etc.

Calcium sources used in learning

The intracellular calcium concentration is separated into distinct pools that are used during the learning process. For learning in glutamatergic synapses, one pool for NMDA-evoked calcium concentration ($[\text{Ca}]_{\text{NMDA}}$) is used, and another for L-type calcium concentration ($[\text{Ca}]_{\text{L-type}}$), to reflect the different synaptic plasticity responses of the SPN's biochemical machinery to these two calcium sources in the corticostriatal synapse (Shen *et al.*, 2008; Fino *et al.*, 2010; Plotkin *et al.*, 2013). Both pools are based on the calcium influx from the corresponding source (NMDA and the L-type channels $\text{Ca}_v1.2$ and $\text{Ca}_v1.3$, respectively). Similarly a third pool, used for inhibitory plasticity, collects the current from all voltage dependent Ca channels (T-, R-, L- and N-type, $[\text{Ca}]_v$). All pools include extrusion mechanisms in the form of a calcium pump as well as a one-dimensional time-decay. The calcium pump follows the implementation in Wolf *et al.*, (2005). The parameters for the

[Ca]_{NMDA} model were manually tuned to match the quantities reported in Dorman et al., (2018). The voltage gated calcium channel conductances in the spines were also manually tuned to match the relative calcium proportions in Carter and Sabatini, (2004) and Higley and Sabatini, (2010), as well as the calcium amplitudes due to stimulation with backpropagating action potentials (Shindou, Ochi-Shindou and Wickens, 2011). Spatial diffusion was not included in the present model.

Glutamatergic synaptic input

In our study, we used a synaptic model that includes both AMPA and NMDA conductances activated on spines. Additionally, we considered extrasynaptic NMDA conductances on dendritic shafts directly under the spines to account for glutamate spillover if the synapses were sufficiently stimulated and spatially clustered. Spillover has been shown to be important for generating robust all-or-none dendritic plateaus (Trpevski *et al.*, 2023).

The conductance of extrasynaptic NMDARs during spillover matches their synaptic counterparts. Glutamate spillover in a set of clustered synapses occurs in the model if a glutamate threshold is met, and in that case a presynaptic spike also activates the corresponding extrasynaptic NMDA conductance (Trpevski *et al.*, 2023). The threshold mechanism involves summing the synaptic weights for each NMDA synapse activated in the cluster, and checking whether their sum is greater than a threshold level. The synaptic weight is represented by a weight parameter (w) which scales the maximal conductance of the synapse (the change of the weights are described by Eq. 1-3). We exemplify this with the equation for calculating the NMDA current I_{NMDA} :

$$I_{NMDA}(t) = g_{syn}(t)(V(t) - E_{NMDA}); g_{syn}(t) = wg_{max}g(t)$$

where $V(t)$ is the membrane voltage, E_{NMDA} is the reversal potential of the NMDA synapse, g_{syn} is the NMDA conductance which has a time varying part, ranging from 0 to 1, scaled by a fixed part comprised of the maximal NMDA conductance and the synaptic weight. In our simulations a cluster needs to reach a summed weight of 2 to reach the glutamate threshold for spillover, at which point the extrasynaptic NMDARs are activated. This is modeled in NEURON by summing the NMDA synaptic weights upon each spike arrival, with the sum normalized so that e.g. ten synapses of 0.2 weight reach the threshold. An integrate-and-fire cell in NEURON generates a spike to extrasynaptic NMDARs when this threshold is met. This model allows for fewer synapses to reach the threshold as synaptic weight increases, a phenomenon linked to LTP and enhanced glutamate spillover due to astroglial process withdrawal (Henneberger *et al.*, 2020).

Extrasynaptic NMDA synapses were not included for feature-unspecific (non-clustered) synapses although we explored the role of a potential branch-specific spillover mechanism, where non-clustered synapses localized on the same branch still could interact through spillover mechanisms, either assuming thresholded or accumulative spillover as detailed in Trpevski *et al.*, (2023). The AMPA and NMDA synapse models were taken from Gao *et al.*, (2021) and are a variation of the saturating synapse models in Destexhe *et al.*, (1994). See Trpevski *et al.*, (2023) for a detailed description of how the synaptic input was set for the SPN model.

Learning rule

Excitatory synaptic plasticity

The learning rule is based on experimental findings showing that striatal LTP depends on NMDA-channel activation and the presence of dopamine, while LTD is dependent on activation of the L-type Ca channel $Ca_v1.3$ and the mGluR5 receptor in the absence of dopamine (low dopaminergic tone) (Shen *et al.*, 2008; Fino *et al.*, 2010; Plotkin *et al.*, 2013; Yagishita *et al.*, 2014; Fisher *et al.*, 2017; Shindou *et al.*, 2019). Under basal dopamine levels, no significant synaptic plasticity is assumed to occur. The rule describes a reward-based learning scheme in which a dopamine peak increases the synaptic weight based on one function (plasticity kernel) while a dopamine pause decreases the weights based on another function. Hence, dopamine acts as a switch selecting which of the LTP or LTD pathways should be activated (as illustrated in Fig. 1D). The plasticity kernels for LTP and LTD depend on NMDA calcium and L-type calcium, respectively, and are described below.

LTP

The LTP process is triggered by increased dopamine levels and is based on calcium influx through the NMDA channel. Here the synaptic strength is updated based on a sliding bell-shaped kernel in such a way that the maximal increase is obtained for peak calcium levels close to the kernel midpoint, while higher or lower peak calcium levels result in a smaller increase. This sets an upper limit on the synaptic strength and guarantees that the weights will not grow beyond a certain limit (Zenke and Gerstner, 2017; Zenke, Gerstner and Ganguli, 2017). Since each synapse has its own calcium level, the sliding of the bell-shaped kernel enables the precise tuning of each synaptic weight separately (see Fig. 7A for an illustration).

Mathematically, the bell-shaped kernel for the LTP process is represented by the derivative of the sigmoid function σ (Eqs. 1 and 2) with a slope that can be adjusted through the β parameter. The LTP learning rule (Eq. 3) describes the increase in synaptic weight (w) of the inputs, as a function of peak calcium level ($[Ca]_{NMDA}$), the learning rate (η_1), and the midpoint (θ_{ltp}) and slope of the sigmoid curve (β_1).

$$\sigma(Ca, \theta, \beta) = \frac{1}{1 + e^{-\beta(Ca - \theta)}} \quad \text{Eq. 1}$$

$$\sigma'(Ca, \theta, \beta) = \beta \cdot \sigma(Ca, \theta, \beta) (1 - \sigma(Ca, \theta, \beta)) \quad \text{Eq. 2}$$

$$dw_{ltp} = \eta_1 \cdot \sigma'(Ca_{NMDA}, \theta_{ltp}, \beta_1) \quad \text{Eq. 3}$$

LTD

The LTD process is triggered by a dopamine pause and is dependent on L-type calcium. The LTD plasticity rule (Eq. 4), describes a threshold level of the calcium level necessary for LTD

to occur (Shindou, Ochi-Shindou and Wickens, 2011), after which the decrease in synaptic weight is linearly proportional to the amplitude of the peak calcium level, and is scaled by the learning rate (η_2 , see left panel of Fig. 7B for an illustration). The peak calcium threshold is implemented with a sigmoid function, whose slope parameter, β_2 , was set to a high-value to make the curve resemble a step function. The midpoint, θ_{ltp} , of the LTP kernel was also increased during LTD, as described in the next section (Fig. 7B right panel).

$$dw_{ltd} = \eta_2 \cdot Ca_{v1.3} \cdot \sigma(Ca_{v1.3}, \theta_{ltd}, \beta_2) \cdot w_{ltd} \quad \text{Eq. 4}$$

Metaplasticity

Metaplasticity is a form of regulatory mechanism changing the state of the synapses in such a way as to influence subsequent learning. In our model we implemented this as a reward-history-dependent change of the location of the calcium range over which the LTP kernel was operating - a pause in dopamine triggered a shift towards higher calcium concentrations while a peak in dopamine pushed the LTP kernel in the opposite direction. The metaplasticity was also implemented using a kernel with the same midpoint as the LTP kernel, but with a wider calcium dependence. Together the setup with the two dynamically moving kernels allowed for a wide range of calcium levels to induce plasticity. This setup typically gave rise to a few characteristic situations:

1. For peaks in dopamine (see Fig. 7C for an illustration),
 - a. In synapses with low calcium levels, the kernel will be shifted closer to the observed level and thereby eventually enable LTP in synapses that are regularly activated during rewards, despite that they initially don't generate big changes in calcium.
 - b. In synapses with already high calcium levels, the kernel will be shifted away from the observed level, and thereby protect the neuron from excessive LTP and instead stabilize the weight (as illustrated with the red circle moving to the blue one in Fig. 7C).
2. For dopaminergic pauses, the kernel will be shifted away from the observed level and thereby reducing the likelihood of LTP in synapses that are undergoing LTD (see Fig. 7D for an illustration).

The LTP kernel update was further asymmetric with regards to dopamine peaks and pauses in such a way that a pause caused a larger shift of the kernel than a peak. This further reduced the likelihood of inducing LTP in synapses often participating in LTD or in synapses randomly activated with regard to the dopamine feedback signal. The metaplasticity kernel thus adjusts the LTP kernel for these processes based on ongoing neural activity. These changes are described by Eq. 5, where the LTP kernel midpoint is updated based on the learning rate (η_s), the peak NMDA calcium level ($[Ca]_{NMDA}$), and the slope of the sigmoid curve (β_3). The learning rate, η_s , captures the described asymmetry in response to dopamine feedback (see Table 1). This setup can give rise to the following situations:

1. For randomly alternating dopamine peaks and pauses,
 - a. In a few synapses that initially happen to generate very high calcium levels the synapses can become even stronger as their postsynaptic calcium response is still within the shifted LTP kernel. This explains e.g. the increase in strengths of the synapses illustrated with the blue and green traces in the

right panel of Fig. 3C. (However, as the computational performance of the neuron successively improves these synapses stabilize when the neuron starts to receive less dopamine pauses).

- b. In synapses that generate somewhat less postsynaptic calcium responses the general trend is that they decrease as their calcium level becomes lower than the LTP kernel. This is illustrated in Fig. 7D with the red circle shifting to the blue one; also e.g. compare the purple trace in the right panel of Fig. 3C.

In summary, the slower shifting metaplasticity kernel, compared to the weight update, allows the individual synapses to stabilize when the neuron has learned a specific task and thus doesn't receive negative reward feedback that often. Also, the rule allows initially weaker synapses to be recruited to undergo LTP if they are more consistently co-active with a reward than with an omitted reward (an illustration of this can be seen in Fig. 5D, where one of the yellow synapse traces jumps upwards during the middle of the training session) .

$$d\theta_{ltp} = \eta_s \cdot \sigma'(Ca_{NMDA}, \theta_{ltp}, \beta_3) \quad \text{Eq. 5}$$

Inhibitory synaptic plasticity

In contrast to the well studied mechanistic underpinnings of glutamatergic plasticity in e.g. SPNs, much less is known about how inhibitory synapses might be updated during learning. The inhibitory plasticity rule developed here is therefore more phenomenological and exploratory in nature and was developed to enhance nonlinearities in the local dendrite. The rule is based on the Bienenstock-Cooper-Munro (BCM) formalism (Bienenstock, Cooper and Munro, 1982). In the BCM rule there is a threshold level of synaptic activity below which LTD is triggered and above which LTP is triggered (Fig. 7E). Inspired by (Gandolfi *et al.*, 2020; Ravasenga *et al.*, 2022) and, we use calcium from all voltage-gated calcium channels ($[Ca]_V$) as the indicator of excitatory synaptic activity near the dendritic shaft where an inhibitory synapse is located. The inhibitory rule is designed to passively observe and respond to the surrounding excitatory synaptic activity, governed by local calcium influx, but without reliance on dopamine or explicit feedback. It operates at a slower pace to ensure alignment with the excitatory activity levels, thus enhancing the contrast in activity by amplifying local differences in excitatory synaptic efficacy. The weight change is further dynamic and dependent on the weight itself, in such a way that large and small weights get small updates. This stabilizes the weight at the end of learning in the inhibitory synapses, together with the dynamical threshold level meeting the local calcium level, and prevents the weight from taking on large or negative values. Large weights are then close to w_{inh}^{max} and small weights close to 0 (see Fig. 5D).

The rule is described by Eq. 6-9 where Eq. 6 describes the basic plasticity curve used in Eq. 7 to control the change in inhibitory weights (dw_{inh}). In these equations, the parameters a and b set the activity values for the two terms in Eq. 6 while β and T control the steepness and intersection with the x-axis.

The parameter c recalibrates the baseline for T_{min} , ensuring T_{min} and T_{max} intersect at zero, as depicted in Fig. 7E (right panel, indicated by a green asterisk *).

The upper threshold T_{max} gradually shifts toward the highest calcium level observed in each synapse while the lower threshold T_{min} shifts towards a level below the maximum calcium

level (Eqs. 8 and 9; see Fig. 7E right panel for an illustration). Once the thresholds stabilize, the synapse will also reach stability. The goal of this is to enhance the difference in depolarization/calcium between the most active level, and the next one.

Further, depending on whether the synapse was active or inactive during the period of the calcium influx, the learning rate of Eq. 7 (η_{act}) takes on negative or positive values. Because of this, the active synapses are depressed when the calcium concentration is higher than the upper threshold T_{max} and potentiated if the concentration is between the two thresholds, while inactive synapses follow the inverse relationship. No change happens below the lower threshold (see Fig. 7E left panel for an illustration). This framework emphasizes the inhibitory synapse capacity to follow and dynamically amplify the contrast between local excitatory synaptic activity levels.

$$\Omega(Ca, T_{min}, T_{max}, a, b) = \frac{a}{1 + e^{-\beta_4(Ca - 0.004 \cdot T_{min})}} + \frac{b}{1 + e^{-\beta_4(Ca - 0.004 \cdot T_{max})}} \quad \text{Eq. 6}$$

$$dw_{inh} = \eta_{act} \cdot \Omega(Ca_v, T_{min}, T_{max}, a_1, b_1) \cdot w_{inh} \cdot (w_{inh}^{max} - w_{inh}) \quad \text{Eq. 7}$$

$$dT_{max} = \eta_7 \cdot \Omega(Ca_v, T_{min}, T_{max}, a_2, b_2) \quad \text{Eq. 8}$$

$$dT_{min} = \eta_8 \cdot \Omega(Ca_v + c, T_{min}, T_{max}, a_3, b_3) \quad \text{Eq. 9}$$

Training Procedure

We present the neuron with a sequence of 960 stimuli, which, as described in Fig. 1, can come in four varieties (four possible feature combinations). The learning rule is always “on”, i.e. synapses can always be updated, meaning that there are no separate training and testing phases in which synapses are plastic and frozen, respectively. We used two distinct synaptic arrangements for when representing the features: the clustered setup and the distributed setup.

Clustered setup: This setup is based on the assumption of pre-existing clustered synapses for each feature. The features of a stimulus are represented here by the shape and color of bananas and strawberries (Fig. 1A). Features were allocated to two dendritic branches. Each branch had each feature represented with five synapses clustered closely on a single dendrite. Depending on the feature combination, two, three, or four features were represented in clusters on one or both dendritic branches (see Fig. 3A, 4A and 4D for example). Additionally, 108 feature-unspecific synapses were distributed throughout the dendrites, activated concurrently with all stimuli to enhance the probability of spiking (as plateau potentials together with the general background synaptic noise used do not often lead to spikes in SPNs, cf Fig 2A). To demonstrate the inhibitory plasticity rule, four inhibitory synapses near each cluster (in the middle of the dendritic branch), representing each of the

four features, were added (Fig. 4A). With this setup a single feature activates both excitatory and inhibitory synapses. To match the level of depolarization seen in our excitatory-only setup, the number of distributed feature-unspecific synaptic inputs was increased to 144. The initial conductance of inhibitory synapses is set at 0.1 ± 0.01 nS.

Distributed setup: In contrast, this setup examines learning dynamics in neurons without pre-existing synaptic clustering for individual features. A total of 200 excitatory synapses were randomly distributed over 30 dendrites. Each feature was represented by 40 excitatory synapses, and an additional 40 feature unspecific excitatory synapses were used (Fig. 6A). We initiated this experiment with excitatory synaptic weights at 0.3 ± 0.1 (around 0.75 nS). This higher initial weight was chosen to compensate for the reduced efficacy of non-clustered synaptic inputs in producing sufficient depolarization and calcium influx. Extending our investigation, we added 60 inhibitory synapses, 15 for each feature, dispersed randomly over the 30 dendrites. We initiated this extended experiment with excitatory synaptic weights at 0.45 ± 0.1 (around 1.125 nS), aiming to maintain the same baseline voltage activity as in the excitatory-only case. Figure 6A illustrates this setup, highlighting the pre- and post-learning synaptic weights for both excitatory and inhibitory synapses. The four stimuli were presented in random order three times each within a block of 12 stimuli, followed by another reshuffled block of 12, and so on, for 960 trials. Each trial consisted of a stimulus presentation, lasting 20 ms, followed by a reward cue arriving 300 ms after stimulus onset, and lasting for 50 ms (Fig. 3A). During the stimulus presentation, the stimuli-related synapses receive one randomly-timed spike per synapse. We opted not to quantitatively model the dopamine receptor response to dopamine concentrations. Instead, dopamine levels from the reward-based learning system are represented with +1 (dopamine peak) for relevant stimuli if the neuron spikes, and -1 (dopamine pause) for irrelevant stimuli if the neuron spikes, and 0 for baseline levels when the neuron was silent, at which time the calcium amplitudes were measured, and the synaptic weights were updated during the presence of the reward cue. The time between two stimuli is 800 ms, long enough to allow for the voltage, calcium (and all other state variables in the model) to return to their baseline values.

Table 1: Excitatory Plasticity Parameters

Parameter	Description	Value
η_1	Learning rate for $[Ca]_{NMDA}$ -dependent plasticity	$1.5 \times 10^{-5} [\mu S \cdot mM \cdot ms^{-1}]$
η_2	Learning rate for $[Ca]_{L-type}$ -dependent plasticity	$3 \times 10^{-3} [ms^{-1} \cdot mM^{-1}]$
$\eta_{s, LTP}$	Learning rate for θ_{LTP} during Da peaks	$1 \times 10^{-7} [mM^2 \cdot ms^{-1}]$
$\eta_{s, LTD}$	Learning rate for θ_{LTP} during Da pauses	$4 \times 10^{-7} [mM^2 \cdot ms^{-1}]$
θ_{LTP}	Midpoint of LTP kernel (initial value)	$0.02 [mM]$

θ_{LTD}	Constant LTD threshold parameter	$7 \times 10^{-5} \text{ [mM]}$
β_1	Parameter controlling width of LTP kernel	$10^3 \text{ [mM}^{-1}\text{]}$
β_2	Parameter controlling steepness of LTD threshold	$10^5 \text{ [mM}^{-1}\text{]}$
β_3	Parameter controlling width of metaplasticity kernel	$334 \text{ [mM}^{-1}\text{]}$

Table 2: Inhibitory Plasticity Parameters

Parameter	Description	Value
η_{act}	Learning rate for the modification of inhibitory weight for active and inactive synapses	<i>active</i> : -0.055 <i>inactive</i> : 0.055 $[\mu\text{S}^{-1} \cdot \text{ms}^{-1}]$
$w_{\text{inh}}^{\text{max}}$	Maximum synaptic strength	5 [nS]
η_7	Learning rate for modifying T_{max}	$9 \times 10^{-4} \text{ [mM} \cdot \text{ms}^{-1}\text{]}$
η_8	Learning rate for modifying T_{min}	$-5 \times 10^{-5} \text{ [mM} \cdot \text{ms}^{-1}\text{]}$
a_1	Parameter controlling inhibitory weight	-1
b_1	Parameter controlling inhibitory weight	3
a_2	Parameter controlling T_{max}	-1
b_2	Parameter controlling T_{max}	3
a_3	Parameter controlling T_{min}	-2
b_3	Parameter controlling T_{min}	3
c	Calcium concentration offset	$6 \times 10^{-4} \text{ [mM]}$

β_4	Parameter for controlling the steepness of the curves	$2.5 \times 10^3 [mM^{-1}]$
-----------	---	-----------------------------




Figure 7: Synaptic Plasticity Rules: Calcium and Dopamine Interactions in Synaptic Weight Modification

A: Synaptic weight updates during a dopamine peak. (Left) LTP kernel is a bell-shaped curve which determines an optimal $[Ca]_{NMDA}$ region in which synaptic weight is increased. (Right) A wider bell-shaped kernel, i.e. the metaplasticity kernel, determines how the LTP kernel (the optimal region for plasticity) slides along the calcium level ($[Ca]_{NMDA}$) axis following a Da peak.

B: Synaptic weight updates during LTD. (Left) The LTD plasticity kernel. This kernel is constant. The LTD threshold is constant and set at 70 nM (Right). Metaplasticity describes how the LTP kernel slides along the calcium axis following a Da pause.

C: A schematic of how the LTP kernel window is updated following a dopamine peak. As NMDA calcium levels increase following activation of the strengthened synapse due to a dopamine peak (illustrated with the red circle jumping to the blue circle), the LTP kernel slides down and in this example the strengthened synapse (blue circle) stabilizes, and doesn't increase in strength even though additional rewards arrive.

D: A schematic showing the update of the LTP window following a dopamine pause leading to that synaptic weight and calcium response decrease as a function of L-type Ca (illustrated with the red circle jumping to the blue one). Here, the LTP kernel slides up towards higher NMDA calcium levels, and thus the weakened synapse (blue circle) is more unlikely to be recruited into the LTP window (unless it in the future is activated with Da peaks more regularly than with Da dips).

E: Depicts the inhibitory plasticity rule. (Left) Changes in synaptic weight for active (orange) and inactive (blue) synapses based on voltage dependent calcium levels in the dendritic shaft at the location of the inhibitory synapse, and minimum ($T_{\min}$) and maximum ($T_{\max}$) threshold. (Right) Functions for sliding the minimum and maximum thresholds with voltage dependent calcium level. The asterisk denotes the semi-stable zero level where the curves for active and inactive synapses meet.

Acknowledgement

We would like to thank members of the Hellgren Kotaleski laboratory for helpful discussion of various aspects of the manuscript. DT would like to thank Ana Kalajdjieva for illustrating the mouse brain. We acknowledge the use of Fenix Infrastructure resources, which are partially funded from the European Union's Horizon 2020 research and innovation programme through the ICEI project under the grant agreement No. 800858.

Simulations were also performed on resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at PDC KTH partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

The study was supported by the Swedish Research Council (VR-M-2020-01652), Swedish e-Science Research Centre (SeRC), Science for Life Lab, EU/Horizon 2020 no. 945539 (HBP SGA3) and No. 101147319 (EBRAINS 2.0 Project), European Union's Research and Innovation Program Horizon Europe under grant agreement No 101137289 (the Virtual Brain Twin Project), and KTH Digital Futures.

Author contributions

ZK and DT implemented the model, ran the simulations and analyzed the data with support from RL.

JHK and RL supervised all aspects of the study.

All authors contributed to the writing of the manuscript.

References

- Antic, S.D. *et al.* (2010) 'The decade of the dendritic NMDA spike', *Journal of neuroscience research*, 88(14), pp. 2991–3001.
- Bicknell, B.A. and Häusser, M. (2021) 'A synaptic learning rule for exploiting nonlinear dendritic computation', *Neuron*, 109(24), pp. 4001–4017.e10.
- Bienenstock, E.L., Cooper, L.N. and Munro, P.W. (1982) 'Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 2(1), pp. 32–48.
- Branco, T., Clark, B.A. and Häusser, M. (2010) 'Dendritic discrimination of temporal input sequences in cortical neurons', *Science*, 329(5999), pp. 1671–1675.
- Branco, T. and Häusser, M. (2010) 'The single dendritic branch as a fundamental functional unit in the nervous system', *Current opinion in neurobiology*, 20(4), pp. 494–502.
- Bruce, N.J. *et al.* (2019) 'Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals', *PLoS computational biology*, 15(10), p. e1007382.
- Carter, A.G. and Sabatini, B.L. (2004) 'State-dependent calcium signaling in dendritic spines of striatal medium spiny neurons', *Neuron*, 44(3), pp. 483–493.
- Chen, S.X. *et al.* (2015) 'Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning', *Nature neuroscience*, 18(8), pp. 1109–1115.
- Chen, X. *et al.* (2011) 'Functional mapping of single spines in cortical neurons in vivo', *Nature*, 475(7357), pp. 501–505.
- Cichon, J. and Gan, W.-B. (2015) 'Branch-specific dendritic Ca(2+) spikes cause persistent synaptic plasticity', *Nature*, 520(7546), pp. 180–185.
- Day, M. *et al.* (2024) 'GABAergic regulation of striatal spiny projection neurons depends upon their activity state', *PLoS biology*, 22(1), p. e3002483.
- Destexhe, A., Mainen, Z.F. and Sejnowski, T.J. (1994) 'An efficient method for computing synaptic conductances based on a kinetic model of receptor binding', *Neural computation*, 6(1), pp. 14–18.
- Dorman, D.B., Jędrzejewska-Szmek, J. and Blackwell, K.T. (2018) 'Inhibition enhances spatially-specific calcium encoding of synaptic input patterns in a biologically constrained model', *eLife*, 7. Available at: <https://doi.org/10.7554/eLife.38588>.
- Doron, M. *et al.* (2017) 'Timed Synaptic Inhibition Shapes NMDA Spikes, Influencing Local Dendritic Processing and Global I/O Properties of Cortical Neurons', *Cell reports*, 21(6), pp. 1550–1561.
- Du, K. *et al.* (2017) 'Cell-type-specific inhibition of the dendritic plateau potential in striatal spiny projection neurons', *Proceedings of the National Academy of Sciences of the United States of America*, 114(36), pp. E7612–E7621.
- Fino, E. *et al.* (2010) 'Distinct coincidence detectors govern the corticostriatal spike timing-dependent plasticity', *The Journal of physiology*, 588(Pt 16), pp. 3045–3062.

Fisher, S.D. *et al.* (2017) 'Reinforcement determines the timing dependence of corticostriatal synaptic plasticity in vivo', *Nature communications*, 8(1), p. 334.

Gandolfi, D. *et al.* (2020) 'Inhibitory Plasticity: From Molecules to Computation and Beyond', *International journal of molecular sciences*, 21(5). Available at: <https://doi.org/10.3390/ijms21051805>.

Gao, P.P. *et al.* (2021) 'Local glutamate-mediated dendritic plateau potentials change the state of the cortical pyramidal neuron', *Journal of neurophysiology*, 125(1), pp. 23–42.

Gidon, A. *et al.* (2020) 'Dendritic action potentials and computation in human layer 2/3 cortical neurons', *Science*, 367(6473), pp. 83–87.

Henneberger, C. *et al.* (2020) 'LTP Induction Boosts Glutamate Spillover by Driving Withdrawal of Perisynaptic Astroglia', *Neuron*, 108(5), pp. 919–936.e11.

van den Heuvel, M.P. and Sporns, O. (2013) 'Network hubs in the human brain', *Trends in cognitive sciences*, 17(12), pp. 683–696.

Higley, M.J. and Sabatini, B.L. (2010) 'Competitive regulation of synaptic Ca²⁺ influx by D2 dopamine and A2A adenosine receptors', *Nature neuroscience*, 13(8), pp. 958–966.

Hwang, F.-J. *et al.* (2022) 'Motor learning selectively strengthens cortical and striatal synapses of motor engram neurons', *Neuron*, 110(17), pp. 2790–2801.e5.

Iacuruso, M.F., Gasler, I.T. and Hofer, S.B. (2017) 'Synaptic organization of visual space in primary visual cortex', *Nature*, 547(7664), pp. 449–452.

Jia, H. *et al.* (2010) 'Dendritic organization of sensory input to cortical neurons in vivo', *Nature*, 464(7293), pp. 1307–1312.

Johansson, Y. and Silberberg, G. (2020) 'The Functional Organization of Cortical and Thalamic Inputs onto Five Types of Striatal Neurons Is Determined by Source and Target Cell Identities', *Cell reports*, 30(4), pp. 1178–1194.e3.

Ju, N. *et al.* (2020) 'Spatiotemporal functional organization of excitatory synaptic inputs onto macaque V1 neurons', *Nature communications*, 11(1), p. 697.

Kahneman, D. and Tversky, A. (1979) 'Prospect theory: An analysis of decision under risk', *Econometrica: journal of the Econometric Society*, 47(2), p. 263.

Kawai, R. *et al.* (2015) 'Motor cortex is required for learning but not for executing a motor skill', *Neuron*, 86(3), pp. 800–812.

Kerlin, A. *et al.* (2019) 'Functional clustering of dendritic activity during decision-making', *eLife*, 8. Available at: <https://doi.org/10.7554/eLife.46966>.

Kleindienst, T. *et al.* (2011) 'Activity-dependent clustering of functional synaptic inputs on developing hippocampal dendrites', *Neuron*, 72(6), pp. 1012–1024.

Larkum, M.E. *et al.* (2009) 'Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle', *Science*, 325(5941), pp. 756–760.

Lavzin, M. *et al.* (2012) 'Nonlinear dendritic processing determines angular tuning of barrel cortex neurons in vivo', *Nature*, 490(7420), pp. 397–401.

Legenstein, R. and Maass, W. (2011) 'Branch-specific plasticity enables self-organization of

nonlinear computation in single neurons', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 31(30), pp. 10787–10802.

Lindroos, R. *et al.* (2018) 'Basal Ganglia Neuromodulation Over Multiple Temporal and Structural Scales-Simulations of Direct Pathway MSNs Investigate the Fast Onset of Dopaminergic Effects and Predict the Role of Kv4.2', *Frontiers in neural circuits*, 12, p. 3.

Lindroos, R. and Hellgren Kotaleski, J. (2021) 'Predicting complex spikes in striatal projection neurons of the direct pathway following neuromodulation by acetylcholine and dopamine', *The European journal of neuroscience*, 53(7), pp. 2117–2134.

Losonczy, A. and Magee, J.C. (2006) 'Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons', *Neuron*, 50(2), pp. 291–307.

Major, G. *et al.* (2008) 'Spatiotemporally graded NMDA spike/plateau potentials in basal dendrites of neocortical pyramidal neurons', *Journal of neurophysiology*, 99(5), pp. 2584–2601.

von der Malsburg, C. (1999) 'The what and why of binding: the modeler's perspective', *Neuron*, 24(1), pp. 95–104, 111–25.

Matsuda, W. *et al.* (2009) 'Single nigrostriatal dopaminergic neurons form widely spread and highly dense axonal arborizations in the neostriatum', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 29(2), pp. 444–453.

McCulloch, W.S. and Pitts, W. (1943) 'A logical calculus of the ideas immanent in nervous activity', *The Bulletin of mathematical biophysics*, 5(4), pp. 115–133.

Nair, A.G. *et al.* (2015) 'Sensing Positive versus Negative Reward Signals through Adenylyl Cyclase-Coupled GPCRs in Direct and Indirect Pathway Striatal Medium Spiny Neurons', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 35(41), pp. 14017–14030.

Niculescu, D. *et al.* (2018) 'A BDNF-Mediated Push-Pull Plasticity Mechanism for Synaptic Clustering', *Cell reports*, 24(8), pp. 2063–2074.

Oikonomou, K.D. *et al.* (2014) 'Spiny neurons of amygdala, striatum, and cortex use dendritic plateau potentials to detect network UP states', *Frontiers in cellular neuroscience*, 8, p. 292.

Plotkin, J.L. *et al.* (2013) 'Regulation of dendritic calcium release in striatal spiny projection neurons', *Journal of neurophysiology*, 110(10), pp. 2325–2336.

Plotkin, J.L., Day, M. and Surmeier, D.J. (2011) 'Synaptically driven state transitions in distal dendrites of striatal spiny neurons', *Nature neuroscience*, 14(7), pp. 881–888.

Poirazi, P., Brannon, T. and Mel, B.W. (2003) 'Pyramidal neuron as two-layer neural network', *Neuron*, 37(6), pp. 989–999.

Polsky, A., Mel, B.W. and Schiller, J. (2004) 'Computational subunits in thin dendrites of pyramidal cells', *Nature neuroscience*, 7(6), pp. 621–627.

Ravasenga, T. *et al.* (2022) 'Spatial regulation of coordinated excitatory and inhibitory synaptic plasticity at dendritic synapses', *Cell reports*, 38(6), p. 110347.

Reig, R. and Silberberg, G. (2014) 'Multisensory integration in the mouse striatum', *Neuron*, 83(5), pp. 1200–1212.

- Reynolds, J.N.J. *et al.* (2022) 'Coincidence of cholinergic pauses, dopaminergic activation and depolarisation of spiny projection neurons drives synaptic plasticity in the striatum', *Nature communications*, 13(1), p. 1296.
- Roskies, A.L. (1999) 'The binding problem', *Neuron*, 24(1), pp. 7–9, 111–25.
- Sanabria, B.D. *et al.* (2024) 'Cell-Type Specific Connectivity of Whisker-Related Sensory and Motor Cortical Input to Dorsal Striatum', *eNeuro*, 11(1). Available at: <https://doi.org/10.1523/ENEURO.0503-23.2023>.
- Santos, F.J. *et al.* (2015) 'Cortico-striatal dynamics encode the refinement of specific behavioral variability during skill learning', *eLife*, 4, p. e09423.
- Schiess, M., Urbanczik, R. and Senn, W. (2016) 'Somato-dendritic Synaptic Plasticity and Error-backpropagation in Active Dendrites', *PLoS computational biology*, 12(2), p. e1004638.
- Schiller, J. *et al.* (2000) 'NMDA spikes in basal dendrites of cortical pyramidal neurons', *Nature*, 404(6775), pp. 285–289.
- Scholl, B., Wilson, D.E. and Fitzpatrick, D. (2017) 'Local Order within Global Disorder: Synaptic Architecture of Visual Space', *Neuron*, 96(5), pp. 1127–1138.e4.
- Schultz, W. (2007) 'Multiple dopamine functions at different time courses', *Annual review of neuroscience*, 30, pp. 259–288.
- Schultz, W., Dayan, P. and Montague, P.R. (1997) 'A neural substrate of prediction and reward', *Science*, 275(5306), pp. 1593–1599.
- Shen, W. *et al.* (2008) 'Dichotomous dopaminergic control of striatal synaptic plasticity', *Science*, 321(5890), pp. 848–851.
- Shindou, T. *et al.* (2019) 'A silent eligibility trace enables dopamine-dependent synaptic plasticity for reinforcement learning in the mouse striatum', *The European journal of neuroscience*, 49(5), pp. 726–736.
- Shindou, T., Ochi-Shindou, M. and Wickens, J.R. (2011) 'A Ca(2+) threshold for induction of spike-timing-dependent depression in the mouse striatum', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 31(36), pp. 13015–13022.
- Silver, D. *et al.* (2018) 'A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play', *Science*, 362(6419), pp. 1140–1144.
- Stauffer, W.R. *et al.* (2016) 'Components and characteristics of the dopamine reward utility signal', *The Journal of comparative neurology*, 524(8), pp. 1699–1711.
- Surmeier, D.J. *et al.* (2010) 'The role of dopamine in modulating the structure and function of striatal circuits', *Progress in brain research*, 183, pp. 149–167.
- Takahashi, N. *et al.* (2012) 'Locally synchronized synaptic inputs', *Science*, 335(6066), pp. 353–356.
- Tran-Van-Minh, A. *et al.* (2015) 'Contribution of sublinear and supralinear dendritic integration to neuronal computations', *Frontiers in cellular neuroscience*, 9, p. 67.
- Trpevski, D. *et al.* (2023) 'Glutamate spillover drives robust all-or-none dendritic plateau potentials-an investigation using models of striatal projection neurons', *Frontiers in cellular neuroscience*, 17, p. 1196182.

Varga, Z. *et al.* (2011) 'Dendritic coding of multiple sensory inputs in single cortical neurons in vivo', *Proceedings of the National Academy of Sciences of the United States of America*, 108(37), pp. 15420–15425.

Wilson, C.J. *et al.* (1983) 'Three-dimensional structure of dendritic spines in the rat neostriatum', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 3(2), pp. 383–388.

Wilson, D.E. *et al.* (2016) 'Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex', *Nature neuroscience*, 19(8), pp. 1003–1009.

Winnubst, J. *et al.* (2015) 'Spontaneous Activity Drives Local Synaptic Plasticity In Vivo', *Neuron*, 87(2), pp. 399–410.

Wolf, J.A. *et al.* (2005) 'NMDA/AMPA ratio impacts state transitions and entrainment to oscillations in a computational model of the nucleus accumbens medium spiny projection neuron', *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 25(40), pp. 9080–9095.

Xu, N.-L. *et al.* (2012) 'Nonlinear dendritic integration of sensory and motor input during an active sensing task', *Nature*, 492(7428), pp. 247–251.

Yagishita, S. *et al.* (2014) 'A critical time window for dopamine actions on the structural plasticity of dendritic spines', *Science*, 345(6204), pp. 1616–1620.

Zenke, F. and Gerstner, W. (2017) 'Hebbian plasticity requires compensatory processes on multiple timescales', *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 372(1715). Available at: <https://doi.org/10.1098/rstb.2016.0259>.

Zenke, F., Gerstner, W. and Ganguli, S. (2017) 'The temporal paradox of Hebbian learning and homeostatic plasticity', *Current opinion in neurobiology*, 43, pp. 166–176.