The role of temporal sensitivity of synaptic plasticity in representation learning

Patricia Rubisch[1,2,4,*], Melanie I Stefan[2], and Matthias H Hennig[1,3]

[1]Institute for Adaptive and Neural Computation, School of Informatics, University of Edinburgh, EH8 9AB Edinburgh, UK
[2]Department of Human Medicine, Medical School Berlin, 14197 Berlin, Germany
[3]Senior author
[4]Lead contact
[*]Correspondence: patricia.rubisch@medicalschool-berlin.de

Abstract

Synaptic plasticity, the process by which synapses change in an activity-dependent manner, is assumed to be the basis of learning. Experiments demonstrated that synaptic plasticity not only depends on excitatory activity but also on the rate and timing of inhibitory events. Hypothesising that the regulatory effect of inhibition is mediated by membrane potential hyperpolarisation, we identify fast fluctuation sensitivity as a contributing factor to the inhibitory regulation of plasticity. Fast fluctuation sensitivity characterises the influence of fast changes in the membrane potential on the plasticity model predictions, including (short) hyperpolarising events. Furthermore, by introducing a novel plasticity model, the Voltage-Dependent Pathway model, we show that fast fluctuation sensitivity enables a precise temporal regulation of plasticity via inhibition tightly locked to stimulus presentation. In recurrent networks, receptive fields develop independent of the degree of fluctuation sensitivity, yet the temporal precision of inhibitory regulation is critical for the heterogeneity and quality of the resulting representation.

Author Summary

Synaptic plasticity plays a central role in learning and enables us to interact within an ever-changing environment flexibly. The term governs all processes which leads to a change in connection strength between two neurons. Previously, studies of synaptic plasticity strongly focused on the role of excitatory activity. However, experiments excluding GABA blockers implicate inhibitory neurons to regulate synaptic plasticity. We further investigate the role of inhibitory neurons as regulators of synaptic plasticity in a model study. Specifically, we focus on the decrease in membrane potential due to inhibitory activity as a key actor in plasticity regulation. This comparative modelling study considers two well-established models and a novel model which differ in the order in which filtering and thresholding is applied to the membrane potential. We find that while all models develop connectivity structures resembling receptive fields characteristics of the visual cortex, the quality and heterogeneity of the developed feature detectors is model-specific. Heterogeneity and representation specificity is promoted by a high degree of fast fluctuation sensitivity of the model enabling temporally precise inhibitory regulation of plasticity determined by the model's filter and threshold.

Introduction

Long Term Potentiation (LTP) and Long Term Depression (LTD), forms of activity-dependent plasticity, are the foundational processes by which an organism learns. Standard plasticity protocols describe conditions under which LTP and LTD can be experimentally observed. These

include the precise timing between pre- and postsynaptic spikes called Spike-Timing Dependent Plasticity (STDP)[1–3] and the description of pre-post spike pairing frequency dependency[2,4] as well as the dependency of LTP/LTD on the presynaptic stimulation frequency[5–7]. Often these experiments are performed in the absence of inhibition[1–3,5,8] and therefore exclude any effects mediated by inhibitory activity. However, multiple studies show that inhibition influences and changes the established experimental results of the standard protocols[9–13]. Specifically, Hayama et al.[14] and Paille et al.[15] found that STDP outcomes differ when inhibition in the form of GABAergic activity is present at the synapse. In detail, Hayama et al.[14] observed a promotion of LTD when a pre-post spike pair is preceded by inhibition. Moreover, Paille et al.[15] demonstrated the reversal of the causal relationship of standard STDP in the presence of inhibition. A direct link between hyperpolarisation due to inhibition and the regulation of plasticity at excitatory synapses is described by Person and Raman[13].

Models of synaptic plasticity allow us to explore the functional role of inhibition in plasticity regulation within networks of neurons. In rate- and STDP models inhibition and its influence on plasticity promotes functional development and stability[16–18]. However, they require a-priori assumptions about the relationship between Excitatory-Inhibitory (EI) balance and plasticity updates. We hypothesise that inhibitory regulation effects are already accounted for in voltage-dependent plasticity models avoiding any assumptions on the mechanisms themselves. In voltage-dependent plasticity models, plasticity induction is dependent on the relationship between the membrane potential and (arbitrary) plasticity thresholds. Thus, any hyperpolarisation of the postsynaptic membrane potential caused by inhibitory activity directly impacts the plasticity updates. Therefore, they provide the opportunity to investigate inhibitory regulation of plasticity without specifying the interaction between inhibition and the plasticity update explicitly. Moreover, a detailed study of inhibitory regulation in voltage-dependent plasticity models is currently missing.

We address this gap by introducing a novel plasticity model, the Voltage-Dependent Pathway model (VDP), and by systematically comparing it against a well-established model introduced by Clopath et al.[19] (CM) and its variation as presented by Meissner-Bernard et al.[20] (MB) with regard to their sensitivity to fast membrane potential fluctuation. We find that the fluctuation sensitivity determines the effect of inhibition through the extent to which hyperpolarising events contribute to plasticity updates. The higher fluctuation sensitivity of the VDP enables a tighter temporal regulation of plasticity by inhibition. We show that this drives synaptic development linked to input and inhibition onset. Overall, high sensitivity to fast fluctuation of the membrane potential promotes receptive field heterogeneity improving the overall quality of the developed input representation.

Results

In voltage-dependent plasticity models, we assume that the direction of plasticity updates depends on the postsynaptic membrane potential. Thus, the models are not only sensitive to postsynaptic spike timing, but also predict plasticity updates in the absence of postsynaptic spikes[19–21]. In the following we assume that inhibition affects plasticity updates through its hyperpolarising effect on the membrane potential, and we investigate this effect in voltage-dependent plasticity models.

Model	update	presynaptic	postsynaptic
MB[20]	LTP & LTD	low-pass filtered	threshold on low-pass filtered
CM[19]	LTP	low-pass filtered	threshold on low-pass filtered & instantaneous
	LTD	spike time (instantaneous)	low-pass filtered
VDP	LTP & LTD	low-pass filtered	threshold on instantaneous followed by down-stream dynamics

Table 1: Summary of the treatment of the pre- and postsynaptic activity by the different voltage-dependent models in the LTP and LTD induction conditions. All models assume individual thresholds for LTP and LTD induction for which $\theta^+ > \theta^-$.

Hyperpolarisation is sufficient to model inhibitory regulation of STDP

Voltage-dependent plasticity models are based on electrophysiological experiments demonstrating that the sign of the plasticity update is related to the depolarisation level of the synaptic membrane[8,22–30]. Although well-established models have demonstrated voltage-dependent plasticity models widen the range of reproducible experimental results[20,21], they only consider standard plasticity protocols, during which inhibitory signalling is often blocked. However, Paille et al.[15] and Hayama et al.[14] report that GABAergic signalling within a typical STDP pairing changes the plasticity outcome. We used an STDP protocol with precisely timed inhibitory inputs to assess the relationship between membrane potential, inhibitory timing and predicted plasticity updates.

Here, we compare three models of this type: The model presented in[19](CM), which is the foundational voltage-dependent model, its variation introduced in[20](MB) and a novel model called the Voltage-Dependent Pathway model (VDP). All of them implement Hebbian plasticity where the sign and amplitude of the update are dependent on the difference between the postsynaptic trace(s) and the plasticity thresholds $\theta^{+,-}$. Based on their formulation (summarised in Table 1, mathematical details provided in Methods), they can be divided into two subclasses: filter-first-threshold-second (MB, CM) and threshold-first-filter-second (VDP).

In STDP simulations, including a precisely timed inhibitory event relative to the postsynaptic event, the effect of filtering the membrane potential before thresholding in the MB and CM becomes apparent. Filtering the membrane potential results in a prolonged effect of the inhibitory event. Even inhibitory activity delivered after the postsynaptic spike influences the STDP kernel and shifts its LTP-LTD ratio, see Figure 1B and C. In comparison, the effective time range in the VDP is restricted to negative inhibition-post intervals.

Paille et al.[15] reported that the LTP reverts to LTD for pre-post pairings. The MB model is the only model out of the three which shows such a polarity change of the plasticity update. Pre-inhib-post pairings induce LTD while the pre-post pairings without inhibition result in LTP. In the CM and VDP, the inhibitory activity decreases the amplitude of both LTP and LTD. However, computed the LTP-LTD ratio reveals that LTP is more strongly attenuated. Overall, the LTP-LTD ratio shifts towards LTD-dominance (Figure 1C).

Another difference between the filter-first-threshold-second models and the VDP is the effect of the hyperpolarisation duration(Figure 1D). Both the MB and CM are sensitive to the length of the inhibitory event. The longer the duration, the stronger the shift of the LTP-LTD ratio. The ceiling of the effect is caused by the leaky integration of the neuron model and the filtering of the membrane potential by the plasticity models. In contrast, the VDP is insensitive to the duration of the inhibitory stimulus. Only the strength of the hyperpolarisation influences the plasticity update.

In summary, the inhibitory STDP experiment offers first insights into the differences between the models with regard to inhibitory regulation of plasticity. In filter-first threshold-second models inhibition has a broader effective window and larger effect amplitude. A model which applies the

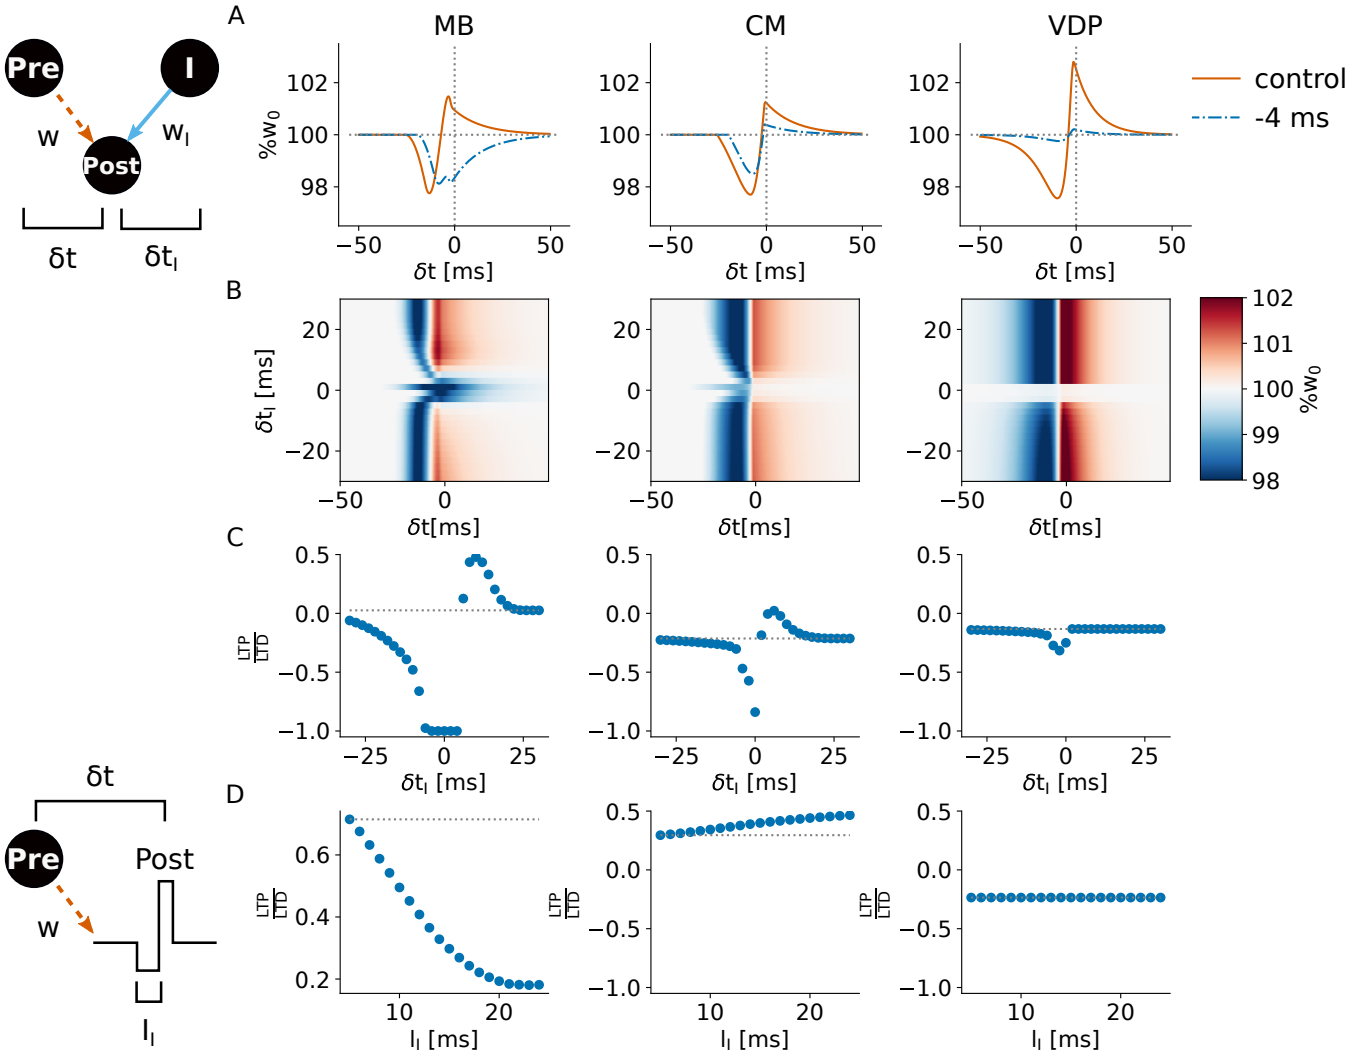

Figure 1: Inhibition mediated via hyperpolarisation in voltage-dependent plasticity models regulates synaptic change in standard STDP experiments. A standard STDP protocol is extended by an inhibitory presynaptic neuron as indicated by the circuit diagram. The inhibitory spike time δt_I is measured relative to the postsynaptic spike. Models are indicated by the abbreviations on top. **A:** Example STDP kernels of the standard protocol without inhibition (orange) and with inhibitory spiking at $\delta t_I = -4$ms. **B**: STDP kernels over several different inhibitory spike timings shown as a heatmap. Blue areas indicate LTD, while red denotes LTP. The MB and CM allow inhibition to affect the plasticity update over a broader range of inhibitory spike timings. **C**: LTP-LTD ratio of the kernels shown in **C**. Pre-inhib-post pairings lead to promotion of LTD. Pre-post-inhib pairings lead to promotion of LTP. **D**: Inhibitory STDP protocol but the inhibitory spike is replaced by a hyperpolarisation of length I_I as shown in the visualisation to the left. δ_t denotes the interval between the presynaptic spike and postsynaptic depolarisation. Filter-first models show a length-dependency, while the VDP's plasticity results are independent of the length of inhibitory event.

thresholds first, followed by downstream dynamics, like the VDP, allows only for inhibitory activity preceding the postsynaptic spike to effect synaptic plasticity disregarding of the duration of the hyperpolarisation.

The order of thresholding and filtering determines fast fluctuation sensitivity of the model

The results above illustrate how plasticity predictions for isolated spike pairs are affected by inhibition and differ between the models. These differences are caused by the order of filtering and thresholding in the model definitions. Low-pass filtering dampens fast components of the signal. Thus, filtering first reduces the sensitivity to fast fluctuations by removing them. To investigate fast fluctuation sensitivity in the models, we replace the postsynaptic spikes of the previous protocol with a fluctuating membrane potential, where we manipulate the fluctuation speed and strength.

To simulate the postsynaptic membrane potential of a neuron receiving input from many neurons, we simulate it as an Ornstein-Uhlenbeck (OU) noise. It is defined by its fluctuation speed τ and standard deviation σ which determines the amplitude of the fluctuations (see example traces in legend of Figure 2A).

The results confirm that weight changes in the VDP model have a higher sensitivity to fast fluctuations than the MB and CM (Fig. 2)A. Filtering reduces high-frequency components, which are further attenuated by thresholding the filtered signal. In contrast, thresholding first preserves the overall frequency structure of the original signal (see Supplementary for illustrations). As the MB and CM are filter-first-threshold-second models, they show a reduced sensitivity to fast fluctuations. Therefore, their plasticity updates are driven by slow and high-amplitude events.

These observations are not unique to the parameters used in these simulations. The high sensitivity to fast fluctuations of the VDP plasticity updates is preserved for a range of parameter variation (Fig. 2B). Both the MB and CM also show an increase in fast fluctuation sensitivity when the postsynaptic filter time constant is reduced. Reducing the filter time constant, decreases the time window over which plasticity occurs shrinking STDP considerably (Fig. 2C). In the VDP, the LTD window also shrinks, but to a much lesser extent than in the CM and MB models.

Increasing the postsynaptic filter constants leads to further reduction in sensitivity to fast fluctuations for the MB and CM, and the models no longer show STDP kernels. These can be recovered by reducing the threshold of both LTP and LTD induction as spikes are strongly attenuated due to the strong filtering. In contrast, the VDP model predicts STDP kernels despite increasing the postsynaptic filter constants. The filter constants do not affect the membrane potential or plasticity threshold. Therefore, the the temporal relationship between membrane potential threshold crossings and presynaptic spikes is preserved accross parameter variation in the VDP. However, the downstream dynamics are either prolonged (”stronger”) or shortened (”weakened”) which leads to an wider/shorted LTD induction window.

To recap, filter-first-threshold-second models more strongly attenuate the postsynaptic signals, which reduces the sensitivity to fast fluctuations and results in a parameter-dependent sensitivity. In contrast, the fast fluctuation sensitivity of the VDP is a model feature, preserving realistic STDP kernels across postsynaptic filter strengths across parameters.

Slow averaged signals drive rate-dependency

The simulations above demonstrate how single spikes and fast signals are picked up by the different plasticity models. Next we analyse the relationship between fluctuation sensitivity and plasticity at different average firing rates. Experiments have shown that LTD is induced at low

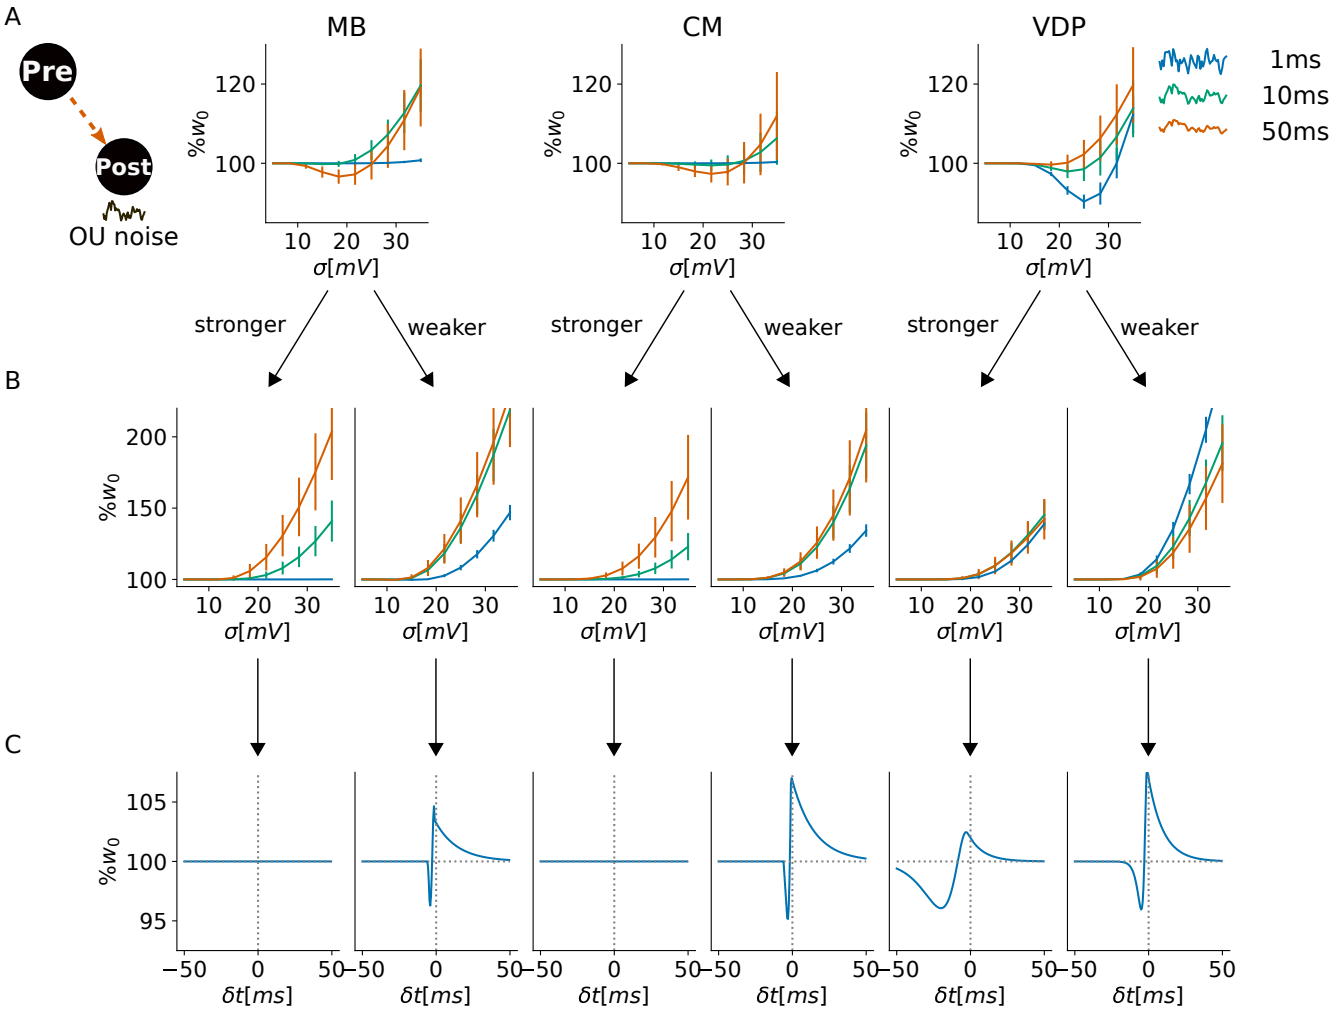

Figure 2: Fast fluctuation sensitivity is a model feature of the VDP while the established models show a parameter-dependent sensitivity. **A:** The postsynaptic membrane potential v is replaced by an Ornstein-Uhlenbeck (OU) noise. Results are obtained over a 10s simulation and averaged across 100 instantiations. To control the fluctuation speed, the time constant τ of the OU noise is varied (colours) while increasing the standard deviation σ results in growing fluctuation amplitude. Plasticity models are indicated by their abbreviations as in previous figures. **B:** Parameter changes which increase filter strength ("stronger", synaptic time constants increased) or decrease filter strength ("weaker", synaptic time constants decreased) demonstrate that fast fluctuation sensitivity can be achieved in the MB and CM with a weak filter. However, the VDP's fast fluctuation sensitivity is independent of parametrisation. **C:** Corresponding STDP kernels to the parameter setting above. Strong filters abolish STDP kernels by reducing the magnitude of the filtered signal to below threshold values. Weaker kernels reduce the LTD window in all models. In the MB and CM, LTD induction is reduced to near instantaneous pairings, while the VDP model retains a wider window in comparison.

frequencies while the synapse strengthens at high presynaptic frequencies[5–7]. To simulate biologically realistic postsynaptic activity, a neuron receives activity from an excitatory and inhibitory Poisson population (Fig. 3A). This allows us to not only examine the frequency dependency of plasticity, but also the influence of the EI balance of the synaptic input.

In voltage-dependent models, postsynaptic frequency is an indirect factor in determining the plasticity update. Prolonged depolarisations which do not generate a spike can induce plasticity. Therefore, the mean membrane potential and its variance is more informative than postsynaptic frequency alone. While counter-intuitive, the mean membrane potential (Fig. 3B) is higher when the synaptic input is dominated by inhibition. Simultaneously, the standard deviation decreases with stronger inhibition at the same postsynaptic frequency (3C). These statistics are similar across neuron models and not limited to the Leaky-Integrate-and-Fire model used here (see Appendix). It is important to note that the standard deviation does not measure fluctuation speed, but the fluctuation amplitude around the mean membrane potential.

The interactions between frequency-dependency and EI balance (Fig. 3D) are driven by the differences in membrane potential statistics. At similar frequencies, inhibition-dominated synaptic input results in a higher mean membrane potential compared to the postsynaptic activity elicited by excitation-dominated synaptic input. For models which are sensitive to the mean of the membrane potential, this causes LTP to be more strongly promoted by inhibition-dominated synaptic input. However, this is not true for all three models. The CM model predicts predominantly LTD for inhibitory-dominated synaptic input ($I > E$) and LTP for excitatory-driven activity ($I < E$). In contrast, the MB and the VDP predict LTP for synaptic input characterised as $I > E$ and LTD for excitatory-dominated input. For an excitatory synapse, LTP updates increase the EI balance while LTD updates decrease it, pushing it further towards $I > E$. Thus, the CM progressively pushes excitatory synapses to the extreme points: to the minimal weight for inhibitory-dominated input and to the maximal weight for excitatory-dominated input. In contrast, the MB and VDP push excitatory synapses towards their starting value.

However, the CM and VDP models show similarities at the extreme condition where inhibitory activity is low. Both models demonstrate a positive feedback between excitatory weight update and postsynaptic frequency. This positive feedback results in exploding weights and has been studied as a feature of Hebbian plasticity[31].

Overall, all models predict a non-linear relationship between the postsynaptic frequency and the resulting plasticity updates. However, investigating the effect of EI balance variation of the synaptic input demonstrates qualitative different behaviour between the models.

Fast fluctuation sensitivity enables tight inhibitory control during learning of structured input

The previous simulations focused on scenarios where postsynaptic neurons either emits single spikes or are driven by constant activity. However, real-world statistics of neural activity are time-dependent and change systematically with different stimuli. Synaptic plasticity allows neurons to learn these patterns, with both single spikes and rates influencing synaptic development.

In the following, we aim to gain an understanding of the interaction between the rate-dependency, STDP effects and difference in fluctuation sensitivity during learning of structured input. We specifically focus on timed inhibitory signalling to investigate inhibitory regulation during stimulus changes.

Three postsynaptic neurons are driven by an input population which is separated into three distinct groups. At each point in time two out of the three input groups are active. Additional input to each postsynaptic neuron is provided by a dedicated inhibitory population. The activity pattern of the input groups and the inhibitory population are inverse of each other: When input population

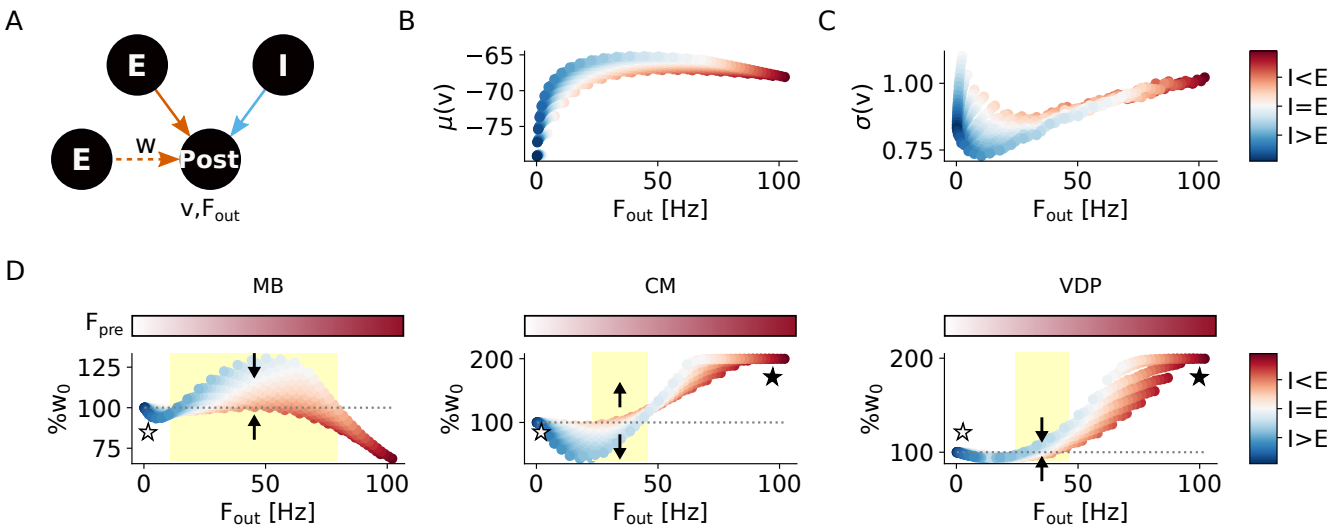

Figure 3: The nonlinearity of frequency-dependency is driven by slow and averaged signals, resulting in small differences between the plasticity models. EI balance variation at constant postsynaptic frequency can shift LTP to LTD and vice versa. **A** A postsynaptic neuron is driven by 100 excitatory and 100 inhibitory Poisson neurons with firing rates $F_E, F_I \in [0, 100)$Hz (5Hz resolution). The colour-coding of the data points indicates the EI balance of the input computed as $\frac{F_E}{F_I}$ with $\frac{F_E}{F_I} > 1$ coloured red and $\frac{F_E}{F_I} < 1$ coloured blue . We measure the postsynaptic membrane potential, output frequency, as well as the weight change of an excitatory (orange,dashed) synapse. The presynaptic neuron of that synapse fires regularly with the same frequency as the excitatory Poisson population. All results are calculated over 10s simulation time. **B**: The mean membrane potential increases when the synaptic input is inhibition-dominated. **C**: The standard deviation of the membrane potential measuring the fluctuation amplitude is increasing when the synaptic input is dominated by excitation. **D**: Plasticity predictions of the plasticity models of the excitatory weights. Model identity is indicated by the abbreviations on top. All results are averaged over 100 runs. The color bar on top indicates the frequency of the regular spiking presynaptic neuron. The open-faced star indicates the regime of pathological behaviour at low postsynaptic frequencies (negative plasticity-activity feedback). The black star indicates the regime of pathological behaviour at high postsynaptic frequencies (positive plasticity-activity feedback). Yellow-backed regions indicate intermediate frequency prediction of special interest. Arrows indicate the direction of EI balance shift due to the predicted synaptic plasticity update.

1 and 2 are active and population 3 is silent, inhibitory population 3 is active while 1 and 2 are in-active. We are interested in the development in the input synapses in order to investigate how the synaptic development is driven by the timed inhibition. In order to differentiate we divide the synapses into to groups: The encoding weights and the non-encoding weights. The encoding weights connect the input population with its corresponding postsynaptic neuron (E1 to 1, E2 to 2, E3 to 3), while the non-encoding weights include the remaining synapses (E2/E3 to 1, E1/E3 to 2 and E1/E2 to 3). Figure 4A and B show an illustration of the set-up and activity.

We consider two key factors when comparing the weight development between plasticity models: the amplitude difference between encoding and non-encoding weight change and the temporal pattern of the update predictions. The update size of non-encoding and encoding weights differ strongly between the models. The MB's and CM's maximal non-encoding synaptic change is significantly larger than the synaptic change of the encoding weights. In contrast, the results of the VDP model show similar amplitudes for both the non-encoding and encoding synapses.

The temporal pattern of the encoding weights are fairly similar between the models. The weight updates reaches their maximum shortly after the stimulus changes and the first presentation does not show any updates for the encoding weights of neuron 3 due to no presynaptic activity. However, the non-encoding weight updates show strong differences. The MB and CM weight changes ramp up towards a plateau over a presentation period. In contrast, the VDP's weight change can be described as a short-timed bump activity. It occurs after the stimulus changes, but the second half of the presentation period is characterised by a constant size of the weight changes. While the time point at which the maximal amplitude is reached is fairly similar for the development of the encoding weights between models, the development of the non-encoding weights differ visibly. The MB model results reach their peak synaptic change at the end of the stimulus presentation. It slowly increases over the presentation time window. The CM model results show a prolonged plateau of at the maximal weight change. This plateau is quickly reached. Therefore, the synaptic update size is constant for half the duration of a single stimulus presentation. In contrast, to the ramping behaviour of the CM and MB model, the VDP updates reach their maximum shortly after the stimulus changes which is followed by a decrease in update size. Similarly to the CM results, the VDP results also show a constant weight update size in the second half of the presentation window.

The higher amplitude of weight updates of non-encoding synapses as well as the ramping behaviour demonstrate that synaptic development in the MB and CM model is strongly driven by slow signals. The weight update increases when the neuron receives inhibition, which is driven by the increase in the mean membrane potential. Visualising the instantaneous activity of the postsynaptic neurons, results in a similar, superposable pattern. In contrast, the VDP weight updates resembles an indicator of change in the postsynaptic activity rather than the development of the activity itself. Weight updates are maximal when the postsynaptic activity changes, but decrease to a constant level quickly after.

The results presented here demonstrate the importance of the fast fluctuation sensitivity during learning of structured input. The MB weights are driven by slow signals due to both LTP and LTD induction relying on filtered signals. In the CM, LTP is sensitive to the instantaneous membrane potential which drives synaptic development after presentation changes. However, LTD is still driven by the filtered membrane potential, a slow signals. Only the VDP, a model in which LTP and LTD take into account the instantaneous membrane potential shows synaptic development tied to sudden changes in postsynaptic activity driven by inhibition onset. Overall, this highlights the importance of fast fluctuation sensitivity of the plasticity to enable a tighter temporal link between activity changes due to inhibition and weight development.

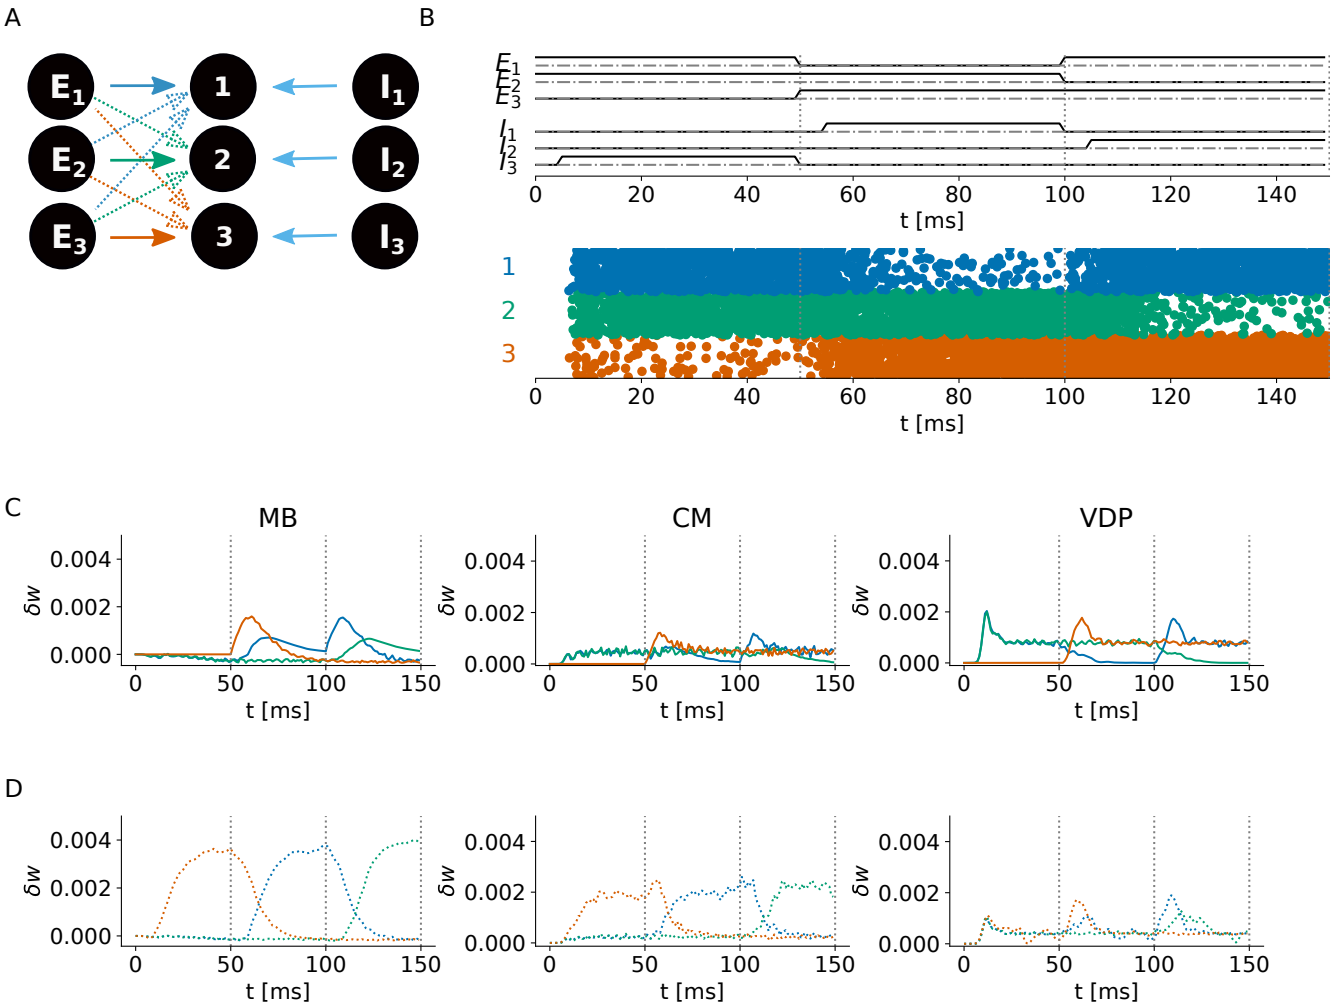

Figure 4: Sensitivity to fast fluctuations promotes stimulus-locked plasticity updates. The higher temporal precision of the inhibitory regulation increases the VDP's sensitivity and reaction speed to changes in the received synaptic input. **A**: Three distinct excitatory Poisson neuron populations (N=10) drive neural activity of 3 LIF neurons. Each LIF neuron receives inhibition at a predefined time from its dedicated inhibitory Poisson neuron population (N=10). Solid arrows indicate encoding weights which connect a specific Poisson population with the neuron which is inhibited when the input population is in-active. Dotted lines indicate non-encoding weights. All results are means of 100 independent runs. **B**: Postsynaptic firing rate of the excitatory neurons and the activation pattern of the input and inhibitory populations. **C** Synaptic update over time of the encoding weights. **D**: Synaptic updates of the non-encoding weights.

High fluctuation sensitivity promotes heterogeneity of learned representations

The results above demonstrate the fluctuation sensitivity and prevalence of fast and/or slow signals in determining the plasticity outcomes predicted by the different models. Based on these observations, we presume that the sensitivity to fast signals also plays a pivotal role in the development of representations in recurrent neural networks. In the following, we use two established experiments, the Bars-and-Stripes problem[32] and receptive field development from natural input statistics[33,34], to investigate differences in representation learning between the models and connect the developed representations to to well-known properties of the visual cortex.

Brito and Gerstner[35] demonstrated that non-linear Hebbian plasticity implements a general principle for receptive field development and unified the biologically motivated plasticity models with principled approaches such as Independent Component Analysis (ICA). The non-linearity specifically refers to the frequency dependency aka the dependence on slow signals, a characteristic all models demonstrate. Thus, they should be able to perform ICA as well.

To test this, networks are presented with the bars-and-stripes problem or Földiak's task[32,36]. This problem consists of randomly sampled bars and stripes which are presented at the same time, thus forming crosses (Fig. 5A). A network which successfully performs ICA on this input develops receptive fields resembling the independent components, a bar or a stripe, and not the presented crosses. As the non-linear frequency-dependency of all models suggest, every network develops single component receptive fields. However, the results vary in their overall statistics and characteristics of non-single component receptive fields(Fig. 5B-C). The MB network develops predominantly multi-factor cells, showing not only one bar or stripe but up to four independent factors encoded by a single excitatory neuron. Interestingly, these cells do not necessarily resemble the presented input but might consist of two stripes instead of a cross. The CM network is the only network which develops "copying" cells. These cells do not show any discernible pattern in their receptive fields. All their synapses reach the maximal weight causing the cell to always be active. It also develops neurons with opposite connectivity patter: uniform synapses with 0 weights. Neurons with this connectivity are non-functional, as their activity is not driven by feed-forward input but only recurrent activity. The VDP network also develops such 0 receptive fields.

The separation score S_i measures the quality and quantity of the receptive fields (Fig. 5C). Based on its definition, we can define lower bounds for single-factor receptive fields (green line) and multi-factor cells (red). Single-factor cells can either encode a bar or a stripe. Their score is the same. The histograms plotted in Figure 5C show that the VDP develops the highest number of single-factor cells and overall heterogeneous receptive fields. The separation score distribution is bimodal with one peak at single-factor scores and a second one at the transition from random receptive field scores to multi-factor cells. The tight temporal window of the inhibitory regulation in the VDP results in a higher degree of representational separability. This confirms our presumption that tight temporal inhibitory regulation promotes functional development in recurrent networks. Furthermore, the development of 0-weight synapses is a form of pruning and shows hallmarks of sparse coding. The network consists of more neurons than the input has factors. Thus, not all neurons are needed to represent the input.

The bars-and-stripes problem illustrates that the variation in fluctuation sensitivity drives differences between learned representations. However, as an experiment it is far removed from biological relevant input. When the concept of ICA is applied to naturalistic input, it results in oriented edge-detectors similar to receptive fields in V1[33–35,37,38]. Therefore, we present the MB, CM and, VDP network with CIFAR10 input patches[39] in order to learn receptive fields resembling those in the visual cortex. In the following we investigate whether the variation in fluctuation sensitivity also promotes differences in the quality and heterogeneity of receptive fields learned

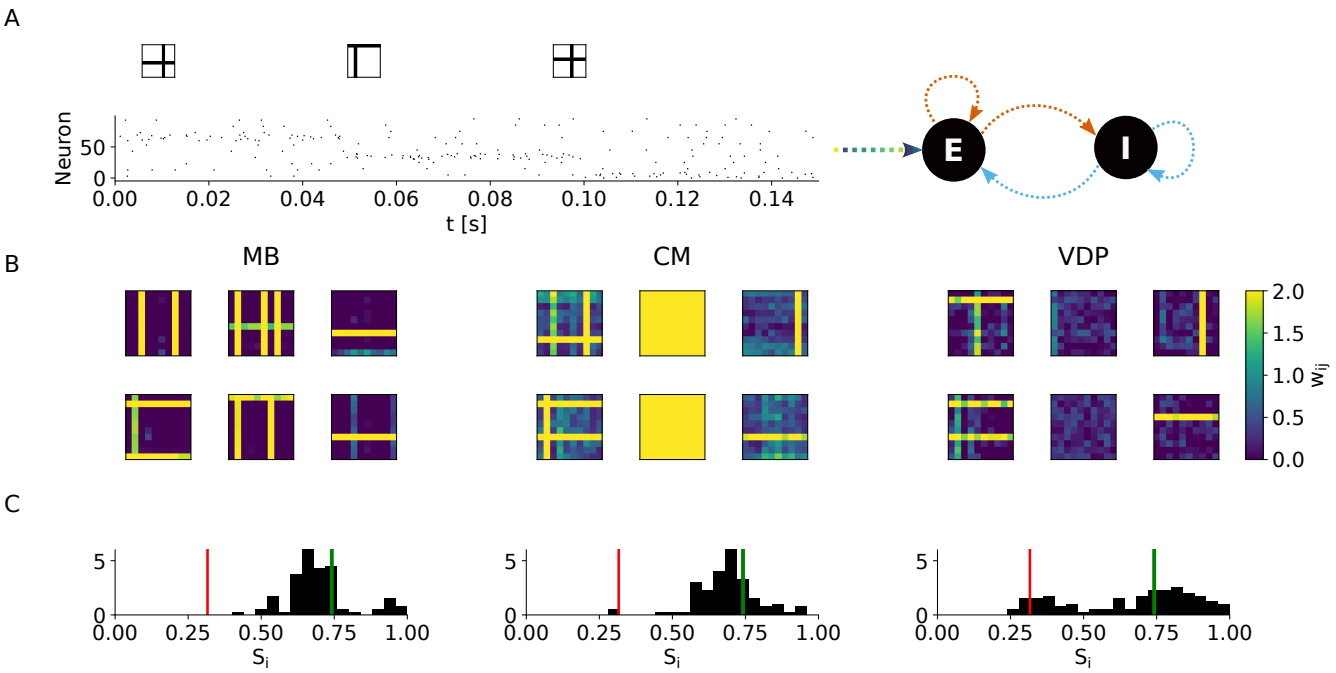

Figure 5: The overlap of learned representations of independent components is reduced in plasticity models with high fast fluctuation sensitivity. **A**: An EI network is presented with the bars-and-stripes task[32]. **B**: Learned receptive fields. **C**: Separation scores of the excitatory neurons. A neuron which encodes a single bar perfectly has $S_i = 1.0$. The green line indicates the upper bound for a multi-factor cell. The red line indicates the corresponding lower bound. Any score below this characterises random receptive fields with no discernible pattern. These scores are computed with cleaned-up receptive fields to reduce the contribution of spurious patterns present in near-0 weights. All weights $w_{ij} < \mu(w)$ are set to 0. The mean is computed over all receptive fields of the network.

from naturalistic input.

The CIFAR-10 input patches are presented by two separate input populations, an On- and Off-population, similar to On- and Off-cells in the retina. Neurons in the On-population fire proportionally to the positive part of the presented image while the value of the negative pixels scales the Off-neuron's firing rate (Fig. 6A). The receptive field of an excitatory neuron is equal to the difference between its Off-weights and its On-weights. It is important to note that homogeneous receptive fields with $w_i = 0 \forall i$ such as those present in the MB network (see Fig. 6B) do not necessarily indicate all synapses to be silent as in the previous experiments. While this can be the cause, perfectly aligned, equally strong On-and Off-synapses cancel out also resulting in a receptive field with 0 weights.

All networks develop a mixture of (uniform) mono-polar and oriented receptive fields (Fig. 6B). Mono-polar receptive fields are characterised by a small variance and the mean μ is either -1 or 1 (Fig. 6C). These receptive fields copy the input from either the On- or Off-population. The MB network receptive fields show the strongest clustering of their statistics (μ and σ) in comparison to the others. They are dominated by mono-polar and 0-mean-receptive fields. The CM receptive fields also shows some clustering at $\mu = [1, -1]$ in their statistics. However the standard deviation scores are covering the value range more uniformly, leading to less clear clusters compared to the MB statistics. The statistics of the VDP receptive fields show even less clustering. This suggest that the VDP network developed the most diverse representations while the MB receptive fields include a lot overlapping patterns.

The other characteristic which is of interest when talking about receptive fields is their orientation. The visual cortex has several orientation columns in which neurons are sensitive to the same orientation and can detect edges in images[40]. To compute the orientation, the receptive fields are convolved with predefined Gabor patches or varying orientation between $-180°$ and $180\circ$. The relative response patterns as shown in Figure 6D indicates how closely the different network fit the distribution of orientations of the input data.

Horizontal edges are dominant in the input data. In comparison to diagonal orientations, vertical features show a stronger response, but not to the same extent as horizontal orientations. Overall, this leads to a "w"-shape of the orientation profile of the input data where the bump of the w is lower than its tips. This distinct "w"-shape is only present in the response profile of the VDP network receptive fields. However, the difference between the response contribution of horizontal and vertical edges is bigger than compared to the input data response profile. Neither the MB nor the CM profile can be desribed as a "w"-shape. The, the middle bump is absent in both. This indicates that neither network learned vertical orientations well. Interestingly, the MB receptive fields underrepresent orientations with a small positive angle and show the highest relative response to horizontal edges. In contrast, edges with small negative angles are less present in the receptive fields developed by the CM. Furthermore, the difference between the response to horizontal stimuli and vertical/diagonal oriented gabor patches is bigger compared to the VDP response and the input data. Therefore, the receptive fields developed by the VDP represent the input data the closest. Even though, the absolute contributions do not exactly match, the overall shape of the response profile is well approximated.

In summary, the characteristics of the developed connectivity structures in these two functional experiments suggest that the higher sensitivity to fast signals of the VDP enabling a tight temporal inhibitory plasticity regulation, promotes the separability of input representations. The receptive fields show a higher degree of heterogeneity and resemble the orientation biases present in the input data more closely. Therefore, the degree of fast fluctuation sensitivity directly affects the quality of the developed representations in a recurrent network.

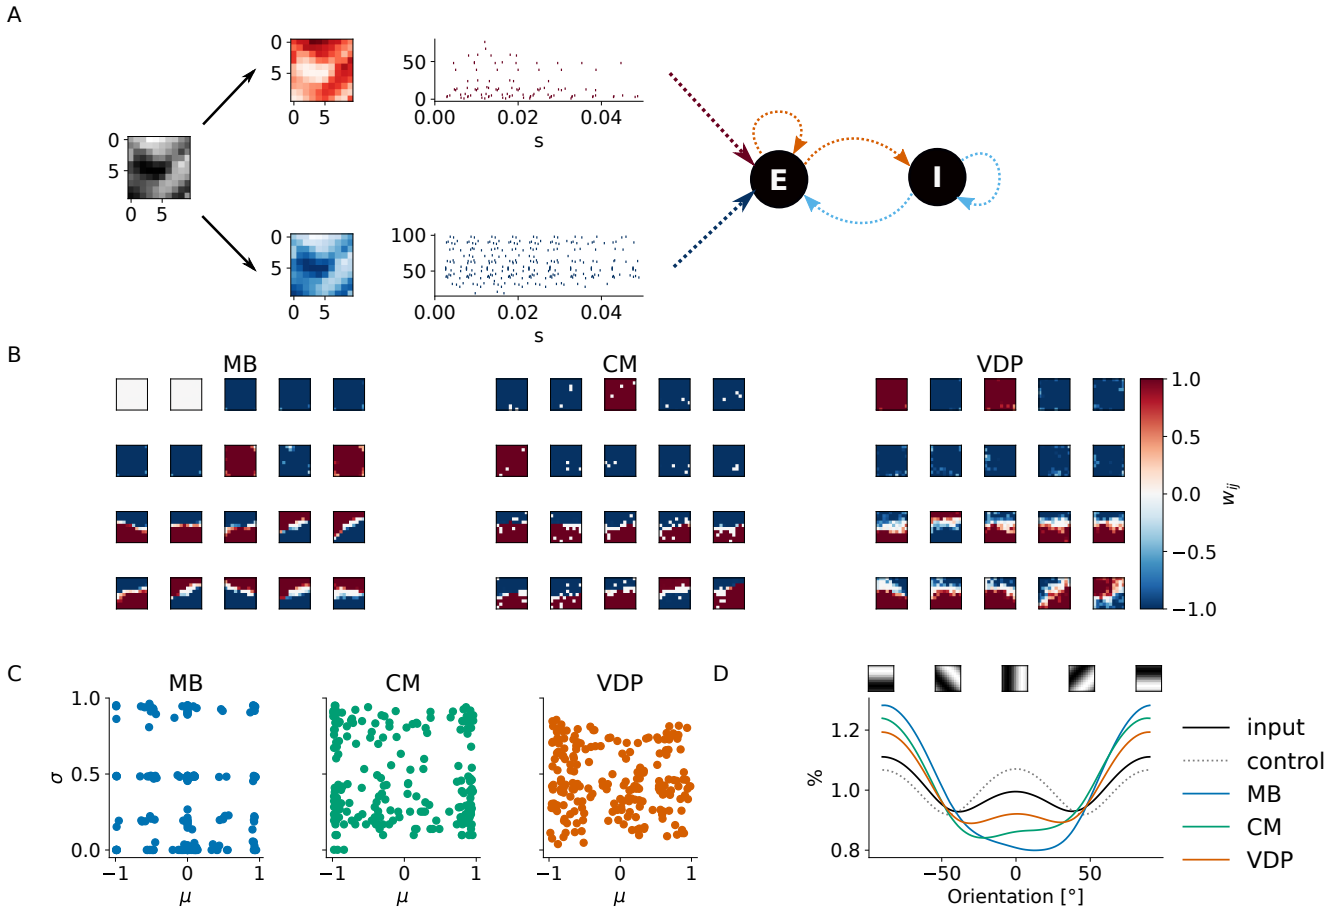

Figure 6: Receptive field quality and heterogeneity are promoted by plasticity models with high sensitivity to fast signals. **A**: A fully-plastic recurrent EI network is presented with CIFAR10 input patches using an On- and Off-population for input encoding. **B**: The developed receptive fields show specific characteristics depending on the plasticity model. Shown are 10 receptive fields with the smallest standard deviation and 10 receptive fields with the highest standard deviation. Blue indicates a strong connection to the Off-population, while red indicates a strong synapses to the On-population. **C**: Mean and standard deviation of the receptive fields. **D**: Responsiveness of oriented Gabor-patches to the input data (black), random patches (grey-dotted), receptive fields of the MB (blue), CM (green) and VDP-network (orange). The colour scheme for the receptive fields is consistent between **C** and **D**.

Discussion

These results suggest that the inhibitory regulation through hyperpolarisation of the membrane potential is a plausible model for developing functional connectivity patterns. A detailed analysis of three voltage-dependent plasticity models highlights the importance of the sensitivity to fast signals to enable precise inhibitory regulation. The stronger the filtering of the postsynaptic membrane potential, the weaker the sensitivity to fast fluctuations. In recurrent networks, the sensitivity to fast fluctuation is important to enable accurate input representation due to inhibitory signalling.

Biologically, these results suggest that synapses located on distal dendrites can rely more on the local, fast-changing information[41–43], if plasticity has higher temporal sensitivity, such as proposed in the VDP. This implies that distal synapses might develop more precise and distinct representations, as they are less influenced by longer temporal integration. In contrast, at proximal synapses, the voltage signal corresponds to information integrated over longer periods of time and over many synapses. For this case our simulations suggest plasticity will yield representations with more overlaps. As the models by Clopath et al.[19] and Meissner-Bernard et al.[20] do not have a high temporal sensitivity, they are predicted to behave similarly at distal and proximal synapses. In contrast, the VDP predicts a difference in representation learning at distal and proximal locations. This prediction is testable in biophysical neuron models, and differences in distal and proximal tuning properties should be detectable in imaging experiments.

Neurons and synapses are highly heterogeneous: from genetic markers to protein and channel sub-type composition and local distribution[43–45]. Thus it is unlikely that a single rule can account for all behaviours. Here, we only considered synapses reproducing the standard STDP kernel, while several different STDP kernel forms have been identified experimentally[1]. These models cannot reproduce effects such as the the reversal from LTD to LTP described by Paille et al.[15], as it is incompatible with voltage-dependent plasticity models where the LTD threshold is always below the LTP threshold[8].

Interestingly, all models show a non-linear frequency-dependency. While we can argue that the frequency-dependency is driven by slow signals such as the mean membrane potential, the models also differ in their level of sensitivity to this slow signal. This is detectable in the reported effect of the synaptic EI balance driving the postsynaptic frequency. We want to stress, that while the experiment approximates a mean-field analysis of the model's behaviour, it is not sufficient to make claims about the stability of the models. A detailed exploration of the activity development over long simulations of recurrent networks, including parameter sweeps, would be necessary for an accurate investigation of model stability. Unfortunately, this is outside the scope of this work which focuses on the role of model characteristics on representation learning.

Similarly, the models implement Hebbian plasticity without considering homeostatic components. Homeostatic plasticity increases network stability and can be implemented by a plethora of mechanisms[31]. Thus, future research investigating network stability in voltage-dependent plasticity rules should also include a detailed exploration of the interactions with homeostatic principles acting on biologically plausible time scales.

This work exclusively focuses on the inhibitory regulation in voltage-dependent plasticity models. While this family of plasticity models is well positioned to investigate hyperpolarisation as a mean of inhibitory regulation, it is not the only one possible to address this. Calcium-dependent plasticity models can also incorporate a direct link between the membrane potential and calcium influx and are able to reproduce experimental data incorporating neural architecture[46,47]. Specifically, Hiratani and Fukai[47] demonstrates the development of dendritic EI balance in their work. However, the work is not investigating the effects of representation learning. Furthermore, the insights derived here and the categorisation of models can be extended to them. Both models

would be classified as filter-first-threshold-second and therefore would show a decreased sensitivity to fast signal components. What these components might signal and their contribution to the activity of downstream processes in biological processes remains to be further investigated.

In summary, this work investigates the role of inhibition in voltage-dependent plasticity models. Inhibitory regulation of plasticity in these models is determined by their sensitivity to fast fluctuations. The higher the fast fluctuation sensitivity, the tighter the temporal inhibitory regulation of plasticity. By studying representation learning and stimulus/inhibition-onset sensitivity, we identified that a high sensitivity to fast signals is important for recurrent inhibition to effectively drive plasticity and functional development. The novel plasticity model, the VDP, demonstrates that a high degree of fast fluctuation sensitivity results in a more heterogeneous receptive fields which accurately represent the input data. This work predicts that synapses whose plasticity is driven by local, fast-fluctuating information like those on distal dendrites develop more distinct and precise input representations due to the precisely timed inhibitory plasticity regulation. Investigating the relationship and interplay between local EI balance and LTP/LTD induction at different dendritic sites warrants further investigations to accurately describe the functional role of inhibitory regulation of plasticity in biological neurons.

Methods

Neuron model

All experiments, except the STDP experiments, are performed using a Leaky Integrate-and-Fire (LIF) neuron with conductance-based synapses. They follow the equation:

$$\tau \frac{dv}{dt} = -(v - E_L) - g_I \cdot (v - E_I) - g_E \cdot (v - E_E) \tag{1}$$

$$\frac{dg_x}{dt} = -\frac{g_x}{\tau_{syn}} + \sum_j s_j \cdot w_j \cdot g_x^{max} \tag{2}$$

$$\text{if} \quad v(t) > \theta : \; v(t) = E_L \tag{3}$$

with $x = [E, I]$ and E_L denoting the resting potential of the neuron. g_I and g_E are the synaptic conductances of the inhibitory and excitatory synapses. Whenever a spike arrives at the synapse g_x grows proportionally to the current synaptic weight w_x. We use g_x^{max} to denote the maximal conductance. We generally assume the notation of w_{ij} denoting the synapse connection between presynaptic neuron j and postsynaptic neuron i. Therefore, j in Equation 2 indicates that the synaptic current is summed over all presynaptic neurons. The standard parametrisation of the LIF is: $E_L = 79mV, E_E = 0mV, E_I = -80mV$ and $\tau = \tau_E = \tau_I = 10ms$.

In order to enable a fair comparison of isolated events as occurring in STDP experiments at low repetition frequencies, the results presented in Figure 1 and Figure 2C use a Generalised LIF (GLIF) model of class 5 as introduced in Teeter et al.[48]. The following equations define this

model:

$$C\frac{dv}{dt} = (-(v - E_L) - g_I \cdot (v - E_I) - g_E \cdot (v - E_E))\frac{1}{R} + \sum_m I_m + I_e \tag{4}$$

$$\frac{dg_x}{dt} = -\frac{g_x}{\tau_{syn}} + \sum_j s_j \cdot w_j \cdot g_x^{max} \tag{5}$$

$$\frac{dI_m}{dt} = -k_m I_m \tag{6}$$

$$\frac{d\theta_s}{dt} = -b_s \theta_s \tag{7}$$

$$\frac{d\theta_v}{dt} = a_v(v - E_L) - b_v \theta_v \tag{8}$$

$$\text{if} \quad v(t) > \theta + \theta_s(t) + \theta_v(t) : \tag{9}$$

$$v(t^+) = E_L + f_v(v(t^-) - E_L) + \delta v \tag{10}$$

$$I_m(t^+) = f_m I_m(t^-) + \delta I_m \tag{11}$$

$$\theta_s(t^+) = \theta_s(t^-) + \delta\theta_s \tag{12}$$

t^+ denotes the time step after spike generation, while t^- denotes the timestep before spike generation. The synaptic conductance g_x follows the same dynamics as in the LIF model. Teeter et al.[48] include a full description of the parameters and variables, including explanations of biological equivalence. The parameter setting used here is simulating a neuron of cluster 1[48], producing excitatory neuronal behaviour. I_e denotes the external current applied to force a spike at precise times.

The GLIF5 neuron model includes After-Spike-Currents (Eq.6) which approximate the spike shape. Thus, the depolarisation is prolonged in the GLIF5 model after spike generation in comparison to a LIF neuron model. This enables models which use postsynaptic membrane filters before thresholding (MB and CM), to predict updates in STDP experiments despite isolated postsynaptic events. The GLIF5 model is therefore similar to the AdaptiveExponential (AdEx) neuron model which was used by Clopath et al.[19] to introduce the CM.

Voltage-dependent plasticity models

All three voltage-dependent plasticity models implement Hebbian plasticity. They can be divided into their presynaptic component and their postsynaptic component. The presynaptic component follows the same dynamics in all models:

$$\tau_s \frac{d\bar{s}}{dt} = -\bar{s} + a^s \cdot s(t) \tag{13}$$

The models either use the presynaptic spike train $s(t)$ or the presynaptic trace $\bar{s}$. The spike train $s(t)$ is 1 if a spike was generated at time t and is 0 otherwise. $\bar{s}$ is computed by applying a low-pass filter to $s(t)$. The scaling parameter a^s is set to $\frac{1}{\tau_s}$ if not indicated otherwise.

A further commonality between the models is the use of an explicit potentiation and depression threshold, θ^+ and θ^- respectively. $\theta^+ > \theta^-$ is true for all parameter settings used in this work. This assumption follows the results presented in Ngezahayo et al.[8].

Voltage-Dependent Pathway model The Voltage-Dependent Pathway (VDP) model is a novel plasticity model and can be summarised as a threshold-first-filter-second model. The membrane potential is first passed through a sigmoidal function $\sigma(v(t) - \theta, k)$ with the plasticity threshold θ

and k = 0.4. The outcome sets the update amplitude of the pathway activity $p(t)$. The pathway activity implements the postsynaptic filter whose properties are demonstrated in the Appendix. It is defined as bi-exponential differential equations:

$$\tau_f \frac{p}{dt} = -p + o \cdot p_r \tag{14}$$

$$\tau_r \frac{p_r}{dt} = -p_r + [[\sigma(v-\theta, k) - p]_+ - p_r]_+ \tag{15}$$

$$o = \frac{\tau_r}{\tau_f}^{\frac{\tau_f}{\tau_r - \tau_f}} \tag{16}$$

where $[\cdot]_+$ denotes a linear rectifier: $[x]_+ = \begin{cases} 0 & if \;\; x < 0 \\ 1 & else \end{cases}$ The plasticity update is computed by the difference between the potentiation and depression pathway, p^+ and p^- respectively. We omitted the superscripts $^{+,-}$ in the definition for the pathway activity p above to improve legibility. This means that all parameters in the bi-exponential definition are pathway specific. Therefore, the dynamics of the potentiation pathway p^+ is determined by the values of τ_r^+, τ_f^+ and θ^+. The difference between the potentiation and depression pathway dynamics is multiplied by the presynaptic trace to compute the overall plasticity update. This ensures that updates only occur if presynaptic and postsynaptic activity coincides within a specified time window determined by the pathway specific parameters.

$$\Delta w = \eta \cdot c(t) \cdot \left(a^+ \bar{s}(t) p_+(t) - a^- \bar{s}(t) p^-(t)\right) \tag{17}$$

η denotes the learning rate, while a^+ and a^- are pathway specific amplitudes. $c(t)$ is a scaling factor which introduces a dependence on the external calcium concentration Ca^{2+} and is based on the results reported in Inglebert et al.[4]. This factor is dependent on both the activity of $p^+ \cdot [\bar{s}]_+$ and p^-:

$$c(t) = p^-(t)^{20 \cdot exp(-Ca^{2+})} + p^+(t)^{20 \cdot exp(-Ca^{2+})} \cdot [\bar{s(t)}]_+$$

This uncommon formulation causes a stronger expression of depression at low calcium levels. However, the results presented here are independent of this scaling factor because it does not effect the fluctuation sensitivity of the plasticity model. A step-by-step example of the induction mechanism of the VDP is illustrated in the Appendix. As standard parametrisation the VDP was fit to a standard STDP kernel[3] (see Appendix). We refer to the values of time constants and threshold as the standard parametrisation as these determine the dynamics of the model. These are $\tau_r^+ = 1ms, \tau_f^+ = 5ms, \theta^+ = -40.75mV, \tau_r^- = 1ms, \tau_f^- = 20ms, \theta^- = -43.75mV$ and $\tau_s = 10ms$. The values of scaling factors are provided in the Appendix which lists all used parametrisations and their corresponding results.

The Clopath et al.[19] and Meissner-Bernard et al.[20] models The CM and MB model are in contrast filter-first-threshold-second models. The membrane potential v is low-pass filtered to compute the synaptic trace $\bar{v}$. $\bar{v}$ is than directly compared against the plasticity thresholds θ in the weight update equation. Potentiation and depression are dependent on their respective postsynaptic traces $\bar{v}^{+,-}$ and plasticity thresholds $\theta^{+,-}$.

$$\tau \frac{d\bar{v^{+,-}}}{dt} = -\bar{v^{+,-}} + v \tag{18}$$

The difference between the CM and MB model lies in the exact combination of the postsynaptic traces and the presynaptic trace. The CM defines the weight-update as follows:

$$\Delta w = \eta \left(a^+ \bar{s}(t) [\bar{v}^+(t) - \theta^-]_+ [v(t) - \theta^+]_+ - a^- s(t) [\bar{v}^-(t) - \theta^-]_+\right) \tag{19}$$

It uses the spike train $s(t)$ for the induction of LTD. Thus, depression can only occur at the time of a presynaptic spike. Potentiation is conditional on the instantaneous membrane potential. Therefore, both the membrane potential as well as the postsynaptic potentiation trace need to be above a certain threshold.

In comparison the MB model updates synaptic weights following this definition:

$$\Delta w = \eta \left(a^+ \bar{s}(t)[\bar{v}^+(t) - \theta^+]_+ - a^- \bar{s}(t)[\bar{v}^-(t) - \theta^-]_+ \right) \tag{20}$$

Updates only occur depending on the pre- and postsynaptic traces. Thus, any synaptic weight change is determined by filtered signals, decreasing the contribution of fast signals.

Due to the similar definition of the models, their standard parametrisations are equivalent. They follow the parametrisation described in Clopath et al.[19] (thresholds have been adapted due to difference in neuron model): $\tau^+ = 7ms, \tau^- = 10ms$ and $\tau_s = 15ms$. A full list of all parameters per experiment is given in the Appendix.

Simulation Set-up

All simulations are carried out using the Euler integration method with an integration time step of $0.1ms$

Spike-Timing-Dependent-Plasticity with Inhibition The synapse of interest is the connection between a regular spiking presynaptic neuron ($F = 0.1Hz$) and postsynaptic neuron. The postsynaptic neuron is simulated using the GLIF5 neuron model with an input current of $I_e = 5nA$ which is applied for $0.8ms$. The input current is delivered with a frequency of $F = 0.1Hz$ with a spike time interval δt relative to the presynaptic neurons spike. δt takes values in the range of $-50ms$ to $50ms$ with a spacing of $1ms$. To include inhibition in the STDP protocol, the postsynaptic neuron also received input from an inhibitory neuron, whose spikes are timed relative to the postsynaptic neuron. This timing difference is denoted as δt_I and ranges from $-20ms$ to $5ms$ (resolution of $1ms$).

To test the effects of the length of an inhibitory event at constant peak membrane potentials, we approximate the postsynaptic hyperpolarisation and spike as a step-wise membrane. Inhibitory events hyperpolarise the membrane potential by 5mV for a duration of $I_I ms$ and followed by a depolarisation to $E_L + 30mV$ for $1ms$.

The LTP-LTD ratio r is computed via:

$$r = \frac{\sum_{i=-t}^{t} (\delta w^{\delta i} \cdot \delta w^{\delta i} > 0) - \sum_{i=-t}^{t} (\delta w^{\delta i} \cdot \delta w^{\delta i} < 0)}{\sum_{i=-t}^{t} |\delta w^{\delta i}|} \tag{21}$$

It computes the absolute sum of all positive and negative weight updates over all measured δt, subtract them from each other and scales it by the sum of all absolute weight changes. Therefore, r is restricted to be between -1 and 1.

Fast fluctuation sensitivity using Ornstein-Uhlenbeck membrane potential dynamics To investigate the varying degree of fluctuation sensitivity of the plasticity models (Figure 2), the postsynaptic membrane potential is simulated as Ornstein-Uhlenbeck noise. It follows the dynamics governed by:

$$\frac{dv}{dt} = -(v - E_L)/\tau + \sigma \delta W$$

δW is a Wiener process. The resting potential E_L is equivalent to the resting potential of the LIF neurons of other experiments with $E_L = -79mV$. Conditions include a variation of the fluctuation speed $\tau = [1, 10, 50]ms$ and amplitude $\sigma = (5, 35)mV$. Presynaptic activity is simulated using a regular spiking neuron with frequency $F = 10Hz$. Each condition is simulated for $10s$ with 100 repetitions. The presented results are averages over the repetitions computed for each condition defined by τ and σ.

Frequency-dependency To approximate the frequency dependency, two separate Poisson populations drive the postsynaptic LIF neuron. The simulation includes an explicit excitatory and inhibitory presynaptic population to account for the variation in sub-threshold membrane potential dependent on the synaptic input composition. Their activity is set by F_E and F_I which take values between 1 and $100Hz$. The synaptic EI balance which is visualised by the colour of each data point is defined as the ratio $\frac{F_E}{F_I}$. The presynaptic neuron which connects the plastic synapse to the postsynaptic neuron fires regularly with the same frequency as their corresponding Poisson population. Thus, the presynaptic frequency in Figure 3D is F_E. Results showing the inhibitory synaptic development with presynaptic frequency F_I are illustrated in the Appendix. The simulation is run for $10s$. Each condition is repeated 100 times and results are averages computed across repetitions.

Inhibition/Input-Onset specificity of synaptic updates To investigate the effect of rapidly changing input, three postsynaptic LIF neurons receive activity from an input population($N_{input} = 30$) and an inhibitory populations ($N_I = 30$). To structure the network, the input population is divided into three groups based on their activity and the inhibitory population is subdivided by their activity and synaptic connections. The input groups are all-to-all connected to the postsynaptic neurons. In contrast, each of the three inhibitory sub-population is connected to one dedicated postsynaptic neuron. Both inhibitory and input population consists of Poisson neurons. When active, the input neurons' firing rate is set to $F_{input} = 100Hz$ and the inhibitory neurons fire with a rate of $F_I = 50Hz$. The activity pattern during one presentation is the following: Two out of the three input-groups are active. The inhibitory sub-population of the inactive input-population inhibits their corresponding postsynaptic neuron. Therefore, when E1 and E2 are active, I3 is active (see Figure 4A). Each presentation lasts for $50ms$ and each possible pattern is presented once. Results are shown as the synaptic updates averaged over 100 instantiations of this simulation.

Bars-and-Stripes The Bars-and-Stripes experiment of Földiak's bars[32] consists of bars and stripes which are simultaneously presented. Each pixel of the 10x10 pixel input patches is represented by a Poisson neuron in the input popualtion. When the pixel is part of the presented overlapping strip and bar, the neurons firing rate is set to $F_{input} = 100Hz$. This input is presented to a fully-connected recurrent network with $N = 120$ ($N_I = 0.2 \cdot N_E$). The initial input weights are all-to-all connected to the excitatory network. All synapses are plastic.

To assess the quality of the developed representations, we compute a separation score S_i based on the receptive field r_i of each excitatory neuron i. The receptive field r_i is defined as all weights w_{ij} which connect the postsynaptic (excitatory neuron i) with the input population. Therefore, j denotes the presynaptic neuron for which$j \in N_{input}$. The separation score S_i measures the maximal cosine-similarity c_S between the receptive field r_i and the ideal receptive fields of all bars $0°$ and the ideal receptive fields of all stripes $90°$:

$$S_i = max(c_S(r_i, [0°, 90°])) \tag{22}$$

Prior to computing S_i, we clean the receptive fields r_i by discarding any weights which are smaller than the overall average postsynaptic connection. This decreases the contribution of residual patterns in non-functional, small synapses. When the receptive field encodes a single factor the cosine similarity c_S to this factor is 1.0, while the similarity to all remaining factors is smaller than 1. The presence of multiple factors in the receptive fields, disregarding whether in the form of a cross or two non-intersecting factors, decreases the cosine similarities between the receptive field and any given single factor. Based on this, we can compute an upper bound for multi-factor receptive fields, which simultaneously provides a lower bound for a single-factor cells ($S_i = 0.74$). Additionally, we compute an upper bound for non-informative random fields ($S_i = 0.32$).

Receptive field development from naturalistic input 1000 Patches from the CIFAR10 data set are presented to a fully recurrent EI network with $N = 250, N_I = 0.2 \cdot N_E$. Each of the centered 10x10 pixel input image ($\mu = 0$) is encoded by two dedicated neurons. One neuron is part of the On-population and the other belongs to the Off-population. Neurons in the On-population are active relative to the positive of the image patch, while the Off-population neuron's frequency is set according to the negative of the image. The neurons in either population follow the dynamics of:

$$\tau_I \frac{dI}{dt} = -I + I_{in} - 2 \cdot I_2 \tag{23}$$

$$\tau_{I_2} \frac{dI_2}{dt} = -I_2 + I \tag{24}$$

$$\tau_v \frac{dv}{dt} = (-v + I) \tag{25}$$

$$\tau_\theta \frac{d\theta}{dt} = (0.2 - \theta) \tag{26}$$

$$\text{if } v > \theta : v = 0 \tag{27}$$

where I_{in} is the corresponding pixel value (On-population) or the pixel value multiplied by -1 (Off-population).

This experiment also adds an adaptive threshold to the neuronal model.

$$30ms \frac{d\theta}{dt} = (-60mV - \theta) \tag{28}$$

with $\theta(0) = -69mV$. This increases the activity during the first stimuli presentations. Furthermore, the resting potential was set to -70mV and upon reset the membrane potential jumped to -75mV.

To characterise receptive fields, we compute the mean and standard deviation of the receptive field r_i of each excitatory neuron. To quantify the distribution of orientations of the receptive field of each network, we convolved each with differently oriented Gabor patches. The Gabor patches are defined as:

$$G(\theta, f) = \cos\left(-2\pi f \cos(\theta) x + \pi f \sin(\theta) y + \frac{1}{2}\pi\right) \tag{29}$$

x and y are the pixel values relative to centre of the image, ($x, z \in (-5, 5)$). θ is the orientation of the Gabor patch and f the spatial frequency. Orientations are linearly spaced from $-90°$ to $90°$. Where $0°$ denotes a vertical oriented Gabor patch.

Reproducibility notice and code availability

All code is available at `https://github.com/prubisch/inhibitory_plasticity_regulation_in_voltage_dependent_plasticity_models`. Any simulations which include randomised variables have a set random seed. Following python packages were used for simulation and analysis: numpy (version 1.23.5)[49], matplotlib (version 3.8.4)[50], scikit-learn (version 1.2.2)[51] and the publicly available BRIAN2 package (version 2.5.1)[52].

ACKNOWLEDGMENTS

We would like to thank the members of the " Neurons and Systems" group for helpful discussions and comments. We would like to thank Isabel Cornacchia and the anonymous reviewers for the helpful comments.

AUTHOR CONTRIBUTIONS

Conceptualization, P.R. and M.H.H.; methodology, P.R.; investigation, P.R.; writing. P.R., M.I.S., M.H.H.; supervision, M.H.H.

DECLARATION OF INTERESTS

All authors declare no competing interests.

References

1. Abbott, L. F., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. Nature Neuroscience *3*, 1178–1183.

2. Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron *32*, 1149–1164.

3. Froemke, R. C., and Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. Nature *416*, 433–438.

4. Inglebert, Y., Aljadeff, J., Brunel, N., and Debanne, D. (2020). Synaptic plasticity rules with physiological calcium levels. Proceedings of the National Academy of Sciences *117*, 33639–33648.

5. Dudek, S. M., and Bear, M. F. (1993). Bidirectional long-term modification of synaptic effectiveness in the adult and immature hippocampus. Journal of Neuroscience *13*, 2910–2918.

6. Dudek, S. M., and Bear, M. F. Homosynaptic long-term depression in area ca1 of hippocampus and effects of n-methyl-d-aspartate receptor blockade. In: How We Learn; How We Remember: Toward An Understanding Of Brain And Neural Systems: Selected Papers of Leon N Cooper (200–204). World Scientific (1995):(200–204).

7. Malenka, R. C., and Bear, M. F. (2004). Ltp and ltd: an embarrassment of riches. Neuron *44*, 5–21.

8. Ngezahayo, A., Schachner, M., and Artola, A. (2000). Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. Journal of Neuroscience *20*, 2451–2458.

9. Steele, P. M., and Mauk, M. D. (1999). Inhibitory control of ltp and ltd: stability of synapse strength. Journal of Neurophysiology *81*, 1559–1566.

10. Bazelot, M., Bocchio, M., Kasugai, Y., Fischer, D., Dodson, P. D., Ferraguti, F., and Capogna, M. (2015). Hippocampal theta input to the amygdala shapes feedforward inhibition to gate heterosynaptic plasticity. Neuron *87*, 1290–1303.

11. Tsukada, M., Aihara, T., Kobayashi, Y., and Shimazaki, H. (2005). Spatial analysis of spike-timing-dependent ltp and ltd in the ca1 area of hippocampal slices using optical imaging. Hippocampus *15*, 104–109.

12. Nishiyama, M., Togashi, K., Aihara, T., and Hong, K. (2010). Gabaergic activities control spike timing-and frequency-dependent long-term depression at hippocampal excitatory synapses. Frontiers in Synaptic Neuroscience *2*, 1628.

13. Person, A. L., and Raman, I. M. (2010). Deactivation of l-type ca current by inhibition controls ltp at excitatory synapses in the cerebellar nuclei. Neuron *66*, 550–559.

14. Hayama, T., Noguchi, J., Watanabe, S., Takahashi, N., Hayashi-Takagi, A., Ellis-Davies, G. C., Matsuzaki, M., and Kasai, H. (2013). Gaba promotes the competitive selection of dendritic spines by controlling local ca2+ signaling. Nature Neuroscience *16*, 1409–1416.

15. Paille, V., Fino, E., Du, K., Morera-Herreras, T., Perez, S., Kotaleski, J. H., and Venance, L. (2013). Gabaergic circuits control spike-timing-dependent plasticity. Journal of Neuroscience *33*, 9353–9363.

16. Miehl, C., and Gjorgjieva, J. (2022). Stability and learning in excitatory synapses by nonlinear inhibitory plasticity. PLOS Computational Biology *18*, e1010682.

17. Rossbroich, J., and Zenke, F. (2023). Dis-inhibitory neuronal circuits can control the sign of synaptic plasticity. arXiv.

18. Agnes, E. J., and Vogels, T. P. (2024). Co-dependent excitatory and inhibitory plasticity accounts for quick, stable and long-lasting memories in biological networks. Nature Neuroscience *27*, 964–974.

19. Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based stdp with homeostasis. Nature Neuroscience *13*, 344–352.

20. Meissner-Bernard, C., Tsai, M. C., Logiaco, L., and Gerstner, W. (2020). Dendritic voltage recordings explain paradoxical synaptic plasticity: a modeling study. Frontiers in Synaptic Neuroscience *12*, 585539.

21. Clopath, C., and Gerstner, W. (2010). Voltage and spike timing interact in stdp–a unified model. Frontiers in Synaptic Neuroscience *2*, 25.

22. Brandalise, F., Carta, S., Helmchen, F., Lisman, J., and Gerber, U. (2016). Dendritic nmda spikes are necessary for timing-dependent associative ltp in ca3 pyramidal cells. Nature Communications *7*, 13480.

23. Lisman, J., and Spruston, N. (2005). Postsynaptic depolarization requirements for ltp and ltd: a critique of spike timing-dependent plasticity. Nature Neuroscience *8*, 839–841.

24. Golding, N. L., Staff, N. P., and Spruston, N. (2002). Dendritic spikes as a mechanism for cooperative long-term potentiation. Nature *418*, 326–331.

25. Stiefel, K. M., Tennigkeit, F., and Singer, W. (2005). Synaptic plasticity in the absence of backpropagating spikes of layer ii inputs to layer v pyramidal cells in rat visual cortex. European Journal of Neuroscience *21*, 2605–2610.

26. Hardie, J., and Spruston, N. (2009). Synaptic depolarization is more effective than back-propagating action potentials during induction of associative long-term potentiation in hippocampal pyramidal neurons. Journal of Neuroscience *29*, 3233–3241.

27. Magee, J. C., and Johnston, D. (1997). A synaptically controlled, associative signal for hebbian plasticity in hippocampal neurons. Science *275*, 209–213.

28. Christie, B. R., Magee, J. C., and Johnston, D. (1996). The role of dendritic action potentials and ca2+ influx in the induction of homosynaptic long-term depression in hippocampal ca1 pyramidal neurons. Learning & Memory *3*, 160–169.

29. Sjöström, P. J., and Häusser, M. (2006). A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. Neuron *51*, 227–238.

30. Fuenzalida, M., Fernandez de Sevilla, D., Couve, A., and Buno, W. (2010). Role of ampa and nmda receptors and back-propagating action potentials in spike timing–dependent plasticity. Journal of Neurophysiology *103*, 47–54.

31. Keck, T., Toyoizumi, T., Chen, L., Doiron, B., Feldman, D. E., Fox, K., Gerstner, W., Haydon, P. G., Hübener, M., Lee, H.-K. et al. (2017). Integrating hebbian and homeostatic plasticity: the current state of the field and future research directions. Philosophical Transactions of the Royal Society B: Biological Sciences *372*, 20160158.

32. Földiak, P. (1990). Forming sparse representations by local anti-hebbian learning. Biological Cybernetics *64*, 165–170.

33. Van Hateren, J. H., and van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. Proceedings of the Royal Society of London. Series B: Biological Sciences *265*, 359–366.

34. Bell, A. J., and Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. Vision Research *37*, 3327–3338.

35. Brito, C. S., and Gerstner, W. (2016). Nonlinear hebbian learning as a unifying principle in receptive field formation. PLoS Computational Biology *12*, e1005070.

36. Savin, C., Joshi, P., and Triesch, J. (2010). Independent component analysis in spiking neurons. PLoS Computational Biology *6*, e1000757.

37. Zylberberg, J., Murphy, J. T., and DeWeese, M. R. (2011). A sparse coding model with synaptically local plasticity and spiking neurons can account for the diverse shapes of v1 simple cell receptive fields. PLoS Computational Biology *7*, e1002250.

38. Zhu, M., and Rozell, C. J. (2013). Visual nonclassical receptive field effects emerge from sparse coding in a dynamical system. PLoS Computational Biology *9*, e1003191.

39. Krizhevsky, A., Hinton, G. et al. (2009). Learning multiple layers of features from tiny images.

40. Kandel, E. R., Schwartz, J. H., Jessell, T. M., Siegelbaum, S., Hudspeth, A. J., and Mack, S. Principles of neural science vol. 5. McGraw-hill New York (2013).

41. Branco, T., and Häusser, M. (2011). Synaptic integration gradients in single cortical pyramidal cell dendrites. Neuron *69*, 885–892.

42. Oviedo, H., and Reyes, A. D. (2005). Variation of input-output properties along the somatodendritic axis of pyramidal neurons. Journal of Neuroscience *25*, 4985–4995.

43. Magee, J., Hoffman, D., Colbert, C., and Johnston, D. (1998). Electrical and calcium signaling in dendrites of hippocampal pyramidal neurons. Annual Review of Physiology *60*, 327–346.

44. Cizeron, M., Qiu, Z., Koniaris, B., Gokhale, R., Komiyama, N. H., Fransén, E., and Grant, S. G. (2020). A brainwide atlas of synapses across the mouse life span. Science *369*, 270–275.

45. Zhu, F., Collins, M. O., Harmse, J., Choudhary, J. S., Grant, S. G., and Komiyama, N. H. (2020). Cell-type-specific visualisation and biochemical isolation of endogenous synaptic proteins in mice. European Journal of Neuroscience *51*, 793–805.

46. Chindemi, G., Abdellah, M., Amsalem, O., Benavides-Piccione, R., Delattre, V., Doron, M., Ecker, A., Jaquier, A. T., King, J., Kumbhar, P. et al. (2022). A calcium-based plasticity model for predicting long-term potentiation and depression in the neocortex. Nature Communications *13*, 3038.

47. Hiratani, N., and Fukai, T. (2017). Detailed dendritic excitatory/inhibitory balance through heterosynaptic spike-timing-dependent plasticity. Journal of Neuroscience *37*, 12106–12122.

48. Teeter, C., Iyer, R., Menon, V., Gouwens, N., Feng, D., Berg, J., Szafer, A., Cain, N., Zeng, H., Hawrylycz, M. et al. (2018). Generalized leaky integrate-and-fire models classify multiple neuron types. Nature Communications *9*, 1–15.

49. Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. (2020). Array programming with NumPy. Nature *585*, 357–362. URL: `https://doi.org/10.1038/s41586-020-2649-2`. doi:`10.1038/s41586-020-2649-2`.

50. Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. Computing in Science & Engineering *9*, 90–95. doi:`10.1109/MCSE.2007.55`.

51. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research *12*, 2825–2830.

52. Stimberg, M., Brette, R., and Goodman, D. F. (2019). Brian 2, an intuitive and efficient neural simulator. eLife *8*, e47314.

53. Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. The Journal of Physiology *117*, 500–544.

Appendix

Illustration of the Voltage-Dependent Pathway model The VDP models the downstream processes of plasticity induction as a simplified pathway-dynamic which is activated by depolarisation of the postsynaptic membrane. Therefore, the overall weight update from a single pairing can only be determined after several milliseconds when the pathway activations for LTP and LTD have both returned to 0. It emphasizes the temporal competition between LTP and LTD pathway activation as the central mechanism determining the sign and amplitude of the synaptic change. The illustration provided in Figure 7 visualises the mechanism used in the VDP.

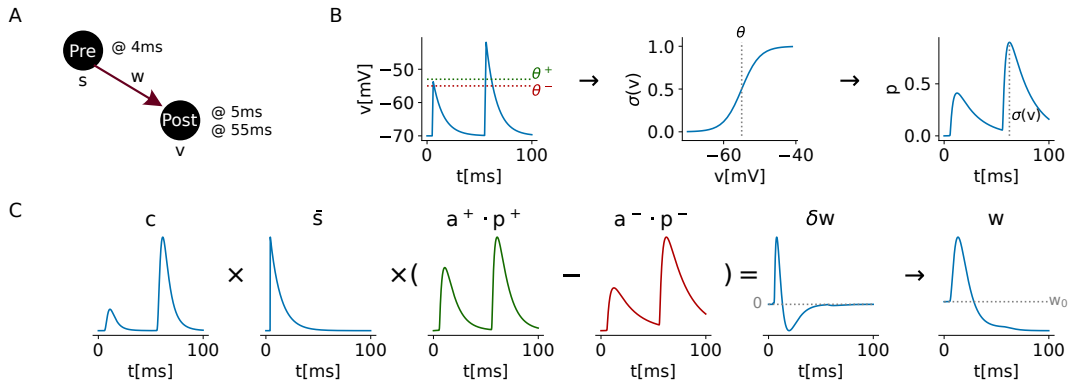

Figure 7: Illustration of the VDP model. **A**: Stimulation schema. The presynaptic spike train s and the postsynaptic membrane potential v are measured to compute the synaptic weight update. The time points denote spike and postsynaptic stimulation time to elicit the activation of the traces in B and C. **B**: Visualisation of the pathway induction process. The instantaneous membrane potential (left) is applied to a sigmoid function (middle) as a soft threshold with the middle points θ^+ and θ^- for the potentiation and depression pathway. The result of the sigmoid operation is used as the maximum amplitude or offset for the biexponential pathway activation p (right). **C**: Illustration of the computation and functions in the synaptic weight update (equation 17). The component is a scaling factor mimicking calcium concentration dependent on the potentiation and depression pathway. This has been omitted in the equations to allow a direct comparison of the models. The scaling does not qualitatively influence the inhibitory effect. From left to right: The calcium level c is multiplied with the presynaptic trace $\bar{s}$ and the difference between the potentiation p^+ and depression pathway p^- in order to calculate the momentary weight w. Due to the amplitude parameter $a_+ < 1$, the illustrative example results overall in LTD, despite the weight increasing at first after the induction mechanism.

Reproduction of central experimental results with the VDP To test the validity of the VDP, we tested whether core results of the experimental literature focusing on excitatory synaptic plasticity can be reproduced. The results which show that the VDP can be considered a valid synaptic plasticity model are shown in Figure 8.

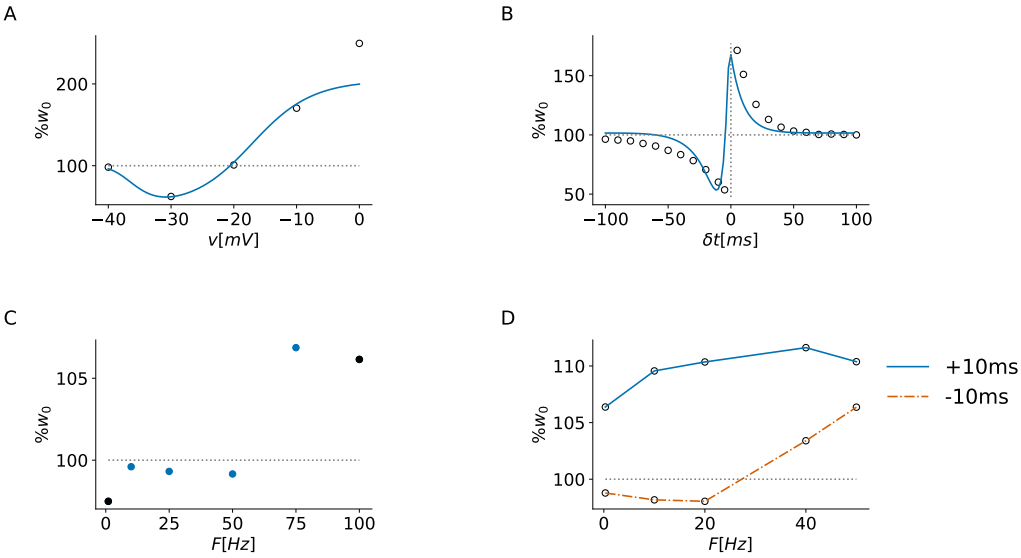

Figure 8: Predictions of the VDP for several experimental protocols. **A**: The voltage dependency follows the relationship as described by Ngezahayo et al.[8] closely. **B**: The VDP can reproduce the standard STDP kernel as measured by Froemke and Dan[3]. The open-faced circles indicate the experimental measurements. **C**: Low Frequency Stimulation (LFS) and Theta Burst Stimulation effects (TBS) are reproduced. The solid circles indicate conditions measured in Dudek and Bear[5]. **C**: The qualitative effect of the Pairing Frequency as described in Sjöström et al.[2] can be reproduced by the VDP. Open-faced circled here indicate the conditions tested by Sjöström et al.[2]. They do not represent their measurements, as fitting the data quantitatively is limited by the description of the experimental protocol and neuron model used.

Furthermore, by changing the relative dynamics of the potentiation and depression pathway the VDP can reproduce different STDP kernel shapes as described in Abbott and Nelson[1]. To characterise the different kernel shapes we measure the shape s defined as :

$$s = \frac{\sum_{i=0}^{t} \delta w^{\delta i} - \delta w^{-\delta i}}{\sum_{i=-t}^{t} |\delta w^{\delta i}|} \tag{30}$$

Figure 9 shows four distinct types of STDP kernel shapes alongside their s and r (LTP-LTD ratio) scores.

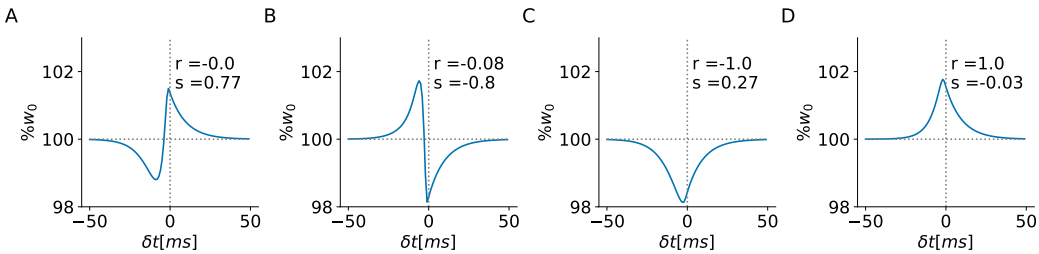

Figure 9: Different STDP kernel shapes reproduced by the VDP. The differences can be measured by calculating the kernel shape s and LTP-LTD ratio r. **A**: Balanced Hebbian STDP kernel with r = 0.0 and s = 0.77. **B**: Balanced anti-Hebbian STDP kernel with r = -0.08 and s = -0.8. **C**: LTD-dominated nearly-symmetric kernel. r = -1.0, s = 0.27. **D**: LTP-dominated nearly symmetric STDP kernel. r = 1.0 and s = -0.03.

EI balance and membrane potential statistics in a Hodgkin-Huxley regular spiking neuron

All simulations use a LIF neuron as described in Section Methods. As this is a simplified neuron model without an explicit spike shape estimation, it is central to test whether the relationship between EI balance and membrane potential statistics holds in neuron models which approximate the spike shape. The results of the same simulation with a regular spiking Hodgin-Huxley model[53](HH) are illustrated in Figure 10. It confirms that inhibitory-dominated synaptic input results in a higher mean membrane potential at similar firing rates in comparison to excitatory-dominated synaptic input. However, the dip in the variance of the membrane potential at low firing rates is not present in this simulation, indicating that this is an artefact caused by the (missing) approximation of the spike shape of the LIF neuron.

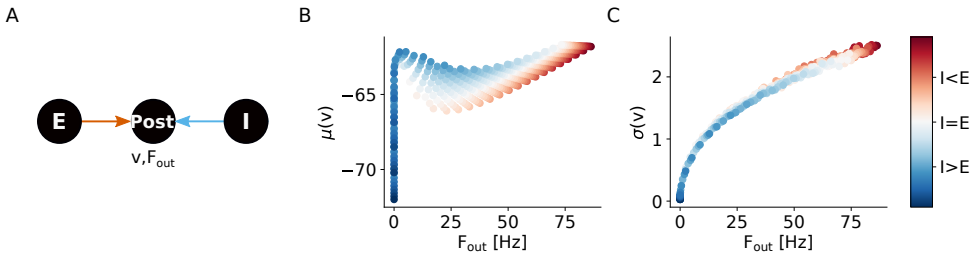

Figure 10: EI balance influences on mean and standard deviation of the membrane potential is similar in LIF neurons and HH neuron.**A** A postsynaptic HH neuron is driven by 100 excitatory and 100 inhibitory Poisson neurons with firing rates $F_E, F_I \in [0, 100)$Hz (5Hz resolution). The colour-coding of the data points indicates the EI balance of the input computed as $\frac{F_E}{F_I}$. We measure the postsynaptic membrane potential, output frequency, as well as the weight change of an excitatory (orange,dashed) synapse. The presynaptic neuron of that synapse fires regularly with the same frequency as the excitatory Poisson population. All results are calculated over 10s simulation time. **B**: The mean membrane potential increases when the synaptic input is inhibition-dominated. **C**: The standard deviation of the membrane potential measuring the fluctuation amplitude is increasing when the synaptic input is dominated by excitation.

Inhibitory frequency dependency in voltage-dependent plasticity models The simulation results portrayed in Figure 3 also enable us to investigate the development of inhibitory weights. The presynaptic frequency is set to be equivalent to the frequency of the inhibitory Poisson population. The results for this simulation are shown in Figure 11. They show that the inhibitory weight development is strongly driven by the presynaptic firing rate. At high postsynaptic frequencies the update amplitude decreases progressively due to the decrease in presynaptic frequency. The models differ only slightly in their inhibitory frequency dependency and show similar qualitative behaviour.

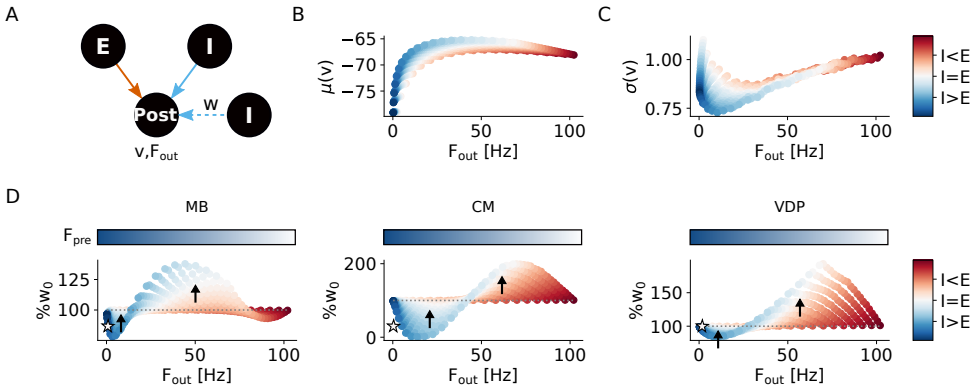

Figure 11: The frequency dependency of inhibitory synapses is driven by the presynaptic activity leading to similar patterns in all plasticity models. **A** A postsynaptic neuron is driven by 100 excitatory and 100 inhibitory Poisson neurons with firing rates $F_E, F_I \in [0, 100)$Hz (5Hz resolution). The colour-coding of the data points indicates the EI balance of the input computed as $\frac{F_E}{F_I}$. We measure the postsynaptic membrane potential, output frequency, as well as the weight change of an inhibitory (blue,dashed) synapse. The presynaptic neuron of that synapse fires regularly with the same frequency as the excitatory Poisson population. All results are calculated over 10s simulation time. **B**: The mean membrane potential increases when the synaptic input is inhibition-dominated. **C**: The standard deviation of the membrane potential measuring the fluctuation amplitude is increasing when the synaptic input is dominated by excitation. **D**: Plasticity predictions of the plasticity models of the inhibitory weights. Model identity is indicated by the abbreviations on top. All results are averaged over 100 simulation runs. The color bar on top indicates the frequency of the regular spiking presynaptic neuron. The open-faced star indicates the regime of pathological behaviour at low postsynaptic frequencies (negative plasticity-activity feedback). Arrows indicate the direction of EI balance shift due to the predicted synaptic plasticity update.

Filter properties of the plasticity models The postsynaptic filter of the MB and CM are low-pass filters by definition. However, adding the threshold and in case of the CM combining a filtered and an unfiltered condition in the plasticity induction introduces conditional filtering. The VDPs down-stream filter in form of the pathway activities can neither be described as a low-pass nor high-pass filter by definition. Therefore, we simulated a burst of 5 spikes (200Hz) repeated with either 5Hz (Figure 12a) or 50Hz (Figure 12b) to visualise the differences in the filter properties of the models.

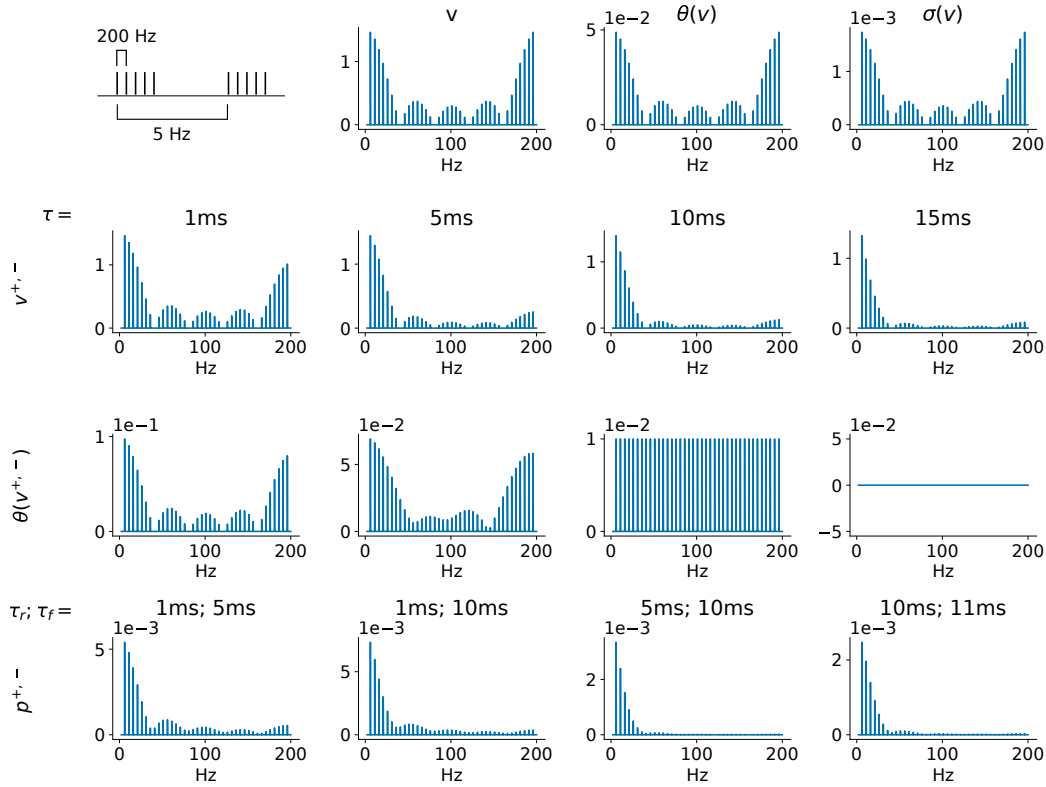

(a) Filter properties of the plasticity model when a burst is repeated with 5Hz. The results show the frequency spectra calculated using a Fourier-Transform of the membrane potential v, membrane potential after applying the threshold $\theta(v), \sigma(v)$, the filtered membrane potential $v^{+,-}$with different filter strengths as indicated above the plot. Further the frequency spectra of the applied induction conditions in the MB and CM $\theta(v^{+,-})$, as well as the frequency spectra of the pathway activities $p^{+,-}$ is shown. The applied time constants (τ_r, τ_f) are shown in the sub-figure titles.

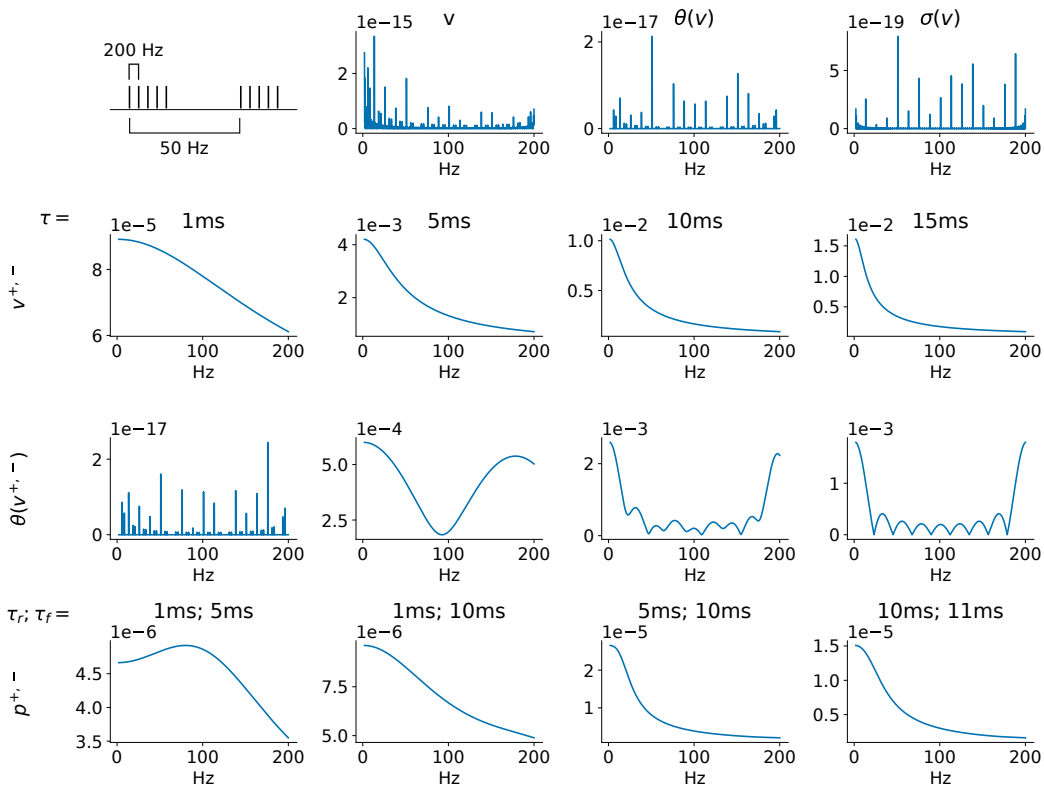

(b) Filter properties of the plasticity model when a burst is repeated with 50Hz.

Influence of time constants on the sensitivity to inhibition onset A central question which arises, when comparing different synaptic plasticity models is whether similar results can be achieved simply by adapting model parameters. In this work we argue that the temporal extent of the inhibitory regulation as well as the sensitivity is dependent on the model definition. In the MB and CM models show a parameter-sensitive fast fluctuation sensitivity, while the VDP's fast fluctuation sensitivity is a model feature and independent of its parametrisaiton. To demonstrate the parameter influence on the plasticity patterns elicit by changing stimuli, we vary the time constants and measure how they influence the slope of weight change of the mechanistic simulation illustrated in Figure 4. Figure 13 demonstrates that the presynaptic temporal constant has the biggest effect in the MB anc CM model behaviour by allowing a longer effective pairing window in these models. The VDP has slower postsynaptic dynamics which is why the presynaptic constant has a less strong effect in this model while the postsynaptic constants influence the slope by prolonging pathway activity.

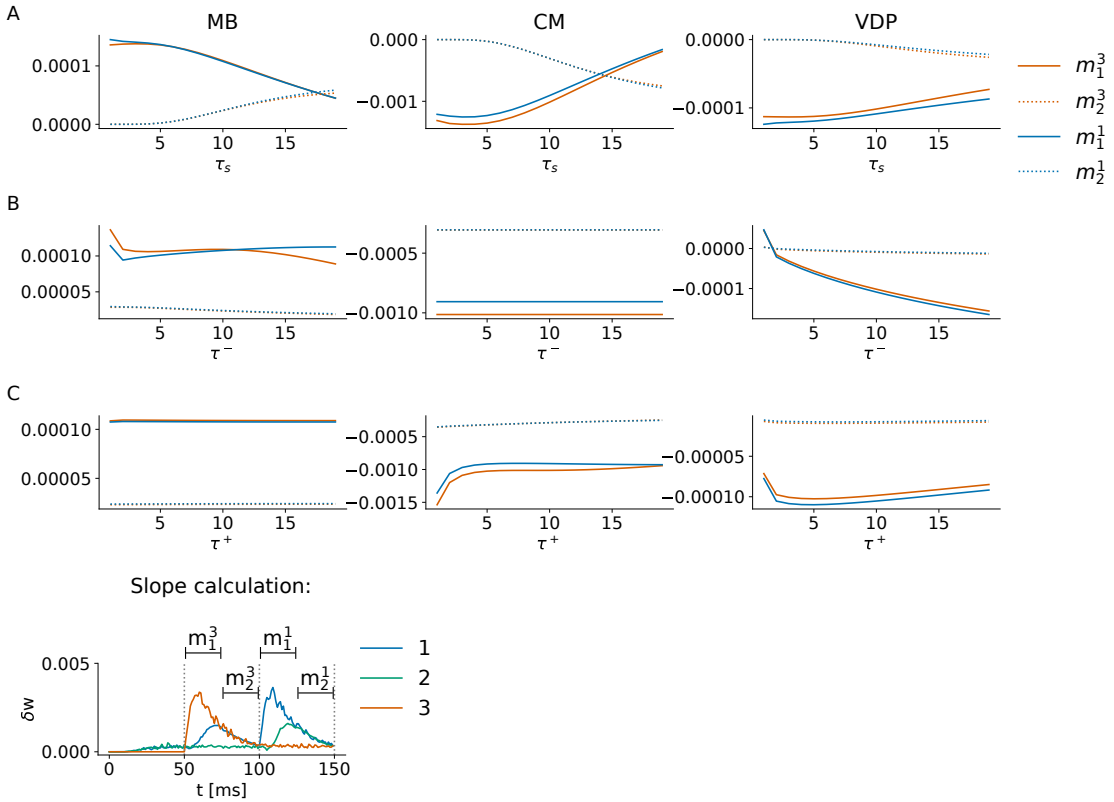

Figure 13: Different STDP kernel shapes reproduced by the VDP. The differences can be measured by calculating the kernel shape s and LTP-LTD ratio r. **A**: Balanced Hebbian STDP kernel with r = 0.0 and s = 0.77. **B**: Balanced anti-Hebbian STDP kernel with r = -0.08 and s = -0.8. **C**: LTD-dominated nearly-symmetric kernel. r = -1.0, s = 0.27. **D**: LTP-dominated nearly symmetric STDP kernel. r = 1.0 and s = -0.03.

Within-network receptive field similarity Figure 14A shows further results on the clustering of the receptive field statistics. To investigate the robustness of the clusters, the statistics were randomly shuffled (grey data). The similarity between receptive fields of the same network and its relationship to the variance of the receptive fields, Figure 14B-D, show the gradual change of receptive fields in the CM and VDP network, while the MB network builds clusters of receptive field with similar structures which is also captured by the clustering in their descriptive statistics.

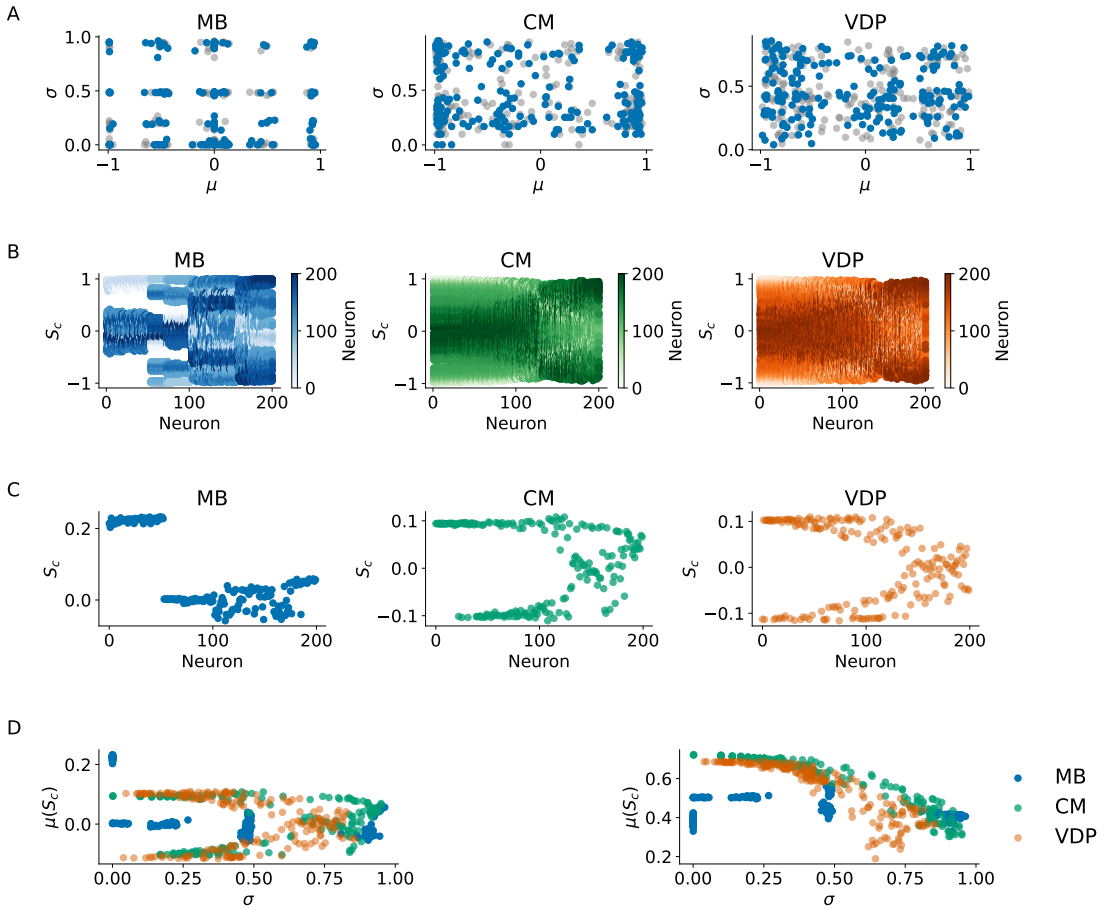

Figure 14: Clustering and further similarity analysis of the receptive fields. **A**:Original pairings of mean μ and variance σ (blue) and reshuffled pairs (gray). **B**: Cosine similarity $S_c(i,j)$ of the neurons i and j receptive field. **C**: Mean cosine similarity $\mu(S_c(i,j))$ or neuron i. **D**: Relationship between the standard deviation σ and $\mu(S_c(i,j))$ (left) and $\mu|(S_c(i,j))|$ (right).

Parametrisation details

type	τ_r^+	τ_f^+	θ^+	τ_r^-	τ_f^-	θ^-	τ_s	a^+	a^-	a_s	η	Ca^{2+}	experiments
standard	1 ms	5 ms	-40.75 mV	1 ms	20 ms	-43.75 mV	10ms	1	0.18	0.1	1	3.0	Figure 1
								1	0.5	0.1	.1	3.0	Figure 2A
								1	0.17	0.1	100	3.0	Figure 3, Figure 5
								1	0.2	0.1	1	3.0	Figure 4
stronger	10 ms	5 ms	-40.75 mV	10 ms	20 ms	-43.75 mV	10ms	1	0.25	0.1	0.2	3.0	Figure 2B & C
weaker	1 ms	2 ms	-40.75 mV	1 ms	5 ms	-40.75 mV	10ms	1	0.2	0.1	0.2	3.0	Figure 2B & C
CIFAR10*	1ms	10ms	-63mV	12ms	30ms	-66mV	10ms	1	0.95	0.03	1	1.9	Figure 6

*:This condition also deviates from the standard slope parameter of sigmoid function in the pathway dynamics, $k^+ = .7$ and $k^- = 1$.

Table 2: Parametrisation of the VDP

type	τ^+	θ^+	τ^-	θ^-	τ_s	a^+	a^-	a_s	η	experiments
standard	7 ms	-68 mV	10 ms	-75mV	15 ms	2.e-4	1.e-2	$\frac{1}{15}$	.1	Figure1
		-58 mV		-61 mV		8.e-5	12.e-3		1	Figure 2A
		-58 mV		-69 mV		8.e-4	1.e-2		1	Figures3, 4
		-50 mV		-70 mV		8.e-2	2.e-4		.01	Figure5
		-66 mV		-70 mV		8.e-4	2.e-2		.1	Figure6
stronger	14 ms	-58 mV	20 ms	-61mV	15 ms	16.e-5	8.e-4	$\frac{1}{15}$	1	Figure 2B-C
weaker	1 ms	-58 mV	1 ms	-61mV	15 ms	16.e-5	8.e-4	$\frac{1}{15}$	.1	Figure 2B-C

Table 3: Parametrisation of the CM

type	τ^+	θ^+	τ^-	θ^-	τ_s	a^+	a^-	a_s	η	experiments
standard	7 ms	-70 mV	10 ms	-75mV	15 ms	2.e-3	3.e-4	$\frac{1}{15}$	10	Figure1
		-58 mV		-61 mV		8.e-3	6.e-3		1	Figure2A
		-65 mV		-75 mV		8.e-4	8.e-5		1	Figure3
		-65 mV		-75 mV		8.e-4	8.e-5		10	Figure4
		-55 mV		-75 mV		2.e-2	8.e-5		.1	Figure5
		-68 mV		-70 mV		1.5.e-4	8.e-5		1	Figure6
stronger	14 ms	-58 mV	20 ms	-61mV	15 ms	10.e-3	4.e-3	$\frac{1}{15}$	1	Figure 2B-C
weaker	1 ms	-58 mV	1 ms	-61mV	15 ms	10.e-3	4.e-3	$\frac{1}{15}$	1	Figure 2B-C

Table 4: Parametrisation of the MB