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

Zahra Khodadadi[1*], **Daniel Trpevski**[1], **Robert Lindroos**[1†], **Jeanette Hellgren Kotaleski**[1,2†]

***For correspondence:** zahra.khodadadi@scilifelab.se (ZK)

[†]These authors contributed equally to this work

[1]Science for Life Laboratory, Department of Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden; [2]Department of Neuroscience, Karolinska Institutet, Stockholm, Sweden

Abstract This study investigates 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 dependent on 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 order to devise a self-adjusting learning rule, which ensures stability for individual synaptic weights, metaplasticity is also used. We demonstrate that this rule allows single neurons to solve the nonlinear feature binding problem, a task traditionally attributed to neuronal networks. We also detail an inhibitory plasticity mechanism that contributes to dendritic compartmentalization, further enhancing computational efficiency in dendrites. This in silico study highlights the computational potential of single neurons, providing deeper insights into neuronal information processing and the mechanisms by which the brain executes complex computations.

Introduction

Classically, single neurons in the nervous system have been thought to operate as simple linear integrators such that the nonlinearity of dendrites can be neglected (***McCulloch and Pitts, 1943***). Based on this simplification, powerful artificial neural systems have been created that outperform humans on multiple tasks (***Silver et al., 2018***). However, in recent decades it has been shown that active dendritic properties shape neuronal output and that dendrites display nonlinear integration of input signals (***Antic et al., 2010***). These dendritic nonlinearities enable a neuron to perform sophisticated computations (***Tran-Van-Minh et al., 2015***; ***Gidon et al., 2020***), expanding its computational power beyond what is available with the somatic voltage threshold and making it similar to a multilayer artificial neural network (***Poirazi et al., 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 spatiotemporally clustered glutamatergic input (***Schiller et al., 2000***; ***Polsky et al., 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 spatial requirement is somewhat loose as even single dendritic branches have been proposed to act as computational units (*Losonczy and Magee, 2006*; *Branco and Häusser, 2010*). Nevertheless, multiple so-called hotspots, 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*; *Iacaruso et al., 2017*; *Scholl et al., 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 its most basic form, the NFBP involves discriminating between two groups of feature combinations. This problem is nonlinear because the neuron must learn to respond only to specific feature combinations, even though all features contribute equally in terms of synaptic input. A commonly used example involves two different shapes combined with two different colors, resulting in four total combinations. Out of these, the neuron should respond only to two specific feature combinations (exemplified in *Figure 1*A and B).

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 *Figure 1*A (*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. Here, however, we will continue to illustrate the NFBP with the more intuitive features borrowed from the visual field, although for the 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 et al., 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 converging 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 signify rewarding outcomes and pauses in dopamine represent the omission of expected rewards (*Schultz et al., 1997*). Dopamine signals further control the synaptic plasticity of corticostriatal synapses on the SPNs (*Figure 1*C). In direct pathway SPNs (dSPN) expressing the D1 receptor, a dopamine peak together with significant calcium influx through NMDA receptors triggers synaptic strengthening (long-term potentiation – LTP). Conversely, when little or 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, see *Figure 1*D) (*Shen et al., 2008*; *Fino et al., 2010*; *Plotkin et al., 2013*).

The ability to undergo LTP or LTD is itself regulated (*Huang et al., 1992*), a concept termed metaplasticity (*Abraham and Bear, 1996*). Metaplasticity refers to changes in synaptic plasticity driven by prior synaptic activity (*Frey et al., 1995*) or by neuromodulators (*Moody et al., 1999*), effectively making plasticity itself adaptable. Metaplasticity can further regulate synaptic physiology to shape future plasticity without directly altering synaptic efficacy, acting as a homeostatic mechanism to keep synapses within an optimal dynamic range (*Abraham, 2008*).Previous theoretical studies (*Bienenstock et al., 1982*; *Fusi et al., 2005*; *Clopath et al., 2010*; *Benna and Fusi, 2016*; *Zenke and Gerstner, 2017*) demonstrate the essential role of metaplasticity in maintaining stability in synaptic

weight distributions.

If dopamine peaks are associated with the relevant feature combinations in the NFBP and dopamine pauses with the irrelevant ones, they 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 (***Figure 1***B) (***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 et al., 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 precalculated learning rule – dependent on somatic spike timing and high local dendritic voltage – was 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 study, we ask whether—and under what conditions—the theoretical solution to the nonlinear feature-binding problem (NFBP) can be achieved in a biophysically detailed model of a striatal projection neuron (SPN) using only local, biologically grounded learning rules. We frame this paper around two questions: First, can a single dSPN equipped with only calcium- and dopamine-dependent excitatory plasticity solve the NFBP when the relevant features are pre-clustered on one dendritic branch? Second, if that mechanism is insufficient—as with randomly distributed or very distal inputs—does adding branch-specific inhibitory plasticity restore plateau-based nonlinear computation and spiking?

To answer these questions we adopt an approach that relies on the following key mechanisms:

A local learning rule: We develop a learning rule driven by local calcium dynamics in the synapse and by reward signals from the neuromodulator dopamine. This plasticity rule is based on the known synaptic machinery for triggering LTP or LTD at the corticostriatal synapse onto dSPNs (***Shen et al., 2008***). Importantly, the rule does not rely on supervised learning paradigms, and no separate training and testing phase is required.

Robust dendritic nonlinearities: According to ***Tran-Van-Minh et al.*** (***2015***), sufficient supralinear integration is needed to ensure that, e.g., two inputs (one feature combination in the NFBP, ***Figure 1***A) on the same dendrite generate greater somatic depolarization than if those inputs were distributed across different dendrites. To accomplish this, we generate sufficiently robust dendritic plateau potentials using the approach of ***Trpevski et al.*** (***2023***).

Metaplasticity: Our simulations demonstrate that metaplasticity is necessary for synaptic weights to remain stable and within physiologically realistic ranges, regardless of their initial values.

We first demonstrate the effectiveness of the proposed learning rule under the assumption of pre-existing clustered synapses for each individual feature, as suggested by ***Tran-Van-Minh et al.*** (***2015***). These clustered synapses are trained, to a degree where they can reliably evoke robust plateau potentials for the relevant feature combinations required to solve the NFBP, while synapses representing irrelevant features are weakened (illustrated in ***Figure 1***B).

We then extend the analysis by applying the learning rule to more randomly distributed synapses, which initially exhibit minimal local supralinear integration. However, when incorporating the assumption that branch-specific plasticity mechanisms are at play, supralinear integration emerges within distinct dendritic branches. This suggests that branch-specific plasticity could play a critical role in enabling single neurons to solve nonlinear problems. Furthermore, we explore an activity-dependent rule for GABAergic plasticity and demonstrate its potential importance

in shaping dendritic nonlinearities. This mechanism may thus further enhance computational capabilities by refining and stabilizing the integration of inputs across dendritic branches.

Although brain systems like the striatum, which integrate multimodal inputs, somehow solve nonlinear problems at the network or systems level, it remains unclear whether individual neurons in the brain regularly solve the NFBP. Our investigation suggests, however, that single SPNs possess the computational capacity to address linearly non-separable tasks. This is achieved by leveraging the organism's performance feedback, represented by dopamine peaks (success) and dopamine pauses (failure), in combination with their ability to generate robust dendritic plateau potentials. Since the mechanisms used in the rule are general to the brain, this capability may also extend to other projection neurons capable of producing dendritic plateaus, such as pyramidal neurons. However, the specific feedback mechanisms, represented here by dopamine, would need to be associated with alternative neuromodulatory signals depending on the type of neuron and synapse involved.

Results

Characterization of the dendritic nonlinearities in the model

The nonlinear sigmoidal voltage sensitivity of NMDA receptors is a crucial element for forming dendritic plateau potentials. We used a model to generate 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, the model includes glutamate spillover from the synaptic cleft that activates extrasynaptic NMDA receptors. Glutamate spillover occurs when the total synaptic weight of the nearby activated synapses reaches a threshold value (see Methods). The threshold value here is set to be equivalent to the total weight of eight clustered synapses with weights of 0.25 each (weights of 0.25 correspond to a maximal conductance of 0.625 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 as well as the NMDA and L-type calcium accumulated in a single spine (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 eventually leading to glutamate spillover). Reaching the spillover threshold produces a sudden and prolonged increase in NMDA conductance, caused by the activation of extrasynaptic NMDA receptors (where the clearance of glutamate is assumed to be slower). The result is a more robust plateau potential that better resembles experimentally generated plateau potentials, as investigated in more detail 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 ***Figure 2***B, and gives the voltage and calcium responses as exemplified in ***Figure 2***A. ***Figure 2***C shows the maximal amplitudes of the somatic and spine voltage together with NMDA and L-type calcium signals in the dendritic spines, averaged over 10 trials and over 10 dendrites. The "baseline" results in ***Figure 2***C 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 (***Figure 2***C, spine voltage panel). In the simulation using "strengthened" synapses the synaptic weights in the cluster are 60% greater than in the "baseline" case, and hence need fewer synapses to trigger plateau potentials. Conversely, with weaker synapses, where weights are 20% 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. Further, clusters of strengthened synapses can reliably generate robust plateau potentials, which produces a long lasting somatic depolarization and thereby increases the likelihood for somatic spiking (in accordance with (***Du et al., 2017***) and also illustrated in (***Trpevski et al., 2023***)). Conversely, clusters of weakened synapses will most likely not generate plateau potentials, and

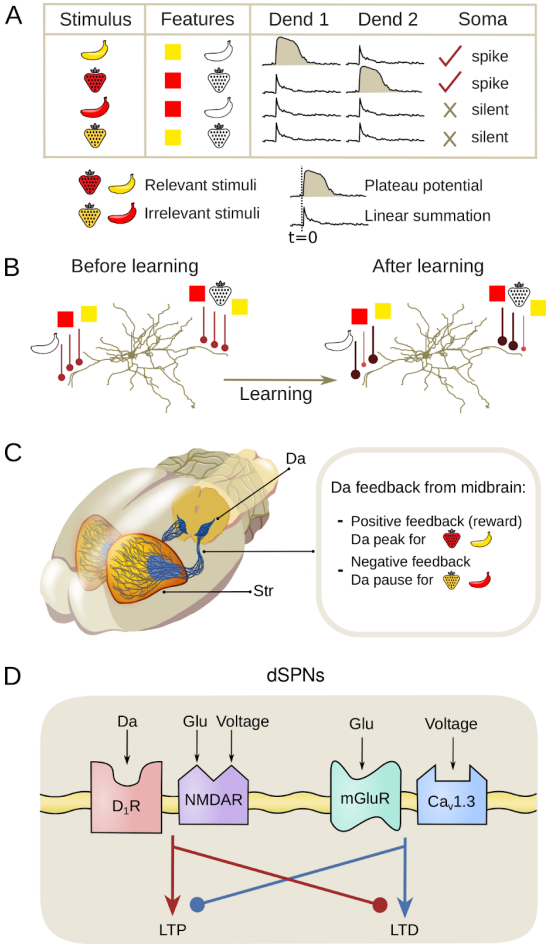

Figure 1. Learning Mechanisms in direct pathway Striatal Projection Neurons (dSPNs) for the Nonlinear Feature Binding Problem (NFBP)

(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. In the NFBP the neuron should learn to respond by spiking to two of the feature combinations, representing the relevant stimuli (red strawberry and yellow banana), while 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, corresponding to a relevant stimulus, evoke a plateau potential, thus superlinearly exciting the soma. Conversely, co-activation of synaptic clusters for the irrelevant combinations should not evoke plateau potentials.

(B): Synaptic clustering in dendrites: 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: dopaminergic inputs 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 dopamine pause is necessary for LTD.

(D): Signaling pathways underlying synaptic plasticity in dSPNs: Illustrations of signaling components at the corticostriatal synapse that modify synaptic strength (redrawn from Shen et al., 2008). 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 dopamine triggers LTD (while counteracting the LTP cascade).

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

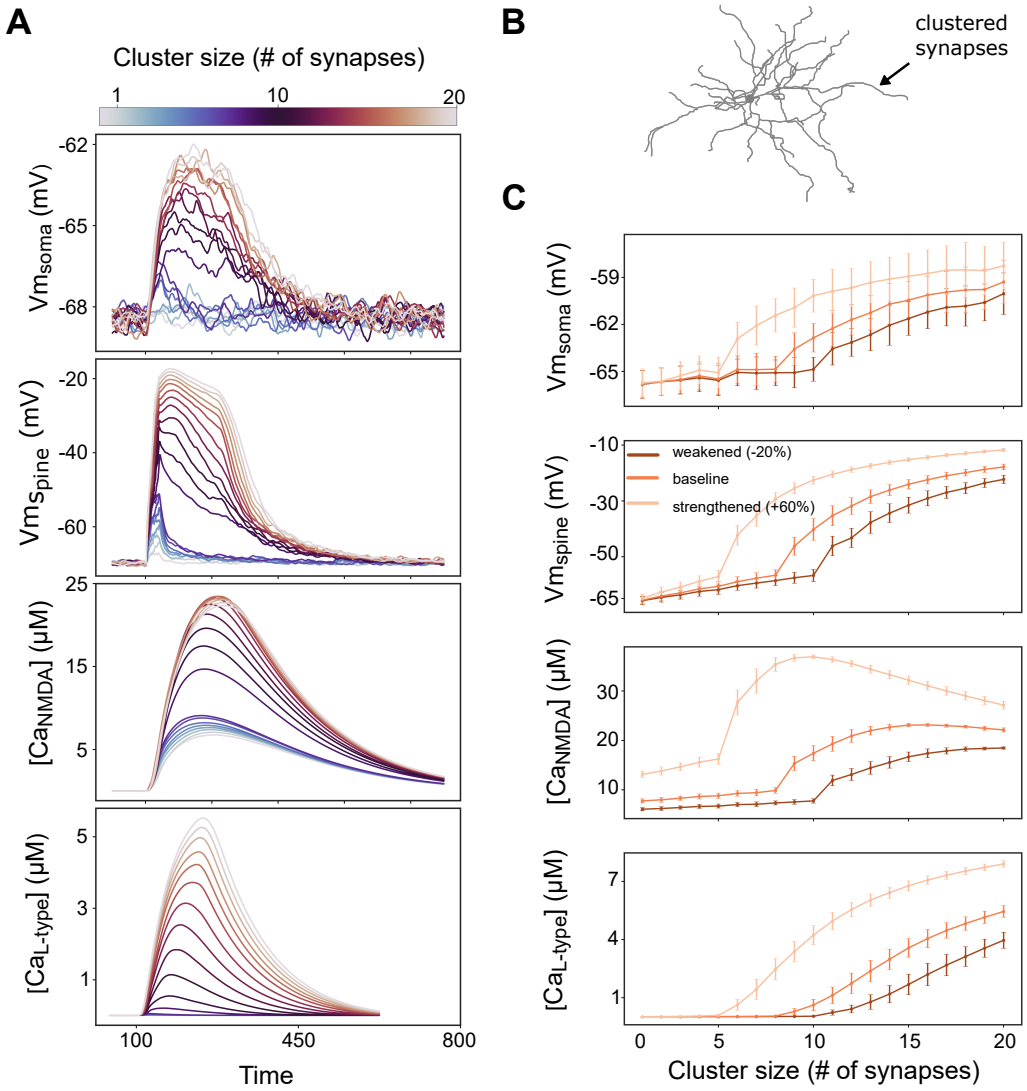

Figure 2. Characterization of dendritic plateau potentials in the model.
(A): Somatic voltage, spine voltage, NMDA calcium, and L-type calcium evoked by a cluster varying in size from 1 to 20 synapses. A plateau potential is evoked when glutamate spillover is activated, here triggered when 8 synapses with a weight of 0.25 each are coactivated (corresponding to the "baseline" weights in C).
(B): Schematic of the neuron morphology with an arrow indicating the stimulated dendritic branch shown in panel A.
(C): The mean maximal amplitude, with standard deviation shown in bars, of the measures in A averaged over 10 different dendrites, and 10 trials per dendrite. The curves represent clusters with different synaptic weights: baseline (0.25, 0.625 nS), strengthened (0.4, 1 nS) and weakened weights (0.2, 0.5 nS). 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 synaptic plasticity rule

To characterize the learning rule we started with a simple setup where three features, each illustrating either a color or a shape, were distributed onto two dendritic branches. The setup was such that one relevant and one irrelevant feature combination were represented in each dendrite and the relevant–irrelevant feature combination was unique to each dendrite (see illustration in ***Figure 3***A). Each feature was represented with 5 synapses, and we start with the assumption that

those synapses are already organized in pre-existing clusters. In addition to background synaptic noise inputs, 108 randomly distributed glutamatergic synapses were also added. These synapses 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 allowed 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 learning rule modified synaptic weights based on local calcium levels at each synapse, employing distinct plasticity kernels for LTP and LTD, each associated with different calcium sources—NMDA and L-type channels, respectively (***Figure 1***; schematic in ***Figure 3***B; see also *Methods*). The LTD kernel was governed by a calcium threshold, above which the change of synaptic weight was proportional to the amplitude of L-type calcium. That is, higher calcium levels trigger a larger reduction of weights than low concentrations. Conversely, the LTP kernel was represented by a bell-shaped function over a range of NMDA calcium concentrations, allowing LTP in the presence of dopamine when calcium levels fell within this range, while calcium outside this range did not elicit LTP (see *Methods*).

An additional feature of the learning rule was metaplasticity (***Abraham, 2008***), which dynamically adjusted the calcium dependence of the LTP kernel. That is, the LTP kernel's calcium dependence was adjusted over time (see ***Figure 3***B–D). Metaplasticity enabled, e.g., previously weak synapses that were repeatedly co-activated with rewards to eventually participate in LTP, as their calcium threshold was gradually lowered when the LTP kernel was shifted. The same mechanism also prevented excessive strengthening of already potent synapses as shifting the LTP kernel downwards also reduced the potentiation of synapses with large calcium transients. This behavior of the LTP kernel ensured stability in synaptic weights and allowed for continuous refinement of synaptic efficacy based on activation pattern and presence/absence of dopamine.

Initially, the synaptic weight of each synapse was set to a value chosen uniformly at random from a range of 0.25 ± 0.05 (corresponding to conductances of 0.625 ± 0.125 nS). Initially, the synapses therefore experienced different NMDA and L-type calcium levels following activation, and hence their weights were also updated differently based on where the calcium level fell within their individual plasticity kernels.

During training, the neuron model was activated with a sequence of 1400 feature combinations, including equal amounts of relevant (i.e. 'red strawberry' and 'yellow banana') and irrelevant feature combinations (i.e. 'yellow strawberry' and 'red banana'). Thus, dopamine peaks and dopamine pauses occurred equally often at the beginning of the learning phase. When the neuron spiked for the relevant feature combinations, dopamine rewards were delivered, triggering LTP in the active synapses with NMDA calcium levels within the LTP kernel. Conversely, spiking for the irrelevant feature combinations elicited a dopamine pause as feedback, triggering LTD as a function of L-type calcium. 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 the model (***Figure 3***E, gray lines). After learning, however, the neuron could differentiate between the two sets of stimuli—the relevant feature combinations associated with a reward continued to evoke a plateau potential and elicit somatic spiking, while in contrast, the neuron's response to irrelevant feature combinations was notably decreased (***Figure 3***E, black lines).

Figure 3C shows the evolution of synaptic conductances during the learning process for dendrite 1 (dendrite 2 is not shown, but has a similar behavior for relevant and irrelevant stimuli). 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 during the learning process LTD could also occur in some synapses representing 'yellow' and 'banana' because these features were also part of the irrelevant stimuli ('yellow strawberry' and 'red

banana', respectively). As a result, a small number of these synapses, in particular those whose initial weights are low, have been weakened and are not recruited in the clusters to encode the 'yellow banana' stimulus. 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 (as some by chance may be outside of the plastic region of the LTP kernel). Depending on the initial local calcium response in the synapse to relevant and irrelevant stimuli, individual synapses might either be preferentially recruited into the LTP or LTD process. An animated rendering of the first 400 training epochs for this panel, illustrating the frame-by-frame evolution of conductance and calcium in the example with clustered synapses, is provided in ***Figure 3—video 1***.

The evolution of synaptic weights for clustered and feature-unspecific synapses is illustrated in the top panels of ***Figure 3***C and D. Correspondingly, the calcium dynamics for three color-coded synapses in each case (indicated with arrows in the top panels) are shown in the bottom panels, demonstrating how the dynamics evolve under the influence of LTP, LTD and metaplasticity. This highlights the process by which synapses stabilize at specific conductances/weights during learning.

The red synapses (exemplified in ***Figure 3***C, second row) represent an irrelevant feature combination ('red banana'), and as learning progresses, LTD becomes more frequent due to repeated dopamine pauses. These pauses shift the dynamic LTP and metaplasticity kernels towards higher calcium levels, moving the red synapses outside the range for LTP. This shift ensures that the red synapses weaken over time, preventing them from encoding irrelevant stimuli and reducing the dendrite's response to these inputs. In setups lacking metaplasticity this does not happen, leading to an inability of the neuron to separate the response to relevant and irrelevant feature combinations (***Figure 3—figure Supplement 1***, red synapse).

Most synapses representing the relevant feature combination ('yellow banana', exemplified in ***Figure 3***C, bottom rows) initially undergo LTP as repeated dopamine rewards strengthen them. However, synapses with weak initial weights may fail to potentiate further due to the slow adaptation of the LTP kernel and their inability to reach the required calcium levels. Specifically, for the synapses representing yellow, "Yellow syn 1" exhibits an increase in synaptic conductance, whereas "Yellow syn 2" does not undergo potentiation (see arrows in ***Figure 3***C, top).

Feature-unspecific synapses, like Example 1 and 3 in ***Figure 3***D, typically weaken over time. These synapses are co-activated with irrelevant stimuli and experience frequent dopamine pauses, causing the metaplasticity kernel to shift upward and away from the calcium levels required for LTP. As a result, LTD dominates in these synapses, leading to their gradual weakening. Initially, LTP and LTD typically occur equally often, but as learning progresses, both the frequency of dopamine pauses and thus LTD decrease, allowing the weights to stabilize over time.

In contrast, Example 2, though feature-unspecific, behaves differently due to its proximity to a cluster. It benefits from high local voltage and calcium during relevant stimuli (e.g. 'yellow banana'), keeping calcium levels within the high end of the LTP kernel. As LTP dominates in this synapse, the LTP kernel gradually shifts down towards lower calcium concentrations. Eventually, this causes the calcium level associated with the other relevant stimulus ('red strawberry') to also fall within the plastic region of the LTP kernel, and thereby further strengthening the synapse.

In summary, synaptic weights are dynamically regulated by the interplay of LTD and LTP, modulated by the metaplasticity kernel. This process ensures that synapses adjust appropriately based on their relevance to feature combinations, stabilizing in a manner that reflects their functional roles and initial conditions.

Clustering enables excitatory-only learning, but overlapping features make performance less reliable

After demonstrating that the SPN can differentiate between relevant and irrelevant stimuli in the simplified example in ***Figure 3***, we generalized this setup by varying the innervation of the four features across the two dendrites and recorded the SPN's performance on the NFBP as learning

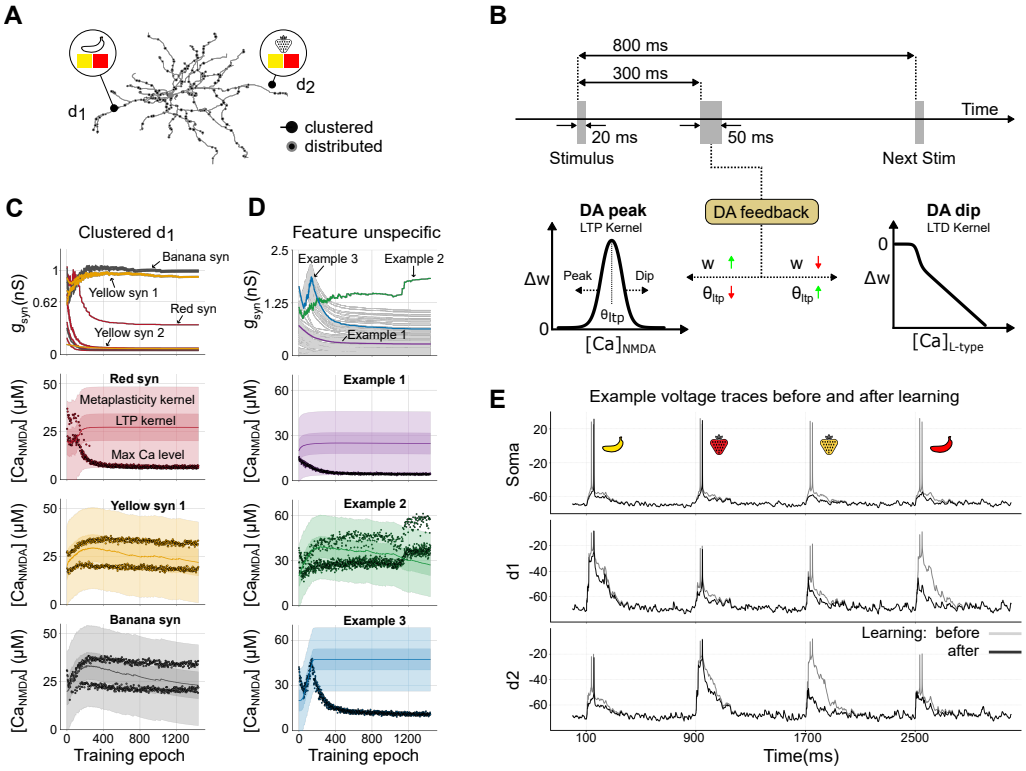

Figure 3. Example of setup and learning-induced synaptic plasticity

(A): Input configuration. The panel shows how the four stimulus features are distributed across two dendrites (d1, d2). Each dendrite contains pre-existing synaptic clusters for three features (black circles), and distributed, feature-unspecific synapses (shown in grey).

(B): Stimulation timeline and plasticity schematic. The top diagram depicts a 20 ms stimulus followed by a 50 ms dopaminergic feedback pulse delivered with a 300 ms delay, only if the neuron spikes. Successive stimuli are separated by 800 ms to let calcium return to baseline. Lower panels: a dopamine peak (left) gates LTP—synapses whose $[Ca]_{NMDA}$ falls within the bell-shaped window are potentiated (w ↑, green) while the kernel midpoint θ_{LTP} shifts downward (↓, red); a dopamine pause (right) gates LTD—the degree of depression scales with $[Ca]_{L\text{-}type}$ (w ↓, red) and θ_{LTP} shifts upward (↑, green).

(C), (D): Evolution of synaptic conductances during learning (top row) and examples of peak Ca^{2+} levels (black dots) alongside kernel dynamics in single synapses (three lower rows). Panels in (C) depict clustered synapses in dendrite 1 (d1), where 'yellow' and 'banana' generally undergo LTP, while 'red' undergoes LTD. See ***Figure 3—video 1*** for an animation of panel C showing the first 400 training stimuli of the learning sequence. Panels in (D) show distributed, feature-unspecific synapses. Among these, the Example 1 and Example 3 traces represent synapses that are weakened, while the Example 2 trace exemplifies a synapse near one of the clusters that, by chance, achieves a sufficiently high local NMDA calcium level for LTP to dominate. The initial synaptic conductances are drawn from a normal distribution with a mean of 0.625 and a standard deviation of 0.125 nS. The solid line represents the midpoint of the kernels, where LTP is strongest. 'Max' indicates the peak NMDA calcium during a single stimulus. The darker shaded regions represent the LTP kernel and the lighter shaded ones show the wider metaplasticity kernel.

(E): Example voltage in the soma and the middle of dendrite 1 (d1) and dendrite 2 (d2) before and after learning. Each dendrite here stops responding to the respective irrelevant stimuli during learning.

Figure 3—video 1. Animated rendering of the first 400 training stimuli (~400 epochs) showing frame-by-frame changes in conductance and calcium for the clustered synapses shown in panel C.

Figure 3—figure supplement 1. Learning without metaplasticity.

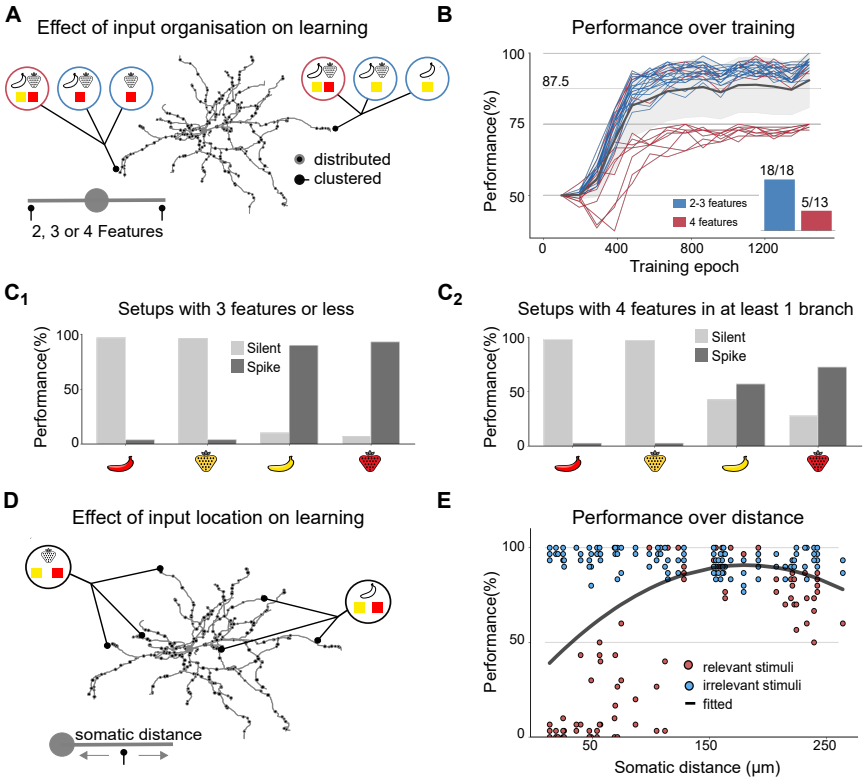

Figure 4. Impact of input configurations and synaptic cluster locations on NFBP learning performance
(A): Illustration of creating different input configurations where two, three (light blue circles), or four features (deep red circles) are placed in two dendritic locations.
(B): Performance trajectories for all 31 unique input configurations over training epochs. Each colored line represents the learning performance of a single configuration, with blue traces indicating setups where no dendrite receives more than three features and red traces representing setups where at least one dendrite contains all four features. The black line and grey shaded area represent the mean and standard deviation over all trials. The inset shows the number of configurations in each group that reached the NFBP learning criterion: 18/18 for the blue group and 5/13 for the deep red group.
(C): Performance on the last 160 stimuli for the two groups of configurations: **(C1)** setups where each dendrite has at most three features and **(C2)** setups where at least one dendrite contains all four features.
(D) and (E): Performance in a three-feature configuration as a function of cluster location. (D) illustrates the distribution of synaptic clusters across dendritic locations, and (E) shows performance over the last 160 stimulus presentations as a function of somatic distance for these configurations. A total of 60 unique configurations were tested, with synapse clusters randomly assigned to different dendritic locations. The solid line represents a quadratic fit to the data.

progressed. We only used feature innervations where both relevant feature combinations were present, with at least one relevant feature on each dendrite. This ensures sufficient innervation to potentially solve the NFBP (for an illustration see ***Figure 4***A). The SPN's performance was assessed during training to determine whether it spiked for relevant feature combinations and remained silent for irrelevant ones. Performance of 100% indicates the SPN spikes exclusively for relevant stimuli, while 50% typically indicates one of two scenarios: (i) the SPN spikes for all four stimuli, or (ii) it remains silent for all four. A performance of 75% reflects intermediate spiking behavior, such as spiking for one relevant combination while remaining silent for the other three, or spiking for three combinations and being silent for one irrelevant combination. The NFBP is considered solved when performance exceeds 87.5%, which for example occurs when the SPN consistently spikes for one relevant combination and spikes at least half the time for the other relevant combination while ignoring irrelevant stimuli.

In cases where both dendrites receive two or three features, ensuring that each dendrite has at least one relevant combination, the mean performance exceeds 90%, indicating that both relevant combinations are learned (***Figure 4***B, blue traces). However, when at least one dendrite is innervated by all four features, learning becomes significantly more challenging. The dendrite with four features struggles to resolve the competition between feature combinations, making it difficult to encode a single relevant combination (***Figure 4***B, deep red traces). However, in some cases, depending on the order of stimuli during training, both relevant combinations can still be learned, solving the NFBP. Here, 5 out of 13 setups in which one dendrite received all four features successfully solved the task (***Figure 4***B, inset).

Figure 4C provides a detailed view of performance during the last 160 training examples for each of the four stimuli. ***Figure 4***C1 corresponds to the blue traces in ***Figure 4***B, where the NFBP is considered solved. ***Figure 4***C2 corresponds to the deep red traces in ***Figure 4***B, where the NFBP is generally not solved, as relevant stimuli elicit spikes only about 75% of the time. These results suggest that when one dendrite is innervated by all four features, the competition between feature combinations creates a bottleneck for learning. In cases where learning fails, the dendrite with four features often either fails to encode any combination due to excessive LTD or instead encodes the same combination already represented in the other dendrite. Addressing this challenge may require mechanisms to "break symmetry" during learning, as noted in previous studies using abstract neuron models. Symmetry-breaking strategies, such as amplifying or attenuating dendritic nonlinearities as they propagate to the soma, can reduce competition and enable successful differentiation of feature combinations (e.g., branch plasticity; ***Legenstein and Maass 2011***).

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

We also investigated the impact of synaptic positioning on learning (***Figure 4***D) when using the same input setup as in ***Figure 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 (***Figure 4***E). From ***Figure 4***E one can infer that after learning the proximal synapses have all been decreased to such an extent that the neuron has stopped spiking for all four combinations (as the total performance is around 50% and relevant feature combinations is at 0%). 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 supralinear input integration 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. Note that in our simulations we allow glutamatergic

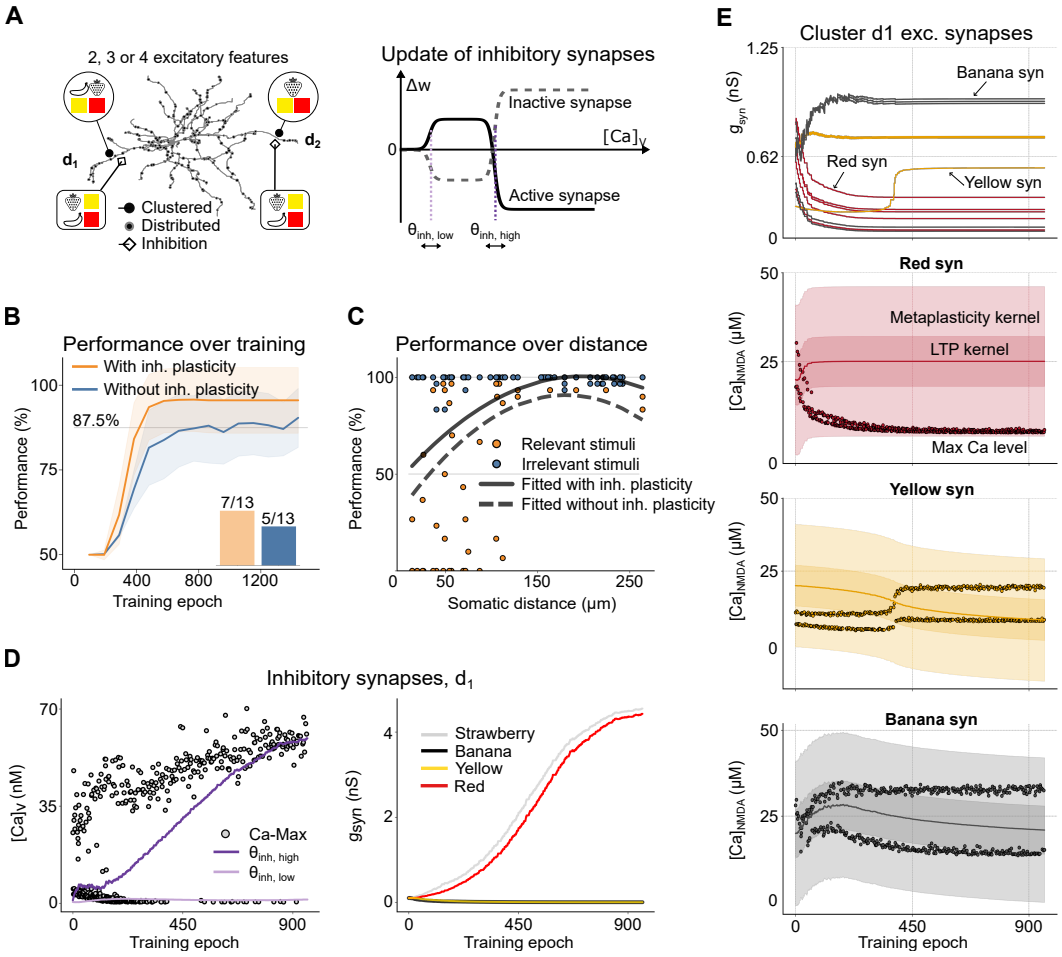

Figure 5. Effects of inhibitory plasticity on performance

(A): Dendritic input configuration with inhibitory synapses added. The setup of the excitatory and feature-unspecific synapses is the same as in ***Figure 4***A. Plastic inhibitory synaptic connections for each of the four features are added in both dendrites, with one synapse per feature.

(B): Displays average performance for 31 unique input configurations of 2, 3, or 4 features on two dendrites as a comparison between the setup with (orange) and without inhibitory plasticity (blue). Shaded areas show standard deviation. The bar plot inset shows the number of configurations, with 4 features in at least one dendrite, that successfully solved the NFBP with and without inhibitory plasticity.

(C): Performance over the last 160 stimulus presentations as a function of somatic distance of the synaptic clusters for the input configuration in ***Figure 4***D, with added inhibition. A total of 60 unique configurations were tested, with synapse clusters randomly assigned to different dendritic locations. The solid line represents a quadratic fit to the data. The dashed line is the corresponding quadratic fit from ***Figure 4***E.

(D): Peak voltage-gated calcium (left panel, dots) and plasticity threshold dynamics (lines), in dendrite 1, over training epochs. The upper threshold ($\theta_{inh,high}$) moves towards peak calcium levels while the lower threshold ($\theta_{inh,low}$) moves towards the next highest level. (Right) Inhibitory synaptic conductances for the synapses in dendrite 1. Strengthened inhibitory synapses prevent excitation by the corresponding excitatory features (here 'red' and 'strawberry'), while weakened inhibitory synapses allow excitation by their corresponding features (here 'yellow' and 'banana').

(E): Excitatory synaptic conductances (top) and calcium levels and kernel dynamics (below) over learning. The specific conductances in the top panel corresponding to the kernel dynamics and calcium levels in the bottom panels are indicated with arrows. Note that kernel and calcium dynamics for the example 'yellow' synapse refer to the only 'yellow' synapse in the top panel which initially weakens and is later strengthened. The solid line shows the midpoint of the LTP kernels. *Max* refers to the peak NMDA calcium during a single stimulus. The darker shaded regions represent the LTP kernel and the lighter shaded ones show the wider metaplasticity kernel.

Figure 5—figure supplement 1. Learning with and without inhibitory plasticity.

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, however, this ultimately results in decreased performance on the NFBP for the following reasons. In distal synaptic clusters, excessive spiking for irrelevant stimuli occurs during the early and middle stages of learning. This leads to frequent punishment signals (negative feedback), which, in turn, reduce the neuron's ability to spike for relevant combinations later in training. Additionally, while distal synapses can generate plateau potentials and NMDA-dependent responses, these signals naturally attenuate as they propagate toward the soma, sometimes failing to elicit somatic spikes for the relevant stimuli. This combination of excessive negative feedback early on and attenuated distal signals reduces the effectiveness of the feedback loop, making it difficult for the neuron to selectively strengthen synapses associated with relevant stimuli and weaken those tied to irrelevant ones.

In our simulations, the ideal learning zone for NFBP thus lies at an intermediate somatic distance, roughly 100–150 μm from the soma, where synapses can effectively contribute to learning the NFBP (***Figure 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.

Inhibitory plasticity can enhance contrast and robustness of dendritic computations

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 ***Figure 4***B. The second is a decrease in performance observed across very distal synapses, as detailed in ***Figure 4***E. These findings prompted us to explore the potential role of inhibitory synapses.

We developed a phenomenological inhibitory plasticity rule to enhance the contrast of dendritic nonlinearities (see ***Figure 5***A, detailed in the *Methods* section). This rule is designed to compartmentalize dendrites, ensuring they predominantly respond to excitatory inputs that cause the strongest activation. Such compartmentalization has been experimentally observed in radial oblique and basal dendrites of CA1 pyramidal neurons, where NMDA spikes and plateau potentials are the main forms of dendritic nonlinearity. Co-located inhibitory synapses in these dendrites regulate whether these nonlinear responses are elicited (***Lovett-Barron et al., 2012***; ***Milstein et al., 2015***; ***Grienberger et al., 2017***).

Unlike excitatory plasticity, which relies on dopaminergic feedback signals, our inhibitory plasticity model follows a rule that passively follows the local calcium level from voltage-gated channels. It reinforces the voltage elevation of the most active excitatory synapses within a dendritic branch by decreasing the inhibition corresponding to the same features there (***Chapman et al., 2022***), and conversely, it strengthens inhibitory synapses for features which generate less excitatory activity (see ***Figure 5***A, right panel, illustrating how the rule updates active and inactive inhibitory synapses).

To demonstrate how the inhibitory plasticity rule works, we use the same excitatory synapse setup as in ***Figure 4*** to which we add four inhibitory synapses near each cluster in the middle of

the dendritic branch, representing each of the four features (***Figure 5***A). 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 number of feature-unspecific inputs from 108 to 144. Alongside this, we began with low inhibitory synaptic weights. These two changes were needed in order to maintain higher baseline activity in the model, as starting with strong inhibitory weights could excessively suppress excitatory activity, since inhibitory inputs close to clustered synapses can effectively counteract the NMDA-dependent nonlinearities (***Doron et al., 2017***; ***Du et al., 2017***; ***Dorman et al., 2018***). Here, as in the example in ***Figure 3***, the conductances of the excitatory inputs representing 'yellow' and 'banana' increase in dendrite 1, strengthening the 'yellow banana' pairing, while the weights of synapses representing the feature 'red' decrease (***Figure 5***E). Conversely, the inhibitory synapses associated with the 'yellow' and 'banana' features in dendrite 1 are weakened, while those linked to the 'red' and 'strawberry' features in the same dendrite are strengthened (***Figure 5***D, right panel). This behaviour of the inhibitory plasticity rule effectively prevents the cell from spiking following activation of the irrelevant stimuli. In this particular example the 'red banana' combination is strongly inhibited, thereby effectively compartmentalizing dendrite 1 to be responsive to 'yellow banana' (since in this case no excitatory 'strawberry' input is present).

We also show the dynamics of peak calcium levels associated with both excitatory and inhibitory synapses for each task (***Figure 5***D, left panel, and ***Figure 5***E). For excitatory synapses, the patterns in peak NMDA calcium levels behave as in the example in ***Figure 3***C (compare to ***Figure 5***E). For inhibitory synapses, the peak calcium levels arising from voltage-gated calcium channels further exemplify how inhibition suppresses calcium activity arising from irrelevant stimuli while allowing that evoked by relevant stimuli (***Figure 5***D, left panel). The high threshold level of calcium for inhibitory synapses ($\theta_{\text{inh, high}}$) followed the highest excitatory synaptic activity, which in this case corresponded to the 'yellow banana' input.

We also compared the performance for different feature configurations with added inhibitory synapses to the results for the excitatory-only setup from ***Figure 4***B. The results show not only that learning with inhibitory plasticity is faster, i.e. requires fewer training examples, but also achieves high performance, nearing 100%. In the configurations with all four features in one dendrite 7/13 learned the NFBP, compared to 5/13 without inhibition (shown in the bar plot inset in ***Figure 5***B).

To understand why inhibitory plasticity improved learning, we examined the dynamics of L-type calcium during early learning. The presence of inhibition resulted in slightly lower initial L-type calcium levels compared to the excitatory-only setup, leading to reduced LTD at the start. This reduction in LTD helped prevent excessive weakening of synapses, increasing the likelihood of selecting one combination over others (***Figure 5—figure Supplement 1***B, showing excitatory synaptic weights without (left) and with inhibitory plasticity (middle): in the setup without inhibitory plasticity all four features weaken over time due to competition, while in contrast, the setup including inhibitory plasticity demonstrates better-regulated learning and more stable feature selection).

Inhibitory plasticity additionally produced a marked improvement in performance when the location of the synaptic clusters was varied, especially prominent at very distal locations (***Figure 5***C). Without inhibition it was difficult for synapses on distal dendrites to differentiate between relevant and irrelevant stimuli. In the middle of the learning, setups without inhibition typically spiked too much compared to setups including inhibition (e.g. spike for 'red banana' in the left panels of ***Figure 5—figure Supplement 1***A). This in turn caused increased LTD and decreased weights in the feature-unspecific synapses, leading to too little spiking at the end of the learning (e.g. no spike for 'yellow banana' in the right panels of ***Figure 5—figure Supplement 1***A). With inhibition, on the other hand, the contrast between relevant and irrelevant stimuli was larger. Larger contrast in turn gave less spiking for the wrong combination in the middle of learning and thereby less reduction of the weights of the feature-unspecific synapses. Therefore, large excitation in a single dendrite, combined with a larger drive of feature-unspecific synapses were enough to cause spiking in relevant stimuli at the end of learning, while spiking in irrelevant stimuli was suppressed (***Figure 5—figure***

Supplement 1A, red traces).

In summary, incorporating inhibitory plasticity suggests that inhibitory synapses may fine-tune dendritic responsiveness and enhance NFBP performance by preventing excitation from irrelevant stimuli, which in turn allows for more rapid, robust, and accurate learning.

Inhibitory plasticity improves performance with randomly distributed inputs

We finally challenged our plasticity rule by relaxing the assumption that single features are represented by pre-clustered synapses on specific dendritic branches. The behaviour of the synaptic plasticity rule was therefore investigated using a setup where 200 excitatory synapses, representing the different features, were randomly distributed across 30 dendrites. In this setup, each feature was represented by 40 synapses and an additional 40 represented the feature-unspecific synapses. ***Figure 6***A 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. To address the reduced effectiveness of non-clustered synaptic inputs in eliciting sufficient depolarization and calcium influx—key for learning and synaptic plasticity—we slightly increased the initial weights compared to the clustered setup to 0.3 ± 0.1 (around 0.75 nS). A branch-specific mechanism for creating non-linearities was also introduced by assuming that glutamate spillover occurs in all co-activated synapses 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 rise in calcium levels, endowing the dendrites with supralinear responses even without closely clustered synapses.

Starting with an excitation-only setup resulted in a moderate mean performance of about 65% (***Figure 6***B, 'Without inhibitory plasticity'). However, on a linear task (e.g. learning only 'yellow banana' or 'red strawberry') with randomly placed synapses the performance was close to 100% (even in so-called subthreshold learning, starting from an initial condition without spiking; see ***Figure 6—figure Supplement 1*** B and C).

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. In the right panel of ***Figure 6***B we show that, with inhibition, most individual distributions (22/31) reach a performance of about 75%, while a subset of distributions (5/31) achieve performance close to 100%. There are also a few distributions that partially solve the NFBP, reaching a performance between 75% and 100% (4/31).

To explore the model's ability to handle increased feature complexity, we extended the task to include two configurations: a subset of 9 feature combinations, including one non-linear combination, and a more challenging subset of 25 feature combinations, including two non-linear combinations (see ***Figure 6—figure Supplement 1***A). In these extended simulations, the model achieved accuracies of 80% and 75% for the 9- and 25-feature tasks, respectively, as shown in ***Figure 6—figure Supplement 1***A.

For a more granular analysis of the successful cases with randomly distributed synapses and assumed branch-specific plasticity, we looked in detail at the conductances of an example that successfully learned the NFBP. In ***Figure 6***C we present the somatic and dendritic voltage responses for four dendritic branches (d1–d4) where relevant stimuli were encoded. The voltage traces display somatic and dendritic responses to all four stimuli after learning. Additionally, the cumulative synaptic conductances for both excitatory and inhibitory inputs were calculated at the midpoint of each dendritic branch. Notably, dendrites 1 and 4 show enhanced excitatory inputs for the 'yellow banana' pattern and increased inhibitory inputs for the 'red strawberry' pattern, while dendrites 2 and 3 exhibit the reverse arrangement. The synapses in dendrites 1 to 4, whose weights are shown in panel C, are marked in ***Figure 6***A with black borders around the markers.

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 supra-linear 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 typically be encoded by the neuron.

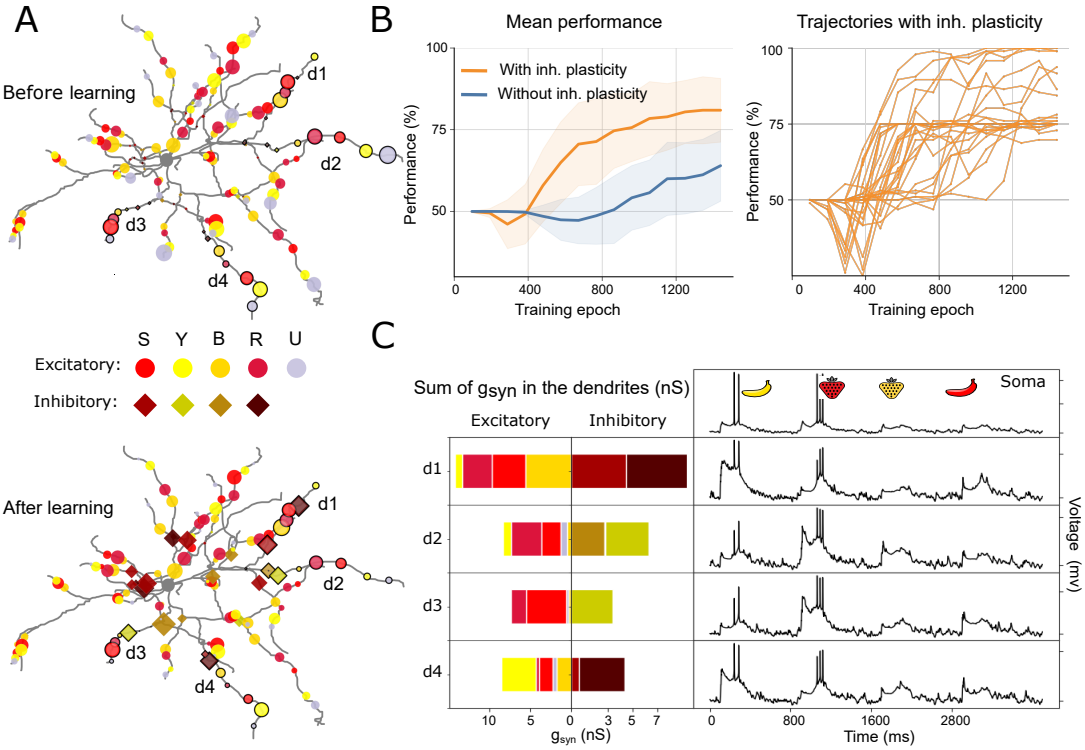

Figure 6. Performance analysis of learning using distributed synaptic inputs
(A): Example illustration of synaptic distribution before (top) and after learning (bottom) of the 200 excitatory and 60 inhibitory inputs. S — strawberry, Y — yellow, B — banana, R — red, U — feature-unspecific inputs. Bordered markers (circles and diamonds) indicate synapses whose conductances are shown in (C).
(B): Performance over training epochs with and without inhibitory plasticity. (Left) Mean performance of the setups with and without inhibitory plasticity. Shaded areas show standard deviation. (Right) Individual traces for the setup with inhibitory plasticity. Each of these individual trials (31 in total) uses a unique random distribution of synapses.
(C): Example of summed synaptic conductances (left) and voltage (right) in the soma, and four example dendrites (d1–d4) of the synaptic distribution in (A) (corresponding to a trial where the NFBP is successfully solved). The sums of both excitatory (left) and inhibitory (right) inputs are shown.

Figure 6—figure supplement 1. Other learning paradigms.

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 multicompartment 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 synap-

tic 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 synaptic input, was largest in the model by ***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 do not control the likelihood of somatic spiking as decisively. As the supralinear dendritic 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 synapses for each feature. In this scenario, only one relevant stimulus is sometimes learned in a dendritic branch by the randomly distributed synapses. In the random setup the glutamate spillover models were updated to include the synapses in the whole dendritic branch, building on the notion that the single branch acts as a single computational unit (***Losonczy and Magee, 2006***; ***Branco and Häusser, 2010***). The realism of this spillover assumption remains open to debate. In actual dendritic branches, however, the diffusion of signaling molecules likely plays a crucial role in both synaptic plasticity at existing synapses and for locally induced structural plasticity (***Nishiyama and Yasuda, 2015***; ***Chater et al., 2024***). These processes may both enhance branch-specific supralinearities in qualitatively similar ways as when using spillover, but the mechanisms are not explored in the current model.

By using a phenomenological inhibitory plasticity rule based on the BCM formalism (***Bienenstock et al., 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 et al., 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 et al., 2010***). The threshold for plateau potential 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.

In calcium-dependent plasticity models, a "no man’s land" has sometimes been proposed as a region where synaptic weights remain largely unchanged within a specific calcium range, reducing sensitivity to fluctuations and improving stability (***Lisman, 2001***; ***Moldwin et al., 2024***). For our inhibitory rule, which builds on a BCM-like mechanism, introduction of such a mechanism in future work could make the system less sensitive to fluctuations in calcium concentration.

Although our learning rule infrequently solves the NFBP when used with only randomly distributed synapses, it can always learn to perform a linearly separable task, such as learning to respond to only one relevant stimulus (e.g.,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, with metaplasticity playing a crucial role in this process. That is, rewards are then seen

very regularly as the neuron has learned to spike for the relevant stimuli, while omitted rewards rarely 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. However, 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***). Specifically 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 (metaplasticity), 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 LTP plastic range (LTP kernel).

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 prevent LTP in synapses that should ultimately undergo LTD. Hence, similarly to the classical loss-aversion tendency described in economic decision theory (***Kahneman and Tversky, 1979***), the model 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***; ***Iacaruso et al., 2017***; ***Scholl et al., 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 dopamine 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 is significantly facilitated by co-regulation of other neuromodulatory systems, such as when there is 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***).

Methods and Materials

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 changes in local calcium concentration resulting from synaptic activation patterns that evoke dendritic non-linearities, such as plateau potentials. Such events affecting the postsynaptic calcium concentration can 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 and Hines, 2006***). Here we will focus on the setup of the learning rule and only give a short summary of the neuron and synapse models, and emphasize changes compared to previously published versions. For a detailed description of the neuron model setup, see ***Lindroos et al.*** (***2018***); ***Lindroos and Hellgren Kotaleski*** (***2021***). The modeling of plateau potentials is explained in ***Trpevski et al.*** (***2023***). See also the section *Code and Software Availability* below.

Neuron model

In short, the dSPN model used here was taken from the library of biophysically detailed models in ***Lindroos and Hellgren Kotaleski*** (***2021***), 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. 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 (Cav2.3), T (Cav3.2 and Cav3.3), and L (Cav1.2 and Cav1.3); 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 ($[Ca]_{NMDA}$) is used, and another for L-type calcium concentration ($[Ca]_{L-type}$), to reflect the different synaptic plasticity responses of the dSPN'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 Cav1.2 and Cav1.3, respectively). Inspired by findings that calcium from various voltage-dependent calcium channels affects GABAergic plasticity in different systems (***Kurotani et al., 2008***; ***Hulme and Connelly, 2014***; ***Udakis et al., 2020***), we also incorporate a third calcium pool based on calcium influx from all voltage gated channels (T-, R-, L- and N-type, $[Ca]_V$). 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 [Ca] amplitudes and durations 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 et al., 2011***). Spatial calcium diffusion was not included in the present model.

Glutamatergic synaptic input

Clustered synapses were modeled following ***Trpevski et al.*** (***2023***) using a version of the saturating synapse model in ***Destexhe et al.*** (***1994***). The synaptic model included AMPA and NMDA conductances activated on spines as well as extrasynaptic NMDA conductances located on the dendritic shafts adjacent to the spine. The strength of the individual synapses on spines were scaled using a weight parameter (w) taking on positive values. More explicitly, in our model, the parameter w represents the percentage of the maximum synaptic conductance g_{max}. For instance, when $w = 0.5$, the synaptic conductance (g) is $0.5 \cdot g_{\text{max}}$. Importantly, changes in w affect both NMDA and

AMPA receptor conductances simultaneously, as both share the same scaling factor. Extrasynaptic conductances were activated following sufficiently large stimulation of nearby spines, resulting in glutamate spillover. Extrasynaptic NMDA synapses were included because they provide robust all-or-none dendritic plateaus (***Trpevski et al., 2023***). The maximal conductance of both synaptic and extrasynaptic NMDARs was set to 2.5 nS, which is lower than the 3.5 nS used in ***Trpevski et al.*** (***2023***). This reduction in conductance was implemented to compensate for the increased number of inputs in the model. In this model, spillover was triggered when the sum of the weights of active synapses in a cluster exceeded a threshold value of 2. For instance, eight synapses with individual weights greater than 0.25 would activate spillover. This threshold was lower than the value of 4 used in ***Trpevski et al.*** (***2023***) in order to facilitate spillover and enhance non-linearity in the context of distributed synapses. However, this scaling only affects the size of the clustered inputs presumed to give rise to plateau potentials, which likely varies across different neuronal subtypes or even between distinct dendritic branches within the same subtype. In the NEURON simulation environment, an integrate-and-fire cell was used to generate a spike in extrasynaptic NMDARs when the spillover threshold is reached. As synaptic weight increases, fewer synapses are required to meet the threshold. This also links long-term potentiation (LTP) to enhanced glutamate spillover. In addition to task-related inputs, the neuron also receives background noise, modeled as excitatory and inhibitory synapses distributed across the dendrites. Since the number of electrical compartments in the model is significantly lower than the number of synapses in a real SPN, these converging synapses were represented by a single non-saturating synapse with input frequency scaled accordingly. This was implemented using a dual exponential synapse model based on the NEURON `exp2syn` mechanism following ***Lindroos and Hellgren Kotaleski*** (***2021***) with time constants and other parameters adapted from ***Hjorth et al.*** (***2020***).

Learning rule

Excitatory synaptic plasticity

The learning rule is grounded in experimental evidence indicating that striatal LTP relies on NMDA receptor activation and the presence of dopamine, whereas LTD depends on the activation of L-type calcium channels (Cav1.3) and mGluR5 receptors in the absence of dopamine (***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. This rule outlines a reward-based learning mechanism where a dopamine peak enhances synaptic weight via an LTP plasticity kernel, while a dopamine pause reduces synaptic weight through an LTD kernel. Thus, dopamine functions as a switch, determining whether LTP or LTD pathways are activated (as illustrated in ***Figure 1***D). Notably, synaptic weights affect both AMPA and NMDA receptors, consistent with findings that the NMDA-to-AMPA ratio remains unchanged following LTP (***Watt et al., 2004***). Additionally, the previously described spillover mechanism is linked to synaptic weights, representing an attempt to associate plasticity with increased glutamate spillover resulting from the retraction of astrocytic processes from the spine (***Henneberger et al., 2020***). The excitatory synaptic plasticity rule encompasses key components such as LTP, LTD, and metaplasticity, each of which will be elaborated upon in the subsequent sections.

LTP

The LTP process is initiated by elevated dopamine levels, following calcium influx through NMDA receptors. Synaptic strength is adjusted using the LTP kernel, a bell-shaped function where the maximum increase occurs when peak NMDA calcium levels are at the kernel's midpoint; deviations above or below this midpoint result in smaller increases (***Figure 7***A, top panel). The left side of the LTP kernel represents the lower calcium threshold necessary for LTP induction, while the right side indicates an upper calcium threshold beyond which LTP does not occur. This framework, in conjunction with metaplasticity, establishes an upper boundary for synaptic strength, ensuring that synaptic weights do not exceed a certain limit (***Zenke et al., 2017***; ***Zenke and Gerstner, 2017***).

Since each synapse experiences unique calcium levels (due to different synaptic conductance, local input resistance, etc), adjusting the calcium dependence of the LTP kernel allows for individualized tuning of synaptic weights (see ***Figure 7***A and C for illustrations).

The update of the synaptic weights follow a rule formulated using the derivative of the sigmoid function (***Equation 2***). The sigmoid function (***Equation 1***) transforms the calcium concentration into a normalized range between 0 and 1:

$$\sigma(Ca, \theta, \beta) = \frac{1}{1 + e^{-\beta(Ca-\theta)}} \tag{1}$$

where θ represents the midpoint of the transition, representing the calcium concentration at which the function increases most rapidly, and β determines the steepness of this transition.

The derivative of the sigmoid function, given by:

$$\sigma'(Ca, \theta, \beta) = \beta \cdot \sigma(Ca, \theta, \beta) \cdot (1 - \sigma(Ca, \theta, \beta)) \tag{2}$$

produces a bell-shaped curve, which forms the LTP kernel. This kernel peaks at the midpoint θ, where synaptic modification is maximized, and decays symmetrically on each side. The parameter β here controls the width of the kernel, thereby regulating the sensitivity of synaptic weight changes to calcium fluctuations.

Based on this general description of the mathematical formulas underlying the LTP kernel, the LTP learning rule for the synaptic weights is formulated as:

$$\Delta w_{\text{ltp}} = \eta_{\text{ltp}} \cdot \sigma'(Ca_{\text{NMDA}}, \theta_{\text{ltp}}, \beta_{\text{ltp}}) \tag{3}$$

where Δw_{ltp} represents the change in synaptic weight, η_{ltp} is the learning rate, and $[Ca]_{\text{NMDA}}$ denotes the calcium concentration mediated by NMDA receptors. The parameters θ_{ltp} and β_{ltp} define the midpoint and slope of the sigmoid curve specific to LTP. The parameter values used in these equations are provided in ***Table 1***).

Table 1. Excitatory plasticity parameters.

Parameter	Description	Value
η_{ltp}	Learning rate for $[\text{Ca}]_{\text{NMDA}}$-dependent plasticity	1.5×10^{-5} µS mM^{-1} ms^{-1}
θ_{ltp}	Midpoint of LTP kernel (initial value)	0.02 mM
β_{ltp}	Controls width of LTP kernel	1.0×10^{3} mM^{-1}
η_{ltd}	Learning rate for $[\text{Ca}]_{\text{L-type}}$-dependent plasticity	3×10^{-3} ms^{-1}mM^{-1}
θ_{ltd}	Constant LTD threshold	7×10^{-5} mM
β_{ltd}	Controls steepness of LTD threshold	1.0×10^{5} mM^{-1}
$\eta_{s,\text{ltp}}$	Rate for shifting θ_{ltp} during dopamine peaks	1×10^{-7} mM^{2}ms^{-1}
$\eta_{s,\text{ltd}}$	Rate for shifting θ_{ltp} during dopamine pauses	4×10^{-7} mM^{2}ms^{-1}
β_{mp}	Controls width of metaplasticity kernel	334 mM^{-1}

LTD

The LTD process is triggered by a dopamine pause and is dependent on L-type calcium. The LTD plasticity kernel (***Equation 4***; see ***Figure 7***B, *top* panel) describes a threshold calcium level necessary for LTD to occur (***Shindou et al., 2011***). Once this threshold is exceeded, the decrease in synaptic weight is linearly proportional to the amplitude of the peak calcium level and is scaled by the learning rate (η_{ltd}). The peak-calcium threshold is implemented with a sigmoid function whose slope parameter (β_{ltd}) is set high so the curve approximates a step function. The threshold is fixed at 70 nM. The NMDA-calcium dependence of the LTP kernel is also increased during LTD, as described in the next section (***Figure 7***B, *bottom* panel, and ***Figure 7***C, right panel).

$$\Delta w_{\text{ltd}} = \eta_{\text{ltd}} \cdot Ca_{\text{v1.3}} \cdot \sigma(Ca_{\text{v1.3}}, \theta_{\text{ltd}}, \beta_{\text{ltd}}) \cdot w_{\text{ltd}} \tag{4}$$

Metaplasticity

Metaplasticity is a form of regulatory mechanism changing the state of the synapses in such a way as to influence subsequent learning (***Abraham, 2008***). In our model we implemented this as a reward-dependent change of the calcium concentration over which the LTP kernel was operating. Specifically, activation of the LTD pathway triggered a shift of the LTP kernel towards higher calcium concentrations while activation of the LTP pathway pushed it in the opposite direction.

Metaplasticity was also implemented using a bell-shaped kernel with the same midpoint as the LTP kernel, but with a wider calcium dependence. The position of the metaplasticity kernel is updated at the same time with the LTP kernel, so that both kernels remain centered at the same calcium level (***Figure 7***A–C) . This setup with the two dynamically moving kernels together allowed for a wide range of calcium levels to induce plasticity.

In addition, using a wider metaplasticity kernel ensured the following:

- Following activation of the LTP pathway (i.e., after a dopamine peak, see ***Figure 7***C ,left panel for an illustration):
  - In synapses with low calcium levels, the kernel will be shifted closer to the observed calcium level and thereby eventually enable LTP in synapses that are regularly activated during rewards, despite that they initially don't generate big elevations in calcium.
  - In synapses with already high calcium levels (above the kernel midpoint), the kernel will be shifted away from the observed calcium level, and thereby protect the synapse from excessive LTP and instead stabilize the weight (as illustrated with the red circle moving to the blue one in ***Figure 7***C,left panel).
- Following activation of the LTD pathway (i.e., after a dopamine pause), the kernel will be shifted towards higher calcium levels and thereby reduce the likelihood of LTP in these synapses that are weakened following LTD (see ***Figure 7***C, right panel for an illustration).

The update of the LTP kernel was further asymmetric following activation of the LTD or LTP pathways in such a way that the LTD pathway caused a larger shift of the metaplasticity kernel than the LTP pathway. 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.

It also allows initially weaker synapses to be recruited for LTP if they are more consistently co-active with a reward than with an omitted reward (an illustration of this can be seen in ***Figure 5***E, where one of the yellow synapses is strengthened in the middle of the training session).

The LTP kernel was updated according to ***Equation 5***, based on the learning rate (η_s), the peak NMDA calcium level ($[\text{Ca}]_{\text{NMDA}}$), and the slope of the sigmoid curve (β_{mp}). The value of the slope β_{mp} creates a wider bell-shaped function in ***Equation 5***. The rate η_s captures the described asymmetry in how the metaplasticity kernel is shifted in response to dopamine feedback, where the specific rates following dopamine peaks and pauses are denoted $\eta_{s,\text{ltp}}$ and $\eta_{s,\text{ltd}}$, respectively. Here, $\eta_{s,\text{ltd}}$ is four times larger than $\eta_{s,\text{ltp}}$ (see ***Table 1***).

$$\Delta\theta_{\text{ltp}} = \eta_s \cdot \sigma'(Ca_{\text{NMDA}}, \theta_{\text{ltp}}, \beta_{\text{mp}}) \tag{5}$$

Pseudocode for the Excitatory Plasticity Rule

Inputs:

- $[\text{Ca}]_{\text{NMDA}}$: Calcium level associated with the NMDA channel
- $[\text{Ca}]_{\text{Ltype}}$: Calcium level associated with the L-type channel
- `dopamine_level`: Dopamine level (1 for peak, -1 for pause, 0 for basal)

Output:

- Δw: Change in synaptic weight

Algorithm:

- `begin`
  - `if dopamine_level == 1 then`
  - `# Compute LTP weight change and adjust kernel midpoint for metaplasticity`
    - $\Delta w =$ solve ***Equation 3*** with $[\text{Ca}]_{\text{NMDA}}$, η_{ltp}, θ_{ltp}, and β_{ltp}
    - update θ_{ltp} using ***Equation 5*** with $[\text{Ca}]_{\text{NMDA}}$, $\eta_{s,\text{ltp}}$, β_{mp}
  - `else if dopamine_level == -1 then`
  - `# Compute LTD weight change and adjust LTP kernel midpoint for metaplasticity`
    - $\Delta w =$ solve ***Equation 4*** with $[\text{Ca}]_{\text{Ltype}}$, η_{ltd}, θ_{ltd}, and β_{ltd}
    - update θ_{ltp} using ***Equation 5*** with $[\text{Ca}]_{\text{NMDA}}$, $\eta_{s,\text{ltd}}$, β_{mp}
  - `else if dopamine_level == 0 then`
  - `# No plasticity under basal dopamine`
    - $\Delta w = 0$
  - `return` Δw
- `end`

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 et al., 1982***). In the BCM rule there is a threshold level of synaptic activity below which LTD is triggered and above which LTP is triggered (***Figure 7***D). Inspired by ***Gandolfi et al.*** (***2020***); ***Chapman et al.*** (***2022***); ***Ravasenga et al.*** (***2022***), we use calcium from all voltage-gated calcium channels ($[\text{Ca}]_V$) 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 voltage-gated calcium influx, but without reliance on dopamine or explicit feedback. It operates at a slower pace to ensure that it follows the average excitatory activity levels, and is meant to enhance the contrast in activity by amplifying local differences in excitatory synaptic efficacy.

Equation 6 describes how the inhibitory synaptic weights change based on the calcium concentration (calculated from calcium influx from all voltage-gated ion channels close to the synapse).

$$\Omega(Ca, \theta_{\text{inh,low}}, \theta_{\text{inh,high}}, a, b) = \frac{a}{1 + e^{-\beta_{\text{inh}}(Ca - 0.004 \cdot \theta_{\text{inh,low}})}} + \frac{b}{1 + e^{-\beta_{\text{inh}}(Ca - 0.004 \cdot \theta_{\text{inh,high}})}} \tag{6}$$

It consists of two sigmoidal terms, where the parameters a and b determine the contribution of each component, β_{inh} determines the steepness of the transitions and $\theta_{\text{inh,low}}$ and $\theta_{\text{inh,high}}$ determines the calcium concentration at which each sigmoid is half activated (see ***Figure 7***D). Based on this, and modulated by a rate constant, η_{act}, ***Equation 7*** describes the specific rule used to update the weights of the inhibitory synapses.

$$\Delta w_{\text{inh}} = \eta_{\text{act}} \cdot \Omega(Ca_v, \theta_{\text{inh,low}}, \theta_{\text{inh,high}}, a_{\text{inh}}, b_{\text{inh}}) \cdot w_{\text{inh}} \cdot (w_{\text{inh}}^{\text{max}} - w_{\text{inh}}) \tag{7}$$

The weight change also depends on the magnitude of the weights themselves, in such a way that intermediate weights update faster than small or large weights. This helps stabilize the weights of the inhibitory synapses at the end of learning, and prevents the weight from taking on large or

negative values. Large weights are then close to $w_{\text{inh}}^{\text{max}}$ at the end of learning and small weights close to 0 (see ***Figure 5***D for an example and ***Table 2*** for parameter values).

The learning rule further differentiates between active and inactive synapses. For active synapses: i) if the calcium concentration exceeds $\theta_{\text{inh,high}}$, the weight is depressed, ii) if calcium falls between $\theta_{\text{inh,low}}$ and $\theta_{\text{inh,high}}$, the weight is potentiated and iii) when calcium is below $\theta_{\text{inh,low}}$, little or no weight change occurs. For inactive synapses, the rule is reversed: i) when calcium exceeds $\theta_{\text{inh,high}}$, the weight is potentiated, ii) if calcium falls between $\theta_{\text{inh,low}}$ and $\theta_{\text{inh,high}}$, the weight is depressed and iii) no changes occur when calcium is below $\theta_{\text{inh,low}}$ (as in the case of active synapses). The schematic in ***Figure 7***D (top panel) illustrates these weight updates, showing how the learning rule differentially affects active and inactive synapses. The black curves represent the relationship between calcium concentration and synaptic weight change, with the beige and blue shaded regions indicating potentiation and depression zones for each case.

The half activation thresholds, $\theta_{\text{inh,low}}$, and $\theta_{\text{inh,high}}$, are further dynamically updated to track the calcium levels of the synapse. These updates of the thresholds ensure that the plasticity rule remains sensitive to changes in calcium concentration and stabilize synaptic learning.***Equation 8*** and ***Equation 9*** describe these adaptive dynamics. This is, in fact, metaplasticity of the inhibitory synapses.

$$\Delta\theta_{\text{inh,high}} = \eta_{\text{inh,high}} \cdot \Omega(Ca_v, \theta_{\text{inh,low}}, \theta_{\text{inh,high}}, a_{\text{inh,high}}, b_{\text{inh,high}}) \tag{8}$$

$$\Delta\theta_{\text{inh,low}} = \eta_{\text{inh,low}} \cdot \Omega(Ca_v + c, \theta_{\text{inh,low}}, \theta_{\text{inh,high}}, a_{\text{inh,low}}, b_{\text{inh,low}}) \tag{9}$$

Table 2. Inhibitory plasticity parameters.

Parameter	Description	Value
β_{inh}	Controls steepness of the curves	2.5×10^3 mM^{-1}
η_{act}	Learning rate for active/inactive inhibitory weight	active: -0.055, inactive: 0.055 µs^{-1} ms^{-1}
a_{inh}	Controls inhibitory weight	-1
b_{inh}	Controls inhibitory weight	3
$w_{\text{inh}}^{\text{max}}$	Maximum synaptic strength	0.005 µS
$\eta_{\text{inh,high}}$	Learning rate for modifying $\theta_{\text{inh,high}}$	9×10^{-4} mM ms^{-1}
$a_{\text{inh,high}}$	Controls $\theta_{\text{inh,high}}$	-1
$b_{\text{inh,high}}$	Controls $\theta_{\text{inh,high}}$	3
$\eta_{\text{inh,low}}$	Learning rate for modifying $\theta_{\text{inh,low}}$	-5×10^{-5} mM ms^{-1}
$a_{\text{inh,low}}$	Controls $\theta_{\text{inh,low}}$	-2
$b_{\text{inh,low}}$	Controls $\theta_{\text{inh,low}}$	3
c	Calcium concentration offset	6×10^{-4} mM

The upper threshold $\theta_{\text{inh,high}}$ shifts towards the highest observed calcium level in the synapse, while the lower threshold $\theta_{\text{inh,low}}$ moves to a level below the maximum. The rate of these adjustments is determined by the parameters $\eta_{\text{inh,high}}$ and $\eta_{\text{inh,low}}$, which scale the influence of calcium activity on the threshold dynamics. The parameters $a_{\text{inh,high}}, b_{\text{inh,high}}$ determine the rate of change of $\theta_{\text{inh,high}}$, while $a_{\text{inh,low}}, b_{\text{inh,low}}$ define the rate of change of $\theta_{\text{inh,low}}$ (see ***Table 2*** for parameter values). The parameter c in ***Equation 9*** is used to make the two curves for the high and low thresholds ($\theta_{\text{inh,low}}, C_{\text{inh,high}}$) intersect at zero which minimizes the fluctuations around the transition point (green asterisk in ***Figure 7***D, bottom panel).

Pseudocode for the Inhibitory Plasticity Rule

Inputs:

- $[\text{Ca}]_V$: Ca concentration from voltage-gated calcium channels

Output:

- `return` Δw_{inh}: Change in synaptic weight

Algorithm:

- `begin`
  - `if synapse_active then`
  - `# Active synapse: Compute weight change using negative learning rate` (η_{act})
    * Δw_{inh} = solve ***Equation 7*** with $[\text{Ca}]_V$, η_{act}, $\theta_{\text{inh,high}}$, $\theta_{\text{inh,low}}$, a_{inh} , b_{inh}, $W_{\text{inh}}^{\text{max}}$
    * update $\theta_{\text{inh,high}}$ using ***Equation 8*** with $\theta_{\text{inh,high}}$, $\theta_{\text{inh,low}}$, $\eta_{\text{inh,high}}$, $a_{\text{inh,high}}$,$b_{\text{inh,high}}$
    * update $\theta_{\text{inh,low}}$ using ***Equation 9*** with $\theta_{\text{inh,high}}$, $\theta_{\text{inh,low}}$, $\eta_{\text{inh,low}}$,$a_{\text{inh,low}}$, $b_{\text{inh,low}}$, c
  - `else`
  - `# Inactive synapse: Compute weight change using positive learning rate` (η_{act})
    * Δw_{inh} = solve ***Equation 7*** with $[\text{Ca}]_V$, η_{act}, $\theta_{\text{inh,high}}$, $\theta_{\text{inh,low}}$, a_{inh} , b_{inh}, $W_{\text{inh}}^{\text{max}}$
    * update $\theta_{\text{inh,high}}$ using ***Equation 8*** with $\theta_{\text{inh,high}}$, $\theta_{\text{inh,low}}$, $\eta_{\text{inh,high}}$, $a_{\text{inh,high}}$, $b_{\text{inh,high}}$
    * update $\theta_{\text{inh,low}}$ using ***Equation 9*** with $\theta_{\text{inh,high}}$, $\theta_{\text{inh,low}}$, $\eta_{\text{inh,low}}$, $a_{\text{inh,low}}$, $b_{\text{inh,low}}$, c
  - `return` Δw_{inh}
- `end`

Training Procedure

The features of a stimulus were represented by the shape and color of bananas and strawberries. We presented the neuron with a sequence of stimuli, typically around 960, each belonging to one of four possible feature combinations, as illustrated in ***Figure 1***. The four stimuli: 'red strawberry', 'yellow banana', 'red banana' and 'yellow strawberry', were presented in random order three times each within a block of 12 stimuli, followed by another reshuffled block of 12, and so on. Each stimulus presentation lasted 20 ms, during which time all the stimuli-related synapses receive one randomly-timed spike per synapse(see ***Figure 3***B). Whenever relevant, inhibitory synapses were activated concurrently with the excitatory synapses and the stimuli lasted for 100 ms. The stimuli presentation was followed by a reward cue lasting 50 ms that arrived 300 ms after the stimulus onset. The reward cue was represented with +1 for relevant stimuli if the neuron spiked (representing a peak in dopamine), with -1 for irrelevant stimuli if the neuron spiked (representing a dopamine pause) and 0 for baseline levels or when the neuron was silent. However, in an additional subthreshold learning task, the reward cue was delivered even in the absence of somatic spiking for active glutamate synapses, such that if they were active following the relevant stimuli ('red strawberry' and 'yellow banana') they got a dopamine peak, and a dopamine pause for irrelevant stimuli (***Figure 6—figure Supplement 1***B and C). The time between two stimuli was 800 ms, long enough to allow for the voltage, calcium (and all other state variables in the model) to return to their baseline values (***Figure 3***B). During the stimulus presentation, the stimuli-related synapses receive one randomly-timed spike per synapse. Further, the learning rule was on during the whole procedure, such that synapses were continuously updated throughout the simulation. Hence, there were no separate training and testing phases in which synapses were plastic and frozen, respectively.

Clustered setup

This setup was based on the assumption of pre-existing clustered synapses for each feature. 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 ***Figure 3***A, ***Figure 4***A, and ***Figure 4***D for examples). 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 ***Figure 2***A). The features representing a stimulus (clustered synapses) were active within 20 ms, while feature-unspecific input was activated over a 50 ms window. In simulations with inhibition, four inhibitory synapses, each one representing a feature, were placed near each cluster (***Figure 5***A). Within this setup a single stimulus activated both excitatory and inhibitory synapses. To match the level of depolarization seen in our excitatory-only setup, the number of feature-unspecific synapses were increased to 144 in the simulations including inhibitory plasticity. The initial conductance of inhibitory synapses was set to 0.1 ± 0.01 nS.

Distributed setup

In contrast, this setup examined 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 (***Figure 6***A). 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 6***A illustrates this setup, exemplifying the pre- and post-learning synaptic weights for both excitatory and inhibitory synapses.

Code and Data Availability

The modeling code used in this study is structured for reproducibility and is available at https://github.com/zahradd/dSPN-learning-rule. This repository includes the simulations underlying the figures in the manuscript, along with a README file providing detailed instructions for environment setup and result reproduction. Additionally, model components adapted from ***Lindroos and Hellgren Kotaleski*** (***2021***) and ***Trpevski et al.*** (***2023***) are available on ModelDB (accession numbers: 266775 and 2017143).

Acknowledgments

We would like to thank the members of the Hellgren Kotaleski laboratory for their helpful discussions on various aspects of the manuscript. DT extends gratitude to Ana Kalajdjieva for illustrating the mouse brain.

We acknowledge the use of Fenix Infrastructure resources, which are partially funded by the European Union's Horizon 2020 research and innovation programme through the ICEI project under 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.

This study was supported by the Swedish Research Council (VR-M-2020-01652), the Swedish e-Science Research Centre (SeRC), Science for Life Laboratory, EU/Horizon 2020 No. 945539 (HBP SGA3) and No. 101147319 (EBRAINS 2.0 Project), the 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.

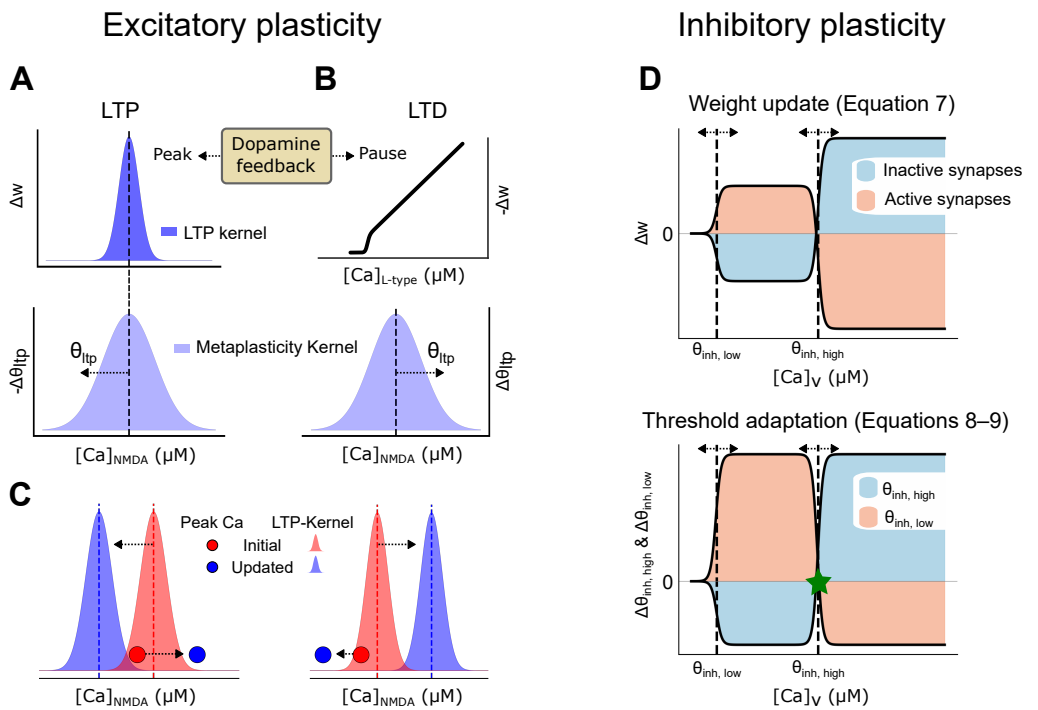

Figure 7. Synaptic Plasticity Rules: Calcium and Dopamine Interactions in Synaptic Weight Modification

(A): Synaptic weight updates following a dopamine peak. (Top) The LTP kernel is a bell-shaped curve describing the amount of weight increase, which happens over a region of $[Ca]_{NMDA}$. (Bottom) A wider bell-shaped kernel, i.e., the metaplasticity kernel, determines how the LTP kernel is shifted along the calcium level ($[Ca]_{NMDA}$) axis following a peak in dopamine.

(B): Synaptic weight updates following a dopamine pause. (Top) The LTD plasticity kernel. The LTD threshold is constant and set at 70 nM. (Bottom) Metaplasticity describing how the LTP kernel shifts along the calcium axis following a dopamine pause.

(C): A schematic of how the LTP kernel is shifted following a dopamine peak (left) and pause (right), together with examples of NMDA calcium levels before and after the shift. The NMDA calcium levels change following the potentiation or suppression of the synapse (illustrated with the red circle jumping to the blue circle). The LTP kernel also moves as indicated by the arrows and colors following a peak or a pause, respectively.

(D): Illustration of the inhibitory plasticity rule. (Top) Changes in synaptic weight for active (beige) and inactive (blue) synapses based on voltage-dependent calcium levels in the dendritic shaft at the location of the inhibitory synapse. The dashed lines show the minimum/lower ($\theta_{inh,low}$) and maximum/higher ($\theta_{inh,high}$) thresholds. (Bottom) Functions for updating the thresholds of the upper panel, depending on the voltage-dependent calcium level. The asterisk denotes the calcium level where the curves for active and inactive synapses meet, which is a point of zero update.

References

Abraham WC, Bear MF. Metaplasticity: the plasticity of synaptic plasticity. Trends Neurosci. 1996 Apr; 19(4):126–130.

Abraham WC. Metaplasticity: tuning synapses and networks for plasticity. Nat Rev Neurosci. 2008 May; 9(5):387.

Antic SD, Zhou WL, Moore AR, Short SM, Ikonomu KD. The decade of the dendritic NMDA spike. J Neurosci Res. 2010 Nov; 88(14):2991–3001.

Benna MK, Fusi S. Computational Principles of Synaptic Memory Consolidation. Nature Neuroscience. 2016; 19(12):1697–1706. doi: 10.1038/nn.4401.

Bicknell BA, Häusser M. A synaptic learning rule for exploiting nonlinear dendritic computation. Neuron. 2021 Dec; 109(24):4001–4017.e10.

Bienenstock EL, Cooper LN, Munro PW. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. J Neurosci. 1982 Jan; 2(1):32–48.

Branco T, Clark BA, Häusser M. Dendritic discrimination of temporal input sequences in cortical neurons. Science. 2010 Sep; 329(5999):1671–1675.

Branco T, Häusser M. The single dendritic branch as a fundamental functional unit in the nervous system. Curr Opin Neurobiol. 2010 Aug; 20(4):494–502.

Bruce NJ, Narzi D, Trpevski D, van Keulen SC, Nair AG, Röthlisberger U, Wade RC, Carloni P, Hellgren Kotaleski J. Regulation of adenylyl cyclase 5 in striatal neurons confers the ability to detect coincident neuromodulatory signals. PLoS Comput Biol. 2019 Oct; 15(10):e1007382.

Carnevale NT, Hines ML. The NEURON Book. Cambridge University Press; 2006.

Carter AG, Sabatini BL. State-dependent calcium signaling in dendritic spines of striatal medium spiny neurons. Neuron. 2004 Oct; 44(3):483–493.

Chapman CA, Nuwer JL, Jacob TC. The Yin and Yang of GABAergic and Glutamatergic Synaptic Plasticity: Opposites in Balance by Crosstalking Mechanisms. Front Synaptic Neurosci. 2022 May; 14:911020.

Chater TE, Eggl MF, Goda Y, Tchumatchenko T. Competitive processes shape multi-synapse plasticity along dendritic segments. Nat Commun. 2024 Aug; 15(1):7572.

Chen SX, Kim AN, Peters AJ, Komiyama T. Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning. Nat Neurosci. 2015 Aug; 18(8):1109–1115.

Chen X, Leischner U, Rochefort NL, Nelken I, Konnerth A. Functional mapping of single spines in cortical neurons in vivo. Nature. 2011 Jun; 475(7357):501–505.

Cichon J, Gan WB. Branch-specific dendritic Ca(2+) spikes cause persistent synaptic plasticity. Nature. 2015 Apr; 520(7546):180–185.

Clopath C, Büsing L, Vasilaki E, Gerstner W. Connectivity Reflects Coding: A Model of Voltage-Based STDP with Homeostasis. Nature Neuroscience. 2010; 13(3):344–352. doi: 10.1038/nn.2479.

Day M, Belal M, Surmeier WC, Melendez A, Wokosin D, Tkatch T, Clarke VRJ, Surmeier DJ. GABAergic regulation of striatal spiny projection neurons depends upon their activity state. PLoS Biol. 2024 Jan; 22(1):e3002483.

Destexhe A, Mainen ZF, Sejnowski TJ. An efficient method for computing synaptic conductances based on a kinetic model of receptor binding. Neural Comput. 1994 Jan; 6(1):14–18.

Dorman DB, Jędrzejewska-Szmek J, Blackwell KT. Inhibition enhances spatially-specific calcium encoding of synaptic input patterns in a biologically constrained model. Elife. 2018 Oct; 7.

Doron M, Chindemi G, Muller E, Markram H, Segev I. Timed Synaptic Inhibition Shapes NMDA Spikes, Influencing Local Dendritic Processing and Global I/O Properties of Cortical Neurons. Cell Rep. 2017 Nov; 21(6):1550–1561.

Du K, Wu YW, Lindroos R, Liu Y, Rózsa B, Katona G, Ding JB, Kotaleski JH. Cell-type-specific inhibition of the dendritic plateau potential in striatal spiny projection neurons. Proc Natl Acad Sci U S A. 2017 Sep; 114(36):E7612–E7621.

Fino E, Paille V, Cui Y, Morera-Herreras T, Deniau JM, Venance L. Distinct coincidence detectors govern the corticostriatal spike timing-dependent plasticity. J Physiol. 2010 Aug; 588(Pt 16):3045–3062.

Fisher SD, Robertson PB, Black MJ, Redgrave P, Sagar MA, Abraham WC, Reynolds JNJ. Reinforcement determines the timing dependence of corticostriatal synaptic plasticity in vivo. Nat Commun. 2017 Aug; 8(1):334.

Frey U, Schollmeier K, Reymann KG, Seidenbecher T. Asymptotic hippocampal long-term potentiation in rats does not preclude additional potentiation at later phases. Neuroscience. 1995 Aug; 67(4):799–807.

Fusi S, Drew PJ, Abbott LF. Cascade Models of Synaptically Stored Memories. Neuron. 2005; 45(4):599–611. doi: 10.1016/j.neuron.2005.02.001.

Gandolfi D, Bigiani A, Porro CA, Mapelli J. Inhibitory Plasticity: From Molecules to Computation and Beyond. Int J Mol Sci. 2020 Mar; 21(5).

Gao PP, Graham JW, Zhou WL, Jang J, Angulo S, Dura-Bernal S, Hines M, Lytton WW, Antic SD. Local glutamate-mediated dendritic plateau potentials change the state of the cortical pyramidal neuron. J Neurophysiol. 2021 Jan; 125(1):23–42.

Gidon A, Zolnik TA, Fidzinski P, Bolduan F, Papoutsi A, Poirazi P, Holtkamp M, Vida I, Larkum ME. Dendritic action potentials and computation in human layer 2/3 cortical neurons. Science. 2020 Jan; 367(6473):83–87.

Grienberger C, Chen X, Konnerth A. Inhibitory Plasticity Enables Branch-Specific Integration of Inputs. Nature Neuroscience. 2017; 20(10):1382–1390. doi: 10.1038/nn.4631.

Henneberger C, Bard L, Panatier A, Reynolds JP, Kopach O, Medvedev NI, Minge D, Herde MK, Anders S, Kraev I, Heller JP, Rama S, Zheng K, Jensen TP, Sanchez-Romero I, Jackson CJ, Janovjak H, Ottersen OP, Nagelhus EA, Oliet SHR, et al. LTP Induction Boosts Glutamate Spillover by Driving Withdrawal of Perisynaptic Astroglia. Neuron. 2020 Dec; 108(5):919–936.e11.

van den Heuvel MP, Sporns O. Network hubs in the human brain. Trends Cogn Sci. 2013 Dec; 17(12):683–696.

Higley MJ, Sabatini BL. Competitive regulation of synaptic Ca2+ influx by D2 dopamine and A2A adenosine receptors. Nat Neurosci. 2010 Aug; 13(8):958–966.

Hjorth JJJ, Kozlov A, Carannante I, Frost Nylén J, Lindroos R, Johansson Y, Tokarska A, Dorst MC, Suryanarayana SM, Silberberg G, Hellgren Kotaleski J, Grillner S. The microcircuits of striatum in silico. Proc Natl Acad Sci U S A. 2020 Apr; 117(17):9554–9565.

Huang YY, Colino A, Selig DK, Malenka RC. The influence of prior synaptic activity on the induction of long-term potentiation. Science. 1992 Feb; 255(5045):730–733.

Hulme SR, Connelly WM. L-type calcium channel-dependent inhibitory plasticity in the thalamus. J Neurophysiol. 2014 Nov; 112(9):2037–2039.

Hwang FJ, Roth RH, Wu YW, Sun Y, Kwon DK, Liu Y, Ding JB. Motor learning selectively strengthens cortical and striatal synapses of motor engram neurons. Neuron. 2022 Sep; 110(17):2790–2801.e5.

Iacaruso MF, Gasler IT, Hofer SB. Synaptic organization of visual space in primary visual cortex. Nature. 2017 Jul; 547(7664):449–452.

Jia H, Rochefort NL, Chen X, Konnerth A. Dendritic organization of sensory input to cortical neurons in vivo. Nature. 2010 Apr; 464(7293):1307–1312.

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

Ju N, Li Y, Liu F, Jiang H, Macknik SL, Martinez-Conde S, Tang S. Spatiotemporal functional organization of excitatory synaptic inputs onto macaque V1 neurons. Nat Commun. 2020 Feb; 11(1):697.

Kahneman D, Tversky A. Prospect theory: An analysis of decision under risk. Econometrica. 1979 Mar; 47(2):263.

Kawai R, Markman T, Poddar R, Ko R, Fantana AL, Dhawale AK, Kampff AR, Ölveczky BP. Motor cortex is required for learning but not for executing a motor skill. Neuron. 2015 May; 86(3):800–812.

Kerlin A, Mohar B, Flickinger D, MacLennan BJ, Dean MB, Davis C, Spruston N, Svoboda K. Functional clustering of dendritic activity during decision-making. Elife. 2019 Oct; 8.

Kleindienst T, Winnubst J, Roth-Alpermann C, Bonhoeffer T, Lohmann C. Activity-dependent clustering of functional synaptic inputs on developing hippocampal dendrites. Neuron. 2011 Dec; 72(6):1012–1024.

Kurotani T, Yamada K, Yoshimura Y, Crair MC, Komatsu Y. State-dependent bidirectional modification of somatic inhibition in neocortical pyramidal cells. Neuron. 2008 Mar; 57(6):905–916.

Larkum ME, Nevian T, Sandler M, Polsky A, Schiller J. Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. Science. 2009 Aug; 325(5941):756–760.

Lavzin M, Rapoport S, Polsky A, Garion L, Schiller J. Nonlinear dendritic processing determines angular tuning of barrel cortex neurons in vivo. Nature. 2012 Oct; 490(7420):397–401.

Legenstein R, Maass W. Branch-specific plasticity enables self-organization of nonlinear computation in single neurons. J Neurosci. 2011 Jul; 31(30):10787–10802.

Lindroos R, Dorst MC, Du K, Filipović M, Keller D, Ketzef M, Kozlov AK, Kumar A, Lindahl M, Nair AG, Pérez-Fernández J, Grillner S, Silberberg G, Hellgren Kotaleski J. 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. Front Neural Circuits. 2018 Feb; 12:3.

Lindroos R, Hellgren Kotaleski J. Predicting complex spikes in striatal projection neurons of the direct pathway following neuromodulation by acetylcholine and dopamine. Eur J Neurosci. 2021 Apr; 53(7):2117–2134.

Lisman JE. Three Ca2+ levels affect plasticity differently: the LTP zone, the LTD zone and no man's land. J Physiol. 2001 Apr; 532(Pt 2):285.

Losonczy A, Magee JC. Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons. Neuron. 2006 Apr; 50(2):291–307.

Lovett-Barron M, Turi GF, Kaifosh P, Lee PH, Bolze F, Sun XH, Nicoud JF, Zemelman BV, Sternson SM, Losonczy A. Regulation of neuronal input transformations by tunable dendritic inhibition. Nat Neurosci. 2012 Jan; 15(3):423–30, S1–3.

Major G, Polsky A, Denk W, Schiller J, Tank DW. Spatiotemporally graded NMDA spike/plateau potentials in basal dendrites of neocortical pyramidal neurons. J Neurophysiol. 2008 May; 99(5):2584–2601.

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

Matsuda W, Furuta T, Nakamura KC, Hioki H, Fujiyama F, Arai R, Kaneko T. Single nigrostriatal dopaminergic neurons form widely spread and highly dense axonal arborizations in the neostriatum. J Neurosci. 2009 Jan; 29(2):444–453.

McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity. Bull Math Biophys. 1943 Dec; 5(4):115–133.

Milstein AD, Li Y, Bittner SR, Grinvald A, Spruston N. Inhibitory Gating of Input Comparison in the Caudate Nucleus. Nature Neuroscience. 2015; 18(11):1591–1600. doi: 10.1038/nn.4121.

Moldwin T, Azran LS, Segev I. A Generalized Framework for the Calcium Control Hypothesis Describes Weight-Dependent Synaptic Plasticity. bioRxiv. 2024 Mar; https://www.biorxiv.org/content/10.1101/2023.07.13.548837v1, preprint.

Moody TD, Carlisle HJ, O'Dell TJ. A nitric oxide-independent and beta-adrenergic receptor-sensitive form of metaplasticity limits theta-frequency stimulation-induced LTP in the hippocampal CA1 region. Learn Mem. 1999; 6(6):619–633.

Nair AG, Gutierrez-Arenas O, Eriksson O, Vincent P, Hellgren Kotaleski J. Sensing Positive versus Negative Reward Signals through Adenylyl Cyclase-Coupled GPCRs in Direct and Indirect Pathway Striatal Medium Spiny Neurons. J Neurosci. 2015 Oct; 35(41):14017–14030.

Niculescu D, Michaelsen-Preusse K, Güner Ü, van Dorland R, Wierenga CJ, Lohmann C. A BDNF-Mediated Push-Pull Plasticity Mechanism for Synaptic Clustering. Cell Rep. 2018 Aug; 24(8):2063–2074.

Nishiyama J, Yasuda R. Biochemical Computation for Spine Structural Plasticity. Neuron. 2015 Jul; 87(1):63–75.

Oikonomou KD, Singh MB, Sterjanaj EV, Antic SD. Spiny neurons of amygdala, striatum, and cortex use dendritic plateau potentials to detect network UP states. Front Cell Neurosci. 2014 Sep; 8:292.

Plotkin JL, Day M, Surmeier DJ. Synaptically driven state transitions in distal dendrites of striatal spiny neurons. Nat Neurosci. 2011 Jun; 14(7):881–888.

Plotkin JL, Shen W, Rafalovich I, Sebel LE, Day M, Chan CS, Surmeier DJ. Regulation of dendritic calcium release in striatal spiny projection neurons. J Neurophysiol. 2013 Nov; 110(10):2325–2336.

Poirazi P, Brannon T, Mel BW. Pyramidal neuron as two-layer neural network. Neuron. 2003 Mar; 37(6):989–999.

Polsky A, Mel BW, Schiller J. Computational subunits in thin dendrites of pyramidal cells. Nat Neurosci. 2004 Jun; 7(6):621–627.

Ravasenga T, Ruben M, Regio V, Polenghi A, Petrini EM, Barberis A. Spatial regulation of coordinated excitatory and inhibitory synaptic plasticity at dendritic synapses. Cell Rep. 2022 Feb; 38(6):110347.

Reig R, Silberberg G. Multisensory integration in the mouse striatum. Neuron. 2014 Sep; 83(5):1200–1212.

Reynolds JNJ, Avvisati R, Dodson PD, Fisher SD, Oswald MJ, Wickens JR, Zhang YF. Coincidence of cholinergic pauses, dopaminergic activation and depolarisation of spiny projection neurons drives synaptic plasticity in the striatum. Nat Commun. 2022 Mar; 13(1):1296.

Roskies AL. The binding problem. Neuron. 1999 Sep; 24(1):7–9, 111–25.

Sanabria BD, Baskar SS, Yonk AJ, Linares-Garcia I, Abraira VE, Lee CR, Margolis DJ. Cell-Type Specific Connectivity of Whisker-Related Sensory and Motor Cortical Input to Dorsal Striatum. eNeuro. 2024 Jan; 11(1).

Santos FJ, Oliveira RF, Jin X, Costa RM. Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning. Elife. 2015 Sep; 4:e09423.

Schiess M, Urbanczik R, Senn W. Somato-dendritic Synaptic Plasticity and Error-backpropagation in Active Dendrites. PLoS Comput Biol. 2016 Feb; 12(2):e1004638.

Schiller J, Major G, Koester HJ, Schiller Y. NMDA spikes in basal dendrites of cortical pyramidal neurons. Nature. 2000 Mar; 404(6775):285–289.

Scholl B, Wilson DE, Fitzpatrick D. Local Order within Global Disorder: Synaptic Architecture of Visual Space. Neuron. 2017 Dec; 96(5):1127–1138.e4.

Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. Science. 1997 Mar; 275(5306):1593–1599.

Schultz W. Multiple dopamine functions at different time courses. Annu Rev Neurosci. 2007; 30:259–288.

Shen W, Flajolet M, Greengard P, Surmeier DJ. Dichotomous dopaminergic control of striatal synaptic plasticity. Science. 2008 Aug; 321(5890):848–851.

Shindou T, Ochi-Shindou M, Wickens JR. A Ca(2+) threshold for induction of spike-timing-dependent depression in the mouse striatum. J Neurosci. 2011 Sep; 31(36):13015–13022.

Shindou T, Shindou M, Watanabe S, Wickens J. A silent eligibility trace enables dopamine-dependent synaptic plasticity for reinforcement learning in the mouse striatum. Eur J Neurosci. 2019 Mar; 49(5):726–736.

Silver D, Hubert T, Schrittwieser J, Antonoglou I, Lai M, Guez A, Lanctot M, Sifre L, Kumaran D, Graepel T, Lillicrap T, Simonyan K, Hassabis D. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. Science. 2018 Dec; 362(6419):1140–1144.

Stauffer WR, Lak A, Kobayashi S, Schultz W. Components and characteristics of the dopamine reward utility signal. J Comp Neurol. 2016 Jun; 524(8):1699–1711.

Surmeier DJ, Shen W, Day M, Gertler T, Chan S, Tian X, Plotkin JL. The role of dopamine in modulating the structure and function of striatal circuits. Prog Brain Res. 2010; 183:149–167.

Takahashi N, Kitamura K, Matsuo N, Mayford M, Kano M, Matsuki N, Ikegaya Y. Locally synchronized synaptic inputs. Science. 2012 Jan; 335(6066):353–356.

Tran-Van-Minh A, Cazé RD, Abrahamsson T, Cathala L, Gutkin BS, DiGregorio DA. Contribution of sublinear and supralinear dendritic integration to neuronal computations. Front Cell Neurosci. 2015 Mar; 9:67.

Trpevski D, Khodadadi Z, Carannante I, Hellgren Kotaleski J. Glutamate spillover drives robust all-or-none dendritic plateau potentials-an investigation using models of striatal projection neurons. Front Cell Neurosci. 2023 Jun; 17:1196182.

Udakis M, Pedrosa V, Chamberlain SEL, Clopath C, Mellor JR. Interneuron-specific plasticity at parvalbumin and somatostatin inhibitory synapses onto CA1 pyramidal neurons shapes hippocampal output. Nat Commun. 2020 Sep; 11(1):4395.

Varga Z, Jia H, Sakmann B, Konnerth A. Dendritic coding of multiple sensory inputs in single cortical neurons in vivo. Proc Natl Acad Sci U S A. 2011 Sep; 108(37):15420–15425.

Watt AJ, Sjöström PJ, Häusser M, Nelson SB, Turrigiano GG. A proportional but slower NMDA potentiation follows AMPA potentiation in LTP. Nat Neurosci. 2004 May; 7(5):518–524.

Wilson CJ, Groves PM, Kitai ST, Linder JC. Three-dimensional structure of dendritic spines in the rat neostriatum. J Neurosci. 1983 Feb; 3(2):383–388.

Wilson DE, Whitney DE, Scholl B, Fitzpatrick D. Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. Nat Neurosci. 2016 Aug; 19(8):1003–1009.

Winnubst J, Cheyne JE, Niculescu D, Lohmann C. Spontaneous Activity Drives Local Synaptic Plasticity In Vivo. Neuron. 2015 Jul; 87(2):399–410.

Wolf JA, Moyer JT, Lazarewicz MT, Contreras D, Benoit-Marand M, O'Donnell P, Finkel LH. NMDA/AMPA ratio impacts state transitions and entrainment to oscillations in a computational model of the nucleus accumbens medium spiny projection neuron. J Neurosci. 2005 Oct; 25(40):9080–9095.

Xu NL, Harnett MT, Williams SR, Huber D, O'Connor DH, Svoboda K, Magee JC. Nonlinear dendritic integration of sensory and motor input during an active sensing task. Nature. 2012 Dec; 492(7428):247–251.

Yagishita S, Hayashi-Takagi A, Ellis-Davies GCR, Urakubo H, Ishii S, Kasai H. A critical time window for dopamine actions on the structural plasticity of dendritic spines. Science. 2014 Sep; 345(6204):1616–1620.

Zenke F, Gerstner W. Hebbian plasticity requires compensatory processes on multiple timescales. Philos Trans R Soc Lond B Biol Sci. 2017 Mar; 372(1715).

Zenke F, Gerstner W, Ganguli S. The temporal paradox of Hebbian learning and homeostatic plasticity. Curr Opin Neurobiol. 2017 Apr; 43:166–176.

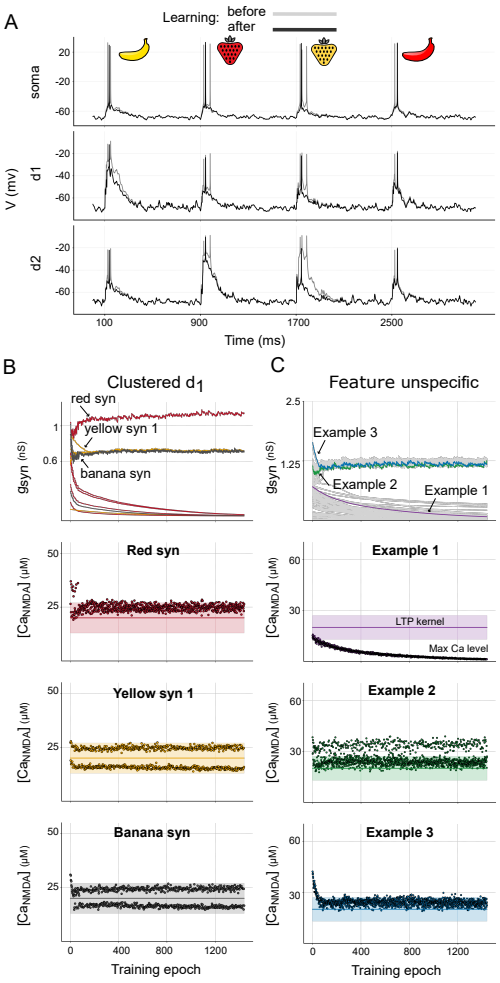

Figure 3—figure supplement 1. Learning-induced synaptic plasticity with metaplasticity turned off

(A): Example voltage traces recorded at the soma and in two dendrites before and after learning. Without metaplasticity, no significant learning occurs—every stimulus triggers spikes (50% performance).

(B, C): Evolution of synaptic conductances throughout learning (top row) and characteristic examples of peak calcium levels (dots), along with the constant kernels in single synapses (shaded areas, bottom rows). The same synaptic arrangement and initial synaptic weights as in ***Figure 3***A is used (for both clustered and feature-unspecific synapses). (B) illustrates changes in clustered synapses on dendrite 1 (d1). (C) shows distributed, feature-unspecific synapses. The final outcome depends on the initial synaptic weights. Synapses whose calcium levels are within or above the LTP region cannot be weakened, because any LTD during irrelevant stimuli is counteracted by LTP occurring for the relevant stimuli. The calcium levels remain "trapped" within the LTP kernel, and the synaptic weights "zig-zag" around a steady state level (seen in the example clustered 'red', 'yellow' and 'banana' synapse, as well as for the Example 2 and Example 3 feature-unspecific synapses). The synapses with small initial weights, whose calcium levels are situated below the LTP kernel are weakened (Example 1 synapse). The solid lines within the LTP kernels represent the midpoint of the kernels where LTP is strongest.

In summary, without metaplasticity learning of NFBP is difficult if initial synaptic weights are not optimally tuned with regard to the postsynaptic calcium concentration. Metaplasticity allows the synapses to fall in and out of the LTP kernel range. They can be strengthened as long as their calcium levels are within the LTP kernel, are stabilized when they are above the LTP kernel, and can be weakened when they are below the LTP kernel.

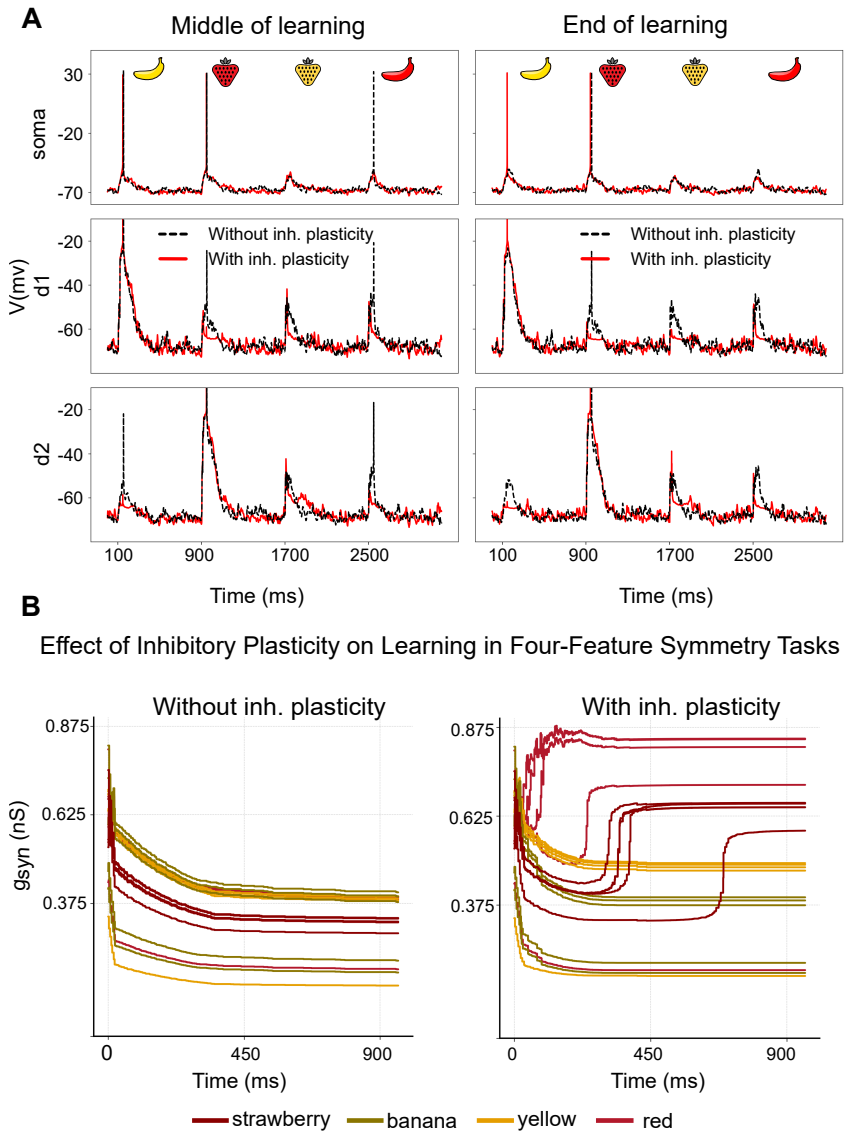

Figure 5—figure supplement 1. Learning with and without inhibitory plasticity in distal dendrites.

(A): The figure illustrates the impact of inhibitory plasticity on learning dynamics in two distal dendrites (d1 and d2). The top panels depict somatic voltage, while the middle and bottom panels show dendritic voltages. The left panels show voltage in the middle of the learning process (approximately 480 training epochs) with and without inhibitory plasticity. The right panels show voltage at the end of the learning process (approximately 960 training epochs). The dashed black line corresponds to the scenario without inhibitory plasticity, and the solid red line indicates the presence of inhibitory plasticity. In the middle of learning, inhibition prevents spiking for the irrelevant stimuli, thus leading to less LTD. This preserves synaptic strengths in both the clustered and feature-unspecific synapses, allowing the neuron to spike for the relevant stimuli at the end of learning. The overall decrease in unspecific input was 39% without inhibitory plasticity and 35% with inhibitory plasticity.

(B): The effect of inhibitory plasticity on learning in a dendrite receiving all four features. The left panel shows excitatory synaptic weights in d1 without inhibitory plasticity, where all features weaken over time due to competition. The right panel illustrates excitatory synapses in the presence of inhibitory plasticity, where learning is better regulated.

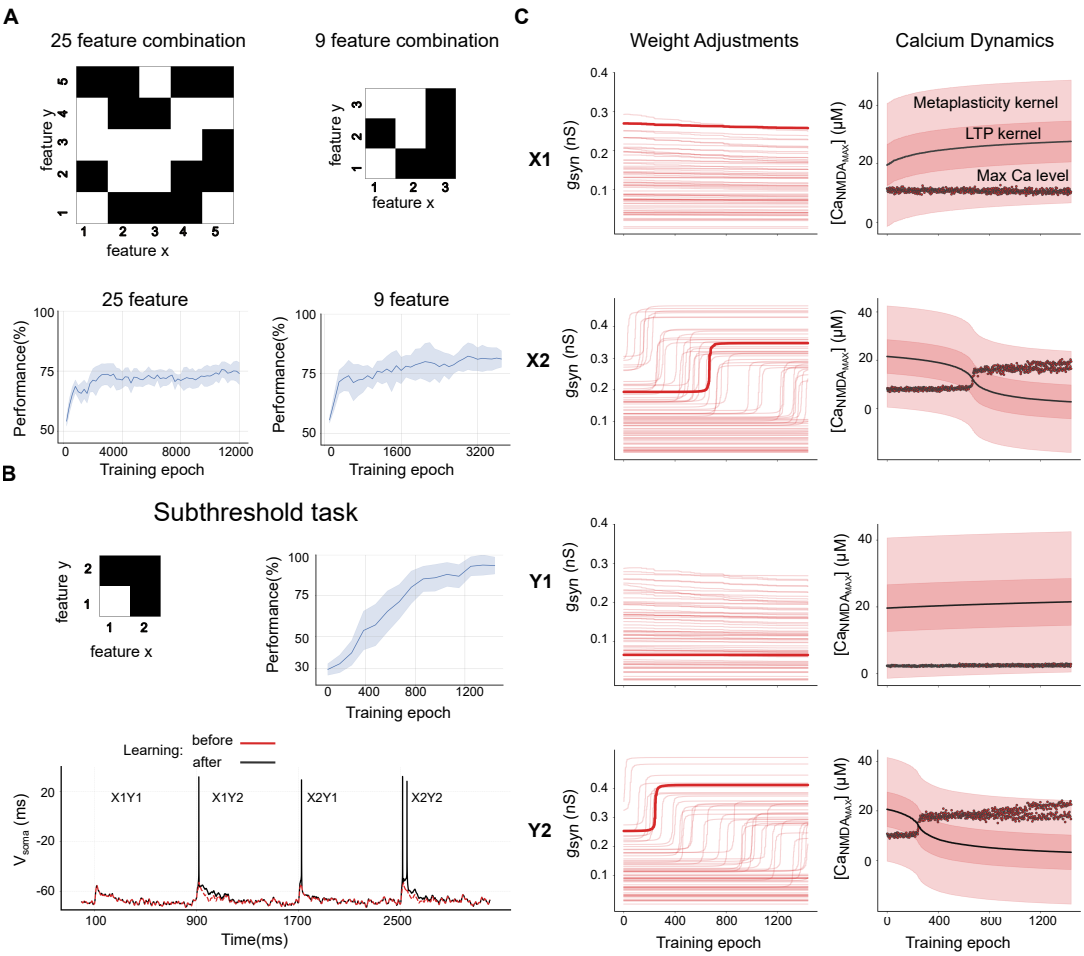

Figure 6—figure supplement 1. Other learning paradigms

(A): Learning of 5×5 and 3×3 feature combination tasks with randomly distributed synapses over 30 dendrites. The neuron is trained to spike for black-marked combinations and remain silent for white-marked ones. For the 25-feature (5×5) combination, the model used 500 excitatory synapses (50 per feature) with initial conductances of 1.125 ± 0.25 nS, along with 100 inhibitory synapses (10 per feature) with initial conductances of 0.1 ± 0.01 nS. For the 9-feature (3×3) combination, the model used 312 excitatory synapses (52 per feature) and 96 inhibitory synapses (16 per feature), with the same initial weights as in the 5×5 case.

(B): Learning of a 2×2 linear feature combination task with randomly distributed synapses. Each feature is represented by 75 synapses distributed across 30 dendrites (300 total synapses), with an initial synaptic weight of 0.375 ± 0.25 nS. The top row shows the feature combinations used, followed by the performance across 30 different synaptic distributions. The bottom row illustrates somatic voltage traces before and after learning, where the neuron starts from a silent state (red trace) and learns to spike through metaplasticity and dopamine-driven weight updates. Additionally, in the subthreshold learning paradigm, dopamine was given without spiking for expected firings and suppressed when unexpected spiking occurred.

(C): Evolution of synaptic conductances (left column) and calcium dynamics with metaplasticity activity (right column) during learning for each feature (X1, X2, Y1, Y2). The left panels display the weight adjustments of all synapses corresponding to each feature, with one synapse per feature bolded to highlight its individual trajectory. The right panels show the corresponding calcium traces and kernel dynamics for the bolded synapse. The solid line represents the midpoint of the kernels where LTP is strongest. 'Max' indicates the peak NMDA calcium level during a single stimulus. The darker shaded regions represent the LTP kernel and the lighter shaded ones show the wider metaplasticity kernel.