Assemblies, synapse clustering and network topology interact with plasticity to explain structure-function relationships of the cortical connectome

András Ecker[1†], Daniela Egas Santander[1†], Marwan Abdellah[1], Jorge Blanco Alonso[1], Sirio Bolaños-Puchet[1], Giuseppe Chindemi[1,2], Dhuruva Priyan Gowri Mariyappan[3,4,5], James B. Isbister[1], James Gonzalo King[1], Pramod Kumbhar[1], Ioannis Magkanaris[1], Eilif B. Muller[1,3,4,5], and Michael W. Reimann[1*]

[1]Blue Brain Project, École polytechnique fédérale de Lausanne (EPFL), Campus Biotech, Geneva, Switzerland
[2]Department of Basic Neurosciences, University of Geneva, Geneva, Switzerland
[3]Department of Neurosciences, Faculty of Medicine, Université de Montréal, Montréal, Canada
[4]CHU Sainte-Justine Research Center, Montréal, Canada
[5]Mila Quebec AI Institute, Montréal, Canada
[†]Co-lead authors
[*]Correspondance: michael.reimann@epfl.ch

Abstract

Synaptic plasticity underlies the brain's ability to learn and adapt. While experiments in brain slices have revealed mechanisms and protocols for the induction of plasticity between pairs of neurons, how these synaptic changes are coordinated in biological neuronal networks to ensure the emergence of learning remains poorly understood. Simulation and modeling have emerged as important tools to study learning in plastic networks, but have yet to achieve a scale that incorporates realistic network structure, active dendrites, and multi-synapse interactions, key determinants of synaptic plasticity. To rise to this challenge, we endowed an existing large-scale cortical network model, incorporating data-constrained dendritic processing and multi-synaptic connections, with a calcium-based model of functional plasticity that captures the diversity of excitatory connections extrapolated to *in vivo*-like conditions. This allowed us to study how dendrites and network structure interact with plasticity to shape stimulus representations at the microcircuit level. In our simulations, plasticity acted sparsely and specifically, firing rates and weight distributions remained stable without additional homeostatic mechanisms. At the circuit level, we found plasticity was driven by co-firing stimulus-evoked functional assemblies, spatial clustering of synapses on dendrites, and the topology of the network connectivity. As a result of the plastic changes, the network became more reliable with more stimulus-specific responses. We confirmed our testable predictions in the MICrONS datasets, an openly available electron microscopic reconstruction of a large volume of cortical tissue. Our results quantify at a large scale how the dendritic architecture and higher-order structure of cortical microcircuits play a central role in functional plasticity and provide a foundation for elucidating their role in learning.

Keywords

Plasticity, simulation, connectomics, cell assembly, dendrites, synapse clustering, network topology

1 Introduction

Learning and memory are orchestrated by synaptic plasticity, the ability of synapses to change their *efficacy* in an activity-dependent manner. Donald O. Hebb's postulate about how synaptic plasticity might manifest was paraphrased with the well known mantra: *"cells that fire together, wire together"* (Hebb, 1949; Shatz, 1992). The first proof of coincident pre- and postsynaptic population activity leading to *potentiation* (an increase in efficacy) came from pathway stimulation in hippocampal slices (Bliss and Lømo, 1973). It was later confirmed at the neuron pair level (Markram et al., 1997; Bi and Poo, 1998), and spike-time dependent plasticity (STDP) became a quintessential protocol to study Hebbian plasticity *in vitro*. In the early 2000s a plethora of cortical pathways were studied and plasticity proved to be synapse location- and therefore pathway-dependent (Sjöström and Häusser, 2006; Letzkus et al., 2006; Froemke et al., 2010). The molecular substrate of Hebbian coincidence detection is the N-methyl-D-aspartate (NMDA) receptor, which upon removal of the Mg^{2+} block by depolarization, conducts Ca^{2+} (Mayer et al., 1984). The calcium-control hypothesis, put forward by Lisman (1989) postulates that prolonged, moderate amounts of Ca^{2+} lead to *depression* (a decrease in efficacy) while large transients of Ca^{2+} lead to potentiation. By putting these together, it became evident that it is not necessarily the timing of the postsynaptic spike, but rather the depolarization induced entrance of calcium into the postsynaptic dendrite that is important to evoke changes in synaptic efficacy (Goldberg et al., 2002; Lisman and Spruston, 2005).

In parallel with slice electrophysiology, Hebbian plasticity was also studied through its effect on behavior via fear conditioning experiments (McKernan and Shinnick-Gallagher, 1997) and this line of research led to numerous new techniques for tagging and re-activating cells that participate in newly formed memories (Tonegawa et al., 2015). While these studies highlighted the need to study plasticity at the network level, most changes are expected to happen at the synapse level. Therefore, high-throughput methods tracking synaptic proteins like PSD95 (Ray et al., 2023) and α-amino-3-hydroxy-5-methyl-4-isoxazolepropionate (AMPA) subunit GluA1 (Graves et al., 2021; Kim et al., 2023) are currently being developed. While readily applicable to monitor synaptic efficacy *in vivo*, currently these techniques cannot be supplemented with recordings of neural activity, thus the reason for the changes in efficacy can only be speculated.

The bridge between *in vitro* pairs of neurons and *in vivo* behavior is often provided by complementary, simulation based-approaches. Early theoretical work explored the potential link between memories and cells that fire and therefore wire together, concentrating on the storage and retrieval of memories in strongly recurrent networks (Hopfield, 1982), which remained an active topic of research (Fusi and Abbott, 2007; Krotov and Hopfield, 2016; Widrich et al., 2020). Inspired by STDP experiments, modelers have postulated diverse *plasticity rules* that fit the most recent experimental findings (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000; Pfister and Gerstner, 2006; Clopath et al., 2010). Models that take into account the crucial role of calcium in driving plasticity outcomes have also been proposed (Shouval et al., 2002; Graupner and Brunel, 2012; Rubin et al., 2005; Jędrzejewska-Szmek et al., 2017; Rodrigues et al., 2023). The calcium-based model of Graupner and Brunel (2012) describes the evolution of intracellular calcium concentration ($[Ca^{2+}]_i$) given the pre- and postsynaptic spike trains and updates the efficacy of the synapse, upon $[Ca^{2+}]_i$ crossing thresholds for depression and potentiation. Bioplausible plasticity rules have been shown to bring about Hebbian *cell assemblies*, i.e., groups

of neurons that fire together to enable long-term memory storage (Litwin-Kumar and Doiron, 2014; Zenke et al., 2015; Fauth and Van Rossum, 2019; Kossio et al., 2021). A common theme in these models is the necessity for homeostatic plasticity to keep the networks stable. While experimental evidence exists for homeostatic plasticity (Turrigiano and Nelson, 2004), this has been shown to be too slow to induce stability (Zenke et al., 2017a) and alternative mechanism have been proposed (Vogels et al., 2011; Delattre et al., 2015). While these studies provide mechanistic explanations of learning and memory, they used point-neuron models, therefore neglecting the structural and functional importance of dendrites in plasticity (with the exception of Bono et al. (2017); Kastellakis and Poirazi (2019)). The compartmentalized nature of dendritic trees gives rise to spatial clustering of synapses (Kastellakis and Poirazi, 2019; Farinella et al., 2014; Iacaruso et al., 2017; Tazerart et al., 2020) and local, non-linear voltage events (Poirazi et al., 2003; Stuart and Spruston, 2015) both of which are thought to contribute to the removal of the Mg^{2+} block from NMDA receptors and therefore gate plasticity.

To investigate the effect of dendrites and multi-synaptic interactions for plasticity at the network level, in this study we equipped the biophysically detailed, large-scale cortical network model of Isbister et al. (2023) with our recently developed model of functional plasticity (Chindemi et al., 2022) between excitatory cells (Figure 1). At this level of detail, a calcium-based plasticity model is the natural choice since calcium dynamics are already part of the simulation. Moreover, it allowed us to model *in vivo*-like conditions, specifically the low extracellular calcium concentration ($[Ca^{2+}]_o$, Chindemi et al., 2022), which has been experimentally shown to reduce plasticity (Inglebert et al., 2020, Figure 1E). As we had access to the pre- and postsynaptic activity and efficacy of millions of synapses including their dendritic location we could characterize the rules governing plasticity at the microcircuit level. During simulated evoked activity, plasticity acted sparsely and specifically, although it was still able to reorganize the network's dynamics, manifesting in more pattern-specificity after plasticity. Potentiation dominated in amplitude and depression counteracted it in frequency, which led to stable firing rates without explicitly introducing any homeostatic terms (Zenke et al., 2017a; Turrigiano and Nelson, 2004). We found plasticity was driven by the network topology beyond pairwise connectivity. Specifically, by co-firing cell assemblies, spatial clustering of synapses on dendrites, and the centrality of a connection in the network. Furthermore, we confirmed our testable predictions in the MICrONS (2021) dataset, an openly available electron microscopic reconstruction.

2 Results

To study how plasticity and various network level features interact to explain structure-function relationships, we used a bio-realistic, large-scale cortical model of the rat non-barrel somatosensory cortex (nbS1). By doing so, we had a continuous readout of both the activity of all neurons and the efficacy of the synapses between them under *in vivo*-like, low $[Ca^{2+}]_o$ conditions ($\sim 1\,\mathrm{mM}$ instead of $2-2.5\,\mathrm{mM}$ *in vitro*). The current model improves on Markram et al. (2015) in terms of both anatomical, e.g., atlas-based cell composition and placement (described in Reimann et al. (2022)), and physiological properties, e.g., improved single cell models, multi-vesicular synaptic release, and layer-wise compensation for missing synapses (described in Isbister et al. (2023)). For this study, we used a seven column subvolume comprising 211,712 neurons in $2.4\,mm^3$ of tissue (Figure 1A). In line with the biological variability, excitatory cells are modeled as a diverse set of morphologies (Reimann et al., 2022; Kanari et al., 2019) (Figure 1B) equipped with conductances distributed across all electrical compartments (Reva et al., 2022) (Supplementary Figure S1A). The connectivity and synaptic physiology of these cells were extensively validated (Isbister et al., 2023; Reimann et al., 2022) (Figure 1C; Supplementary Figure S1C). To deliver input with spatio-temporal precision to the circuit, we used thalamic fibers from the ventral

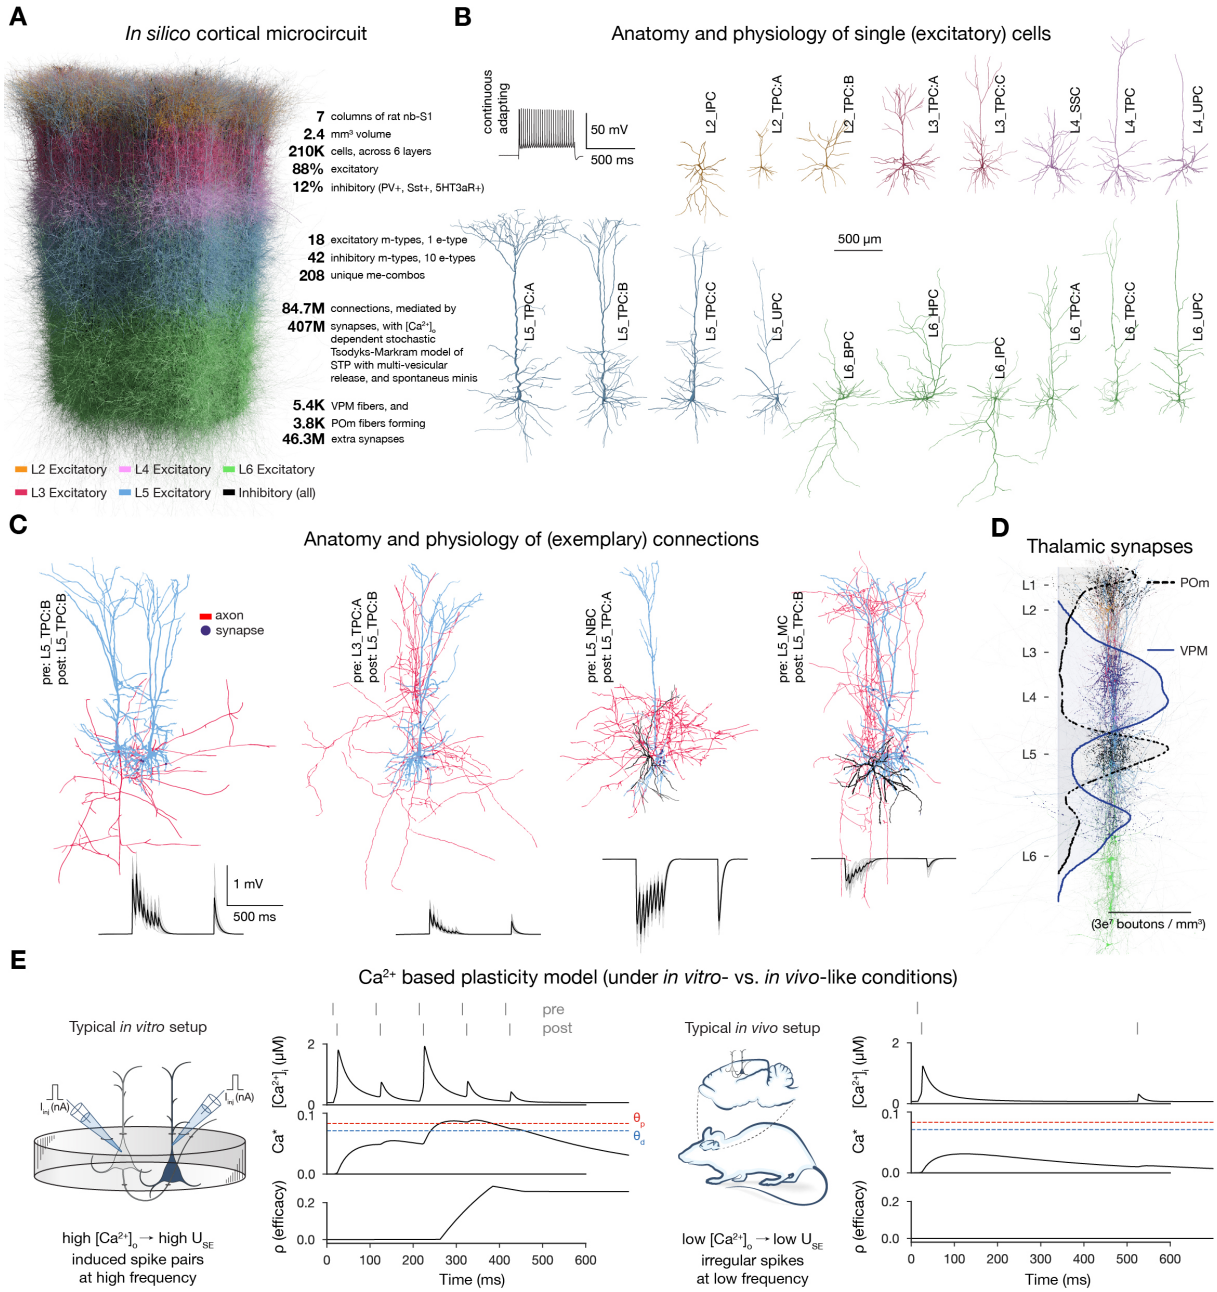

Figure 1: Overview of the cortical network model and calcium-based plasticity rule. A: Rendering of the seven column subvolume of rat nbS1. (10% of the cells are shown). **B:** Representative morphologies for the 18 excitatory morphological types (m-types) and their typical firing pattern (top left). Neurite diameters are scaled (x2) for better resolution. **C:** Exemplary connections to layer 5 (thick-tufted pyramidal cells, L5_TPC:A and L5_TPC:B) L5 TTPCs (top) and their short-term dynamics (bottom). Neurite diameters are uniform for better visibility. On the bottom: thin gray lines represent 20 individual trials, while the thicker black ones their means. **D:** Bouton density profiles of thalamocortical fibers, and locations of VPM (black) and POm (purple) synapses on neurons (in a 5 μm radius subvolume). The spatial scale bar on B applies to the whole figure. Panels adapted from Isbister et al. (2023); Chindemi et al. (2022); Reimann et al. (2022). **E:** Main variables of the plasticity model during coincident activation of the pre- and postsynaptic neurons for an exemplary synapse. Left: under *in vitro*-like conditions (adapted from Chindemi et al. (2022)). Right: same pair of neurons under *in vivo*-like conditions. Schematics on their lefts illustrate the difference between typical *in vitro* and *in vivo* protocols.

posteriomedial nucleus of the thalamus (VPM) and the high-order posteriomedial nucleus of the thalamus (POm; Figure 1D; Meyer et al., 2010).

To simulate long-term plasticity, we integrated our recently published calcium-based plasticity model that was used to describe functional long-term potentiation and depression between pairs of pyramidal cells (PCs) (Chindemi et al., 2022). In short, the model follows the formalism of Graupner and Brunel (2012), where pre- and postsynaptic spikes led to changes in synaptic $[Ca^{2+}]_i$ (Figure 1E). Calcium entering though NMDA receptors and voltage-dependent calcium channels (VDCCs) contributes to $[Ca^{2+}]_i$ (Equation (2) in Methods). When the integrated calcium trace of a synapse crosses the threshold for depression (θ_d) or the higher one for potentiation (θ_p), *synaptic efficacy* (ρ) is updated (Figure 1E left; Equation (1) in Methods). Changes in ρ are then converted into changes in the utilization of synaptic efficacy (U_{SE}), a variable of the Tsodyks-Markram model of short-term plasticity describing the baseline probability of vesicle release (Tsodyks and Markram, 1997), and the peak AMPA receptor conductance ($\hat{g}_{AMPA}$; Equations (7) and (8) in Methods respectively). As a result of updating U_{SE}, short- and long-term plasticity are tightly coupled in the model (Markram and Tsodyks, 1996; Costa et al., 2015; Deperrois and Graupner, 2020). In the absence of $[Ca^{2+}]_i$ influx ρ exhibits bistable dynamics (Lisman, 1985; Graupner and Brunel, 2012) and at initialization, synapses are assumed to be at one of the two fixed points (fully depressed ones at $\rho = 0$ and fully potentiated ones at $\rho = 1$) and their assignment to these states is pathway-specific (Supplementary Figure S1C3).

We calibrated layer-wise spontaneous firing rates and evoked activity to brief VPM inputs matching *in vivo* data from Reyes-Puerta et al. (2015). Spontaneous activity was driven by somatic injection of a layer- and cell type-specific noisy conductance signal (Isbister et al., 2023, see Methods). Evoked activity was driven by a thalamocortical input stream already described in Ecker et al. (2024). In short, 10 VPM input patterns were repeatedly presented in random order with a 500 ms inter-stimulus interval, together with a non-specific POm input. These VPM patterns were defined with varying degrees of overlap in the sets of activated fibers (Figure 2A; see Methods). An exemplary raster plot, characterizing the evoked state of the plastic network is shown in Figure 2B.

2.1 Plasticity changes promote stable network activity

In vivo, plasticity does not seem to create macroscopic network instability. Therefore, it is a critical first step to validate our model in this regard. Previous theoretical work has shown that while embedding simple STDP rules in spiking networks without fast homeostatic plasticity leads to exploding firing rates (Zenke et al., 2017a; Morrison et al., 2007), this is not necessarily the case for calcium-based rules (Higgins et al., 2014; Graupner et al., 2016; Wang and Aljadeff, 2022). To assess the stability of our network while undergoing plasticity, we simulated 10 minutes of biological time of *in vivo*-like, stimulus-evoked network activity and verified that excitatory firing rates remained stable throughout the entire length of the simulation (Chindemi, 2018, Figure 2C).

When comparing the amount of changes in ρ across time steps, we found that most of the plastic changes happened in the first 1-2 minutes of the simulation, after which they stabilized (Figure 2D1). While small changes were still apparent towards the end of the simulation, by visualizing individual synaptic efficacy traces we confirmed that most of them oscillated around a dynamic fixed-point and the amount of changes in the second half of the simulation were negligible (Chindemi, 2018, Supplementary Figure S5). The results were similar (Figure 2E1, black line) when we averaged ρ values within a connection, i.e., all synapses between a pair of neurons (on average 4.1 ± 2.3 synapses per connection; Supplementary Figure S1C1). Considering changes at the connection level allowed us to compare our results against a traditional spike pair-based STDP rule (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000, see Methods) executed

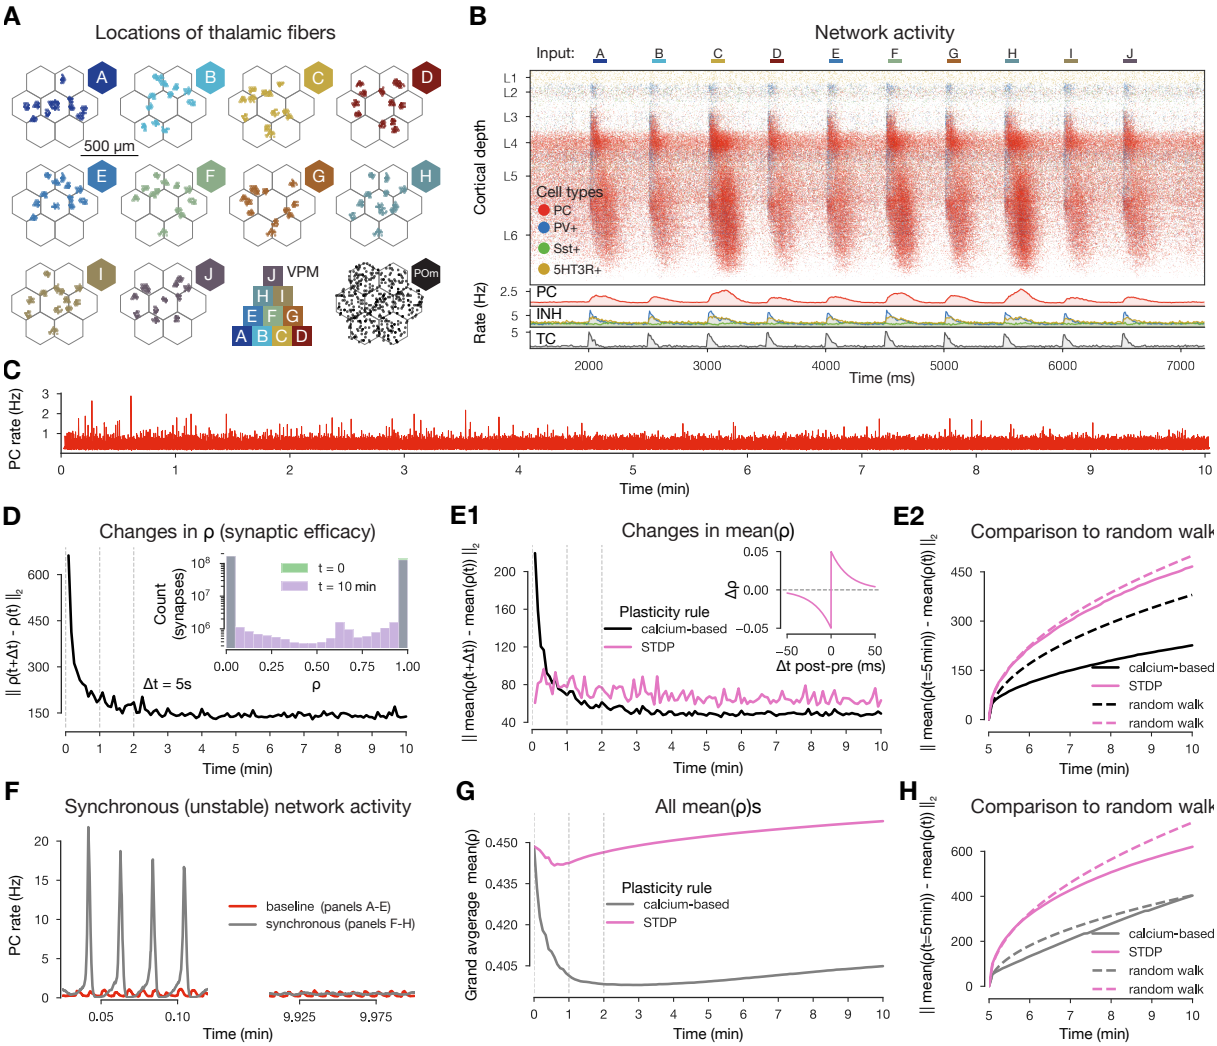

Figure 2: Changes in synaptic efficacy stabilize and they promote macroscale stable dynamics. A: Top-down view of the spatial arrangement of the VPM fiber centers associated with the 10 input patterns. Bottom row 3rd: pyramid-like overlap setup of VPM patterns, 4th (bottom right) POm fibers (high-order thalamic input) associated with all stimuli. **B:** Raster plot of the microcircuit's activity and the population firing rates below. The y-axis shows cortical depth. As cortical layers do not have the same cell density, the visually densest layer is not necessarily the most active. Panel adapted from Ecker et al. (2024). **C:** Firing rate of excitatory cells during the 10-minutes-long simulation. **D:** Plastic evolution of synaptic efficacy (ρ): L2 norm of changes in ρ across time. Inset shows the distribution of ρ values in the beginning (green; strickly 0s and 1s, see Methods) and end (purple) of the simulation. **E:** Evolution of mean ρ (aggregated over multiple synapses within a single connections). **E1:** L2 norm of changes in mean ρ (black) across time against STDP control in pink (inset, see Methods). **E2:** L2 norm of changes in ρ (similar to E1, but) compared to $\rho_{t=5minutes}$ in black, STDP in pink, and their random walk controls (see Methods) with the same colors but dashed lines. **F:** Firing rates of excitatory cells from a simulation with higher $[Ca^{2+}]_o$ and thus population bursts (see Supplementary Figure S6, Markram et al., 2015, and Methods) in gray and the baseline one in red as on panel C. **G:** Grand average mean ρ values across time for the synchronous simulation in gray and for the corresponding STDP control in pink. **H:** Same as E2 but for the synchronous simulation (calcium-based plasticity rule in gray instead of black).

on the 35 M excitatory spikes from our simulation. We observed that the STDP rule kept inducing the same magnitude of changes throughout all 10 minutes of the simulation instead of stabilizing after a transient phase (Figure 2E, pink line). To further assess whether the remaining changes in

mean ρ were in the form of fluctuations around a dynamic fixed-point, we compared the changes in the second half of the simulation to a random walk with the same step size (see Methods). The magnitude of changes were well below the random walk control for our simulation, while the changes from the STDP rule were above both, but below a random walk with a step size fitted to the changes induced by the STDP (Figure 2E2, see Methods). Consequently, the lifetime of synaptic efficacy configurations in the network are longer when using a calcium-based rule than an STPD rule or a random walk. This demonstrates that not only the scale of the simulated network, its biorealistic connection probabilities and the low, *in vivo*-like rates contribute to the stability of changes observed in our simulation, but also the use of a calcium-based plasticity model (Chindemi et al., 2022) in line with previous modeling insights (Higgins et al., 2014).

Next, we went beyond keeping a stable network stable and investigated if an unstable network exhibiting synchronous activity (Figure 2F, left) could be brought back to a stable, asynchronous regime by the calcium-based plasticity rule. The synchronized activity state was the result of simulating a higher $[Ca^{2+}]_o$ and thus higher release probability (U_{SE}, see Methods). Surprisingly, the firing rate flipped from synchronous to asynchronous (Figure 2F, right) and it did so within a minute of biological time (Supplementary Figure S6B). This was likely driven by a sudden decrease in the mean ρ values over synapses in the network, which decrease could not be observed for the STDP rule (Figure 2G). As before, both rules led to smaller changes than their corresponding random walk controls (Figure 2H). Taken together, our results indicate that the calcium-based rule evolves an over-excitable network towards lower excitability, which in this case is a role that a homeostatic rule would play (Turrigiano and Nelson, 2004; Zenke et al., 2017a).

In summary, the calcium-based plasticity rule promotes stability in our network. Indeed, when applied to asynchronous activity, it generates changes in synaptic efficacy that stabilize after a transient phase without leading to exploding firing rates. Furthermore, when applied to an unstable synchronous activity regime, the mean synaptic efficacy of the whole network consistently decreases making the network less excitable and fosters stability in the macroscopic activity.

2.2 Plasticity induces sparse and specific changes driven by the stimulus and the network's topology

We have shown that the changes in synaptic efficacy stabilize after a transient period, which raises the question: What predictions does the model make about the structure of the plastic changes *in vivo*? First of all, the changes were sparse, implying that the configuration of synaptic efficacy of the overall network remained largely unchanged. Indeed, we found that the number of connections changing was below 18% across the whole network and remained low when restricted to individual layers (Figure 3A1). Moreover, the propensity of changes increased as the pairwise firing rates increased (Figure 3A2), in line with previous modeling insights (Litwin-Kumar and Doiron, 2014; Graupner et al., 2016). On the level of individual synapses changes were even sparser (Figure 3B1), and in both cases depression was more common than potentiation. Layer 5 (L5) PCs contributed mostly to changes on the basal dendrites, while apical changes happened mostly on L6 PCs (Figure 3B2).

So far, we characterized plastic change in terms of ρ, as this parameter is bounded to the $[0, 1]$ interval and thus easy to interpret across pathways. Another parameter affecting the function of the synaptic connection is $\hat{g}_{AMPA}$. We found a minimal decrease (-2 pS, Figure 3C1) in its mean value due to plasticity, which was explained by frequent depression and a heavy tail of potentiation amplitudes (Figure 3C2). The distribution of $\hat{g}_{AMPA}$ remained lognormal, in line with biology (Buzsáki and Mizuseki, 2014; Rößler et al., 2023, Figure 3C1, bottom). The plasticity model employed was bistable around $\rho_* = 0.5$, i.e., in the absence of further Ca^{2+}

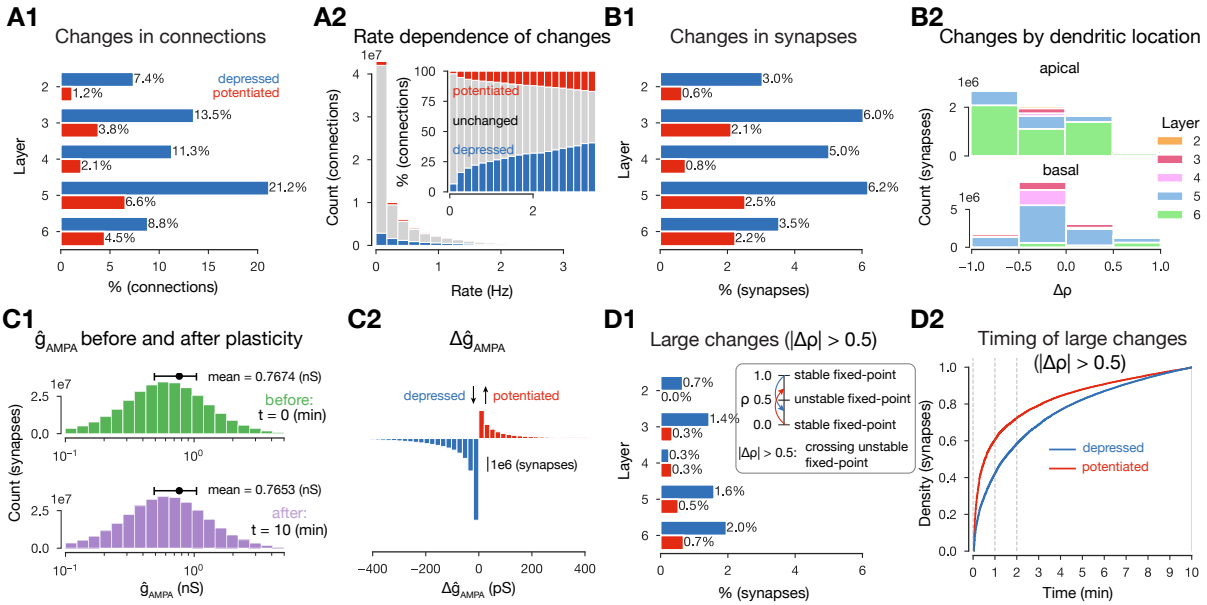

Figure 3: Changes in synaptic efficacy are sparse. A1: Propensity of changes in connections (mean ρ values) across layers. Connections are grouped into layers based on the soma location of the postsynaptic cell. Percentages are calculated within layers, i.e., 7.4% depression (in blue) for layer 2 (L2) means that among all connections that have the postsynaptic cell in L2, 7.4% depressed. **A2:** Plastic changes in mean ρ vs. mean pairwise firing rates of the pre- and postsynaptic cell of a given connection. **B1:** Same as A1 but for ρ (changes in synapses rather than connections). **B2:** Layer- and neurite type-wise (apical and basal dendrites) distribution of $\Delta\rho$. Only synapses that underwent any plastic change ($\sim$ 15M) are shown. **C1**: Distribution of $\hat{g}_{AMPA}$ in the beginning (in green) and end (in purple) of the simulation. **C2:** Plastic changes in $\hat{g}_{AMPA}$ leading to the slight shift in the distribution on C1. **D1:** Same as B1, i.e., propensity of changes at the synapse level, but for synapses that cross the $\rho_* = 0.5$ unstable fixed-point. **D2:** Cumulative histogram of ρ crossing the unstable fixed-point against time for depressing synapses in blue and potentiating ones in red.

influx, values above 0.5 would converge towards 1 and values below towards 0 (see Graupner and Brunel, 2012 and Methods). The fraction of changes in ρ that crossed this threshold was slightly higher for depression (26%) than for potentiation (23%) (compare Figure 3D1 to B1). Furthermore, the vast majority of these crossings occurred early, with 60% of potentiation and 40% of depression happening within the first minute (Figure 3D2).

Additionally, we hypothesized that both the amount, i.e., the level of sparsity, and the structure of plastic changes are non-random and instead shaped by the stimuli and the underlying network's topology. We confirmed our hypothesis first for the amount of plastic changes. We ran control simulations where we delivered random Poisson spikes on the same VPM fibers with the same mean rate, but without the spatio-temporal structure of the stimuli. We found that this reduced the number of connections undergoing plasticity by 25%, demonstrating the importance of stimulus structure over simple firing of pre- and postsynaptic neurons (Supplementary Figure S7). Additionally, without synaptic transmission, plasticity was reduced even further: In simulations of stimulus streams where intrinsic connectivity between the simulated neurons was cut we found occasional changes in ρ, but the number of changing synapses was an order of magnitude below baseline conditions (Supplementary Figure S7). Thus, our plasticity model is not strictly Hebbian, since the effect of postsynaptic firing alone could change synaptic efficacy, although presynaptic release was required for most of the observed changes (Graupner and Brunel, 2012; Graupner et al., 2016). When the external stimuli were also left out, negligible plastic changes occurred (Supplementary Figure S7).

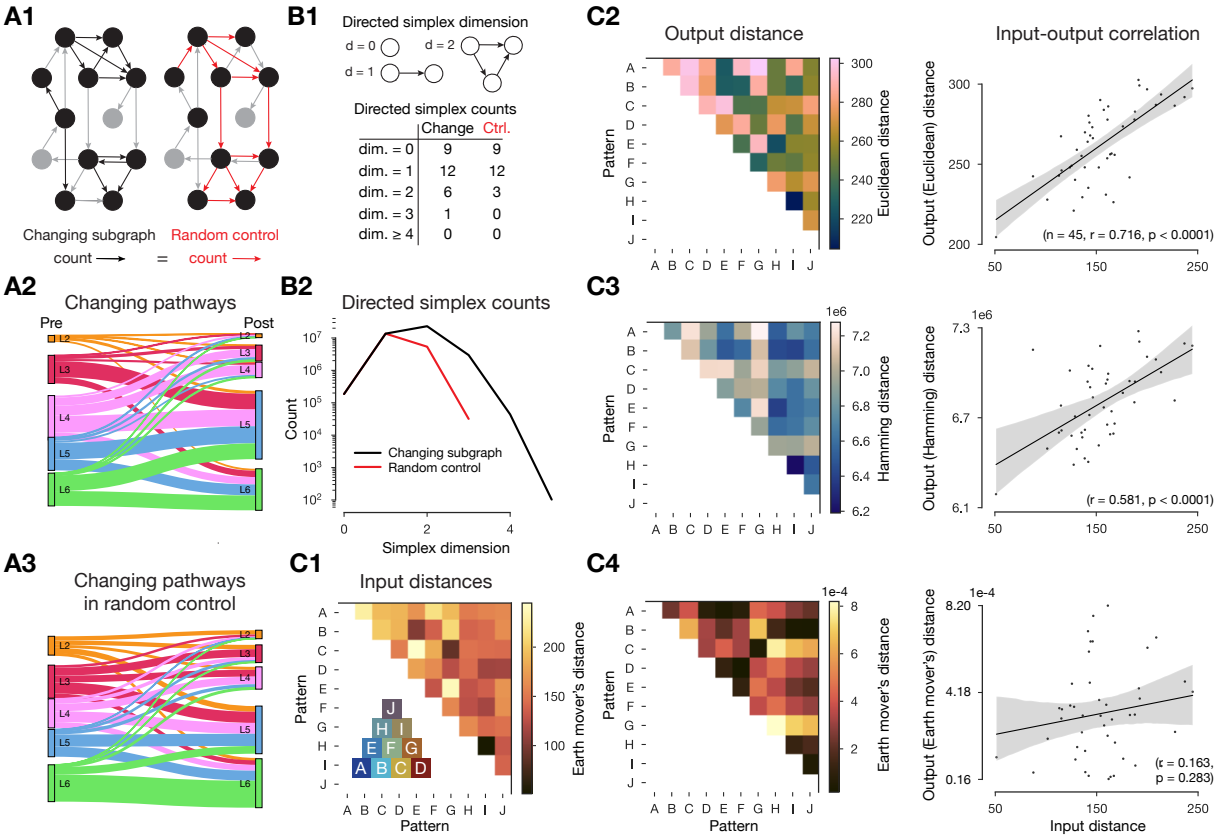

Figure 4: Changes in synaptic efficacy are non-random and shaped by the networks topology and the input stimuli. A: Layer-wise distribution of changing connections. **A1:** Left: schematic of the *changing subgraph*. In black, connections whose efficacies are changing and their pre- and postsynaptic populations. In grey, connections whose efficacies do not change and neurons not partaking in a changing connection. Right: random control of the changing subgraph, generated by randomly selecting the same number of connections as the changing connections (red edges) between the same pre- and postsynaptic populations (black nodes). **A2:** Sankey plot of layer-wise distribution of changing connections. Thickness of lines is proportional to the number of changing connections between pre- and post node populations. **A3:** As A2, but for the random control (A1 right). **B:** Directed simplex counts in the changing subgraph and its random control. **B1:** Schematic of directed k-simplices and their counts in the schematic graphs on A1. Note that by construction the control must have the same number of 0- and 1-simplices (see Methods) which correspond to the number of cells and connections in the subnetwork. **B2:** Directed simplex counts across dimensions in the changing subgraph (black) and its random control (red). **C**: Input - output distance correlations. **C1:** Input distances as the Earth mover's distance of the VPM fiber locations (see Figure 2A). Inset shows the overlap (based on Hamming distance, see Methods) of pattern fibers. **C2:** Output distances as Euclidean distance of mean ρ matrices in the 2-minute-long single pattern simulations. To its right: correlation of input and output distances. **C3**: Same as C2 but with Hamming distance. **C4:** Same as C2 but with Earth mover's distance (on the output side as well).

We then confirmed that the network's topology shaped plastic changes as well. In particular, we found layer-to-layer pathway specificity in the amount of plastic changes observed. That is, the amount of changes between layers differed from those expected in a random control with the same pre- and postsynaptic populations (Figure 4A2 vs. A3). We further quantified the difference in the subnetwork of changing connections from the random control by counting *directed simplices* (Fig 4B). These are motifs that have previously shown to be linked to network function (Reimann et al., 2017) as well as quantify the complexity of the network's topology (Kahle, 2009; Bobrowski and Kahle, 2018). Succinctly, a directed simplex of dimension k, is a

motif on $k+1$ neurons with all-to-all feed-forward connectivity (Figure 4A1 right, see Methods). We found strong overexpression of simplices in the subgraph of changing connections compared to the control (Figure 4B). Furthermore, the maximal simplex dimension found in the subgraph was two higher than expected by chance. This result shows that the connections changing are not determined only by their pre- and postsynaptic populations but also depend on their embedding in the whole network. We will explore this further in Section 2.5.

Finally, we showed that the input stimuli shaped plastic changes by running 2-minute-long plastic simulations in which we only presented a single pattern (several times, with the same 500 ms inter-stimulus interval as before) and compared the changes in mean ρ matrices. By using an input - output distance correlation analysis, we found that the distances between these mean ρ matrices were highly correlated with the input distances of the associated VPM fiber locations ($r = 0.716$, $p < 0.0001$, Figure 4C2). Furthermore, this required using Euclidean distance between the ρ matrices on the output side, implying that both the mean ρ values themselves and their position in the network are relevant. When either Hamming distance (taking only the identity of changing connection into account) or Earth mover's distance (taking only the distribution of mean ρ values into account) were used instead, we found weaker and non-significant correlations respectively (Figure 4C3 and C4 respectively).

In summary, we observed that $\sim 7\%$ of synapses undergo long-term plasticity under realistic *in vivo*-like conditions in 10 minutes of biological time, and most of these synapses are on L5 PC's basal dendrites. Moreover, the changes are not random, but depend not only on the firing rates but also on the recurrent connectivity and input stimuli. Potentiation dominated in amplitude, while depression counteracted it in frequency, keeping the network stable without needing to model homeostatic plasticity (Zenke et al., 2017a; Turrigiano and Nelson, 2004).

2.3 More frequent plastic changes within and across cell assemblies

We have shown that plastic changes are sparse and highly specific. We therefore tried to understand the rules determining which synapses changed. From the parametrization of the plasticity model, we learned that presynaptic spikes contribute orders of magnitude higher $[Ca^{2+}]_i$ than postsynaptic ones if the NMDA receptors are fully unblocked (Supplementary Figure S2B). Thus, in order to effectively depolarize the dendrites and unblock NMDA receptors, spikes at low, *in vivo*-like rates must be synchronized in time. Therefore, we specifically hypothesized that plasticity of connections may be structured by the membership of participating neurons in Hebbian cell assemblies (Hebb, 1949; Harris, 2005). In our previous analysis (Ecker et al., 2024), and in line with experimental results (Harris, 2005; Song et al., 2005; Perin et al., 2011), we observed that the number of afferents from an assembly, is a great predictor of a neuron's membership in an assembly. Strong positive interactions were also found *across assemblies*, but only when the direction of innervation reflected the temporal order of assembly activation. These results, combined with the biophysics of the plasticity model, suggest that connections within an assembly and the ones between temporarily ordered assemblies, are expected to undergo plastic changes with a higher probability.

To test this hypothesis, we detected cell assemblies, based on their co-firing function, from the *in silico* spiking activity of the 10-minute-long plastic simulation using methods established by experimentalists (Carrillo-Reid et al., 2015; Herzog et al., 2021, see Methods). In short, spikes were binned and bins with significantly high firing rates were hierarchically clustered (Figure 5A). These clusters correspond to the functional assemblies of neurons (Figure 5C), with a neuron being considered a member if its spiking activity correlates with the activity of an assembly significantly stronger than chance level. Since time bins and not neurons, were clustered in the first place, this method yields one assembly per time bin and single neurons can be part of several assemblies (Figure 5B, D). Note, that in contrast to modeling studies, where

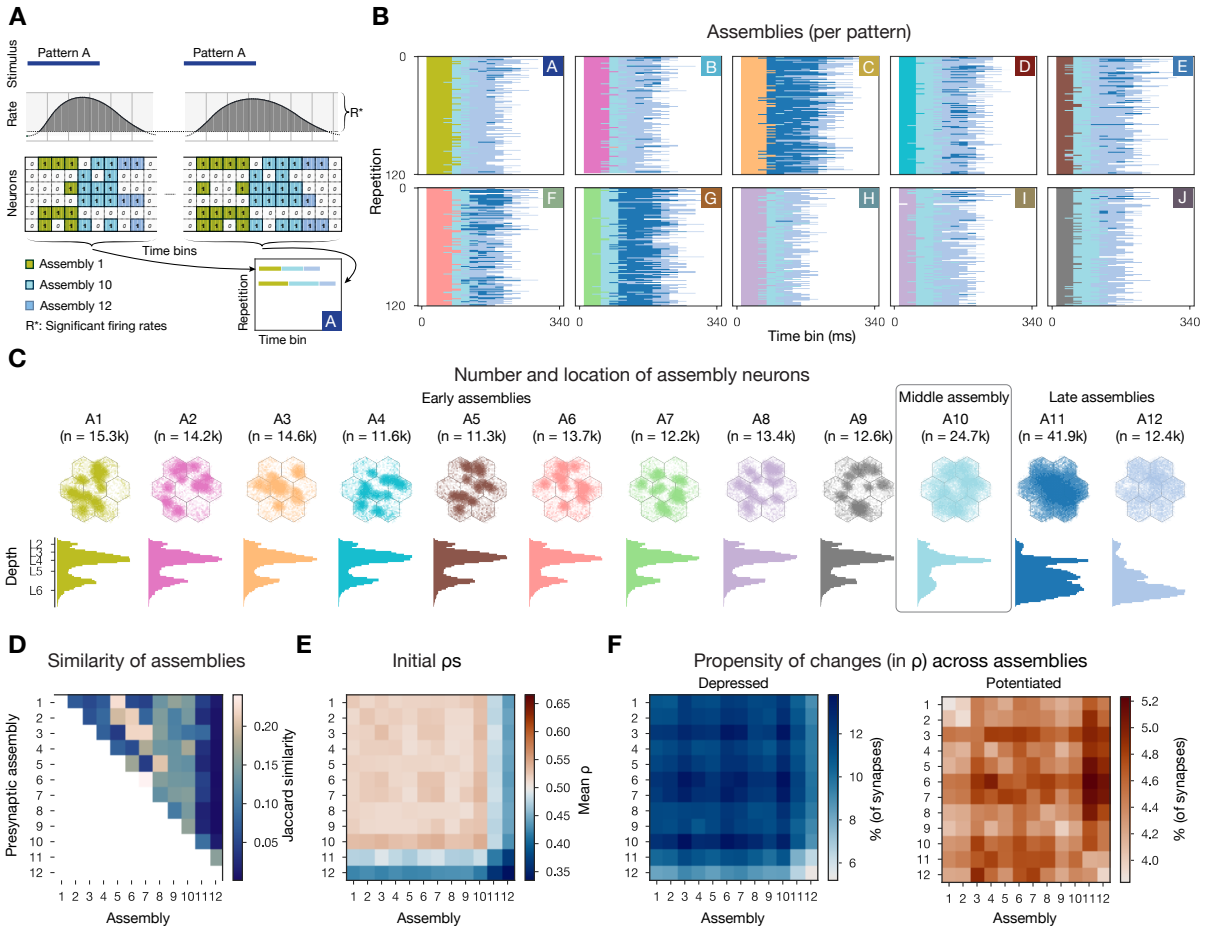

Figure 5: Changes are more frequent within and across cell assemblies. **A:** Schematics of the assembly detection pipeline. Significant time bins are clustered by cosine similarity and these clusters define functional assemblies: olive, light blue, steel blue (see Methods). Bottom right, visual summary of the organization of pattern-specific assembly sequences of the two repetitions of pattern A. **B:** Activation of functional assembly sequences for multiple repetitions of the patterns presented (as in A bottom right). Each row within the 10 matrices corresponds to a single repetition of a given pattern. White: non-significant time bins. **C**: Number and location of neurons in each cell assembly. Constituent neurons are those for which the correlation of their spiking activity within the corresponding time bins is significantly stronger than chance (see Methods). Top: top-down view, bottom: depth-profile. **D:** Jaccard similarity (intersection over union) between cell assemblies. **E:** Initial mean ρ of within- and cross-assembly synapses. **F:** Propensity of depression and potentiation of within- and cross-assembly synapses in blue on the left and in red on the right respectively. Since assemblies are overlapping (see D) some synapses are counted in multiple pre- and postsynaptic assembly pairings.

assemblies are usually defined based on their strong internal connectivity (i.e., their structure Litwin-Kumar and Doiron, 2014; Zenke et al., 2015; Fauth and Van Rossum, 2019; Kossio et al., 2021) our functional assemblies are detected from the activity of the network only, eliminating a potential bias arising from links between structural properties of changing synapses and the assembly detection method.

Assemblies were activated in all stimulus repetitions and a series of three to four assemblies remained active for 190 ± 30 ms (Figure 5B). Pattern C elicited the strongest response, while pattern H the weakest, and the responses of patterns H and I were the most similar to each other, as expected, since they share 66% of the VPM fibers (Figure 2A). Assembly activations had a well-preserved temporal order - with some of them always appearing early during a stimulus,

while others later - and from now on we will refer to them as *early*, *middle*, and *late assemblies*, and will order them in the figures accordingly.

When checking the propensity of changes within and across assemblies, we indeed found more synapses undergoing long-term plasticity (Figure 5F). While only 5% of synapses depressed in the whole dataset, we found up to 13.8% when restricting the analysis to assemblies. Similarly, compared to 2% of all synapses potentiating, we observed up to 5.2% when restricting to assemblies. Interestingly, large values were found in the off-diagonal entries (Figure 5F), i.e., synapses across assemblies underwent more plastic changes than the synapses within these assemblies. In our model, the initial ρ values are pathway-specific and highest in L4 pathways (Brémaud et al., 2007, Supplementary Figure S1C3). Therefore, early assemblies, with large numbers of L4 cells have a higher than average initial ρ (Figure 5C and E respectively), thus their synapses are more likely to depress (Figure 5F left). As early assemblies are stimulus specific, and thus not part of the same assembly sequences, synaptic depression between these cells can be seen as some kind of stimulus separation. On the other hand, late assemblies, that are predominantly composed of cells from the deep layers, have a low initial ρ (Figure 5E; Supplementary Figure S1C3) and synapses towards them are more likely to potentiate (Figure 5F right). These assemblies are mostly non-specific and participate in all assembly-sequences, thus the potentiation of their efferents means a general strengthening of the stimulus response as a whole.

Together these results indicate that, in line with 70-years-old predictions, cells that fire together wire together (Hebb, 1949). Our contribution lies in making the qualitative statement above into a quantitative one: Under *in vivo*-like conditions, cells that fire together more than expected have three times higher chances of changing the efficacy of their connections.

2.4 Synapse clustering contributes to the emergence of cell assemblies, and facilitates plasticity across them

In addition to co-firing, a group of presynaptic neurons is more effective in depolarizing a given dendritic branch if they all innervate the same branch, i.e., they form a spatial synapse cluster (Kastellakis and Poirazi, 2019; Farinella et al., 2014; Iacaruso et al., 2017; Tazerart et al., 2020). To quantify this, we selected the 10 most innervated L5 TTPCs (thick-tufted pyramidal cells) within a cell assembly and then detected spatial clusters of synapses, defined as at least 10 synapses within a 20 μm stretch of a single dendritic branch (Methods). Next, we grouped all synapses on these 10 selected neurons per assembly into four categories based on assembly membership of the presynaptic neuron and whether the synapse was part of a cluster or not (see exemplary clustered assembly synapses in Figure 6A). While some of the ρ values of clustered synapses within an assembly, underwent small constant changes, many of them changed at the same time (vertical stripes on Figure 6B); indicating the importance of spatial co-localization for plasticity.

To systematically quantify this effect, we calculated the likelihood of plastic change in each category: (not) in an assembly and (non-)clustered, by contrasting the conditional probability of observing it in a given category with the probability of observing any change irrespective of the category (see Equation (15) in Methods; Figure 6C). Surprisingly, clustered synapses within an assembly were only likely to undergo plastic changes in late assemblies, but not in early and middle assemblies (Figure 6C black arrow). Furthermore, when comparing the amplitude of changes across conditions with a 2-way ANOVA, we found that clustered within-assembly synapses depress to a smaller degree than the other ones (Supplementary Figure S8A). To explain these trends, we repeated the same analysis on the probability of the initial value of ρ to be either 0 or 1 in each category. Initial ρ values in early and middle assembly synapses, especially the clustered ones, are very likely to be initialized as fully potentiated while synapses within the late assemblies were likely to be initialized in the fully depressed state (Figure 6D black arrow). Thus

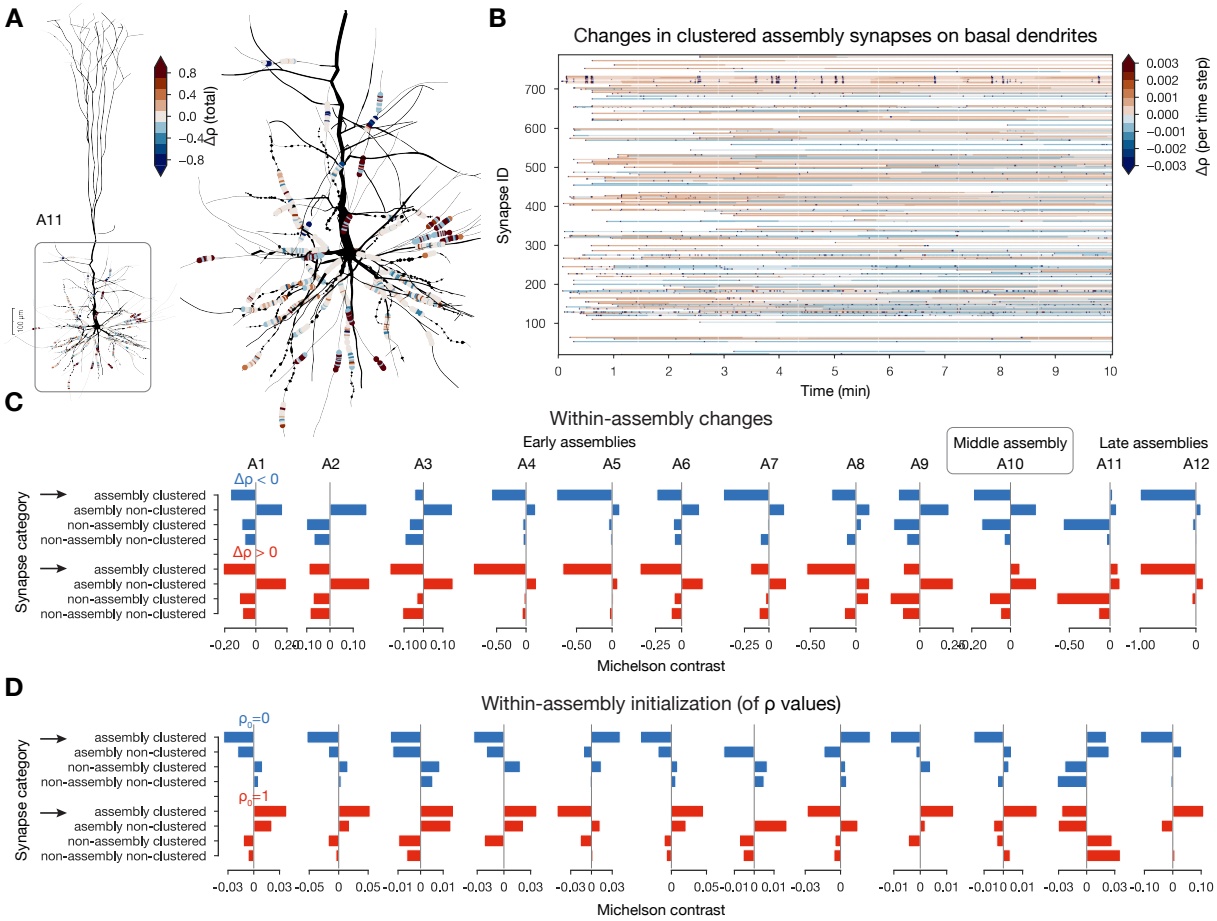

Figure 6: Changes are structured by synaptic clusters. A: Changes in clustered assembly synapse on an exemplary neuron from assembly 11 (A11). Neurite diameters are scaled (x2) for better resolution. (Synapse diameters are arbitrary.) **B:** Temporal evolution of the synapses on basal dendrites shown on A. Synapses are grouped by their dendritic sections (single stretches of non-branching dendrites) and ordered by their mean distance to the soma (closest on the bottom and furthest on top). **C:** Michelson contrast (Equation (15) in Methods) of the likelihood of plastic changes of synapses grouped into four categories: (non)-assembly and (non-)clustered) for each assembly. Negative values indicate that synapses belonging to the given category are unlikely to change. Depression on top (in blue) and potentiation below (in red). Black arrows indicate clustered assembly synapses. **D:** Same as C but for inital ρ values (see Methods).

the picture emerging is as follows: early and middle assemblies are partially defined by clustered (both spatial and functional) synapses that are initialized as fully potentiated. These synapses are unlikely to change, but when they do, they depress less than the others, and would converge back to $\rho = 1$ in absence of activity, as they do not cross the $\rho_* = 0.5$ unstable fixed-point. These stable early assemblies can therefore function as a stable backbone amid ongoing synaptic plasticity.

In our previous investigation, we found that most changes happened across assemblies (Figure 5F), so we extended the analysis described above to cross-assembly synapses (Supplementary Figure S8B). Here, the picture was reversed: cross-assembly synapses that were part of a spatial cluster were likely to be initiated as fully depressed and then had a high chance of undergoing potentiation (Supplementary Figure S8C black arrow). Together with the previous results, this suggests that synapses between assemblies are more likely to change, which is even more pronounced if these synapses form a cluster on the postsynaptic dendrite.

2.5 Network-level metrics predict plasticity and separate potentiation from depression

So far, we have found co-firing and synapse clustering to be good predictors of plastic changes in line with previous work (Zenke et al., 2015; Kastellakis and Poirazi, 2019; Harris, 2005). Additionally, we have shown that changes are also driven by the underlying network's connectivity (Figure 4A and B). We now elaborate on this effect using *edge participation*, a novel network metric of edge centrality that considers how the connections are embedded in the entire network. More precisely, the k-edge participation of a connection is the number of k-simplices that the connection is part of (Figure 7A1-Full, see Methods). Across dimensions, as edge participation increased so did the probability of undergoing plastic change (Figure 7A2). This effect was stronger for potentiation than depression (Figure 7A3). Thus, we hypothesized edge participation to be a good predictor of synaptic strength. We verified this both in our model as well as in the MICrONS (2021) mm^3 dataset, an electron microscopic reconstruction of cortical tissue that combines synaptic resolution with the scale needed to calculate meaningful edge participation. Both networks had similar maximal dimension of simplices as well as ranges of edge participation values (Supplementary Figure S9A), making the results of this analysis comparable. For our model we used the sum of $\hat{g}_{AMPA}$ values as a proxy for connection strength. For MICrONS, we used total synapse volume of a connection (see Methods), as a measure of connection strength (Harris and Stevens, 1989). Indeed, high edge participation led to higher connection strength in MICrONS (2021) as predicted by our modeling insights; moreover, the curves across dimensions are qualitatively the same for both datasets (Figure 7B).

To gain insights into pattern-specificity, we analyzed the 2-minutes-long single pattern simulations and considered instead edge participation within the assembly-specific subgraphs (Figure 7A1-Sub). When doing so, we found an even higher predictive power, which was additionally pattern specific (Figure 7C). More precisely, we found that the probability of changes only increased with edge participation in the assembly sub-graph associated with the presented pattern or to a late assembly and that it was the strongest for the former (see Figure 7C left for pattern A for an example and right for a summary of the maxima attained for all patterns).

Finally, *pattern-indegree*, i.e., the number of VPM fibers belonging to a pattern that innervate a neuron is also a good predictor of the probability of a connection to change i.e., an increase of either the pre- or postsynaptic pattern-indegree of a connection leads to more frequent plastic changes (Supplementary Figure S9B1). However, given the highly non-random structure of the full network, the values of pre- and postsynaptic pattern-indegree are not independent and thus their effect in driving plastic changes is hard to disentangle (Supplementary Figure S9B2, C).

In summary, edge participation in the full network allowed us to delineate potentiation from depression given change, as well as predict connection strength, which we confirmed in the MICrONS (2021) dataset. Moreover, edge-participation within assembly subgraphs is a pattern specific predictor of plastic changes.

2.6 Increased stimulus-specificity and more reliable assembly sequences characterize the network after plasticity

As the changes are sparse, one might wonder if they cause any functional change. To study this, we compared the network's activity *before* and *after* plasticity (see Methods). Firing rates decreased slightly as a result of plasticity (Figure 8A left), while spike correlations remained stable, in line with recent findings (Oby et al., 2019; Feulner et al., 2022, Figure 8A right). Both of these effects were more pronounced in the deeper layers (Supplementary Figure S10A). Correlations of connected pairs of neurons decreased slightly after plasticity following the decrease of firing rate (Supplementary Figure S10A). However, the highest correlations both before and

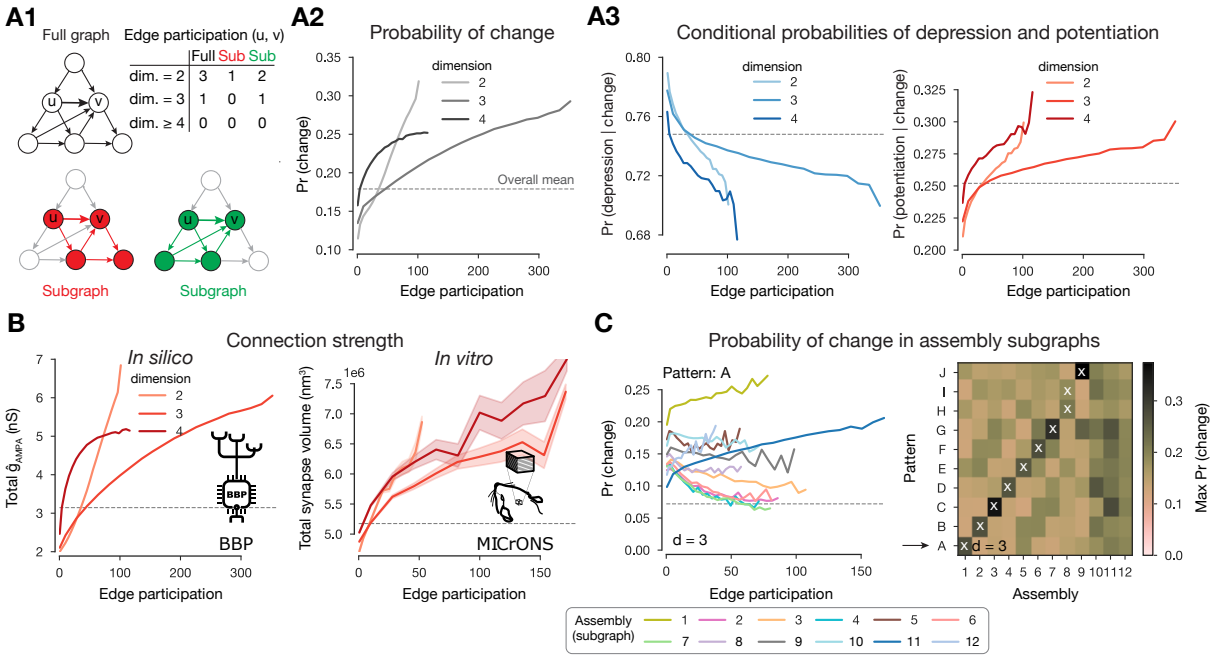

Figure 7: Changes are stronger in central connections in the network. A1: Schematics of edge participation in high-dimensional simplices in an exemplary *"full"* network on the top in black and in (assembly) *"sub"*-graphs on the bottom. Table summarizes edge participation values for the connection from node u to v across dimensions for all three cases. **A2:** Probability of changes (i.e., either depression or potentiation) vs. edge participation in high-dimensional simplices (see A1 and Methods). **A3:** Probability of depression (left) and potentiation (right) conditioned on plastic change (in any direction) vs. edge participation. **B:** Comparison with electron microscopy data. Total AMPA conductance ($\hat{g}_{AMPA}$) against edge participation across dimensions in our nbS1 model on the left. Total synapse volume (see Methods) against edge participation across dimensions in the MICrONS (2021) dataset on the right. Shaded areas indicate SEM. **C:** Probability of changes against edge participation in assembly subgraphs (see A1 bottom). Left, probability of changes in (2-minutes-long) pattern A simulations against edge participation in assembly subgraphs. Right, summary of the maximum probability values across patterns on the right. Arrow indicates the row shown in detail on its left and white crosses indicate the pattern-specific early assemblies (see Figure 5B).

after plasticity were found in the most central edges in the network i.e, those with highest edge participation in high dimensional simplices (Figure 8B1). We further verified this prediction in the MICrONS dataset (Figure 8B2, see Methods). It has been previously shown before that assemblies have many edges that are central in the network, therefore, we hypothesized that they remain functional and reliably active (Fauth and Van Rossum, 2019; Kossio et al., 2021; Pérez-Ortega et al., 2021). We showed this on the level of assembly sequences: plastic changes promote stimulus specific and reliable responses. While we found considerable overlap between assemblies detected before and after plasticity, an early assembly responding to patterns B and E before plasticity, split into two stimulus-specific assemblies (Figure 8C). Note that this is not an artifact of the thresholding applied during assembly detection as forcing the pipeline to yield the same number of assemblies instead splits a middle assembly first (Supplementary Figure S10E). Moreover, the assembly sequences responding to repetitions of a given pattern generally became more reliable after plasticity (Supplementary Figure S10B). To quantify this, we measured the Hamming similarity of assembly sequences between repetitions and found significant increases for most patterns (see Methods, Figure 8D). Interestingly, this increase on the assembly level lied in contrast with a slight decrease of the reliability of the neuronal activity on the single cell level (Figure 8E).

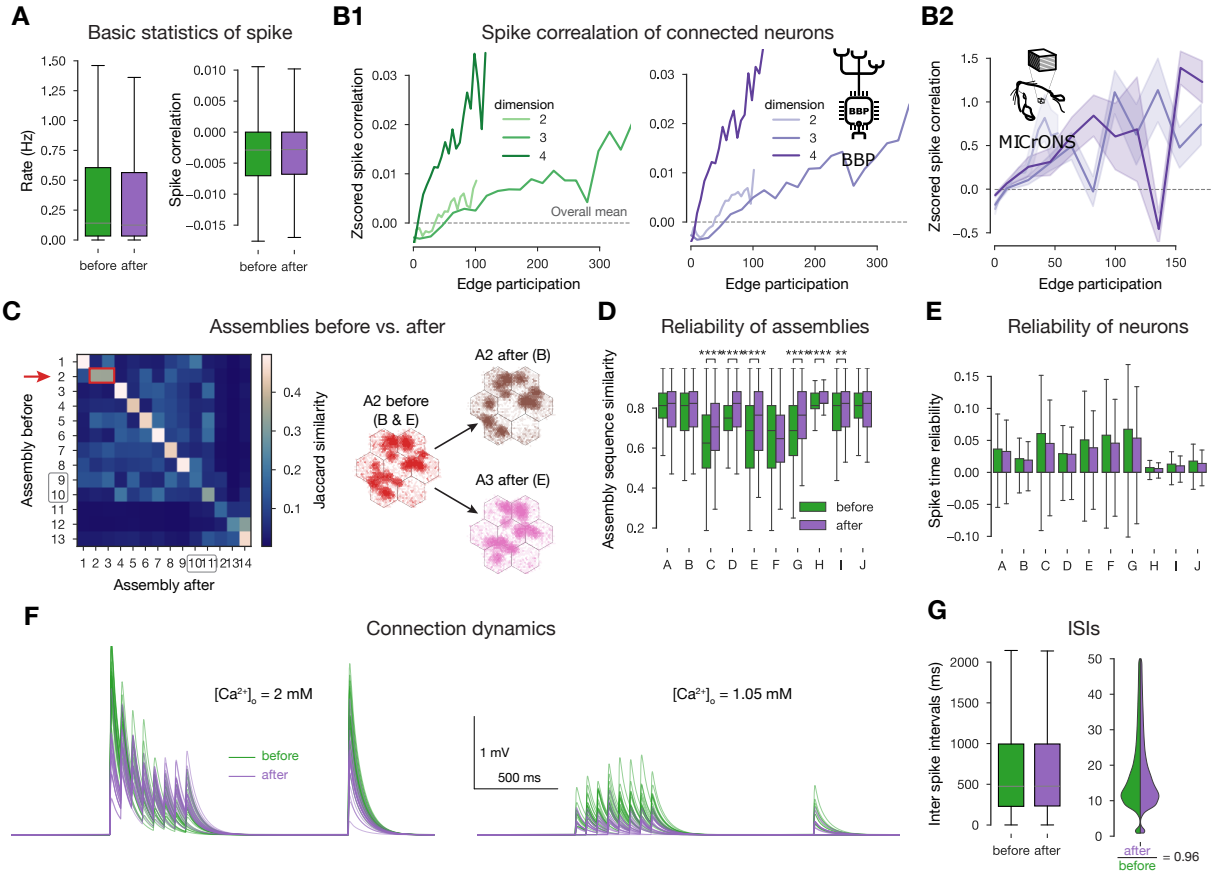

Figure 8: Changes promote stimulus specificity. **A**: Firing rates and pairwise spike correlations extracted from non-plastic simulations *before* plasticity, i.e., in the naive circuit vs. *after* the 10-minute-long plastic simulation. **B1**: Spike correlation of connected cells vs. edge participation in high dimensional simplices (see Methods) before and after plasticity on the left and right respectively. **B2**: Spike correlation of connected cells vs. edge participation in the MICrONS (2021) dataset (see Methods). Shaded areas indicate SEM. **C**: Jaccard similarity of assemblies detected before vs. after plasticity on the left. Shared early assembly of patterns B and E splitting into two stimulus-specific ones after plasticity on the right. (Indicated by red arrow and rectangle on its left). See more detailed plots in Supplementary Figure S10B. **D**: Reliability of assembly sequences (see Methods and Supplementary Figure S10B). Significance of increases was determined with Kruskal-Wallis test: *: $p \leq 0.05$, **: $p \leq 0.01$, ***: $p \leq 0.001$, ****: $p \leq 0.0001$. **E**: Spike time reliability (see Methods) of single cells to the different patterns before and after plasticity. (The lack of stars means no significant increases after plasticity.) **F**: Short-term dynamics of an exemplary potentiated L5 TTPC connection before (in green) and after plasticity (in purple). At *in vitro* $[Ca^{2+}]_o$ on the left, and *in vivo* on the right. Thin lines represent 20 individual trials, while the thicker ones their means. **G**: Interspike interval (ISI) distribution of all excitatory neurons before and after plasticity on the left. Zoom in on low ISIs (≤ 50 ms) on the right.

The plastic changes we observed at the network-level are in contrast with *in vitro* observations at the connection level. In particular, with the redistribution of synaptic efficacy towards earlier spikes during high-frequency firing caused by changes in U_{SE} (Chindemi et al., 2022; Markram and Tsodyks, 1996; Costa et al., 2015; Selig et al., 1999; Sjöström et al., 2003), which under our *in vivo*-like, low firing rates and $[Ca^{2+}]_o$ is not in the ranges where this connection-level redistribution is relevant. Specifically, at *in vitro* levels of $[Ca^{2+}]_o$ the potentiated U_{SE} shifts the connection further into the depressing regime, causing the redistribution (Figure 8F left). *In vivo* the lower $[Ca^{2+}]_o$ counters this by moving the connection into the pseudo-linear regime; additionally, the low firing rates make this phenomenon less relevant in general (Figure 8F

right and G respectively). Thus, while Markram and Tsodyks (1996) showed a redistribution of synaptic efficacy after plasticity at the single connection level *in vitro* (Figure 8F), we found a redistribution at the network level under *in vivo*-like conditions (Figure 8A-D, Supplementary Figure S10A).

In summary, we observed a network-level redistribution of synaptic efficacy that maintained spike correlations globally, while enhancing stimulus-specificity and the reliability of the subplopulations in terms of assembly activation. Furthermore, we derived and confirmed in the MICrONS (2021) dataset a prediction stating that the most central connections of the network drive the correlated activity.

3 Discussion

In this work, we aimed to understand how network structure, function, and dendritic processing interact to drive synaptic plasticity at the population level in the neocortex. Specifically, we used a plasticity model, whose parameters were determined by pairwise STDP experiments but modeled at the level of dendrites. We studied its effect at the population level under *in vivo*-like conditions to determine network-level rules that best predict plastic changes by simulating it in a biophysically detailed model of a cortical microcircuit that was subjected to 10 different stimuli under a variety of protocols. Our principal observations are as follows: (1) Plastic changes were sparse, affecting only 7% of the synapses. A balance between occasional, large-amplitude potentiation and more frequent depression stabilized the network without explicitly modeling homeostatic plasticity. Moreover, the changes were non-random and stimulus specific. (2) Plastic changes were largely determined by the anatomical structure of synaptic connectivity and its relation to functional units, i.e., changes were most likely between co-firing cell assemblies, at clustered synapses, between neurons highly innervated by the same input, and in central edges in the network. (3) While assemblies remained fairly stable, they become more stimulus specific and reliable which has important implications for coding and learning.

The first observation (1) is quite significant considering that we did not design the learning rule to be sparse and stable. In previous models of plastic networks, changes were not sparse and additional mechanisms were needed to keep the network's activity stable (Litwin-Kumar and Doiron, 2014; Delattre et al., 2015; Zenke et al., 2015; Fauth and Van Rossum, 2019; Kossio et al., 2021; Zenke et al., 2017a; Turrigiano and Nelson, 2004). The machine learning community is also aware of the importance of sparse changes, as in continual learning one has to balance plasticity and stability of existing weights to avoid catastrophic forgetting (McCloskey and Cohen, 1989; Ratcliff, 1990). In recent years, they have come up with impressive techniques that mask or slow down weight updates (e.g., by a regularizer on the loss) to improve the performance of deep neural networks (Zenke et al., 2017b; Kirkpatrick et al., 2017; Mallya and Lazebnik, 2018; Frankle and Carbin, 2019), whereas in our model it emerged without any mechanisms explicitly enforcing it. Finally, demonstrating the stimulus-specificity of changes was a crucial validation that our results are not just a description of a biophysical process, i.e., plasticity, but have implications for learning in general.

For our second observation (2), we described three different kinds of plasticity rules that emerge from the underlying plasticity model based on: activity, subcellular structure, and network structure. While all of these are powerful predictors of plastic changes, none of them fully determines them. It is rather the interplay between them (and potentially additional unknown rules) that brings about non-homogeneous changes across the full microcircuit. The first two can be explained from the biophysics of the plasticity model and links our results to the classical work of Hebb (1949) as well as the recent literature on synapse clustering (Kastellakis and Poirazi, 2019; Farinella et al., 2014; Iacaruso et al., 2017; Tazerart et al., 2020). With respect to synapse

clustering, we would highlight that our synapses are stochastic and the release probability between PCs is ~ 0.04 at the simulated low $[Ca^{2+}]_o = 1.05\,\text{mM}$ (Markram et al., 2015; Ecker et al., 2020; Jones and Keep, 1988; Borst, 2010). Therefore, care should be taken when comparing our results with either glutamate uncaging experiments, which bypass the presynaptic machinery (Pettit et al., 1997; Losonczy and Magee, 2006), or with other modeling studies that use deterministic synapses (Farinella et al., 2014; Poirazi et al., 2003; Ujfalussy and Makara, 2020). With respect to network-based rules, previous simulation approaches have characterized the learned connectomes as having an overexpression of reciprocal connectivity and densely connected neuron motifs (Brunel, 2016; Zhang et al., 2019). We expanded on this body of previous work, by introducing cell and connection-specific metrics that directly yield the probability of observing depression or potentiation of any given connection. Most of our predictors behaved similarly with respect to depression and potentiation. Only a metric based on the embedding of individual edges in the global (or assembly-specific) network was able to separate them. This edge-based prediction was confirmed by analyzing an electron microscopic reconstruction of cortical tissue in a comparable way (MICrONS, 2021).

Our third observation (3) that assembly sequence become more reliable and stimulus specific may sound counter-intuitive considering that we observed reduced firing rates and less reliable single neurons. This shows that plasticity mechanism can contribute to the emergence of a robust population code with unreliable neurons. Crucially, this means that plasticity can make circuit responses more distinct (between stimuli by increasing specificity) but also less distinct (for a given stimulus by increasing reliability). This has important computational implications that may be explored in a future study.

To facilitate further research, we are open-sourcing our model alongside detailed instructions to launch simulations and analyze the results (see Data and code availability). Simulating the model requires a performant hardware and software infrastructure (e.g., we needed 12 M core hours to run the simulation presented in this manuscript). With respect to the latter we are continuously improving the performance and efficiency of the simulator (Kumbhar et al., 2019). The model has several assumptions and limitations, which will ideally be improved upon iteratively and in a community-driven manner. First, for our simulations we used the model of Chindemi et al. (2022) with parameters obtained in a previous version of the circuit. Due to various improvements of synapse parameters (e.g. U_{SE} of L5 PCs) since fitting plasticity-specific parameters, the results for pairing protocols at various frequencies shift slightly while maintaining the presence of STDP (Supplementary Figure S2A). Second, the extracellular magnesium concentration of 1 mM used in *in vitro* preparations was assumed to be representative of the *in vivo* level. Third, experimental evidence indicates that firing bursts affects plasticity (Letzkus et al., 2006; Williams and Stuart, 1999; Larkum, 2013). However, L5 TTPC bursts were rare in our simulations, as the model is based on an early developmental stage (P14-16: juvenile rats) and firing bursts only becomes prominent as the animals mature (Zhu, 2000). Furthermore, bursts can be triggered by synaptic input onto apical dendrites (Larkum, 2013; Hay et al., 2011). Inputs to apical dendrites from high-order thalamus have been shown to gate plasticity in L23 PCs *in vivo*, via dis-inhibiting the distal dendrites (Gambino et al., 2014; Williams and Holtmaat, 2019). While POm input was present in our model and targeted apical dendrites, we activated them in random and non-specific ways. On the other hand, top-down inputs represent context or brain state and are thought to serve as an error or target signal for learning and thus are likely very specific (Makino, 2019). Feedback signals from high-order cortical areas also innervate apical dendrites (Harris et al., 2019), but as our model is of a single cortical area only, these inputs were absent. Missing excitatory input from other areas were compensated by somatic conductance injections (Isbister et al., 2023), which could be replaced in the future by simulating additional dendritic synapses instead. Finally, learning in the brain is orchestrated by several mechanisms

in unison and some of these are beyond the scope of this paper, e.g., metabotropic glutamate receptors, neuromodulation, slow homeostatic plasticity, structural plasticity, and calcium induced calcium release (Magee and Grienberger, 2020).

Bottom-up modeling is often criticized for being arbitrarily complex. In our simulations, the mean value of synaptic efficacies of an over excited network reduced, comparable to Turrigiano and Nelson (2004) without introducing rate homeostasis explicitly. In the machine learning field, mechanisms are added to classical algorithms to make updates sparser, overcome catastrophic forgetting, and allow the learning of tasks sequentially (Kirkpatrick et al., 2017; Zenke et al., 2017b). The calcium-based plasticity rule employed here results in sparse updates without additional mechanisms, unlike STDP rules, which update efficacies at each spike arrival. Thus, we showed that adding biological complexity allowed us to reduce the number of mechanisms explicitly modeled. Calcium-based plasticity models existed for a long time (Shouval et al., 2002; Rubin et al., 2005), have been shown to be stable (Higgins et al., 2014; Graupner et al., 2016; Wang and Aljadeff, 2022), have been modified to mimic *in vivo*-like low $[Ca^{2+}]_o$ (Graupner et al., 2016; Chindemi et al., 2022), and versions more complex than ours exist (Mäki-Marttunen et al., 2020; Rodrigues et al., 2023). Our contribution lies in integrating all these different aspects in network simulations at scale with subcellular resolution. Using such simulations allowed us to explore not only the temporal aspects of plasticity but also their link to spatial and network aspects, such as clusters of synapses and cell assemblies. Furthermore, the analysis methods developed on large scale functional and structural data empowered us to make and test novel predictions in the MICrONS (2021) dataset, which while pushing the boundaries of big data neuroscience, was so far only analyzed with single cells in focus instead of the network as a whole (Ding et al., 2023; Wang et al., 2023).

4 Methods

4.1 Calcium-based plasticity model

The calcium-based plasticity model is fully described in Chindemi et al. (2022), but a minimal description of it can be found below. Synaptic efficacy (ρ) is based on the Graupner and Brunel (2012) formalism, which exhibits a bistable dynamics ($\rho = 0$ fully depressed, $\rho = 1$ fully potentiated, and $\rho_* = 0.5$ unstable fixed-point) described as:

$$\tau \frac{d\rho}{dt} = -\rho(1-\rho)(\rho_* - \rho) + \gamma_p(1-\rho)\Theta\big(Ca^* - \theta_p\big) - \gamma_d \rho \Theta\big(Ca^* - \theta_d\big) \tag{1}$$

where $\tau = 70\,\mathrm{s}$ is the time constant of convergence, θ_d and θ_p are depression and potentiation thresholds, $\gamma_d = 101.5$ and $\gamma_p = 216.2$ are depression and potentiation rates and Θ is the Heaviside function. Ca^* is linked to the dynamics of $[Ca^{2+}]_i$ in spines (see below), which was modeled as:

$$\frac{d[Ca^{2+}]_i}{dt} = \big(I^*_{NMDA} + I_{VDCC}\big)\frac{\eta}{2FX} - \frac{[Ca^{2+}]_i - [Ca^{2+}]_i^{(0)}}{\tau_{Ca}} \tag{2}$$

where I^*_{NMDAR} and I_{VDCC} are calcium currents through NMDA receptors and VDCCs, $\eta = 0.04$ is the fraction of unbuffered calcium, F is the Faraday constant, X is the spine volume, $[Ca^{2+}]_i^{(0)} = 70\,\mathrm{pM}$ is the resting value of $[Ca^{2+}]_i$, and $\tau_{Ca} = 12\,\mathrm{ms}$ is the time constant of free (unbuffered) calcium clearance. I^*_{NMDA} depends on the state of the Mg^{2+} block as follows:

$$I_{NMDA}(t) = g_{NMDA} m (V - E_{NMDA}) \tag{3}$$

where g_{NMDA} and $E_{NMDA} = -3\,\mathrm{mV}$ are the conductance and the reversal potential of the NMDA receptor, V is the local voltage, and m describes the nonlinear voltage dependence due

to the Mg^{2+} block following the Jahr and Stevens (1990) formalism:

$$m = \frac{1}{1 + [Mg^{2+}]_o / \theta e^{-\kappa V}} \tag{4}$$

where θ is a scaling factor of the extracellular magnesium concentration $[Mg^{2+}]_o$, and κ is the slope of the voltage dependence. Parameters $\theta = 2.552$ and $\kappa = 0.072$ were obtained by refitting the model to cortical recordings from Vargas-Caballero and Robinson (2003) (as opposed to the original parameters fit to hippocampal ones Jahr and Stevens, 1990). Spines were assumed to be separate biochemical compartments, i.e., $[Ca^{2+}]_i$ of the dendrites does not influence that of the synapses.

Inspired by previous theoretical insights (Rubin et al., 2005), a leaky integrator of $[Ca^{2+}]_i$ was introduced (Ca^*) to slow down its time course instead of modeling enzymes downstream of calcium (e.g. CamKII as others did (Mäki-Marttunen et al., 2020; Rodrigues et al., 2023)):

$$\frac{dCa^*}{dt} = -\frac{Ca^*}{\tau^*} + [Ca^{2+}]_i - [Ca^{2+}]_i^{(0)} \tag{5}$$

where $\tau^* = 278.318\,\text{ms}$ is the time constant of the integrator. Updates in ρ were linked to this Ca^* variable crossing θ_d and/or θ_p (see equation (1)). The two synapse-specific threshold were derived based on peaks in $[Ca^{2+}]_i$ caused by pre- and postsynaptic spikes, c_{pre} and c_{post} respectively. To measure these parameters for all 312,709,576 synapses, simulations of single cells were run, in which either the pre- or the postsynaptic cell was made to fire a single action potential and the local $[Ca^{2+}]_i$ was monitored in each synapse. Presynaptically evoked $[Ca^{2+}]_i$ peaks were three orders of magnitude larger, than the ones evoked by postsynaptic spikes (Supplementary Figure S2B). Postsynaptically evoked $[Ca^{2+}]_i$ peaks had a multimodal distribution in the apical dendrites (Supplementary Figure S2B right), in line with Landau et al. (2022). Since 8% of L6 PCs could not be made to fire a single action potential (only bursts), synapses on those cells (10,995,513 in total) were assumed to be non-plastic, i.e., their thresholds were set to a negative value that could not be crossed. Similarly, as the plasticity of connections between L4 spiny stellate cells was shown to be non-NMDA dependent (Chindemi et al., 2022; Egger et al., 1999) those connections were made non-plastic. For the remainder of cells, θ_d and θ_p were derived as follows:

$$\begin{bmatrix} \theta_d \\ \theta_p \end{bmatrix} = \begin{bmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{bmatrix} \times \begin{bmatrix} c_{pre} \\ c_{post} \end{bmatrix} \text{ or } \begin{bmatrix} b_{00} & b_{01} \\ b_{10} & b_{11} \end{bmatrix} \times \begin{bmatrix} c_{pre} \\ c_{post} \end{bmatrix} \tag{6}$$

where $a_{i,j}$ and $b_{i,j}$ are constants optimized during model fitting for apical and basal dendrites respectively. Changes in ρ were then converted by low-pass filtering into changes U_{SE} and $\hat{g}_{AMPA}$ as follows:

$$\frac{dU_{SE}}{dt} = \frac{\overline{U}_{SE} - U_{SE}}{\tau_{change}} \quad \text{where} \quad \overline{U}_{SE} = U_{SE}^{(d)} + \rho\left(U_{SE}^{(p)} - U_{SE}^{(d)}\right) \tag{7}$$

$$\frac{d\hat{g}_{AMPA}}{dt} = \frac{\overline{g}_{AMPA} - \hat{g}_{AMPA}}{\tau_{change}} \quad \text{where} \quad \overline{g}_{AMPA} = \hat{g}_{AMPA}^{(d)} + \rho\left(\hat{g}_{AMPA}^{(p)} - \hat{g}_{AMPA}^{(d)}\right) \tag{8}$$

where $U_{SE}^{(d)}$, $U_{SE}^{(p)}$, $\hat{g}_{AMPA}^{(d)}$, and $\hat{g}_{AMPA}^{(p)}$ are the fully depressed (d) and fully potentiated (p) values of the given variables in-between which they evolve and $\tau_{change} = 100\,\text{s}$ is the time constant of the conversion of changes in ρ into changes in U_{SE} and $\hat{g}_{AMPA}$.

In the nbS1 model U_{SE} is also modulated by $[Ca^{2+}]_o$, where a reduction in $[Ca^{2+}]_o$ leads to a pathway-specific, non-linear reduction in U_{SE} (Markram et al., 2015; Ecker et al., 2020).

When extracting the network's state after plasticity, not only the U_{SE} and $\hat{g}_{AMPA}$ values, but also the peak NMDA conductances ($\hat{g}_{NMDA}$) were updated according the ρ values in the last time step of the simulation.

Model parameter optimization (Chindemi et al., 2022) and the derivation of thresholds from c_{pre} and c_{post} measurements was done with `plastyfire`.

4.2 *In vivo*-like spontaneous and evoked activity

The calibration process that leads to the *in vivo*-like spontaneous activity is fully described in Isbister et al. (2023), but a minimal description and a list of the parameters used in this article can be found below. As extracellular recordings are known to overestimate firing rates (Wohrer et al., 2013), a spectrum of spontaneous states at fixed percentage of the rates reported in Reyes-Puerta et al. (2015) were calibrated (Isbister et al., 2023). Matching specific firing rates *in silico* was achieved by iterative adjustments of layer and cell-type (excitatory/inhibitory) specific somatic conductance injection (following an Ornstein-Uhlenbeck process (Destexhe et al., 2001)). As in Isbister et al. (2023), extracellular recordings were assumed to have the same bias across layers and neuron populations. Furthermore it was assumed, that different inhibitory subpopulations require the same amount of input compensation. By introducing plasticity at all E to E synapses, an additional depolarizing current from VDCCs was added to the model, which made the network more active than its non-plastic counterpart (Supplementary Figure S3A). This required an algorithmic lowering of the amplitude of injected conductances from Isbister et al. (2023) to achieve the same *in vivo*-like layer-wise spontaneous firing rates (Supplementary Figure S3B). The spontaneous state used in the article is characterized by the parameters: $[Ca^{2+}]_o = 1.05\,\mathrm{mM}$ (Jones and Keep, 1988), percentage of reported firing rates = 40%, the coefficient of variation (CV; std/mean) of the noise process = 0.4.

The thalamic input patterns, and the spike trains delivered on them are fully described in Ecker et al. (2024), but a minimal description, highlighting the changes applied in this study, can be found below. First, the flatmap location (Bolaños-Puchet et al., 2024) of VPM fibers avoiding the boundaries of the network were clustered with k-means to form 100 bundles of fibers. Second, the four base patterns (A, B, C, and D) were formed by randomly selecting four non-overlapping groups of bundles, each containing 12% of them. The remaining six patterns were derived from these base patterns with various degrees of overlap: three patterns as combinations of two of the base ones (E, F, G), two patterns as combinations of three of the base ones (H, I), and one pattern as a combination of all four base ones (J). Third, the input stream was defined as a random presentation of these 10 patterns, in a balanced way. Last, for each pattern presentation, unique spike times were generated for its corresponding fibers following a 100 ms-long inhomogeneous adapting Markov process (Muller et al., 2007). The maximal rate of the VPM fibers was set to 17.5 Hz (compared to 30 Hz for the non-plastic circuit in Ecker et al., 2024) and half of that for POm. The overlap of the patterns is clearly visible in the firing pattern of each group of fibers corresponding to them (Supplementary Figure S4).

4.3 Network simulations

Simulations were run using the `NEURON` simulator as a core engine with the Blue Brain Project's collection of hoc and NMODL templates for parallel execution on supercomputers (Kumbhar et al., 2019; Hines and Carnevale, 1997; Awile et al., 2022) (see Data and code availability). Simulating 10 minutes of biological time with reporting the state of all synapses (in every second) took 2,350,000 core hours ($\sim 4x$ more than the corresponding non-plastic circuit without reporting), on our HPE based supercomputer, installed at CSCS, Lugano. Simulations were always repeated at least three times to assess the consistency of the results.

4.4 Control STDP rule

To compare the amount of changes induced by the calcium-based model of Chindemi et al. (2022) with classical plasticity rules, the 35,264,818 excitatory spikes from the 10-minute-long simulation were evaluated with a pair-based STDP rule (Gerstner et al., 1996; Kempter et al.,

1999; Song et al., 2000). Synaptic weights evolved as follows under the STDP rule:

$$\Delta w_+ = A_+ \exp\Big(-\frac{\Delta t}{\tau_+}\Big) \text{ at } t_{post} \text{ if } t_{pre} < t_{post} \tag{9}$$

$$\Delta w_- = A_- \exp\Big(\frac{\Delta t}{\tau_-}\Big) \text{ at } t_{pre} \text{ if } t_{pre} > t_{post} \tag{10}$$

where t_{pre} and t_{post} are the times of pre- and postsynaptic spikes, $\Delta t = t_{post} - t_{pre}$ is the difference between them; $A_\pm = 0.05$ describe the weight update, which decayed exponentially with time constants $\tau_\pm = 20\,\text{ms}$. The same parameters were used for the comparison with the 10-minute-long synchronous (unstable) simulation (76,025,009 excitatory spikes). The STDP rule was implemented in `Brian2` (Stimberg et al., 2019).

4.5 Random walk control

Changes in ρ induced by the calcium-based model of Chindemi et al. (2022) and the STDP rule described above, were compared to a random control that takes into account global and noise-induced random changes in the network. To that end, a control model was used, where K connections of the network either increase or decrease by a noise value ϵ with equal probability in each time step. Note that, in this control, the size of the change vector at each time step is given by:

$$\|\rho(t + \Delta t) - \rho(t)\|_2 = \epsilon\sqrt{K}. \tag{11}$$

This process was modeled with a *random walk* in K-dimensional space of step size l, which at step N was given by:

$$\vec{R} = \sum_{i=0}^{N} \vec{r_i}, \tag{12}$$

where each $\vec{r_i}$ is a random variable in $\mathbb{R}^K$, where K is the number of changing connections of the network and such that the step size is given by $|\vec{r_i}| = \epsilon\sqrt{K} =: l$. The changes induced by the random walk at step N in the network is given by its distance from the origin at that time step, which for large enough N and K is approximately:

$$\|\vec{R}\| \approx l\sqrt{N}. \tag{13}$$

The step size l was determined as the average step size in the data after the 5 minute mark, when the transient activity has long passed. That is, l is given by:

$$\underset{t>5\text{minutes}}{\text{mean}} \{\|\rho(t + \Delta t) - \rho(t)\|_2\}. \tag{14}$$

4.6 Cell assembly detection

The combination of methods from Carrillo-Reid et al. (2015) and Herzog et al. (2021) yielding the assembly detection pipeline is fully described in Ecker et al. (2024), but a minimal description, highlighting the changes applied in this study, can be found below. First, spikes of excitatory cells were binned using 20 ms time bins (Harris et al., 2003). Second, time bins with significantly high firing rates were determined as crossing a threshold defined as the mean activity level plus the 95th percentile of the standard deviation of 100 shuffled controls. These shuffled controls were less strict than in Ecker et al. (2024). Unlike in the original study, where spikes were only shifted by one time bin forward or backward (Carrillo-Reid et al., 2015), spikes were shifted by any amount. This change was introduced because the network's response to the same patterns was more variable in the plastic simulations, and to not miss any of them, a lower threshold was more

fitting. Third, based on the cosine similarity of activation vectors, i.e., vectors of spike counts of all neurons in the given significant time bins, a similarity matrix was built (Carrillo-Reid et al., 2015). Fourth, this similarity matrix was hierarchically clustered using Ward's linkage (Pérez-Ortega et al., 2021; Montijn et al., 2016). Like for any other unsupervised clustering method, the number of optimal clusters cannot be known beforehand, thus potential number of clusters were scanned between five and twenty. The number of clusters with the lowest Davis-Bouldin index was chosen, which maximizes the similarity within elements of the cluster while minimizing the between cluster similarity (Davies and Bouldin, 1979). Fifth, neurons were associated to these clusters based on their spiking activity, and it was determined whether they formed a cell assembly or not. The correlations between the spike trains of all neurons and the activation sequences of all clusters were computed and the ones with significant correlation selected to be part of the potential assemblies. Significance was determined based on exceeding the 95th percentile of correlations of shuffled controls (1000 controls with spikes of individual cells shifted by any amount as above; Herzog et al., 2021; Montijn et al., 2016). Finally, it was required that the mean pairwise correlation of the spikes of the neurons with significant correlations was higher than the mean pairwise correlation of neurons in the whole dataset (Herzog et al., 2021). Clusters passing this last criterion were considered to be functional assemblies and the neurons with significant correlations their constituent cells.

The reliability of assembly sequences was defined as the Hamming similarity over the repetitions of a single pattern, which attains a value of 1 if two assembly sequences are identical and 0 if they are completely different. Assemblies of neurons were compared using their Jaccard similarity.

Assemblies were detected using the `assemblyfire` package.

4.7 Synapse clusters and likelihood of plastic changes within them

To quantify the importance of co-localization of synapses for plastic changes, clusters of synapses were detected. To be part of a synapse cluster, a synapse was required to have at least nine other synapses on the same dendritic branch, i.e., between two branching points of the dendrite, with $\leq 10\,\mu$m (Euclidean) distance. Significance of spatial clustering was determined similarly to Druckmann et al. (2014). To that end, the distribution of synapse neighbor distances of the 10 selected synapses were compared with a Poisson model, assuming exponentially distributed inter-synapse distances, based on all (same branch) synapse neighbor distances on the given neuron. Clusters were merged in a subsequent step, thus synapse clusters with more than 10 synapses, spanning more than $20\,\mu$ms were also feasible.

Plastic changes in synapse clusters were only analyzed for a small subpopulation of neurons in an assembly (10 L5 TTPCs per assembly), which were selected based on maximizing two connectivity features: assembly-indegree and synaptic clustering coefficient (SCC). The first one is the number of connections from a presynaptic assembly, while the second one quantifies the co-localization of synapses on the dendrites of a neuron from its presynaptic assembly. Maximizing these two metrics, which were introduced (and are described in detail) in Ecker et al. (2024), allows one to select subpopulations with high probability of finding synapse clusters.

Control synapse clusters, originating from non-assembly neurons were also detected on the same postsynaptic neurons.

The normalized likelihood of changes, conditioned on the four *categories* a synapse could fall into (assembly clustered, assembly non-clustered, non-assembly cluster, non-assembly non-clustered) were quantified using the Michaelson contrast, defined as:

$$\frac{Pr(changed \,|\, category) - Pr(changed)}{Pr(changed \,|\, category) + Pr(changed)} \tag{15}$$

where *changed* was split to be either potentiated ($\Delta\rho > 0$) or depressed ($\Delta\rho < 0$). Note that a nonzero value for one category always has to be compensated by a nonzero value with the opposite sign in another. For the normalized likelihood of initial ρ values, the same equation was used but $Pr(\Delta\rho > 0|\, category)$ was replaced with $Pr(\rho_0 = 1|\, category)$ for potentiated and $Pr(\Delta\rho < 0|\, category)$ with $Pr(\rho_0 = 0|\, category)$ for depressed synapses respectively.

4.8 Topological metrics

For quantifying the non-random nature of changing connections, directed simplex counts were used. A k-simplex in a network G is a set of on $k+1$ nodes (neurons) of G that are all-to-all connected in a feedforward fashion. That is, there is a numbering of the nodes $v_0, v_1, \ldots v_k$, such that for every $i < j$ there is an edge from v_i to v_j, and k is called the dimension of the simplex. In particular, 0-simplices are the nodes of the network, 1-simplices directed edges (connections) and 2-simplices are also known as transitive triads. Random controls were defined as the same number of edges between the same nodes, resulting in the same 0- and 1-simplex counts. Given an edge in G, a notion of its edge-centrality in the network is its *k-edge participation*, which is the number of k-simplices the edge belongs to. This extends the classic notion of node participation from nodes to edges. This value can be computed either for the simplices in the entire network or in an specificied subnetwork e.g., the subnetwork on neurons belonging to an assembly.

Simplex counts and edge participation values were computed using the `connalysis` package (Egas Santander et al., 2024), based on a fast C++ implementation from `flagsercount` (Lütgehetmann et al., 2020).

4.9 MICrONS dataset

For the comparison of two of our findings with a rodent electron microscopy dataset with co-registered calcium-imaging traces, the *minnie65_public* release of the MICrONS (2021) dataset was used. Synapses from sources other than one of the 60,048 classified neurons inside the bigger (*minnie65*) volume were discarded. The 117 version of the structural dataset was used for these analysis and made freely available in SONATA format (Dai et al., 2020) at https://doi.org/10.5281/zenodo.8364070. For comparable analysis we restricted the volume to its central part ($650,000 \leq x \leq 950,000$ and $700,000 \leq z \leq 1,000,000$) and considered only E to E connections. Synapse volume was defined as the number of voxels painted by the automatic cleft segmentation (from the correspondings `synapses_pni_2` table) times the 4x4x40 nm voxel size. Total synapse volume was defined as the sum across all synapses mediating a connection.

For the functional dataset, eight calcium-imaging sessions from the 661 version were used. These sessions were selected by the same criteria as in Reimann et al. (2023); Egas Santander et al. (2024), namely at least 1000 neurons were scanned in each session and at least 85% of them were co-registered in the structural data. Neurons with non-unique identifiers were dropped, and the remaining ones restricted to the central part of the structural volume described above. Deconvolved spike trains of neurons (243 neurons/session on average, 1817 unique neurons in total, 2% and 15% of all cells respectively) were used to calculate Pearson correlations. To average spike correlations from different session without any bias, they were first z-scored within each session.

4.10 Spike time reliability

Spike time reliability was defined as the mean of the cosine similarities of a given neuron's mean centered, smoothed spike times across all pairs of repetitions of a single pattern (Schreiber et al., 2003; Cutts and Eglen, 2014). To smooth the spike times, they were first binned to 1 ms time

bins, and then convolved with a Gaussian kernel with a standard deviation of 10 ms as in Ecker et al. (2024); Egas Santander et al. (2024).

4.11 Visualization

Rendering of cells from the cortical network on Figure 1A was done with `Brayns`, while the rendering of selected cells with their thalamo-cortical synapses on D with `BioExplorer`. Morphologies on Figures 1 and 6 were rendered with `NeuroMorphoVis` (Abdellah et al., 2018).

Acknowledgements

The authors thank Nicolas Ninin for his involvement in the early stage of this project, Elvis Boci and Cyrille Favreau for their help with visualizations, Michael Gevaert, Joni Herttuainen and Thomas Delemontex for their assistance with software engineering, and Alberto Antonietti, Christoph Pokorny, Kathryn B. Hess, Ran Levi, Wulfram Gerstner, Roberto Araya and Henry Markram for discussions.

Funding

This study was supported by funding to the Blue Brain Project, a research center of the École polytechnique fédérale de Lausanne (EPFL), from the Swiss government's ETH Board of the Swiss Federal Institutes of Technology.

Author contributions

- Conceptualization: A.E., D.E.S., G.C., E.B.M., M.W.R.
- Formal analysis: A.E., D.E.S., M.W.R.
- Investigation: A.E., S.B.-P., G.C.
- Methodology: A.E., D.E.S., G.C., J.B.I., E.B.M., M.W.R.
- Project administration: M.W.R.
- Software: A.E., D.E.S, S.B.-P., G.C., J.K., P.K., M.W.R.
- Resources: J.B.A, I.M.
- Supervision: M.W.R.
- Validation: A.E., D.E.S., D.P.G.M, S.B.-P., J.B.I., M.W.R.
- Visualization: A.E., M.A.
- Writing - original draft: A.E., D.E.S., M.W.R.
- Writing - review & editing: A.E., D.E.S., S.B.-P., G.C., J.B.I., E.B.M, M.W.R.

Declaration of interest

The authors declare no competing interests.

Data and code availability

The $2.4\,mm^3$ subvolume of the juvenile rat somatosensory cortex, containing 211,712 neurons and 312,709,576 plastic synapses in SONATA format (Dai et al., 2020) is freely available at: https://doi.org/10.5281/zenodo.8158471. It can be loaded and instantiated in CoreNEURON (Kumbhar et al., 2019) with neurodamus. The circuit and the simulations can be analyzed using Blue Brain SNAP and ConnectomeUtilities built on top of it. Cell assemblies were detected and can be analyzed with assemblyfire. Exemplary Jupyter notebooks using the packages above were deposited in the same repository on Zenodo.

References

Abdellah, M., Hernando, J., Eilemann, S., Lapere, S., Antille, N., Markram, H., and Schürmann, F. (2018). NeuroMorphoVis: A collaborative framework for analysis and visualization of neuronal morphology skeletons reconstructed from microscopy stacks. *Bioinformatics*, 34(13):i574–i582.

Awile, O., Kumbhar, P., Cornu, N., Dura-Bernal, S., King, J. G., Lupton, O., Magkanaris, I., McDougal, R. A., Newton, A. J., Pereira, F., Săvulescu, A., Carnevale, N. T., Lytton, W. W., Hines, M. L., and Schürmann, F. (2022). Modernizing the NEURON Simulator for Sustainability, Portability, and Performance. *Frontiers in Neuroinformatics*, 16(884046).

Berger, T., Larkum, M. E., and Lüscher, H. R. (2001). High Ih channel density in the distal apical dendrite of layer V pyramidal cells increases bidirectional attenuation of EPSPs. *Journal of Neurophysiology*, 85(2):855–868.

Bi, G.-q. and Poo, M.-m. (1998). Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *The Journal of Neuroscience*, 18(24):10464–10472.

Bliss, T. V. P. and Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, 232(2):331–356.

Bobrowski, O. and Kahle, M. (2018). Topology of random geometric complexes: a survey. *Journal of Applied and Computational Topology*, 1:331–364.

Bolaños-Puchet, S., Teska, A., Hernando, J. B., Lu, H., Romani, A., Schürmann, F., and Reimann, M. W. (2024). Enhancement of brain atlases with laminar coordinate systems: Flatmaps and barrel column annotations. *bioRxiv*.

Bono, J., Wilmes, K. A., and Clopath, C. (2017). Modelling plasticity in dendrites: from single cells to networks. *Current Opinion in Neurobiology*, 46:136–141.

Borst, J. G. G. (2010). The low synaptic release probability in vivo. *Trends in Neurosciences*, 33(6):259–266.

Brémaud, A., West, D. C., and Thomson, A. M. (2007). Binomial parameters differ across neocortical layers and with different classes of connections in adult rat and cat neocortex. *PNAS*, 104(35):14134–14139.

Brunel, N. (2016). Is cortical connectivity optimized for storing information? *Nature Neuroscience*, 19(5):749–755.

Buzsáki, G. and Mizuseki, K. (2014). The log-dynamic brain: How skewed distributions affect network operations. *Nature Reviews Neuroscience*, 15(4):264–278.

Carrillo-Reid, L., Miller, J. E. K., Hamm, J. P., Jackson, J., and Yuste, R. (2015). Endogenous sequential cortical activity evoked by visual stimuli. *Journal of Neuroscience*, 35(23):8813–8828.

Chindemi, G. (2018). *Towards a unified understanding of synaptic plasticity: parsimonious modeling and simulation of the glutamatergic synapse life-cycle*. PhD thesis, EPFL.

Chindemi, G., Abdellah, M., Amsalem, O., Benavides-Piccione, R., Delattre, V., Doron, M., Ecker, A., Jaquier, A. T., King, J., Kumbhar, P., Monney, C., Perin, R., Rössert, C., Tuncel, M. A., van Geit, W., DeFelipe, J., Graupner, M., Segev, I., Markram, H., and Muller, E. B. (2022). A calcium-based plasticity model predicts long-term potentiation and depression in the neocortex. *Nature Communications*, 13(3038).

Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: A model of voltage-based STDP with homeostasis. *Nature Neuroscience*, 13(3):344–352.

Costa, R. P., Froemke, R. C., Sjöström, P. J., and van Rossum, M. C. (2015). Unified pre- and postsynaptic long-term plasticity enables reliable and flexible learning. *eLife*, 4(e09457).

Cutts, C. S. and Eglen, X. S. J. (2014). Detecting pairwise correlations in spike trains: An objective comparison of methods and application to the study of retinal waves. *Journal of Neuroscience*, 34(43):14288–14303.

Dai, K., Hernando, J., Billeh, Y. N., Gratiy, S. L., Planas, J., Davison, A. P., Dura-Bernal, S., Gleeson, P., Devresse, A., Dichter, B. K., Gevaert, M., King, J. G., van Geit, W. A., Povolotsky, A. V., Muller, E., Courcol, J. D., and Arkhipov, A. (2020). The SONATA data format for efficient description of large-scale network models. *PLoS Computational Biology*, 16(2):e1007696.

Davies, D. L. and Bouldin, D. W. (1979). A Cluster Separation Measure. *IEEE Transactions on Pattern Analysis and Machine Learning*, PAMI-1(2):224–227.

Delattre, V., Keller, D., Perich, M., Markram, H., and Muller, E. B. (2015). Network-timing-dependent plasticity. *Frontiers in Cellular Neuroscience*, 9(220).

Deperrois, N. and Graupner, M. (2020). Short-term depression and long-term plasticity together tune sensitive range of synaptic plasticity. *PLoS Computational Biology*, 16(9):e100826.

Destexhe, A., Rudolph, M., Fellous, J. M., and Sejnowski, T. J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1):13–24.

Ding, Z., Fahey, P. G., Papadopoulos, S., et al. (2023). Functional connectomics reveals general wiring rule in mouse visual cortex. *bioRxiv*.

Druckmann, S., Feng, L., Lee, B., Yook, C., Zhao, T., Magee, J. C., and Kim, J. (2014). Structured Synaptic Connectivity between Hippocampal Regions. *Neuron*, 81(3):629–640.

Ecker, A., Romani, A., Sáray, S., Káli, S., Migliore, M., Falck, J., Lange, S., Mercer, A., Thomson, A. M., Muller, E., Reimann, M. W., and Ramaswamy, S. (2020). Data-driven integration of hippocampal CA1 synaptic physiology in silico. *Hippocampus*, 30(11):1129–1145.

Ecker, A., Santander, D. E., Bolaños-Puchet, S., Isbister, J. B., and Reimann, M. W. (2024). Cortical cell assemblies and their underlying connectivity: an in silico study. *PLoS Computational Biology*, 20(3):e1011891.

Egas Santander, D., Pokorny, C., Ecker, A., Lazovskis, J., Santoro, M., Smith, J. P., Hess, K., Levi, R., and Reimann, M. W. (2024). Efficiency and reliability in biological neural network architectures. *bioRxiv*.

Egger, V., Feldmeyer, D., and Sakmann, B. (1999). Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nature Neuroscience*, 2(12):1098–1105.

Farinella, M., Ruedt, D. T., Gleeson, P., Lanore, F., and Silver, R. A. (2014). Glutamate-Bound NMDARs Arising from In Vivo-like Network Activity Extend Spatio-temporal Integration in a L5 Cortical Pyramidal Cell Model. *PLoS Computational Biology*, 10(4).

Fauth, M. J. and Van Rossum, M. C. (2019). Self-organized reactivation maintains and reinforces memories despite synaptic turnover. *eLife*, 8:e43717.

Feulner, B., Perich, M. G., Chowdhury, R. H., Miller, L. E., Gallego, J. A., and Clopath, C. (2022). Small, correlated changes in synaptic connectivity may facilitate rapid motor learning. *Nature Communications*, 13(5163).

Frankle, J. and Carbin, M. (2019). The lottery ticket hypothesis: Finding sparse, trainable neural networks. *ICLR*.

Froemke, R. C., Letzkus, J. J., Kampa, B. M., Hang, G. B., and Stuart, G. J. (2010). Dendritic synapse location and neocortical spike-timingdependent plasticity. *Frontiers in Synaptic Neuroscience*, 2(29).

Fusi, S. and Abbott, L. F. (2007). Limits on the memory storage capacity of bounded synapses. *Nature Neuroscience*, 10(4):485–493.

Gambino, F., Pagès, S., Kehayas, V., Baptista, D., Tatti, R., Carleton, A., and Holtmaat, A. (2014). Sensory-evoked LTP driven by dendritic plateau potentials in vivo. *Nature*, 515(7525):116–119.

Gerstner, W., Kempter, R., Hemmen, J. L. V., and Wagnert, H. (1996). A neuronal learning rule for sub-milisecond temporal coding. *Nature*, 383:76–78.

Goldberg, J., Holthoff, K., and Yuste, R. (2002). A problem with Hebb and local spikes. *Trends in Neurosciences*, 25(9):433–435.

Graupner, M. and Brunel, N. (2012). Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *PNAS*, 109(10):3991–3996.

Graupner, M., Wallisch, P., and Ostojic, S. (2016). Natural firing patterns imply low sensitivity of synaptic plasticity to spike timing compared with firing rate. *Journal of Neuroscience*, 36(44):11238–11258.

Graves, A. R., Roth, R. H., Tan, H. L., Zhu, Q., Bygrave, A. M., Lopez-Ortega, E., Hong, I., Spiegel, A. C., Johnson, R. C., Vogelstein, J. T., Tward, D. J., Miller, M. I., and Huganir, R. L. (2021). Visualizing synaptic plasticity in vivo by large-scale imaging of endogenous AMPA receptors. *eLife*, 10:e66809.

Harris, J. A., Mihalas, S., et al. (2019). Hierarchical organization of cortical and thalamic connectivity. *Nature*, 575(7781):195–202.

Harris, K. D. (2005). Neural signatures of cell assembly organization. *Nature Reviews Neuroscience*, 6(5):399–407.

Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, 424(6948):552–556.

Harris, K. M. and Stevens, J. K. (1989). Dendritic spines of CA1 pyramidal cells in the rat hippocampus: Serial electron microscopy with reference to their biophysical characteristics. *Journal of Neuroscience*, 9(8):2982–2997.

Hay, E., Hill, S., Schürmann, F., Markram, H., and Segev, I. (2011). Models of Neocortical Layer 5b Pyramidal Cells Capturing a Wide Range of Dendritic and Perisomatic Active Properties. *PLoS Computational Biology*, 7(7):e1002107.

Hebb, D. O. (1949). *The Organization of Behavior; A Neuropsychological Theory.* John Wiley & Sons, Inc., New York.

Herzog, R., Morales, A., Mora, S., Araya, J., Escobar, M. J., Palacios, A. G., and Cofré, R. (2021). Scalable and accurate method for neuronal ensemble detection in spiking neural networks. *PLoS ONE*, 16(7):e0251647.

Higgins, D., Graupner, M., and Brunel, N. (2014). Memory Maintenance in Synapses with Calcium-Based Plasticity in the Presence of Background Activity. *PLoS Computational Biology*, 10(10):e1003834.

Hines, M. L. and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural computation*, 9(6):1179–1209.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *PNAS*, 79(8):2554–2558.

Iacaruso, M. F., Gasler, I. T., and Hofer, S. B. (2017). Synaptic organization of visual space in primary visual cortex. *Nature*, 547(7664):449–452.

Inglebert, Y., Aljadeff, J., Brunel, N., and Debanne, D. (2020). Synaptic plasticity rules with physiological calcium levels. *PNAS*, 117(52):33639–33648.

Isbister, J. B., Ecker, A., Pokorny, C., Bolaños-Puchet, S., Egas Santander, D., et al. (2023). Modeling and Simulation of Neocortical Micro- and Mesocircuitry. Part II: Physiology and Experimentation. *bioRxiv*.

Jahr, C. E. and Stevens, C. F. (1990). Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *The Journal of neuroscience*, 10(9):3178–3182.

Jędrzejewska-Szmek, J., Damodaran, S., Dorman, D. B., and Blackwell, K. T. (2017). Calcium dynamics predict direction of synaptic plasticity in striatal spiny projection neurons. *European Journal of Neuroscience*, 45(8):1044–1056.

Jones, H. and Keep, R. (1988). Brain Fluid Calcium Concentration and Response To Acute Hypercalcaemia During Development in the Rat. *Journal of Physiology*, 402:579–593.

Kahle, M. (2009). Topology of random clique complexes. *Discrete Mathematics*, 309:1658–1671.

Kanari, L., Ramaswamy, S., Shi, Y., Morand, S., Meystre, J., Perin, R., Abdellah, M., Wang, Y., Hess, K., and Markram, H. (2019). Objective Morphological Classification of Neocortical Pyramidal Cells. *Cerebral Cortex*, 29(4):1719–1735.

Kastellakis, G. and Poirazi, P. (2019). Synaptic Clustering and Memory Formation. *Frontiers in Molecular Neuroscience*, 12(300).

Kempter, R., Gerstner, W., and van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Physical Review*, 59(4):4498–4514.

Kim, D., Park, P., Li, X., Wong-Campos, J. D., Tian, H., Moult, E. M., Grimm, J. B., Lavis, L., and Cohen, A. E. (2023). Mapping memories: pulse-chase labeling reveals AMPA receptor dynamics during memory formation. *bioRxiv*.

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., and Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *PNAS*, 114(13):3521–3526.

Kossio, Y. F. K., Goedeke, S., Klos, C., and Memmesheimer, R. M. (2021). Drifting assemblies for persistent memory: Neuron transitions and unsupervised compensation. *PNAS*, 118(46):e2023832118.

Krotov, D. and Hopfield, J. J. (2016). Dense associative memory for pattern recognition. *Advances in Neural Information Processing Systems*, 29:1172–1180.

Kumbhar, P., Hines, M., Fouriaux, J., Ovcharenko, A., King, J., Delalondre, F., and Schürmann, F. (2019). CoreNEURON : An Optimized Compute Engine for the NEURON Simulator. *Frontiers in Neuroinformatics*, 13(63).

Landau, A. T., Park, P., Wong-Campos, J. D., Tian, H., Cohen, A. E., and Sabatini, B. L. (2022). Dendritic branch structure compartmentalizes voltage-dependent calcium influx in cortical layer 2/3 pyramidal cells. *eLife*, 11(e76993).

Larkum, M. E. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151.

Larkum, M. E., Zhu, J. J., and Sakmann, B. (2001). Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. *Journal of Physiology*, 533(2):447–466.

Letzkus, J. J., Kampa, B. M., and Stuart, G. J. (2006). Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *Journal of Neuroscience*, 26(41):10420–10429.

Lisman, J. (1989). A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *PNAS*, 86(23):9574–9578.

Lisman, J. E. (1985). A mechanism for memory storage insensitive to molecular turnover: A bistable autophosphorylating kinase. *PNAS*, 82:3055–3057.

Lisman, J. E. and Spruston, N. (2005). Postsynaptic depolarization requirements for LTP and LTD: A critique of spike timing-dependent plasticity. *Nature neuroscience*, 8(7):839–841.

Litwin-Kumar, A. and Doiron, B. (2014). Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature Communications*, 5(5319).

Losonczy, A. and Magee, J. C. (2006). Integrative Properties of Radial Oblique Dendrites in Hippocampal CA1 Pyramidal Neurons. *Neuron*, 50(2):291–307.

Lütgehetmann, D., Govc, D., Smith, J. P., and Levi, R. (2020). Computing persistent homology of directed flag complexes. *Algorithms*, 13(19).

Magee, J. C. and Grienberger, C. (2020). Synaptic Plasticity Forms and Functions. *Annual Review of Neuroscience*, 43.

Mäki-Marttunen, T., Iannella, N., Edwards, A. G., Einevoll, G. T., and Blackwell, K. T. (2020). A unified computational model for cortical post-synaptic plasticity. *eLife*, 9:e55714.

Makino, H. (2019). Top-down control: A unified principle of cortical learning. *Neuroscience Research*, 141:23–28.

Mallya, A. and Lazebnik, S. (2018). PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science*, 275(5297):213–215.

Markram, H., Muller, E. B., Ramaswamy, S., Reimann, M. W., et al. (2015). Reconstruction and Simulation of Neocortical Microcircuitry. *Cell*, 163:456–492.

Markram, H. and Tsodyks, M. (1996). Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Letters to Nature*, 382:807–810.

Mayer, M. L., Westbrookt, G. L., and Guthriet, P. B. (1984). Voltage-dependent block by Mg2+ of NMDA responses in spinal cord neurones. *Nature*, 309:261–263.

McCloskey, M. and Cohen, N. J. (1989). Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *The Psychology of Learning and Motivation*, 24:109–165.

McKernan, M. G. and Shinnick-Gallagher, P. (1997). Fear conditioning induces a lasting potentiation of synaptic currents in vitro. *Nature*, 390:607–611.

Meyer, H. S., Wimmer, V. C., Hemberger, M., Bruno, R. M., De Kock, C. P., Frick, A., Sakmann, B., and Helmstaedter, M. (2010). Cell type-specific thalamic innervation in a column of rat vibrissal cortex. *Cerebral Cortex*, 20(10):2287–2303.

MICrONS (2021). Functional connectomics spanning multiple areas of mouse visual cortex. *bioRxiv*.

Montijn, X. J. S., Olcese, U., and Pennartz, X. C. M. A. (2016). Visual Stimulus Detection Correlates with the Consistency of Temporal Sequences within Stereotyped Events of V1 Neuronal Population Activity. *The Journal of Neuroscience*, 36(33):8624–8640.

Morrison, A., Aertsen, A., and Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural Computation*, 19(6):1437–1467.

Muller, E., Buesing, L., Schemmel, J., and Meier, K. (2007). Spike-Frequency Adapting Neural Ensembles: Beyond Mean Adaptation and Renewal Theories. *Neural Computation*, 19:2958–3010.

Nevian, T., Larkum, M. E., Polsky, A., and Schiller, J. (2007). Properties of basal dendrites of layer 5 pyramidal neurons: A direct patch-clamp recording study. *Nature Neuroscience*, 10(2):206–214.

Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., Chase, S. M., and Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *PNAS*, 116(30):15210–15215.

Pérez-Ortega, J., Alejandre-García, T., and Yuste, R. (2021). Long-term stability of cortical ensembles. *eLife*, 10:e64449.

Perin, R., Berger, T. K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *PNAS*, 108(13):5419–5424.

Pettit, D. L., Wang, S. S., Gee, K. R., and Augustine, G. J. (1997). Chemical two-photon uncaging: A novel approach to mapping glutamate receptors. *Neuron*, 19(3):465–471.

Pfister, J.-P. and Gerstner, W. (2006). Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. *Journal of Neuroscience*, 26(38):9673–9682.

Poirazi, P., Brannon, T., and Mel, B. W. (2003). Pyramidal Neuron as Two-Layer Neural Network. *Neuron*, 37:989–999.

Ratcliff, R. (1990). Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285–308.

Ray, A., Christian, J. A., Mosso, M. B., Park, E., Wegner, W., Willig, K. I., and Barth, A. L. (2023). Quantitative Fluorescence Analysis Reveals Dendrite-Specific Thalamocortical Plasticity in L5 Pyramidal Neurons during Learning. *Journal of Neuroscience*, 43(4):584–600.

Reimann, M. W., Bolaños-Puchet, S., Courcol, J.-D., Egas Santandre, D., et al. (2022). Modeling and Simulation of Rat Non-Barrel Somatosensory Cortex. Part I: Modeling Anatomy. *bioRxiv*.

Reimann, M. W., Egas Santander, D., Ecker, A., and Muller, E. B. (2023). Specific inhibition and disinhibition in the higher-order structure of a cortical connectome. *bioRxiv*.

Reimann, M. W., Nolte, M., Scolamiero, M., Turner, K., Perin, R., Chindemi, G., Dłotko, P., Levi, R., Hess, K., and Markram, H. (2017). Cliques of neurons bound into cavities provide a missing link between structure and function. *Frontiers in Computational Neuroscience*, 11(48).

Reva, M., Rössert, C., Arnaudon, A., Damart, T., Mandge, D., Tuncel, A., Ramaswamy, S., Markram, H., and Werner, V. G. (2022). A universal workflow for creation, validation and generalization of detailed neuronal models. *bioRxiv*.

Reyes-Puerta, V., Sun, J. J., Kim, S., Kilb, W., and Luhmann, H. J. (2015). Laminar and Columnar Structure of Sensory-Evoked Multineuronal Spike Sequences in Adult Rat Barrel Cortex in Vivo. *Cerebral Cortex*, 25(8):2001–2021.

Rodrigues, Y. E., Tigaret, C. M., Marie, H., O'donnell, C., and Veltz, R. (2023). A stochastic model of hippocampal synaptic plasticity with geometrical readout of enzyme dynamics. *eLife*, 12(e80152).

Rößler, N., Jungenitz, T., Sigler, A., Bird, A., Mittag, M., Rhee, J. S., Deller, T., Cuntz, H., Brose, N., Schwarzacher, S. W., and Jedlicka, P. (2023). Skewed distribution of spines is independent of presynaptic transmitter release and synaptic plasticity and emerges early during adult neurogenesis. *bioRxiv*.

Rubin, J. E., Gerkin, R. C., Bi, G. Q., and Chow, C. C. (2005). Calcium time course as a signal for spike-timing-dependent plasticity. *Journal of Neurophysiology*, 93(5):2600–2613.

Schreiber, S., Fellous, J. M., Whitmer, D., Tiesinga, P., and Sejnowski, T. J. (2003). A new correlation-based measure of spike timing reliability. *Neurocomputing*, 52-54:925–931.

Selig, D. K., Nicoll, R. A., and Malenka, R. C. (1999). Hippocampal long-term potentiation preserves the fidelity of postsynaptic responses to presynaptic bursts. *Journal of Neuroscience*, 19(4):1236–1246.

Shatz, C. J. (1992). The Developing Brain. *Scientific American*, 267(3):60–67.

Shouval, H. Z., Bear, M. F., and Cooper, L. N. (2002). A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *PNAS*, 99(16):10831–10836.

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(2):227–238.

Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2003). Neocortical LTD via Coincident Activation of Presynaptic NMDA and Cannabinoid Receptors. *Neuron*, 39(4):641–654.

Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3(9):919–926.

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology*, 3(3):e68.

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

Stuart, G. J. and Sakmann, B. (1994). Active propogation of somatic action potentials into neocortical pyrimidal cell dendrites. *Nature*, 367(January):69–72.

Stuart, G. J. and Spruston, N. (2015). Dendritic integration: 60 years of progress. *Nature Neuroscience*, 18(12):1713–1721.

Tazerart, S., Mitchell, D. E., Miranda-Rottmann, S., and Araya, R. (2020). A spike-timing-dependent plasticity rule for dendritic spines. *Nature Communications*, 11(4276).

Tonegawa, S., Liu, X., Ramirez, S., and Redondo, R. (2015). Memory Engram Cells Have Come of Age. *Neuron*, 87(5):918–931.

Tsodyks, M. and Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *PNAS*, 94(2):719–723.

Turrigiano, G. G. and Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5(2):97–107.

Ujfalussy, B. B. and Makara, J. K. (2020). Impact of functional synapse clusters on neuronal response selectivity. *Nature Communications*, 11(1413).

Vargas-Caballero, M. and Robinson, H. P. (2003). A slow fraction of Mg2+ unblock of NMDA receptors limits their contribution to spike generation in cortical pyramidal neurons. *Journal of Neurophysiology*, 89(5):2778–2783.

Vogels, T. P., Spekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science*, 334(6062):1569–1573.

Wang, B. and Aljadeff, J. (2022). Multiplicative Shot-Noise: A New Route to Stability of Plastic Networks. *Physical Review Letters*, 129(6):068101.

Wang, E. Y., Fahey, P. G., Ponder, K., Ding, Z., Chang, A., Muhammad, T., Patel, S., Ding, Z., Tran, D., Fu, J., Papadopoulos, S., Franke, K., Ecker, A. S., Reimer, J., Pitkow, X., Sinz, F. H., and Tolias, A. S. (2023). Towards a Foundation Model of the Mouse Visual Cortex. *bioRxiv*.

Widrich, M., Schäfl, B., Pavlovic, M., Ramsauer, H., Gruber, L., Holzleitner, M., Brandstetter, J., Sandve, G. K., Greiff, V., Hochreiter, S., and Klambauer, G. (2020). Modern hopfield networks and attention for immune repertoire classification. *Advances in Neural Information Processing Systems*, 33.

Williams, L. E. and Holtmaat, A. (2019). Higher-Order Thalamocortical Inputs Gate Synaptic Long-Term Potentiation via Disinhibition. *Neuron*, 101(1):91–102.

Williams, S. R. and Stuart, G. J. (1999). Mechanisms and consequences of action potential burst firing in rat neocortical pyramidal neurons. *Journal of Physiology*, 521(2):467–482.

Wohrer, A., Humphries, M. D., and Machens, C. K. (2013). Population-wide distributions of neural activity during perceptual decision-making. *Progress in Neurobiology*, 103:156–193.

Zenke, F., Agnes, E. J., and Gerstner, W. (2015). Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature Communications*, 6(6922).

Zenke, F., Gerstner, W., and Ganguli, S. (2017a). The temporal paradox of Hebbian learning and homeostatic plasticity. *Current Opinion in Neurobiology*, 43:166–176.

Zenke, F., Poole, B., and Ganguli, S. (2017b). Continual learning through synaptic intelligence. *ICML*.

Zhang, D., Zhang, C., and Stepanyants, A. (2019). Robust associative learning is sufficient to explain the structural and dynamical properties of local cortical circuits. *Journal of Neuroscience*, 39(35):6888–6904.

Zhu, J. J. (2000). Maturation of layer 5 neocortical pyramidal neurons: amplifying salient layer 1 and layer 4 inputs by Ca2+ action potentials in adult rat tuft dendrites. *Journal of Physiology*, 526(3):571–587.

Supplementary Material

Supplementary Figures

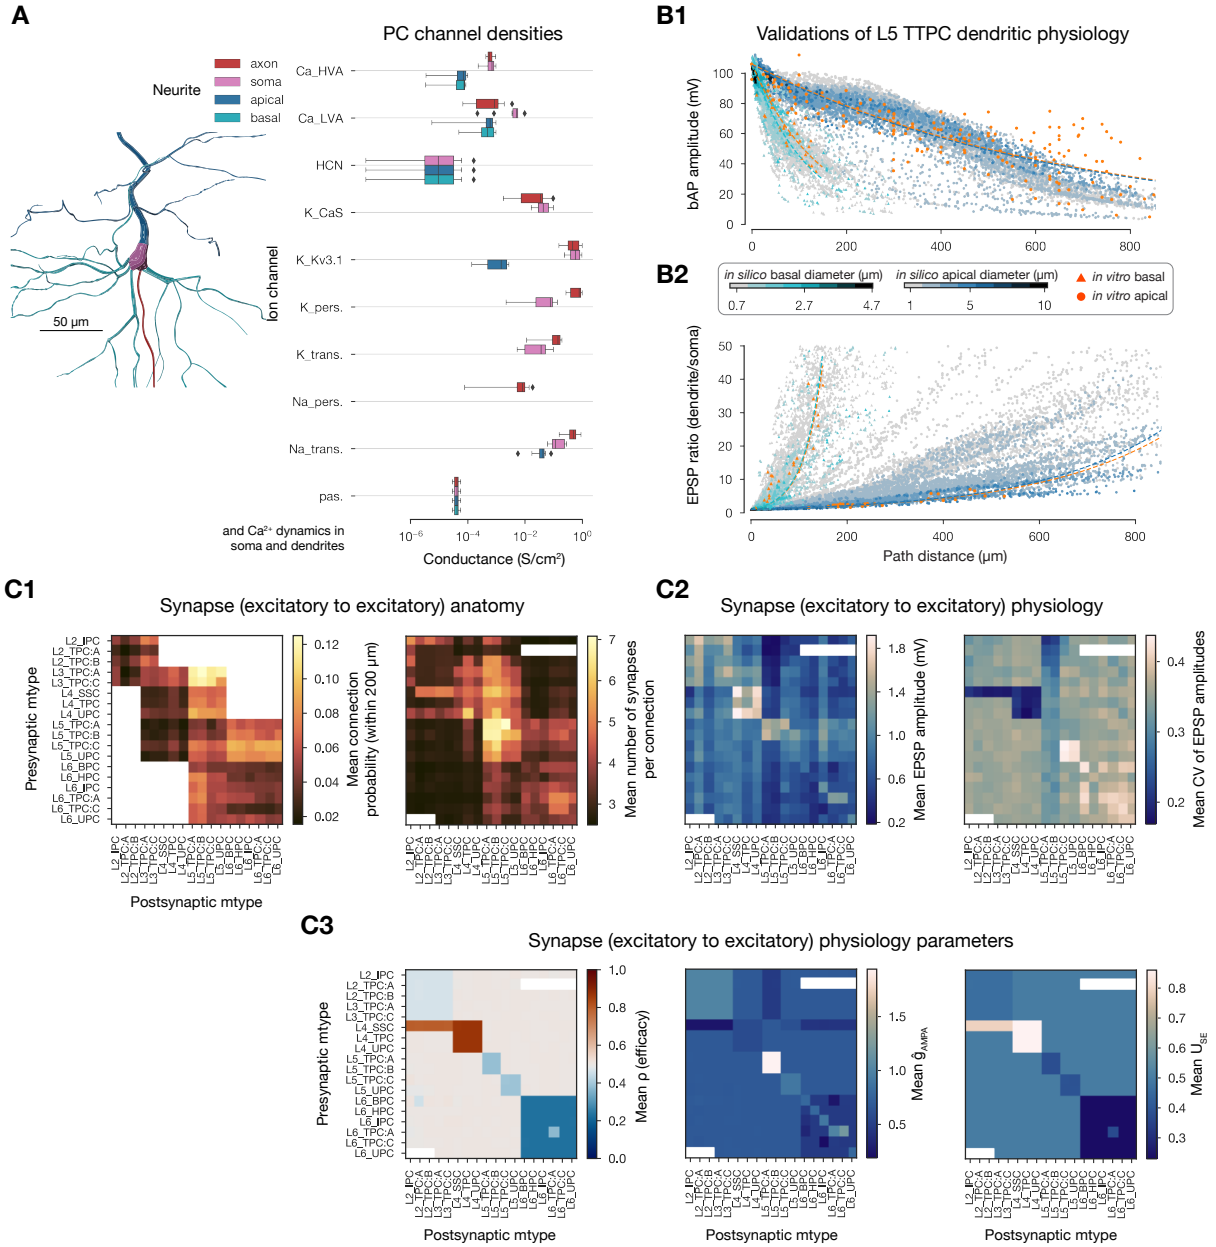

Figure S1: Physiology of excitatory cells and E to E connections. **A**: Distribution of ion-channel densities in the excitatory (cADpyr) electrical type (etype). **B:** Validation of dendritic physiology of the cADpyr etype on L5 TTPC mtypes. **B1**: Validation of back-propagating action potential (bAP) amplitude for basal (teal) and apical (blue) dendrites. Reference data (in orange) comes from Stuart and Sakmann (1994); Larkum et al. (2001) (apical) and Nevian et al. (2007) (basal). Lines show exponential fits for the *in silico* (teal and blue) and *in vitro* (orange) data. Color bar indicates dendritic diameter. **B2:** Validation of EPSP attenuation. Reference data comes from Berger et al. (2001) (apical) and Nevian et al. (2007) (basal). Lines and color bar same as in B2. Data taken from (and partially shown in) Reva et al. (2022). **C:** Anatomy and physiology of E to E connections. **C2**: Connection probability and number of synapses per connections for all E to E connections. White boxes indicate non-feasible connections, or on the left panel: no pairs found within the 200 μm intersomatic distance used. **C2:** Mean (over 100 pairs) PSP amplitude (left) and CV (std/mean on the right) of all E to E connections. Data taken from (and shown in) Isbister et al. (2023). **C3:** Initial synaptic physiology parameters. From left to right: ρ, $\hat{g}_{AMPA}$, and U_{SE}.

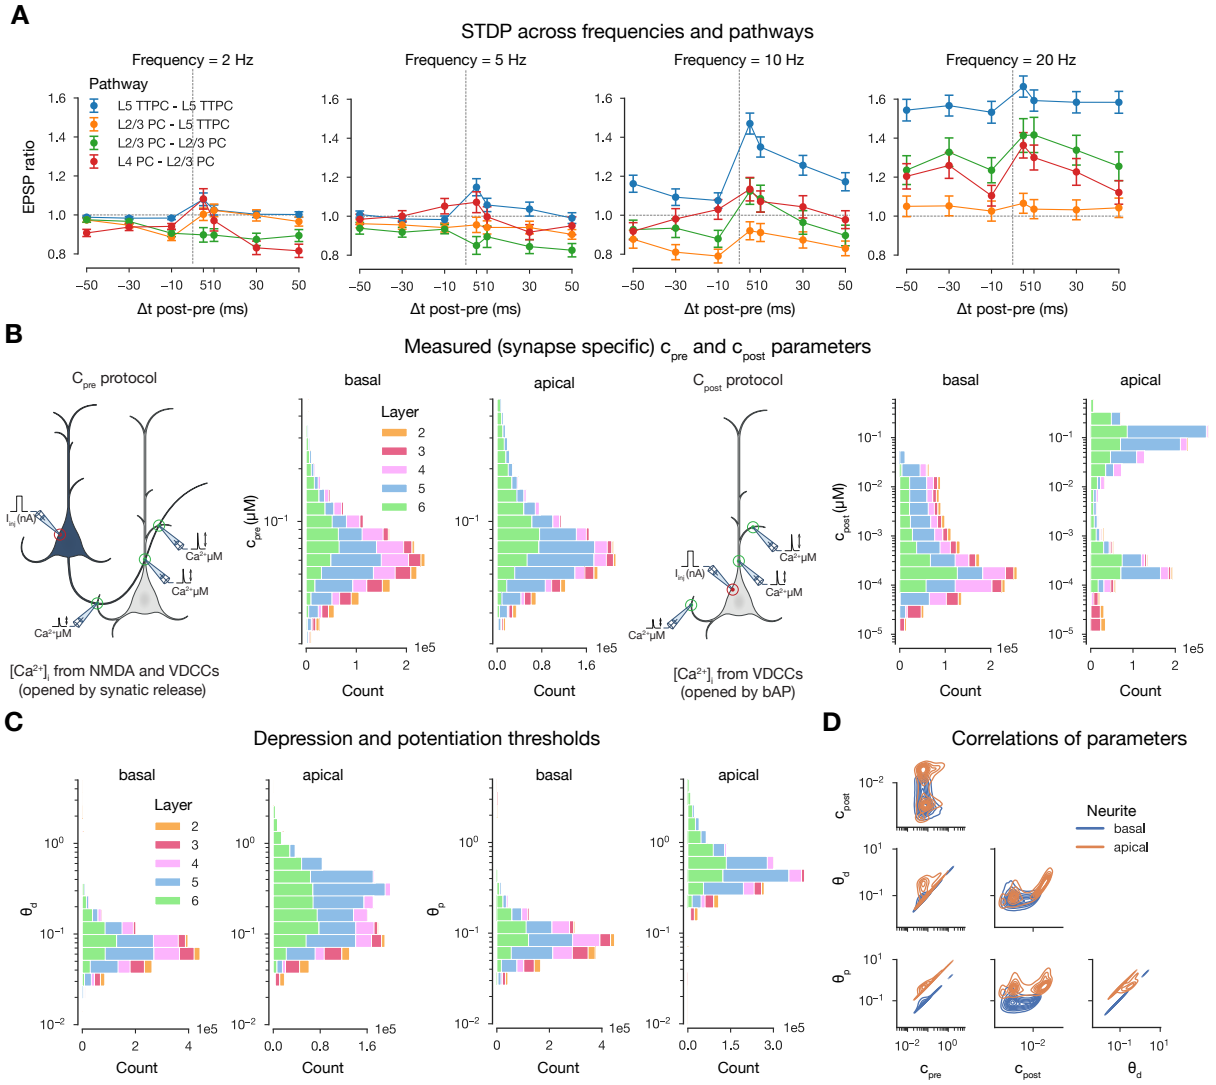

Figure S2: STDP and synapse-specific parameters of the plasticity model. A: Spike timing, pairing frequency, and pathway dependence of plasticity. **B**: Layer- and neurite type-wise distribution of measured $[Ca^{2+}]_i$ peaks (used to derive parameters of the plasticity model shown in C). Synapses are grouped based on the soma location of the postsynaptic cell. 10% of all synapses are shown. Schematics on their lefts illustrate the measurement protocols. **C:** Layer- and neurite type-wise distribution of depression and potentiation thresholds (θ_p and θ_p) of the plasticity model. Synapses grouped and shown as in B. **D**: Correlations of the parameters shown in B and C.

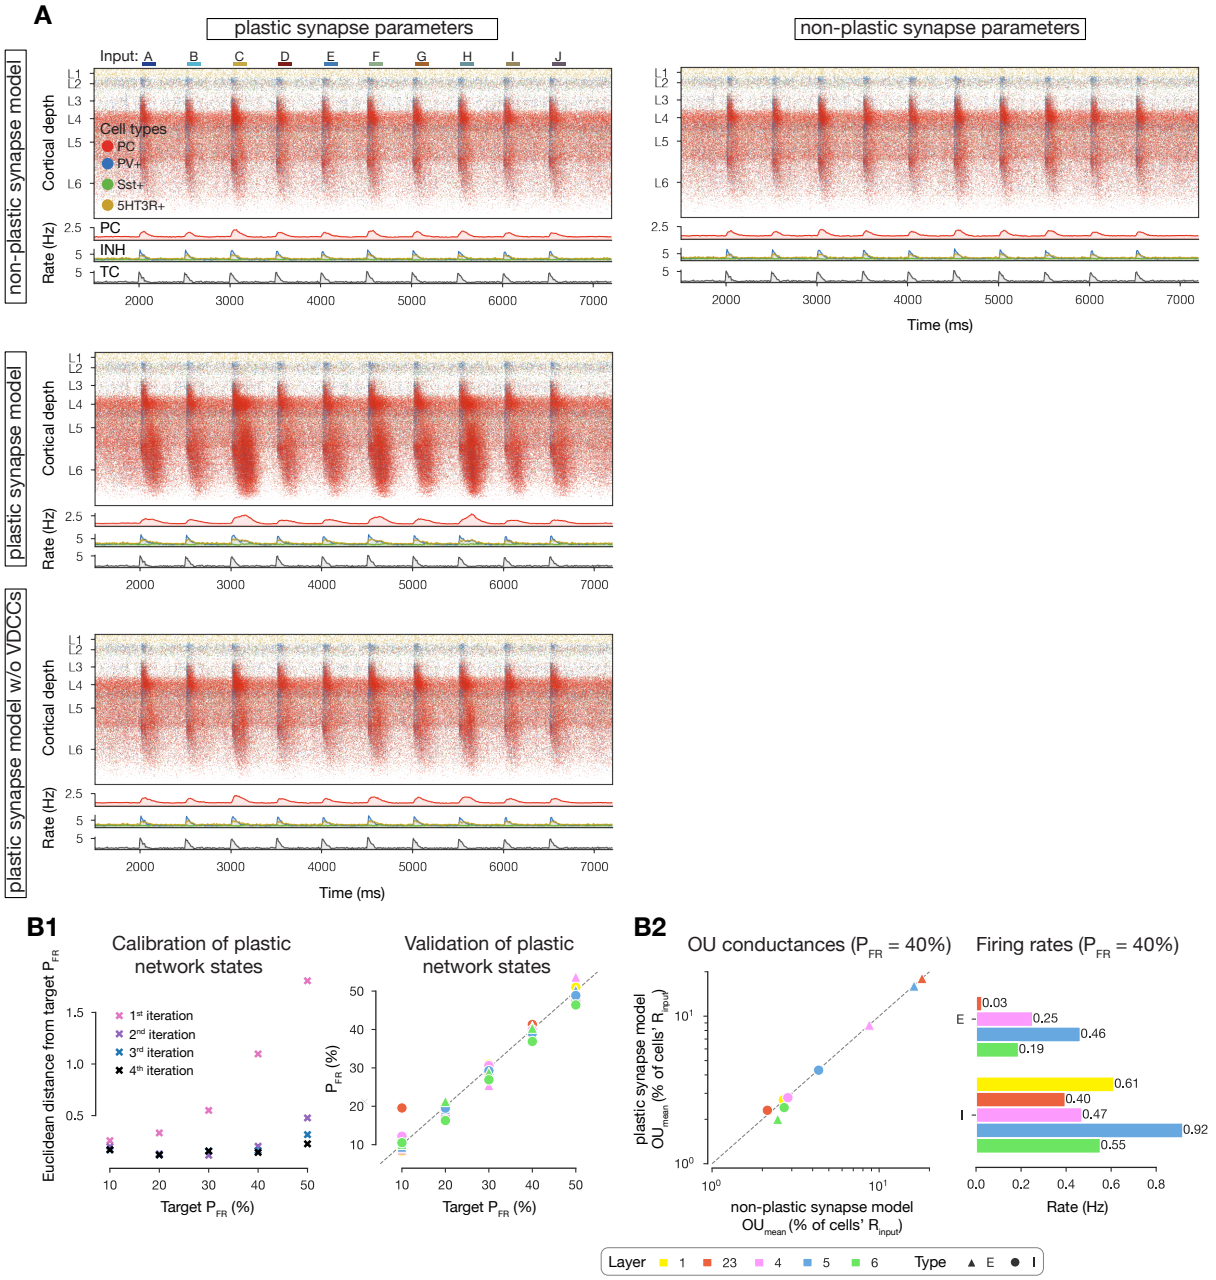

Figure S3: Calibration of the *in vivo*-like network state. A: Same as Figure 2B (i.e., raster plots of the microcircuit's activity) under different synapse setups. The microcircuit equipped with the plasticity model of Chindemi et al. (2022) only resembles that of the non-plastic network's of Isbister et al. (2023) when VDCCs (voltage-dependent calcium channels) are blocked (last row). **B:** Re-calibration of the *in vivo*-like state using the plasticity model. **B1:** Left: Euclidean distance of the measured percentages of firing rates (P_{FR}s) from the target ones in different iterations of the calibration process. Right: Validation of network states after the final (4th) iteration. Dashed gray line along the diagonal indicated perfect match. **B2** Left: Injected Ornstein-Uhlenbeck (OU) conductances in the non-plastic model of Isbister et al. (2023) vs. the plastic one for $P_{FR} = 40\%$ (the state used in the current article). Dashed gray line along the diagonal indicated perfect match. Right: Layer-wise (absolute) firing rates of excitatory (E) and inhibitory (I) subpopulations at $P_{FR} = 40\%$. Legend on the bottom applies to the last three panels in B.

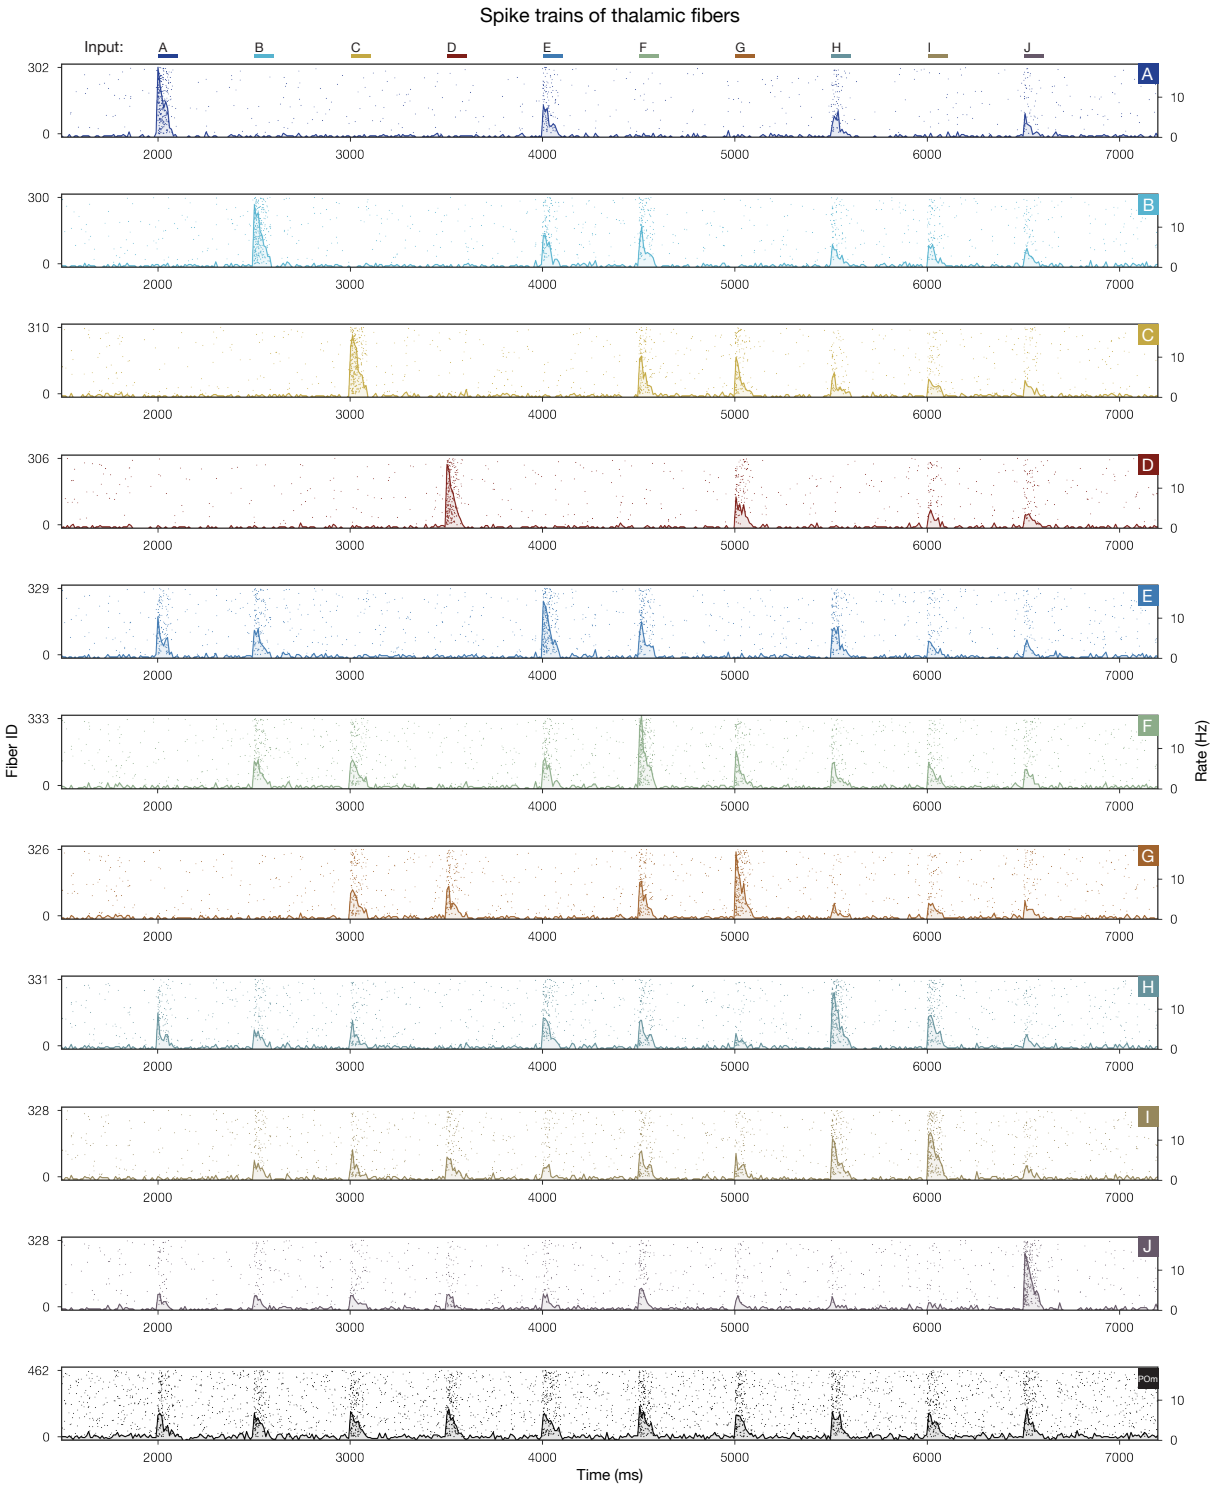

Figure S4: Activity of the thalamic fibers. Raster plots of VPM fibers forming each of the 10 input patterns (Figure 2A) for the stimulus stream (i.e., from pattern A at 2000 ms to pattern J at 6500 ms). Bottom row shows the same for non-specific POm fibers. (Adapted from Ecker et al. (2024).)

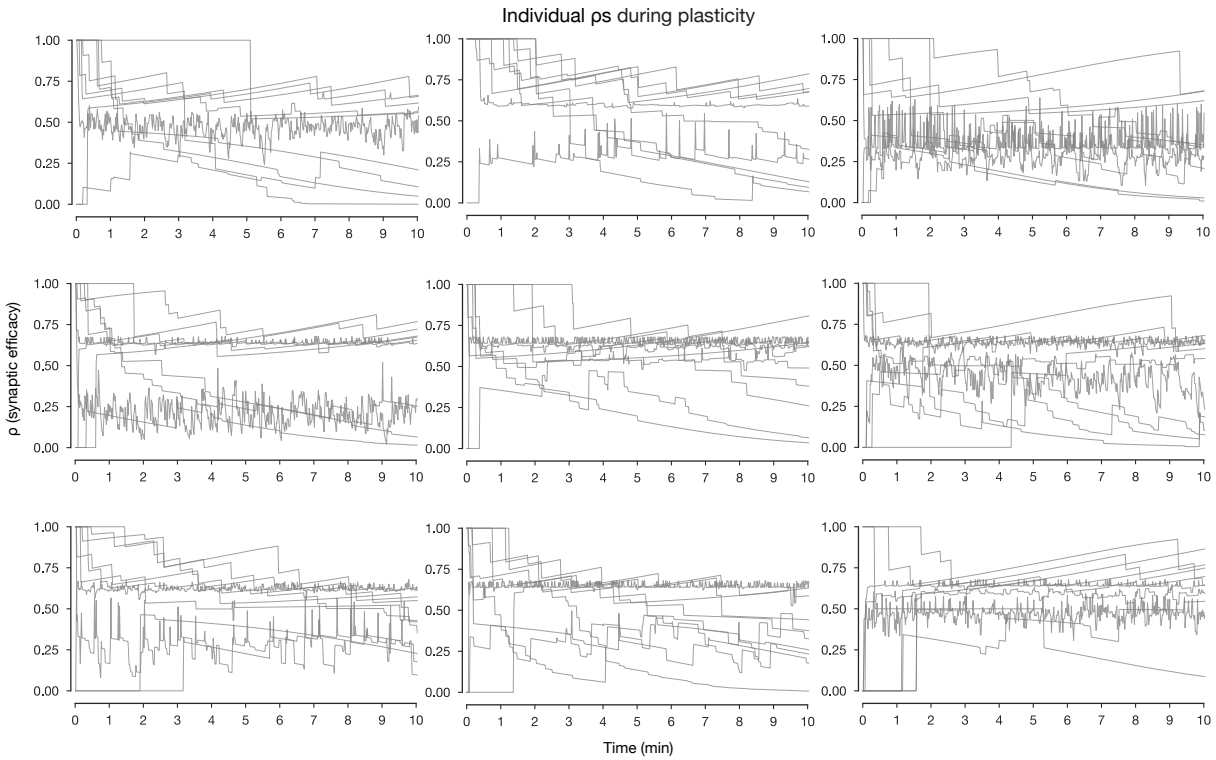

Figure S5: Changes in synaptic efficacy during plasticity. Individual ρ traces (10 per panel) during plasticity. Sampled traces have $|\Delta\rho| \geq 0.1$ and mean over the 10 minutes within $[0.2, 0.8]$.

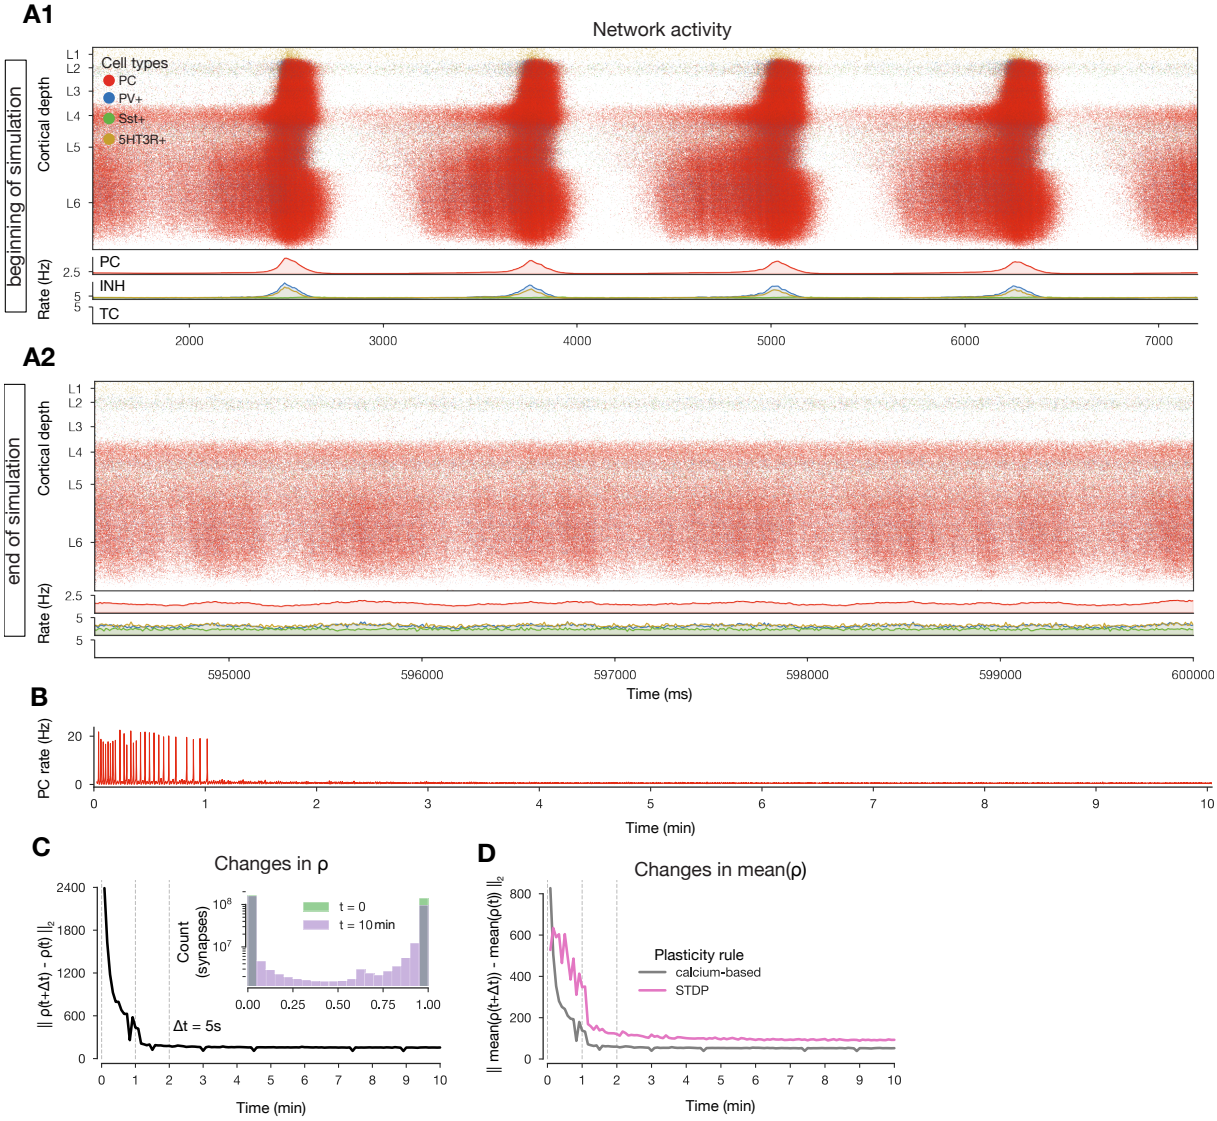

Figure S6: Synchronous (unstable) network activity. A: Same as Figure 2B (i.e., raster plots of the microcircuit's activity) under a different network state, characterized by higher $[Ca^{2+}]_o$ and thus population bursts (see Markram et al. (2015)). **A1** in the beginning of the 10-minute-long simulation, while **A2** at the end. **B:** Firing rate of excitatory cells (zoom-ins on the beginning and end (corresponding to the raster plots above) are shown in Figure 2F). **C** and **D**: Same as Figure 2D and E1 (i.e., L2 norm of changes in ρ (synapses) and mean ρ (connections) across time).

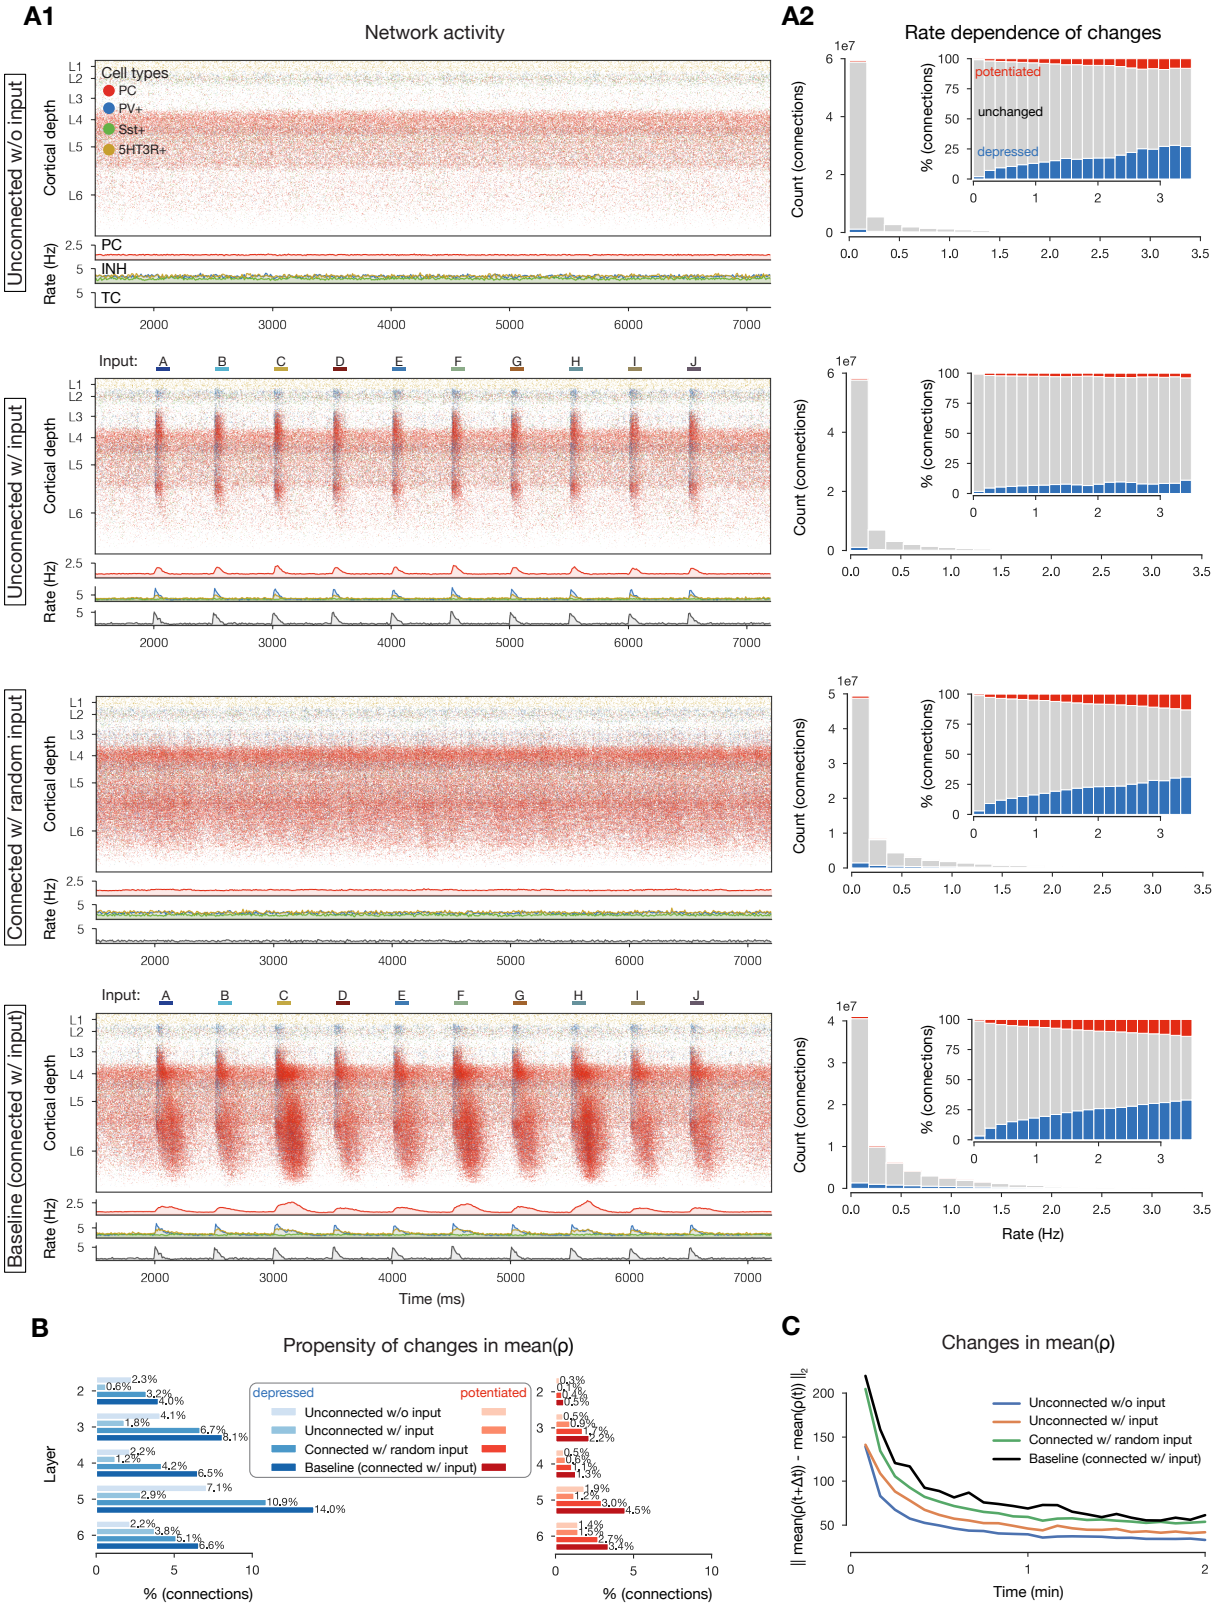

Figure S7: Changing connections in plastic control simulations. A: Same as Figure 2B and 3A2 (i.e., raster plots of the microcircuit's activity and plastic changes in mean ρ vs. firing rates under different conditions). The last row of A2 is not an exact replica of Figure 3A2 as these simulations were run for 2 minutes. **B:** Similar to Figure 3A1 (i.e., layer-wise propensity of changes in mean ρ) but split across conditions. **C:** Similar to Figure 2E1 (i.e., L2 norm of changes in mean ρ values for all conditions.)

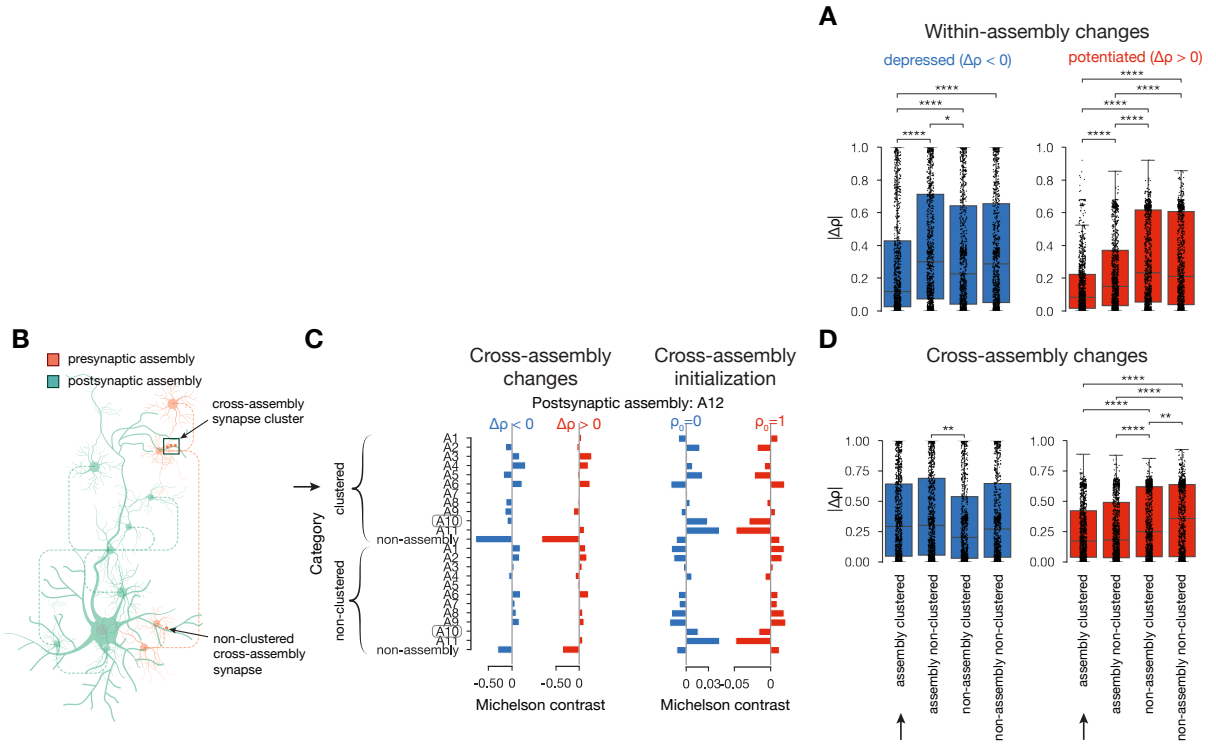

Figure S8: Changes in cross-assembly synapse clusters. A: Distribution of within-assembly $\Delta\rho$s across the four categories introduced in Figure 6 (non)-assembly and (non-)clustered). Black arrow indicates clustered assembly synapses. Boxes show all values, while black dots are 1000 samples from each. Significance test was run on the balanced samples (1000 each): 2-way ANOVA and post-hoc Tukey's test: *: $p \leq 0.05$, **: $p \leq 0.01$, ***: $p \leq 0.001$, ****: $p \leq 0.0001$. **B:** Schematics of cross-assembly synapse clusters. Postsynaptic neuron from the green assembly, presynaptic ones from the orange assembly. **C:** As Figure 6 C and D, i.e., Michelson contrast (Equation (15) in Methods) of the likelihood of plastic changes of synapses grouped into categories (not only four as there are several presynaptic assemblies) on the left, and initialization of ρ values on the right. Depression in blue and potentiation in red. Black arrow indicates clustered assembly synapses. Assembly 12 (A12) was chosen as the postsynaptic assembly as it appeared the latest in the sequence (Figure 5B) and thus had the most presynaptic assemblies which were active before it in time. **D:** Same as A, but for cross-assembly synapses (data from several postsynaptic assemblies, not only A12 shown on C).

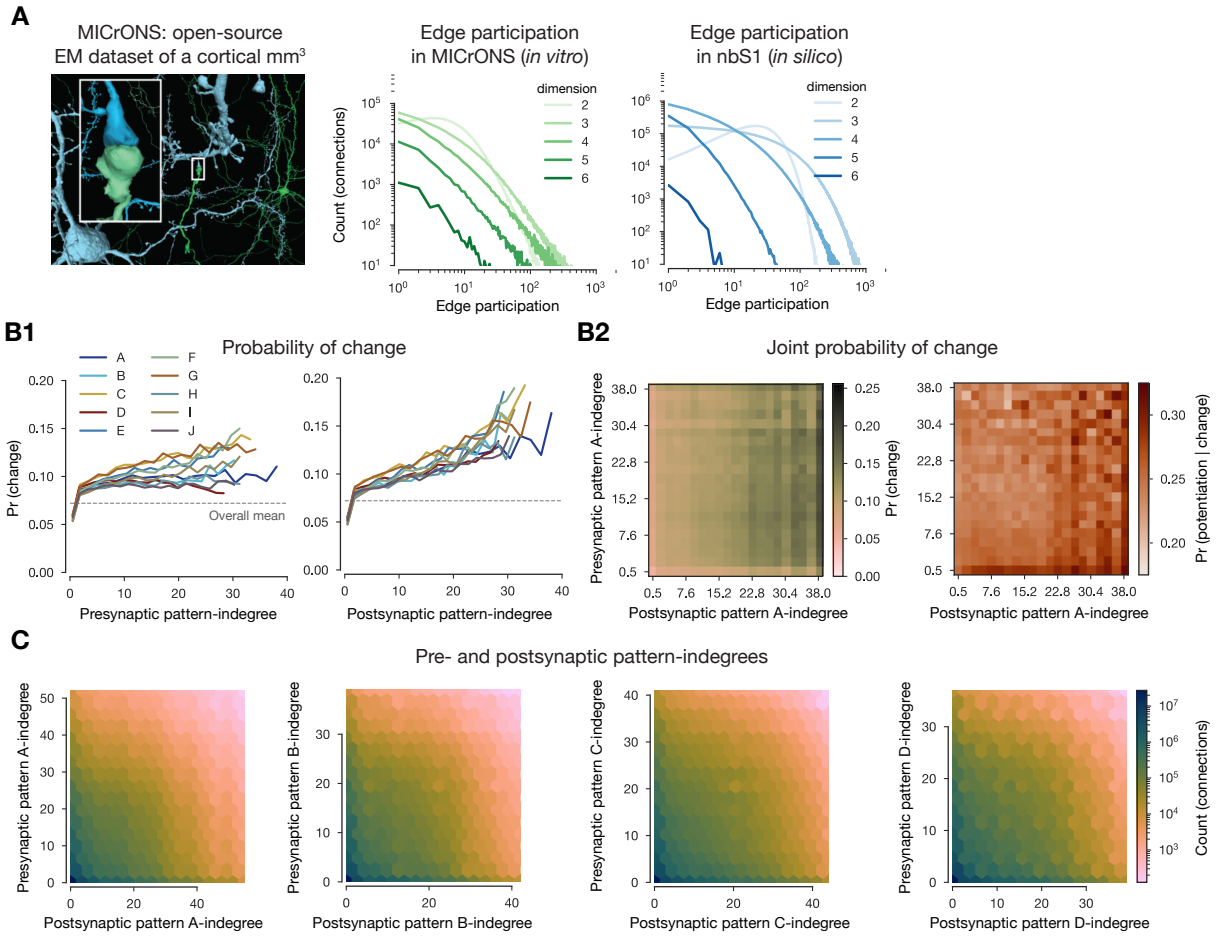

Figure S9: Structural comparison with MICrONS and probability of change vs. pattern-indegree. **A**: Image taken from MICrONS (2021) (under CC-BY 4.0 license). Distributions of edge participation in high dimensional simplices (see Methods) across dimensions in the MICrONS (2021) dataset and in our nbS1 model (Reimann et al., 2022). **B:** Probability of change vs. pattern-indegree (number of connections from VPM patterns). **B1:** Probability of changes (in the 2-minute-long single pattern simulations) vs. pattern-indegree of the presynaptic (left) and postsynaptic (right) neurons. **B2:** Probability of changes (i.e., either depression or potentiation) against the pattern-indegree of both pre- and postsynaptic neurons on the left (for pattern A only). Probability of potentiation conditioned on plastic change (in any direction) against pattern-indegrees on the right. **C:** Pattern-indegree of pre- and postsynaptic cells (for the four base patterns only as the remaining ones are based on these four).

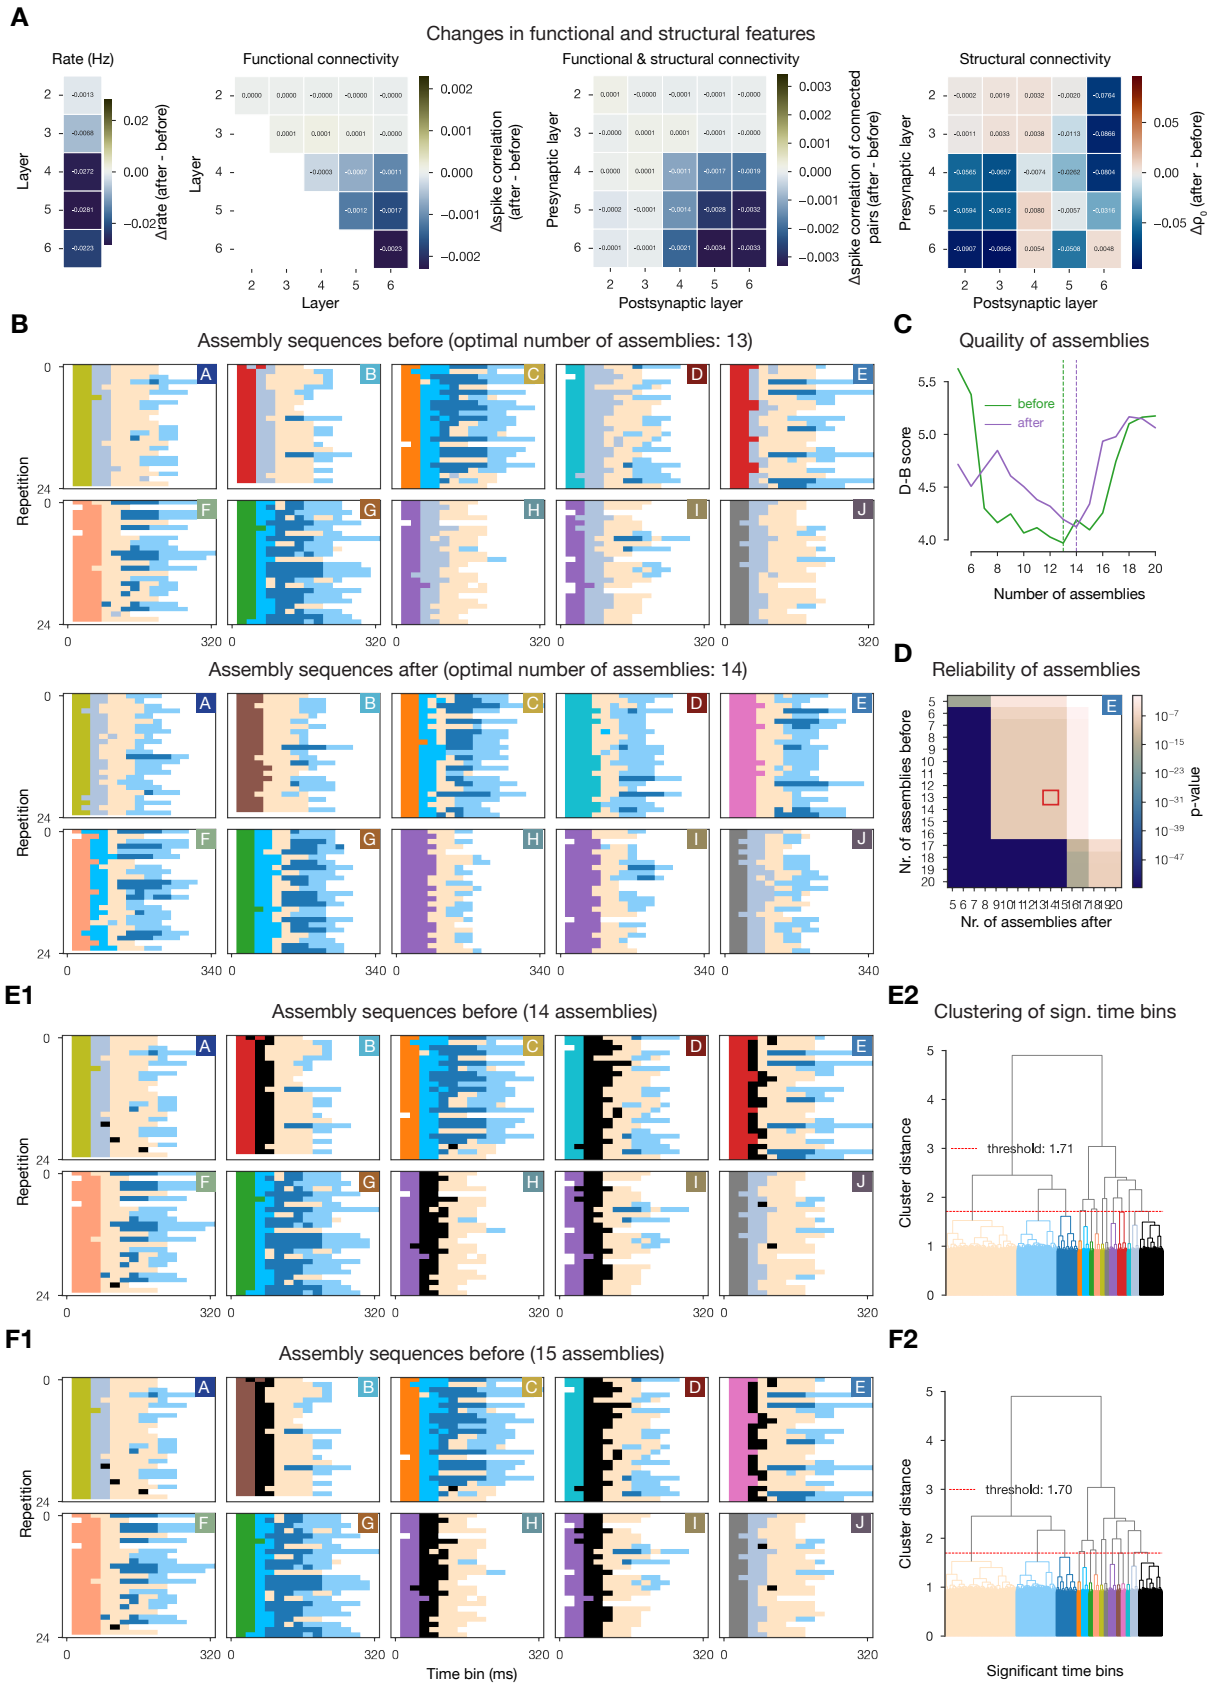

Figure S10: Layer-wise changes and assemblies detected before and after plasticity. Caption on the next page.

A: Layer-wise distribution of changes (after - before) from left to right in: firing rate, pairwise spike correlation, pairwise spike correlation of connected cells, and ρ values. **B**: Same as Figure 5B (i.e., sequential activation of assemblies over repetitions of patterns) but for assemblies detected from 2-minute-long simulations (leading to fewer presentations of patterns) before and after plasticity on top and bottom respectively. **C**: Davis-Bouldin index (Davies and Bouldin, 1979) used to determine optimal number of assemblies before (in green) and after (in purple) plasticity. **D**: Reliability of assembly sequences (see Methods) for selected pattern E (as that is the one with a stimulus specific early assembly after plasticity with increased reliability) for different number of assemblies detected before and after plasticity. Red rectangle indicates data shown on Figure 8D. As on the main Figure p-values from Kruskal-Wallis test. Missing entiries indicate either a decrease or a non-significant (p-value > 0.05) increase in assembly sequence reliability. **E** and **F**: Different number of assemblies before plasticity. **E1**: As B, but for 14 assemblies. **E2**: Ward's linkage tree and the threshold that results in 14 assemblies. **F**: As E, but for 15 assemblies.