Long-term plasticity induces sparse and specific synaptic changes in a biophysically detailed cortical model

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

[1]Blue Brain Project, École polytechnique fédérale de Lausanne (EPFL), Campus Biotech, Geneva, Switzerland
[2]current address: 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
[†]Co-lead authors
[*]Correspondance: michael.reimann@epfl.ch

Abstract

Synaptic plasticity underlies the brain's ability to learn and adapt. While *in vitro* experiments reveal the mechanisms behind plasticity at the level of individual pairs of neurons, they lack the scale to explain how they are coordinated in microcircuits to achieve learning. Conversely, research at the population level still relies on *in silico* approaches of limited generalizability. To overcome these limitations, we embedded a calcium-based model of functional plasticity that captures the diversity of excitatory connections in a thoroughly validated large-scale cortical network model and studied how plasticity shapes stimulus representations at the microcircuit level. We then used an openly available electron microscopic reconstruction of cortical tissue to confirm our testable predictions. We found that in an *in vivo*-like network state, plasticity acted sparsely and specifically, keeping the firing rate stable without additional homeostatic mechanisms. Our results predict that this specificity at the circuit level is governed by co-firing functional assemblies, spatial clustering of synapses on dendrites, and the topology of the whole network's connectivity. These effects cannot be captured with point neuron models, random connectivity and pairwise rules. In summary, our findings elevate descriptions of plasticity rules to the population level, bridging the scales between plasticity and learning in networks.

Keywords

Plasticity, simulation, connectomics, cell assembly, 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 of 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). In parallel with the STDP experiments, modelers developed plenty of *learning rules* that could explain the most recent pathway-specific findings (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000; Pfister and Gerstner, 2006; Clopath et al., 2010). Of particular interest is the calcium-based model of Graupner and Brunel (2012), which models 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. Memory storage and a diverse set of bioplausible learning rules have been linked together to model the formation and maintenance of Hebbian *cell assemblies*, i.e., groups of neurons that fire together (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 of fast homeostatic plasticity, that keeps the networks stable (Zenke et al., 2017a), however experimental evidence for those mechanisms is lacking (Turrigiano and Nelson, 2004). Furthermore, while these studies provided mechanistic explanations of learning and memory, they used point-neuron models, therefore neglecting the structural and functional importance of dendrites and other subcellular components (but see

Bono et al., 2017; Kastellakis and Poirazi, 2019). The compartmentalized nature of dendritic trees gives rise to spatial clustering of synapses (Farinella et al., 2014; Iacaruso et al., 2017; Kastellakis and Poirazi, 2019; 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 go beyond networks of point neurons stabilized with homeostatic plasticity, we equipped the biophysically detailed, large-scale cortical network model of Isbister et al. (2023) with our recently developed, calcium-based model of functional plasticity (Chindemi et al., 2022) between excitatory cells (Figure 1). This way, we had access to more realistic pre- and postsynaptic activity and efficacy of millions of synapses and could characterize the rules governing plasticity at the microcircuit level. To make our predictions more relevant, we calibrated the circuit's activity to mimic an *in vivo* state, characterized by low synaptic release probability and low firing rates. Thanks to the biophysical detail of the model, we could also take the effect of low extracellular calcium concentration ($[Ca^{2+}]_o$) into account (Chindemi et al., 2022), which was experimentally shown to reduce plasticity (Inglebert et al., 2020; Figure 1E). Among 312 million synapses, potentiation dominated in amplitude and depression counteracted it in frequency, which led to stable firing rates without explicitly introducing any homeostatic terms (Turrigiano and Nelson, 2004; Zenke et al., 2017a). As we followed a bottom-up framework and did not model any specific task, we will refer to the effects of plasticity as *changes* in synaptic efficacy instead of learning. Changes in synaptic efficacy were sparse, affecting 5% of all synapses in 10 minutes of biological time. On the other hand, this was still enough to reorganize the network's dynamics, manifesting in more pattern-specificity after plasticity than before. We found an increased likelihood of changes within and across cell assemblies and in synapses forming spatial clusters on postsynaptic dendrites. Additionally, taking the embedding of connections in the whole network connectivity into account, allowed us to predict potentiation even better, and we confirmed this finding in an openly available electron microscopic reconstruction (MICrONS, 2021). To support future, potentially more task-related studies of learning in the cortex, we made the model and the simulator available to the community.

2 Results

To achieve a continuous readout of plastic changes in synaptic efficacy of millions of excitatory synapses, we used a biophysically detailed, large-scale cortical model of the rat non-barrel somatosensory cortex (nbS1). The current model builds upon 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) to keep the complexity of simulation and analysis manageable. In line with the biological variability, excitatory cells are modeled as a diverse set of morphologies (Kanari et al., 2019; Reimann et al., 2022; Figure 1B) equipped with conductances distributed across all compartments (Reva et al., 2022; Supplementary Figure S1A). The connectivity and synaptic physiology of these cells were extensively validated (Reimann et al., 2022; Isbister et al., 2023; Figure 1C; Supplementary Figure S1C). The model is also equipped with fibers from the ventral posteriomedial nucleus of the thalamus (VPM) and the high-order posteriomedial nucleus of the thalamus (POm; Figure 1D; Meyer et al., 2010). We use these fibers to deliver inputs with spatio-temporal precision.

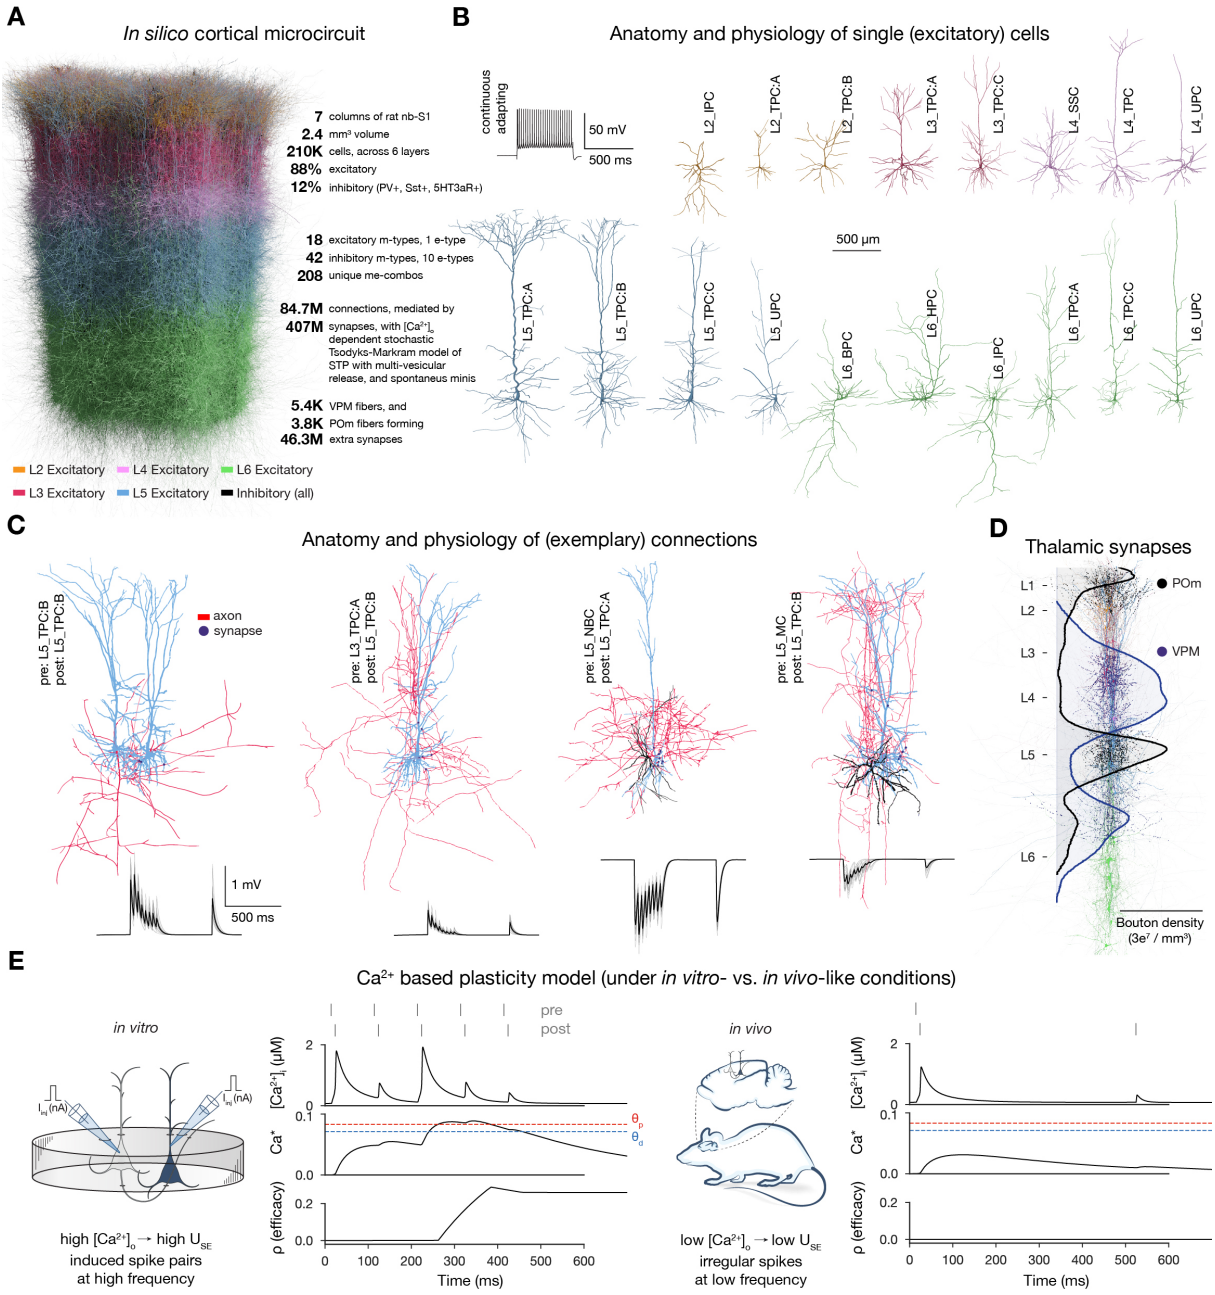

Figure 1: Overview of the network model. **A:** Visualization of the seven column subvolume of rat nbS1. Rendering of 10% of the cells was done with `Brayns`. **B:** Representative morphologies for the 18 excitatory morphological types (m-types) and their typical firing pattern (top left). **C:** Exemplary connections to L5 (thick-tufted pyramidal cells, L5_TPC:A and L5_TPC:B) TTPCs (top) and their STP profiles (bottom). Thin gray represent the 20 individual repetitions, while the thicker black ones their means. Renderings of morphologies (on B as well) were done with `NeuroMorphoVis` (Abdellah et al., 2018). Neurite diameters are scaled (x3) for better resolution. **D:** Bouton density profiles of thalamocortical fibers, and locations of VPM (black) and POm (purple) synapses on neurons (in a 5 μm radius subvolume). Rendering was done with `BioExplorer`. The scale bar on B applies to the whole figure. (Similar panels have been shown in Reimann et al., 2022, Isbister et al., 2023, and Chindemi et al., 2022.) **E:** Variables of the plasticity model during coincident activation of the pre- and postsynaptic neurons. Left: under *in vitro*-like conditions (taken from Chindemi et al., 2022). Right: same pair of neurons under *in vivo*-like conditions. Schematics on their lefts illustrate the difference between *in vitro*- and *in vivo*-like conditions.

2.1 Calcium-based, biophysically detailed model of long-term plasticity

In previous versions of the circuit model, synapses were only equipped with short-term plasticity (STP; Figure 1C). In the remainder of the manuscript we will call this the *non-plastic* version, as our scope here is long-term plasticity. To model long-term plasticity we integrated our recently published calcium-based plasticity model that was used to described 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 lead 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, exhibiting bistable dynamics Lisman, 1985) is updated (Figure 1E left; Equation (1) in Methods).

As Graupner and Brunel (2012) modeled $[Ca^{2+}]_i$ of synapses on point neurons phenomenologically, they had to refit their plasticity model parameters to explain different experimental datasets. Whereas, Chindemi et al. (2022) showed that a generative model, optimized against STDP protocols from only two pathways can explain a large array of other experimentally measured pathways, thanks to the biophysically detailed model of $[Ca^{2+}]_i$ and the morphological complexity of the neurons. The finding of Chindemi et al. (2022), that one unique plasticity rule can rule them all, crucially depends on the location of synapses on the dendrites. The generative model converts location-dependent pre- and postsynaptic $[Ca^{2+}]_i$ peaks into synapse-specific θ_d and θ_p parameters for all excitatory to excitatory (E to E) synapses in the circuit. In our model, we found presynaptically evoked $[Ca^{2+}]_i$ peaks to be three orders of magnitude larger, than the ones evoked by postsynaptic spikes (Supplementary Figure S2A). Postsynaptically evoked $[Ca^{2+}]_i$ peaks had a multimodal distribution in the apical dendrites (Supplementary Figure S2A right), in line with Landau et al. (2022).

Changes in ρ are then converted into changes in the utilization of synaptic efficacy (U_{SE}), a variable of the Tsodyks-Markram model of STP describing the baseline release probability (Tsodyks and Markram, 1997) and the peak AMPA receptor conductance ($\hat{g}_{AMPA}$; Equations (5) and (6) in Methods). As a result of updating U_{SE} as well, short- and long-term plasticity are tightly coupled in the model (Markram and Tsodyks, 1996; Costa et al., 2015; Deperrois and Graupner, 2020). In our network model U_{SE} is also modulated by $[Ca^{2+}]_o$, where a reduction in $[Ca^{2+}]_o$ leads to pathway-specific, non-linear reduction in U_{SE} (Figure 1E right; Markram et al., 2015; Ecker et al., 2020). At initiation, 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).

2.2 Achieving *in vivo*-like network activity

After equipping the circuit with the extra parameters required for long-term plasticity, it was ready to be simulated. To drive network activity, we compensated for missing synaptic input arriving through long-range projections from other brain areas not included in the circuit model (Isbister et al., 2023) and provided inputs through the thalamocortical fibers. Complex phenomena like plasticity are traditionally studied under controlled laboratory conditions *in vitro*, but classical STPD protocols were shown to not induce any plastic changes under *in vivo*-like low $[Ca^{2+}]_o$ (Figure 1E, Inglebert et al., 2020; Chindemi et al., 2022). As our broad interest is understanding the rules governing plasticity in living brains, and our modeling pipeline is capable of taking the effects of low $[Ca^{2+}]_o$ into account (Markram et al., 2015), we calibrated the network's activity to mimic *in vivo* conditions. To that end, 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 (see Isbister et al., 2023 and Methods). 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).

Evoked activity was driven by a thalamocortical input stream already described in Ecker et al. (2023). 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). Spike trains delivered on the pattern fibers followed a 100 ms-long inhomogeneous adapting Markov process (Muller et al., 2007). The overlap of the patterns is clearly visible in the firing pattern of each group of fibers corresponding to them (Supplementary Figure S4). An exemplary raster plot, characterizing the evoked state of the plastic network is shown on Figure 2B.

2.3 Plasticity induces sparse and specific changes while preserving network activity

After achieving *in vivo*-like network activity, we simulated 10 minutes of biological time and assessed the stability of the network activity while it was undergoing plasticity. This is of concern, because previous theoretical work has shown that embedding simple STDP rules in spiking networks without fast homeostatic plasticity leads to exploding firing rates (Morrison et al., 2007; Zenke et al., 2017a). In contrast, with our plasticity rule (Chindemi et al., 2022) excitatory firing rates remained stable throughout the entire length of the simulation (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 fix point and the amount of changes in the second half of the simulation were negligible (Supplementary Figure S5). Moreover, the amount of changes in the second half of the simulation were below a random walk control with the observed step size (Figure 2D2). The results were similar (Figure 2E, black line) when we averaged ρ values of all synapses mediating a connection (4.1 ± 2.3; 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 on the 36 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 stablizing after a transient (Figure 2E, pink line). 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 biophysically detailed plasticity model.

Additionally, for meaningful results the plastic changes had to be non-random and structured by the stimuli. First, we found pathway specificity in the amount of plastic changes observed (Figure 3A). We further quantified this by counting *directed simplices* of dimension k, which are motifs on $k+1$ neurons with all-to-all feed-forward connectivity (see Methods, Figure 3B inset). These motifs have previously shown to be linked to network function (Reimann et al., 2017) as well as quantify complexity of the network's topology (Kahle, 2009; Bobrowski and Kahle, 2018). We found strong overexpression of simplices in the subgraph of changing connections compared

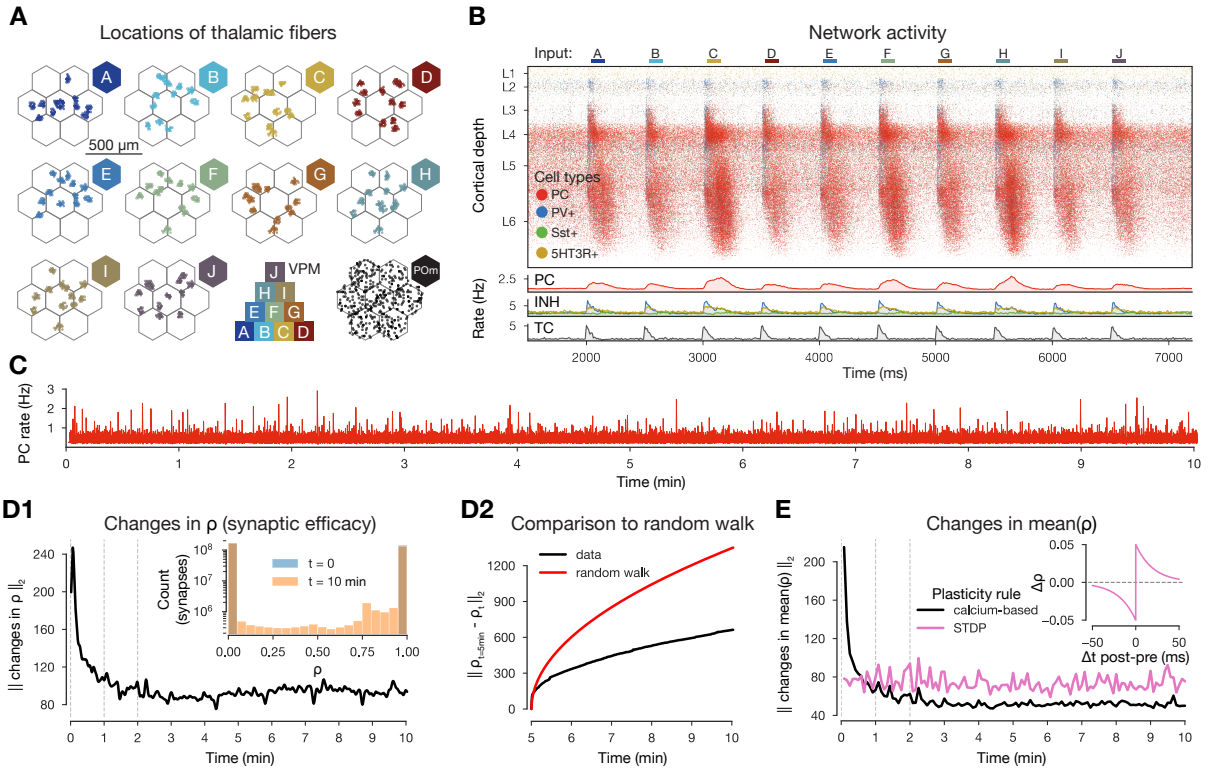

Figure 2: Stable network activity in large-scale plastic simulations. **A:** Centers of the VPM fibers associated with the 10 input patterns in flat map space. Bottom row 3^{rd}: pyramid-like overlap setup of VPM patterns, 4^{th} centers of POm fibers associated with all stimuli. **B:** Raster plot of 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. Similar panels have been shown in Ecker et al. (2023). **C:** Firing rate of excitatory cells.**D1:** Evolution of $\hat{g}_{AMPA}$ during the 10-minute-long simulation. **C1:** Distribution of $\hat{g}_{AMPA}$ in the beginning and end of the plastic simulation. **C2:** Plastic changes that lead to the shift in the distributions shown in C1 (blue: depression, red: potentiation throughout the figure). **C3:** Layer-wise propensity of changes. **D:** Evolution of synaptic efficacy (ρ). **D1:** L2 norm of changes in ρ across time. Inset shows the distribution of ρ values in the beginning (blue) and end (orange) of the simulation. **D2:** L2 norm of changes in ρ (similar to D1, but) compared to $\rho_{t=5minutes}$ in black, random walk control in red. Random walk was defined as $stepsize \cdot \sqrt{2N/\pi}$, where N is the number of steps and $stepsize$ is the average step size of the second half ($t \geq 5minutes$) from D1. **E:** Evolution of mean ρ (aggregated over connections). L2 norm of changes in mean ρ (black) across time against STDP control in pink (inset, see Methods).

to a relevant random control (Figure 3B; see Methods). Furthermore, the maximal simplex dimension found in the subgraph was two higher than expected by chance. By the nature of our control, 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.

Second, we ran 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). We repeated this paradigm three times, for all 10 patterns, 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 ($r = 0.666$, $p < 0.0001$, Figure 3C2). 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

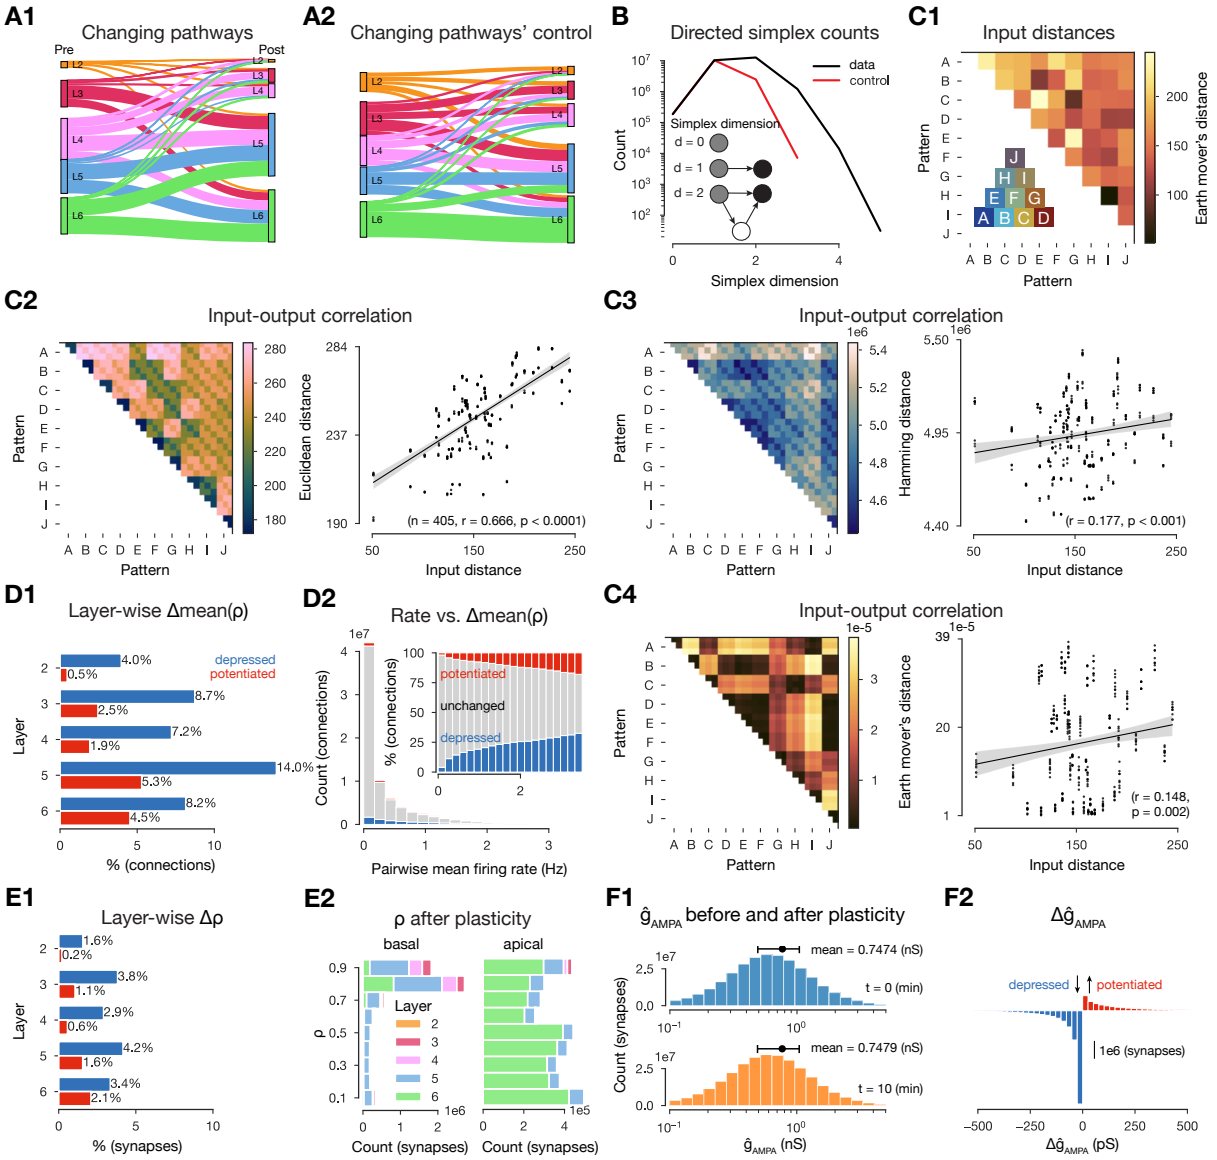

Figure 3: Non-random plastic changes. **A1:** Layer-wise distribution of changing connections. **A2** Control model for A1, generated by taking the same number of connections between the same pre- and postsynaptic populations. **B**: Topology of the changing subnetwork: Directed simplex count in subnetworks from A1 (in black) and A2 (in red). 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. Inset illustrates simplex dimension. **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. Three repetitions (with different random seeds) for each of the 10 input patterns. 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). **D1**: Propensity of layer-wise changes in mean ρ (in the 10-minute-long simulation). **D2:** Plastic changes in mean ρ vs. mean pairwise firing rates. **E1:** Same as D1 but for ρ (changes in synapses rather than connections). **E2:** Layer- and neurite type-wise distribution of non-trivial (neither 0: totally depressed, nor 1: totally potentiated) ρ at the end of the simulation. **F1**: Distribution of $\hat{g}_{AMPA}$ in the beginning (in blue) and end (in orange) of the simulation. **F2:** Plastic changes in $\hat{g}_{AMPA}$ (leading to the slight shift in the distribution on F1).

distribution of mean ρ values into account) were used instead we found only weak correlations (Figure 3C3 and C4 respectively).

Third, we found that the number of connections changing was low (Figure 3D1) and the propensity of changes increased as the pairwise firing rates increased (Figure 3D2), 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 3E1), and in both cases depression was more common than potentiation. Changes occurred an order of magnitude more often on basal compared to apical dendrites although there are roughly the same amount of apical synapses in the circuit (Figure 3E2). Layer 5 (L5) PCs contributed mostly to changes on the basal dendrites, while apical changes happened mostly on L6 PCs.

Finally, we further confirmed stimulus specificity in control simulations where instead of the stimuli we delivered random Poisson spikes on the same VPM fibers at a rate that resulted in the same thalamic spike count. We found that this reduced the number of connections undergoing plasticity by 25% (Supplementary Figure S6), demonstrating the importance of spatiotemporal structure over simple firing of pre- and postsynaptic neurons. Additionally, without synaptic transmission, plasticity was reduced even further: In simulations of stimulus streams where intrinsic connectivity between the simulated neurons was cut (Supplementary Figure S6) we found occasional changes in ρ, but the number of changing synapses was an order of magnitude below baseline conditions. When the external stimuli were also left out, negligible plastic changes occurred.

So far, we characterized plastic change in terms of ρ, as this parameter is bounded to the unit interval and thus easy to interpret. However, the parameter actually affecting the function of the synaptic connection is $\hat{g}_{AMPA}$. We found a minimal increase (0.07%; +0.5 pS, Figure 3F1) in the mean value of this parameter due to plasticity, seemingly contradicting our earlier finding that depression is more frequent (Figure 3D1 and E1). This was explained by a heavier tail of potentiation amplitudes than depression (Figure 3F2). Moreover, the distribution of $\hat{g}_{AMPA}$ remained lognormal, in line with biology (Buzsáki and Mizuseki, 2014; Rößler et al., 2023, Figure 3F1, bottom).

In summary, we observed that $\sim 5\%$ 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 on the recurrent connectivity, the input stimuli, and the firing rates. Potentiation dominated in amplitude, while depression counteracted it in frequency, keeping the network stable without needing to model homeostatic plasticity (Turrigiano and Nelson, 2004; Zenke et al., 2017a). The calcium-based 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.

2.4 More frequent plastic changes within and across cell assemblies

With 95% of synapses remaining unchanged, synaptic plasticity appears to be a highly specific mechanism. 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 S2A). 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*, i.e., groups of neurons that fire together (Hebb, 1949; Harris, 2005). In a previous analysis (Ecker et al., 2023), and in line with experimental results (Harris, 2005; Song et al., 2005; Perin et al., 2011), we observed that *assembly-indegree*, i.e., the number of afferent connections from an assembly, is a great

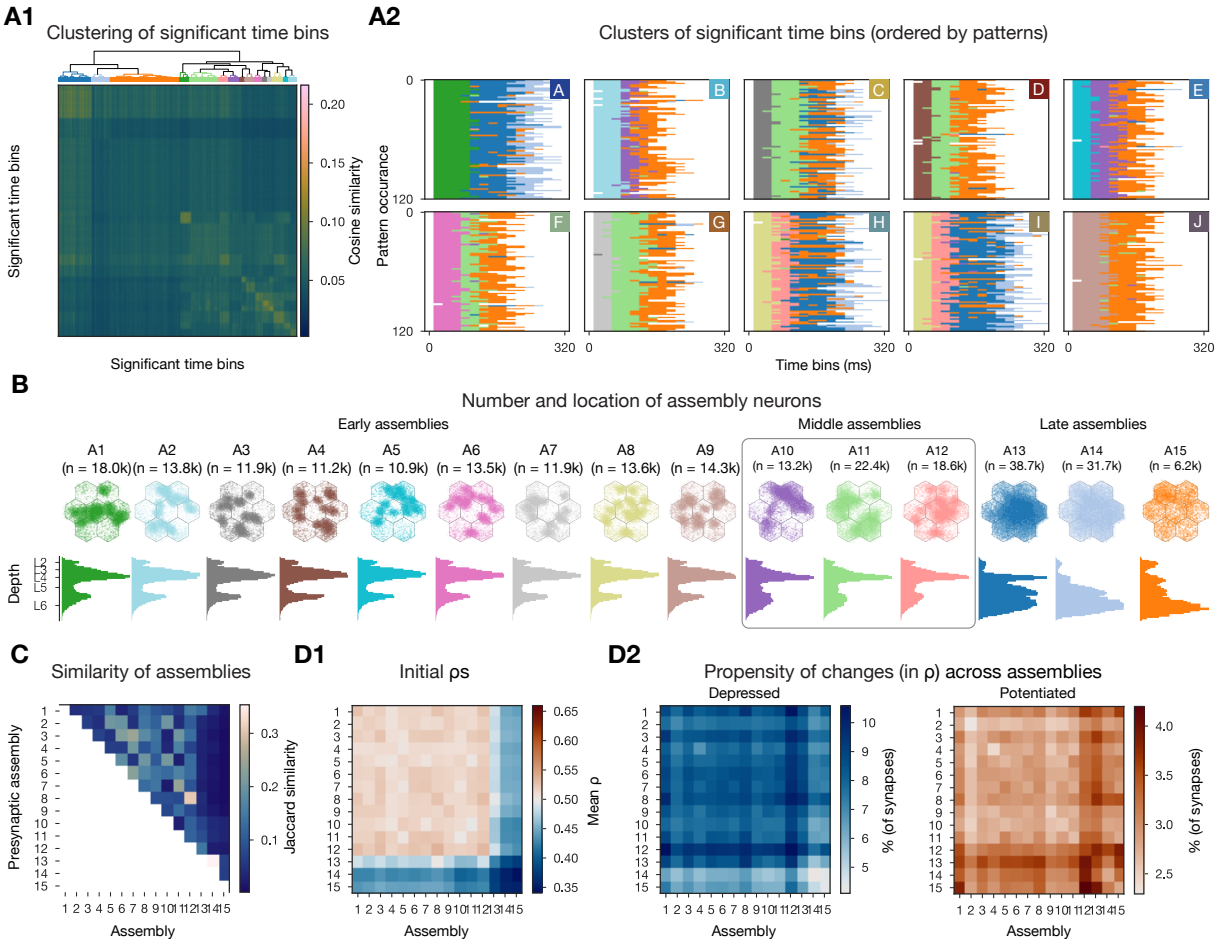

Figure 4: More frequent changes in cell assembly synapses. **A1:** Hierarchical clustering of the cosine similarity matrix of activation vectors of significant time bins (see Methods). **A2:** Clustered significant time bins ordered by the patterns presented. **B**: Number and location of neurons in each cell assembly: flat view on top, depth-profile below. **C:** Jaccard similarity of cell assemblies. **D:** Propensity of changes in cell assemblies. **D1:** Initial mean ρ of within- and cross-assembly synapses. **D2:** Propensity of depression and potentiation of within- and cross-assembly synapses. As assemblies are overlapping (see C) single synapses are taken into account for many different pre- and postsynaptic assembly pairings. (Similar panels (except D) have been shown in Ecker et al., 2023).

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.

Thus, 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 (see Methods; Carrillo-Reid et al., 2015; Herzog et al., 2021). In short, spikes were binned and bins with significantly high firing rates were hierarchically clustered based on the cosine similarity of their activation vector (Figure 4A1). These clusters correspond to the functional assemblies (Figure 4B), 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 4B, C). Note, that in contrast to modeling studies, where

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 4A2). Pattern A elicited the strongest response, while pattern B 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 (Figure 4B-D and 5A, B). We also confirmed our previous findings in this plastic simulation: assembly-indegree is a great predictor of assembly membership with strong cross-assembly interactions; and synapse clustering on postsynaptic dendrites reinforces that trend (Supplementary Figure S7).

When checking the propensity of changes within and across assemblies, we indeed found more synapses undergoing long-term plasticity (Figure 4D2). While only 3.5% of synapses depressed in the whole dataset, we found up to 10.5% when restricting the analysis to assemblies. Similarly, compared to 1.5% of all synapses potentiating, we observed up to 4.2% when restricting to assemblies. Interestingly, large values were found in the off-diagonal entries (Figure 4D2), 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 4B and D1 respectively), thus their synapses are more likely to depress (Figure 4D2). As early assemblies are stimulus specific, and thus not part of the same Hebbian phase-sequences, synaptic depression between these cells can be seen as some kind of orthogonalization of the stimulus responses. On the other hand, late assemblies, that are predominantly composed of cells from the deep layers, have a low initial ρ (Figure 4D1; Supplementary Figure S1C3) and synapses towards them are more likely to potentiate. These assemblies are mostly non-specific and participate in all phase-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.

2.5 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 (Farinella et al., 2014; Iacaruso et al., 2017; Kastellakis and Poirazi, 2019; Tazerart et al., 2020). To quantify this trend, we previously defined the synaptic clustering coefficient (SCC) with respect to an assembly, based on the path distances between synapses from that assembly on to a given neuron (see Ecker et al., 2023 and Methods). For the assemblies detected in this study, we also found SCC to be a good predictor of a neuron's membership in an assembly (Supplementary Figure S7A2), although the effect was less than half as strong as that of assembly-indegree. We used assembly-indegree and SCC to select the 10 most innervated L5 TTPCs (thick-tufted pyramidal cells) within a cell assembly and then explicitly detected spatial clusters of synapses, defined as at least 10 synapses within a 20 μm stretch of a single dendritic branch (see Methods).

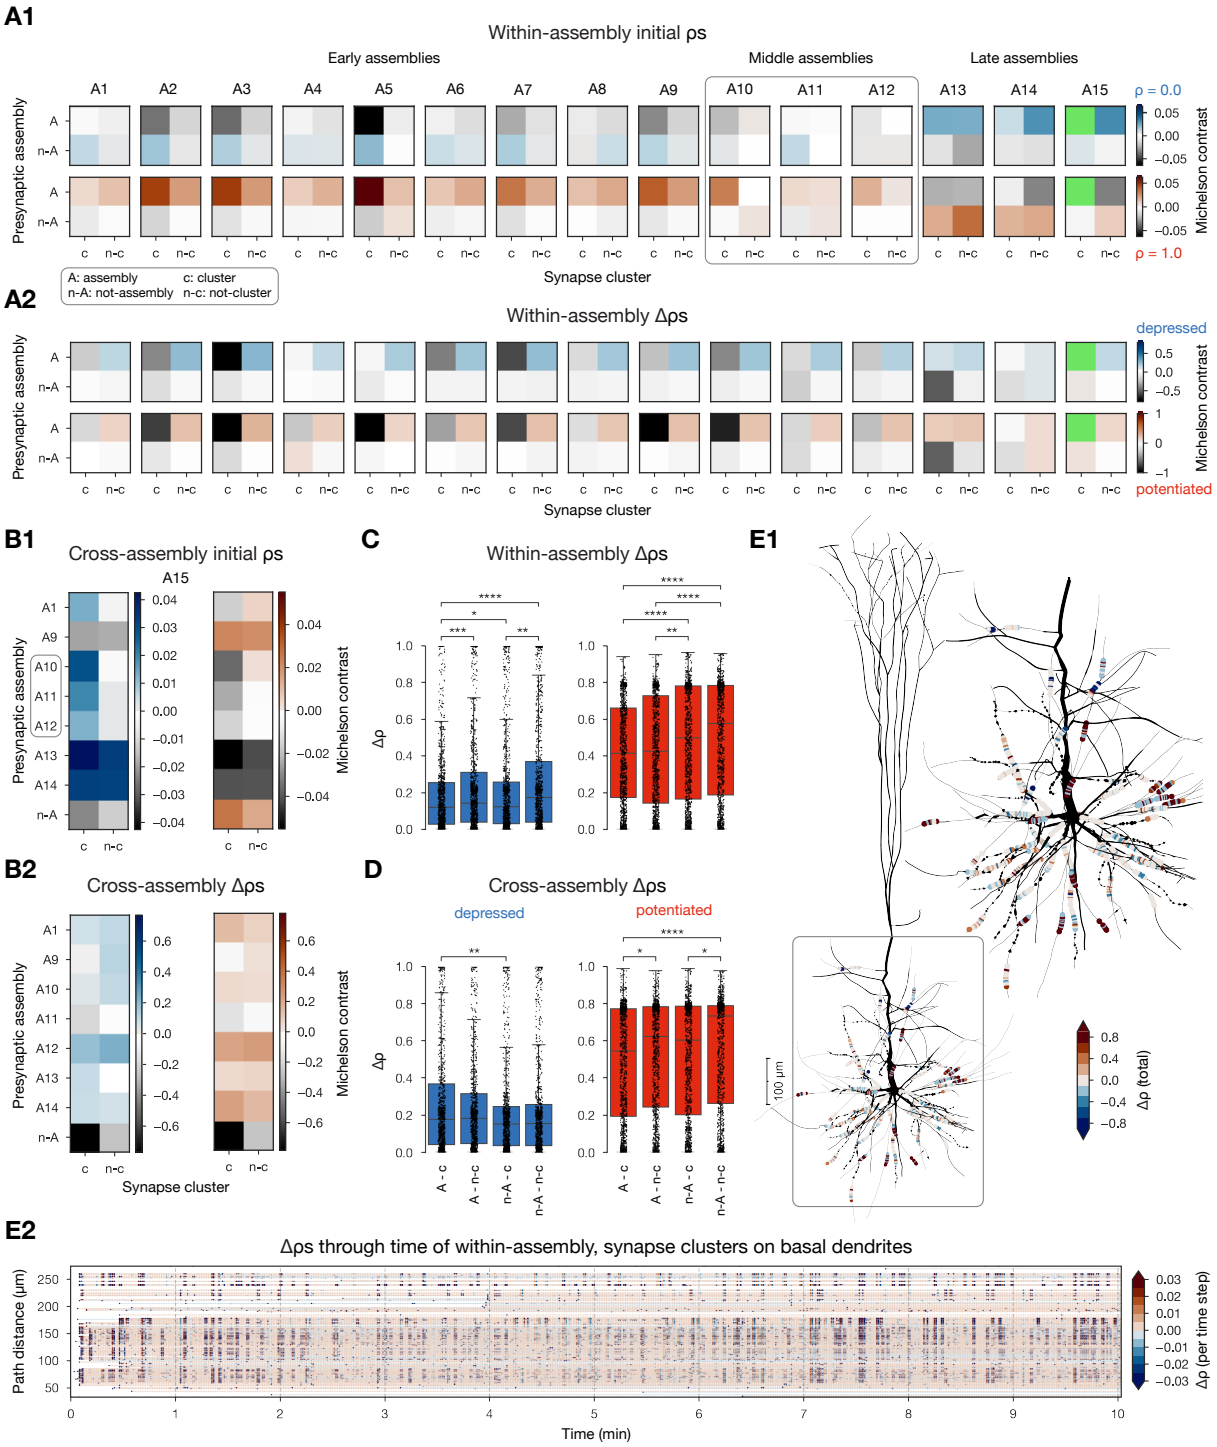

Figure 5: Changes in synapses participating in spatial clusters. **A:** Michelson contrast (Equation (9) in Methods) of probabilities of plastic changes within assembly synapses. Depression on top (blue colormap) and potentiation below (red colormap). Grey parts of the colormaps indicate lower than expected probabilities of observing synapses in a given state, given that it falls into the indicated category, while neon green means no synapses found in the given category. **A1:** Inital ρs, **A2:** plastic changes in ρ. **B:** Same as A, but for cross-assembly synapses (postsynaptic assembly fixed to A15). **C**: Distribution of within-assembly $\Delta\rho$s across the four conditions. 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: *: $p \leq 0.05$, **: $p \leq 0.01$, ***: $p \leq 0.001$, ****: $p \leq 0.0001$. **D:** Same as C, but for cross-assembly synapses (data from several postsynaptic assemblies, not only A15 shown on B). **E1:** Changes in within-assembly, clustered synapse on an examplary A13 neuron. Rendering was done with `NeuroMorphoVis` Abdellah et al. (2018). Neurite diameters are scaled (x2) for better resolution. (Synapse diameters are arbitrary.) **E2:** Temporal evolution of the ($\sim$ 1000) synapses on basal dendrites shown on E1.

For our next analysis, 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 5E1). Then, we quantified the likelihood of plastic change in each category 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 (9) in Methods; Figure 5A2). Note that a nonzero value for one category always has to be compensated by a nonzero value with the opposite sign in another. Surprisingly, clustered within-assembly synapses were not likely to undergo any changes. When we repeated the analysis on the initial ρ values, we found that early and middle assembly synapses, especially the clustered ones, are very likely to be initialized as fully potentiated (Figure 5A1). On the other hand, synapses within the late assemblies were likely to be initialized in the fully depressed state, but were likely to change. 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 (Figure 5C). When we checked the temporal evolution of within-assembly synapse cluster ρ values, we saw that while some of the synapses underwent small constant changes, most of them changed at the same time (vertical stripes on Figure 5E2). Thus 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 $\rho = 1.0$ in absence of activity, as they do not cross the $\rho = 0.5$ unstable fix 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, so we extended the analysis described above to cross-assembly synapses. 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 (Figure 5B). Interestingly, the amplitude of this potentiation was significantly less than that of the other groups' (Figure 5D), but on average, still enough to cross the $\rho = 0.5$ unstable fix point.

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.

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

So far, we have found assembly membership and synapse clustering to be good, neuroscience-based predictors of plastic changes (Harris, 2005; Zenke et al., 2015; Kastellakis and Poirazi, 2019). Here, we extend our predictors to network science-based metrics as well. To do this, as well as to gain insights into pattern-specificity, we analyzed the 2 minutes-long single pattern simulations (Figure 3). We define a connection to be changing for a given pattern if it evolved in the same direction (i.e., potentiation or depression) across all three repetitions of it. This was the case for around 40% of all changing connections for all patterns (Supplementary Figure S8A).

We first studied the propensity of changes against *pattern-indegree*, i.e., the number of VPM fibers belonging to a pattern that innervate a neuron. An increase of either the pre- or postsynaptic pattern-indegree of a connection leads to more frequent plastic changes (Figure 6A). Furthermore, we confirmed that the pre- and postsynaptic effects were not independent by computing the conditional mutual information between them and the probability of their connection to change (see Methods). The mutual information was nonzero (between 0.012 and 0.029) and much larger than the one obtained for corresponding random controls (between 2.2×10^{-7},

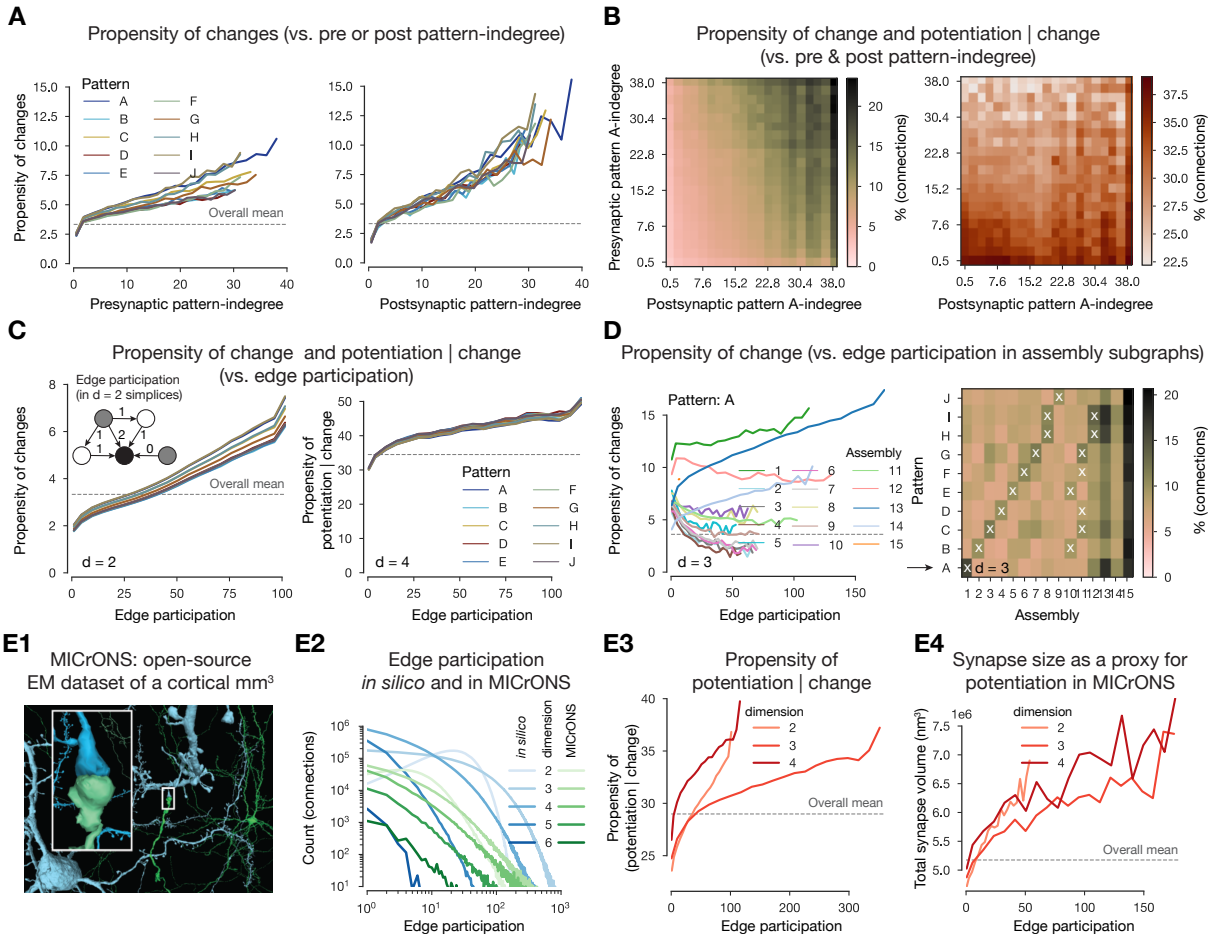

Figure 6: Network-level plasticity rules. **A:** Propensity of changes (in the 2-minute-long single pattern simulations) vs. pattern-indegree of the presynaptic (left) or postsynaptic (right) neurons. **B:** Propensity of changes against the pattern indegree of both pre- and postsynaptic neurons on the left. Propensity of potentiation conditioned on plastic change (in any direction) against pattern indegrees on the right. **C:** Propensity of changes (left) and of potentiation conditioned on plastic change (right) vs. edge participation in high-dimensional simplices (see Methods and inset). **D:** Propensity of changes against edge participation in assembly subgraphs. Propensity of changes in pattern A simulations against edge participation in assembly subgraphs on the left. Summary of the maximum propensity values across patterns on the right. Arrow indicates the row shown in detail on its left and white crosses indicate the pattern-specific early and middle assemblies (see Figure 4A2). **E:** Comparison with electron microscopy data. **E1:** Image taken from MICrONS (2021) (under CC-BY 4.0 license). **E2:** Distributions of edge participation across dimensions in our nbS1 model (Isbister et al., 2023) and in the MICrONS (2021) dataset (see Methods). **E3:** Propensity of potentiation conditioned on plastic change (in the 10-minute-long simulation) against edge participation across dimensions in nbS1. **E4:** Total synapse volume against edge participation across dimensions in MICrONS (2021).

4.5×10^{-7}). Considering their joint distribution showed that even though only 3.3% of the connections changed in the whole network, the propensity of changes rose above 23% when both sides of the connection were highly innervated by the VPM fibers (Figure 6B left). Curiously, pattern-indegree of the postsynaptic side was less important for potentiation, where low pattern-indegree of the presynaptic side was more predictive (Figure 6B right).

Both assembly membership and pattern-indegree are node/neuron-based network metrics. However, due to the highly non-random nature of the network (Reimann et al., 2017, 2022), the value of any node-based metric of the pre- and postsynaptic sides of a connection could be highly

interdependent (Supplementary Figure S8B), making the results difficult to interpret. Thus, we instead used *edge participation*, a metric that considers how the connections are embedded in the entire network. The k-edge participation of a given connection is the number of k-simplices that the connection is part of (see Methods, Figure 6C inset). As edge participation increased so did the propensity of changes (Figure 3C left) and especially of potentiation (Figure 3C right). For overall change, the maximum value reached was lower than for pattern-indegree, indicating lower predictive power; however, for potentiation it was almost twice as high (compare Figure 3B and C). We found an even higher predictive power, when we considered assembly-specific subgraphs (Figure 3D). Additionally, it was pattern specific, i.e., for a given pattern, the strongest predictor of change was the edge participation within the assembly subgraph associated with it (Figure 3D right).

Parts of our network-level plasticity rule were already testable using a freely available dataset. The MICrONS mm^3 dataset, an electron microscopic reconstruction of cortical tissue, combines synaptic resolution (Figure 6E1) with the scale needed to calculate meaningful edge participation. The resulting distributions of edge participation values followed a scaling law that was comparable to our model results (Figure 6 E2). Simultaneously, synaptic resolution allowed calculation of total synapse volume of a connection (see Methods), a measure that is growing with synaptic efficacy (Harris and Stevens, 1989). As *in silico* we found high edge participation to increase the likelihood of potentiation (Figure 6E3), we predicted it to lead to higher total synapse volume in MICrONS (2021). The curves across dimensions are qualitatively the same for both datasets, confirming our prediction (Figure 6E4).

In summary, we found pattern-indegree to be the best predictor of plastic changes. However, considering edge participation restricted to assembly subgraphs resulted in similar predictive power with the additional benefit of pattern specificity. Moreover, edge participation in the full network allowed us to delineate potentiation from depression given change, which we confirmed in the MICrONS (2021) dataset.

2.7 Increased stimulus-specificity characterize the network after plasticity

The evolution of assemblies in terms of their composition and association with stimuli is often used to examine the functional consequences of plasticity and the stability of the neural code in contemporary literature (Fauth and Van Rossum, 2019; Kossio et al., 2021; Pérez-Ortega et al., 2021). To conduct comparable analysis using our setup, we extracted the networks state *before* and *after* our 10-minute-long plastic simulation and ran 2-minute-long non-plastic simulations on each of them (see Methods). On a high level, firing rates increased as a result of plasticity (Supplementary Figure S9A1), while spike correlations remained stable, in line with recent findings (Oby et al., 2019; Feulner et al., 2022, Figure 7A left). The slight increase in spike time reliability (Figure 7A right) could be explained by the increase in firing rates, as the two measures are correlated (Cutts and Eglen, 2014). Plotting pattern-specific peri-stimulus time histograms (PSTHs) before and after plasticity revealed a general lengthening of the late phases of the response and increased amplitudes for selected patterns (Figure 7B).

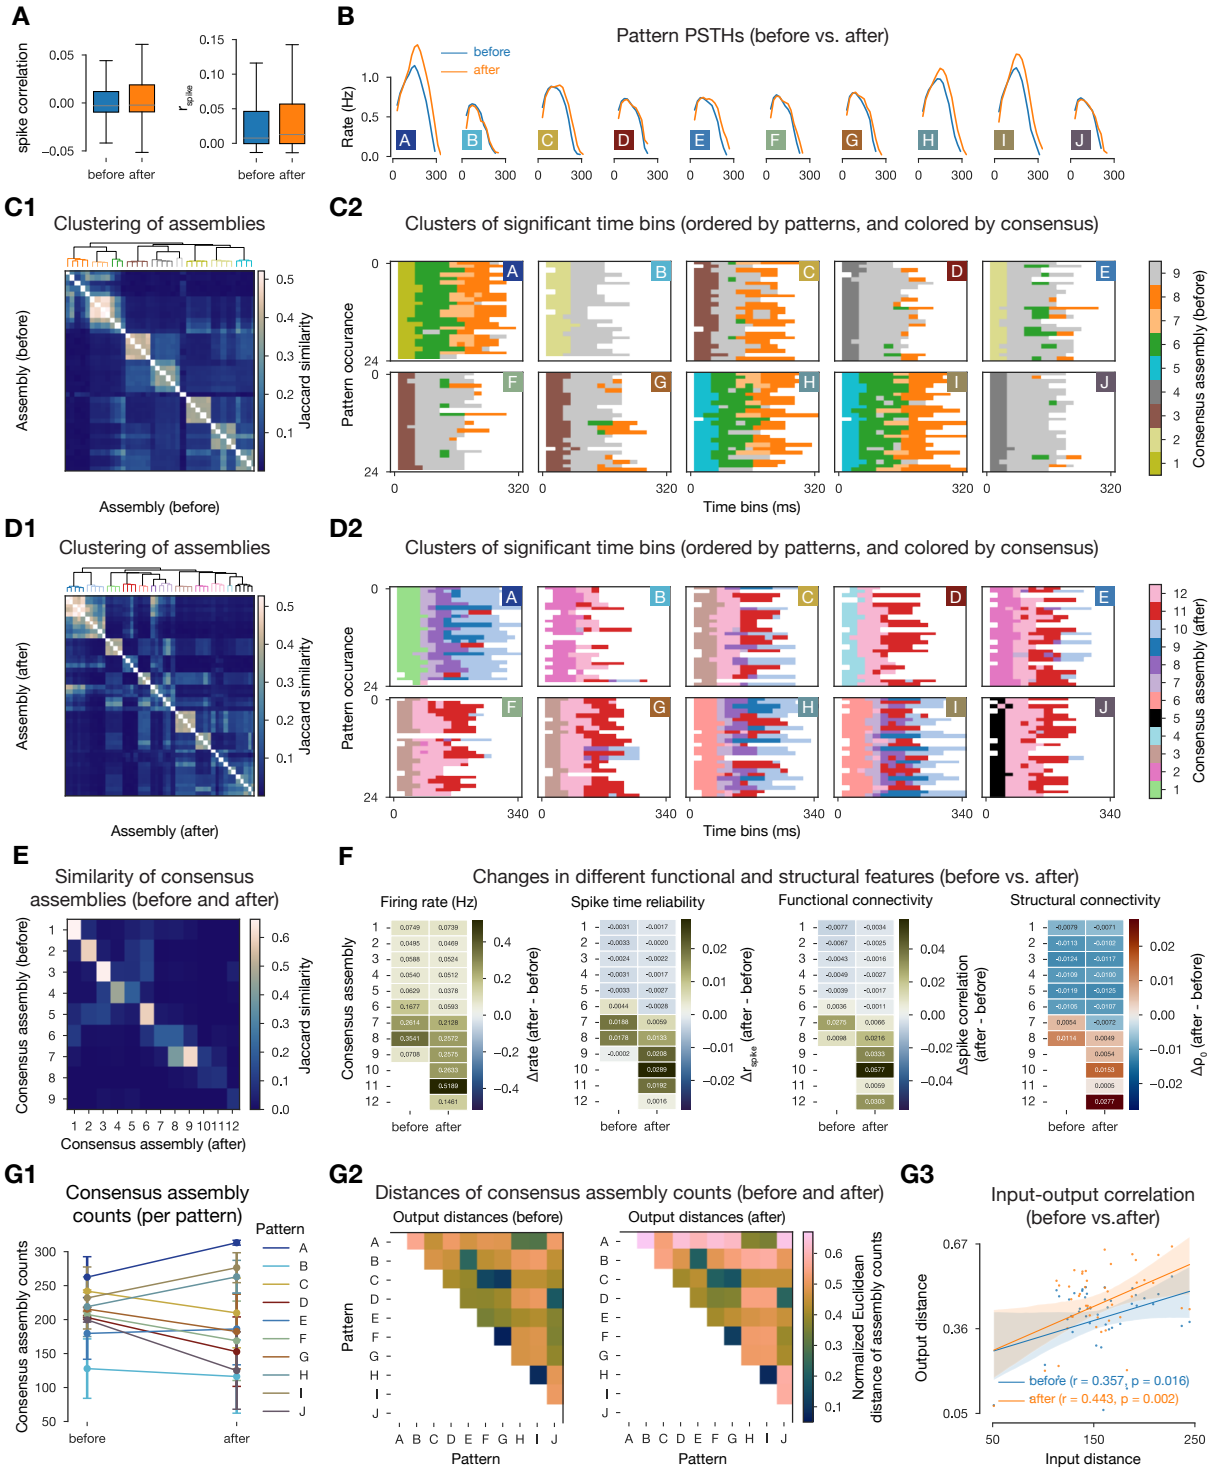

Figure 7: Changes in cell assemblies after plasticity. **A:** Functional network features extracted from spike times of non-plastic simulations *before* plasticity, i.e., in the naive circuit vs. *after* the 10-minute-long plastic simulation. Left: Pairwise spike correlation. Right: Spike time reliability (r_{spike}) measured over five repetitions of the same 2-minute-long simulations with the same input (see Methods). **B:** PSTHs by patterns before vs. after. (Only significant time bins are take into account, see assembly detection in Methods). **C**: Non-plastic consensus assemblies before plasticity. **C1:** Jaccard similarity based hierarchical clustering of assemblies from the five simulation instances. **C2:** Significant time bins from one of the repetitions, ordered by patterns presented, and colored by the consensus assemblies (not the ones detected from that instance). **D:** Same as B, but for non-plastic consensus assemblies after plasticity. (Caption continues on the next page.)

E: Jaccard similarity of consensus assemblies detected before and after plasticity. **F:** From left to right: Changes in firing rate, r_{spike}, pairwise spike correlation, and ρ of within consensus assembly neurons. Colors indicate changes (after - before), while columns indicate at which point the consensus assembly was detected (before or after plasticity). **G:** Input - output correlation: **G1:** Changes in total consensus assembly counts per patter. Error bars are over the five repetitions. **G2:** Output distances as the (normalized) Euclidean distance of pattern evoked consensus assembly cluster over repetitions (see Supplementary Figure S9B). **G3:** Correlation of input distances (see Figure 3C1) and output distances on G2. (Similar panels (except A, B, and F) have been shown in Ecker et al., 2023.)

We detected and compared assemblies before and after plasticity across five repetitions. The number of assemblies in each repetition was objectively determined by minimizing the Davis-Bouldin index (Davies and Bouldin, 1979; see Methods), and the number was consistently higher after plasticity than before, and equal in one repetition (Supplementary Figure S9C). For a robust comparison, we used *consensus assemblies*, i.e., the sets of neurons that were reliably part of a given assembly across repetitions (Figure 7C1 and D1; see Ecker et al., 2023 and Methods). We found more consensus assemblies after plasticity (compare Figure 7C2 and D2), although their sizes remained comparable (Supplementary Figure S9E). Most consensus assemblies before plasticity had a corresponding assembly after plasticity that it shared over 50% of neurons with (Figure 7E).

We observed similar changes in both functional and structural connectivity of consensus assemblies, assessed by firing rate, spike time correlation, reliability, and ρ values (Figure 7F). In particular, we found that through plasticity, early assemblies grew less correlated with weaker internal connectivity and this trend weakened in the subset of neurons that remained part of the consensus assembly (compare columns of Figure 7F). Conversely, in middle and late consensus assemblies correlations and connections grew stronger, especially so in the neurons that were members of the consensus after plasticity.

The activation sequence of consensus assemblies can be seen as a low-dimensional representation of the complex, high-dimensional activity of the network's response. Thus, we analyzed the duration of consensus assembly activations in response to different patterns and their differences. In line with the prolonged PSTHs, we found consensus assemblies to be active longer (190 ± 45 ms before vs. 200 ± 60 ms after plasticity), and could trace it back to selected patterns A, E, H, and I (Figure 7C2 and D2). Following previously established methods (Ecker et al., 2023), we calculated distances between the sequences of consensus assembly activations triggered by different patterns, denoted *output distances* (Figure 7G2). We compared them to the *input distances*, i.e., distances between input pattern fibers (Figure 3C1). We found an overall increase of output distances and of the input - output distance correlation after plasticity ($r = 0.443$ after vs. $r = 0.357$ before plasticity; Figure 7G3).

Taken together, the network became more specific to the patterns it was exposed to. This manifested in assemblies splitting, weakening of early and strengthening of the late assemblies and the consequent prolonged assembly responses to specific patterns.

3 Discussion

Using a detailed, large-scale cortical network model equipped with a calcium-based model of long-term functional plasticity, we have examined changes in synaptic efficacy in response to repeated presentation of 10 different stimuli over 10 minutes of biological time, under *in vivo*-like conditions. Our principal observations in this bottom-up simulation framework are as follows: (1) Plastic changes were sparse, affecting only 5% of the synapses. A balance between occasional large-amplitude potentiation and more frequent depression kept the network stable without ex-

plicitly 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 edges participating in many network motifs. (3) Early-responding cell assemblies were defined by clustered synapses initialized as fully potentiated and remained fairly stable. In contrast, their synapses to late-responding assemblies underwent three times more changes than expected. This resulted in prolonged and more reliable responses to selected patterns after plasticity, and increased stimulus specificity.

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 (Turrigiano and Nelson, 2004; Litwin-Kumar and Doiron, 2014; Zenke et al., 2015, 2017a; Fauth and Van Rossum, 2019; Kossio et al., 2021). The machine learning community is also aware of the importance of sparse changes, as in continual learning one has to balance plasticity and stability to avoid catastrophic forgetting (McCloskey and Cohen, 1989; Ratcliff, 1990). In recent years, they have come up with impressive techniques that mask connections to improve the performance of deep neural networks (Zenke et al., 2017b; Mallya and Lazebnik, 2018; Frankle and Carbin, 2019), whereas in our model it emerged naturally from the high level of biophysical detail. Of course, the amount of data that deep networks are expected to store far exceeds the 10 patterns used here, and it is outside of our scope to find the maximal capacity of our network. On the other hand, we know from theoretical work that for bistable synapses operating on multiple time scales, capacity scales with the square root of the number of synapses (Crick, 1984; Fusi et al., 2005). 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 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 (Farinella et al., 2014; Iacaruso et al., 2017; Kastellakis and Poirazi, 2019; 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 $[Ca^{2+}]_o = 1.05\,\mathrm{mM}$ (Jones and Keep, 1988; Borst, 2010; Markram et al., 2015; Ecker et al., 2020). 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 (Poirazi et al., 2003; Farinella et al., 2014; Ujfalussy and Makara, 2020). With respect to network-based rules, previous approaches used low-level learning rules on point neuron networks and characterized the resulting connectomes as having an overexpression of reciprocal connectivity and densely connected neuron motifs (Brunel, 2016; Zhang et al., 2019). We improved on these by providing more specificity: we developed node and edge-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 finding was validated using an electron microscopic reconstruction of cortical tissue (MICrONS, 2021).

The last observation (3) can be interpreted as a network-level *redistribution of synaptic efficacy* from early assemblies to the middle and late ones. This manifest itself in an increased firing rate for middle and late assemblies and an even larger increase of spontaneous firing (Supplementary Figure S9A1), indicating the reactivation of assemblies, in line with experimental findings (Miller et al., 2014; Carrillo-Reid et al., 2015; Stringer et al., 2019; Herzog et al., 2021; Trägenap et al., 2022). This mirrors *in vitro* observations at the connection level, of a redistribution of synaptic efficacy towards earlier spikes during high-frequency firing caused by changes in U_{SE} (Markram and Tsodyks, 1996; Selig et al., 1999; Sjöström et al., 2003; Costa et al., 2015; Chindemi et al., 2022). However, our *in vivo*-like low firing rates and $[Ca^{2+}]_o$ are not in the ranges where this connection-level redistribution is relevant. Specifically, at *in vitro* levels of $[Ca^{2+}]_o$ the increased U_{SE} shifts the connection further into the depressing regime, causing the redistribution (Supplementary Figure S9B 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 the phenomenon of STP less relevant in general (Supplementary Figure S9B right and A2 respectively). Thus, while Markram and Tsodyks (1996) showed a redistribution of synaptic efficacy after plasticity at the single connection level *in vitro*, we found a redistribution at the network level under *in vivo*-like conditions: efficacy shifted towards synapses targeting the deeper layers of the cortex.

According to the contemporary view of L5 PCs, sensory bottom-up inputs target their basal dendrites, and top-down information arrives at the apical ones, with the coincidence activation of basal and apical inputs encoded by bursts of action potentials (Larkum, 2013; Naud and Sprekeler, 2018). During bursts of action potentials, the action potentials propagate to the distal apical dendrites better (Williams and Stuart, 1999), turning apical depression into potentiation (Letzkus et al., 2006). Therefore, bursts are not only important for coding, but for plasticity as well. L5 TTPC bursts were rare in our simulations, as the model is based on an early developmental stage (P14-16: juvenile rats) and burst firing only becomes prominent as the animals mature (Zhu, 2000). On the other hand, burst firing could probably be rescued with stronger top-down input. As the top-down input represents context/brain state and is thought to serve as an error/target signal for learning, it has to be highly specific (Makino, 2019). Top-down input from POm was shown to gate plasticity in L2/3 PCs *in vivo* via dis-inhibiting the distal dendrites (Gambino et al., 2014; Williams and Holtmaat, 2019). Although POm input was present in our model, we only used randomly distributed fibers to keep our setup simple in this first investigation. For a profound understanding of the role of bursts in apical plasticity, more learning/task-related studies with more precise top-down input would be needed in the future.

We presented here what we believe to be a new way of studying unsupervised learning and plasticity in the cortex by taking the diversity of cell types and morphologies into account, modeling connections as multi-synaptic, validating synapse anatomy and physiology, modeling synapses with bistable dynamics, and simulating the network in an *in vivo*-like state. On the other hand, building a model of this scale and detail required gathering and systematic integration of a large amount of data over several years (Markram et al., 2015; Chindemi et al., 2022; Reimann et al., 2022; Isbister et al., 2023). As the first of its kind, the work presented here did not exhaust all the additional understanding one could possibly gain from the high level of detail. To facilitate that process, we are open-sourcing our model alongside detailed instructions to launch simulations and analyze the results (see Data and code availability). As any other model, it has several assumptions (listed in Table 1) and limitations, which will ideally be improved upon iteratively and in a community-driven manner. Simulating the model requires a performant hardware and software infrastructure (e.g., we needed 16.5M core hours to run the simulation presented in this manuscript). With respect to the second part we are continuously improving the performance and efficiency of the simulator (Kumbhar et al., 2019).

Table 1: List of assumptions

1. As we combined the models of Isbister et al. (2023) and Chindemi et al. (2022) all assumptions therein are inherited. Of particular interest:
2. Extracellular recordings are assumed to have the same bias across layers and neuron populations. Furthermore it is assumed that different inhibitory subpopulations require the same amount of input compensation.
3. The extracellular magnesium concentration of 1 mM used *in vitro* is assumed to be representative of the *in vivo* level.
4. As the plasticity model of Chindemi et al. (2022) is based on $[Ca^{2+}]_i$, by using it we assumed that other factors, like metabotropic glutamate receptors, endocannabinoid release, and BDNF signaling are negligible for the network-level effects of plasticity that we investigated.
5. Spines are assumed to be separate biochemical compartments, i.e., $[Ca^{2+}]_i$ of the dendrites does not influence that of the synapses.
6. By detecting a single set of assemblies in the 10-minute-long plastic simulation, we assumed that assemblies are stable on that time scale.

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 fix point) described as:

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

where τ is the time constant of convergence, θ_d and θ_p are depression and potentiation thresholds, γ_d and γ_p are depression and potentiation rates and Θ is the Heaviside function. The dynamics of $[Ca^{2+}]_i$ in spines was modeled as:

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

where I^*_{NMDAR} and I_{VDCC} are calcium currents through NMDA receptors and VDCCs, η is the fraction of unbuffered calcium, F is the Faraday constant, X is the spine volume, $[Ca^{2+}]_i^{(0)}$ is the resting value of $[Ca^{2+}]_i$, and τ_{Ca} is the time constant of free (unbuffered) calcium clearance. I^*_{NMDAR} depends on the state of the Mg^{2+} block. This nonlinear voltage dependence is described with the Jahr and Stevens (1990) formalism, with parameters fitted to cortical recordings from Vargas-Caballero and Robinson (2003).

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., 2022)):

$$\frac{dCa^*}{dt} = -\frac{Ca^*}{\tau^*} + [Ca^{2+}]_i - [Ca^{2+}]_i^{(0)} \tag{3}$$

where τ^* is the time constant of the integrator. Updates in ρ were done based on 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. 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 (Egger et al., 1999; Chindemi et al., 2022) those connections were made non-plastic. For the remainder of cells, θ_d and θ_p were derived as follows:

$$\begin{bmatrix} \theta_d \\ \theta_p \end{bmatrix} = \begin{bmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{bmatrix} \times \begin{bmatrix} c_{pre} \\ c_{post} \end{bmatrix} \text{ or } \begin{bmatrix} b_{00} & b_{01} \\ b_{10} & b_{11} \end{bmatrix} \times \begin{bmatrix} c_{pre} \\ c_{post} \end{bmatrix} \tag{4}$$

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:

$$\frac{dU_{SE}}{dt} = \frac{\overline{U}_{SE} - U_{SE}}{\tau_{change}} \quad \text{where} \quad \overline{U}_{SE} = U_{SE}^{(d)} + \rho\left(U_{SE}^{(p)} - U_{SE}^{(d)}\right) \tag{5}$$

$$\frac{d\hat{g}_{AMPA}}{dt} = \frac{\overline{g}_{AMPA} - \hat{g}_{AMPA}}{\tau_{change}} \quad \text{where} \quad \overline{g}_{AMPA} = \hat{g}_{AMPA}^{(d)} + \rho\left(\hat{g}_{AMPA}^{(p)} - \hat{g}_{AMPA}^{(d)}\right) \tag{6}$$

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. All values (fixed and optimized alike) are listed in Chindemi et al. (2022). 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. Rough time scales are as follows: $[Ca^{2+}]_i$ evolves at the timescale of tens of ms, Ca^* on the hundreds of ms, while changes in ρ are converted to changes in U_{SE} and $\hat{g}_{AMPA}$ in seconds.

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). The spontaneous state used in the article is characterized by the parameters: $[Ca^{2+}]_o = 1.05\,\text{mM}$ (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. (2023), but a minimal description, highlighting the changes applied in this study, can be found below. First, the flat map location 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., 2023) and half of that for POm.

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 (Hines and Carnevale, 1997; Kumbhar et al., 2019; 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 Chindemi et al. (2022) with classical plasticity rules, the 36,573,737 excitatory spikes from the 10-minute-long simulation were evaluated with pair-based STDP rules (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000). Synaptic weights evolved as follows under the STDP rule:

$$\Delta w_+ = A_+ \exp\left(-\frac{\Delta t}{\tau_+}\right) \text{ at } t_{post} \text{ if } t_{pre} < t_{post} \tag{7}$$

$$\Delta w_- = A_- \exp\left(\frac{\Delta t}{\tau_-}\right) \text{ at } t_{pre} \text{ if } t_{pre} > t_{post} \tag{8}$$

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 $\tau_\pm = 20$ ms. The STDP rule was implemented in `Brian2` (Stimberg et al., 2019).

4.5 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. (2023), 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 95^{th} percentile of the standard deviation of 100 shuffled controls. These shuffled controls were less strict than in Ecker et al. (2023). 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 (Montijn et al., 2016; Pérez-Ortega et al., 2021). 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. In Ecker et al. (2023), the one with the lowest Davis-Bouldin index was chosen, which maximizes the similarity within elements of the cluster while maximizing the between cluster similarity (Davies and Bouldin, 1979). For assemblies

detected over the 10 minutes-long plastic simulation, this optimal value was overwritten, to have at least one pattern-specific assembly for all ten patterns. For the assemblies detected over the 2 minutes-long non-plastic simulation, the optimal value was chosen, to avoid biasing the before vs. after assembly comparisons. 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 95^{th} percentile of correlations of shuffled controls (1000 controls with spikes of individual cells shifted by any amount as above; Montijn et al., 2016; Herzog et al., 2021). 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. Assemblies of neurons were compared using their Jaccard distances. The `assemblyfire` package, developed for Ecker et al. (2023) is publicly available on GitHub.

4.6 Determination of consensus assemblies

Consensus assemblies, resulting from the hierarchical clustering of the Jaccard similarity matrix of assemblies across repetitions of the same input stream, are fully described in (Ecker et al., 2023), but a minimal description of them can be found below. It was ensured that assemblies from the same repetition did not cluster together, first by setting their distances to twice the maximum, and second, by cutting the tree in a way that resulted in the lowest number of cluster in which two assemblies from the same repetition did not cluster together. Membership of neurons in these consensus assemblies was based on the fraction of assembly instances they were part of, normalized by a binomial control and thresholded. As shown in Ecker et al. (2023), consensus assemblies are similar to assemblies detected over the average spike trains across repetitions, but with the added benefit of the membership threshold. In rough terms, this threshold can be understood as follows: if a neuron was part of 80% of assembly instances that made up the consensus, then it was also a member of the consensus assembly.

In order to assess the functional connectivity of consensus assemblies before and after plasticity, the spike trains of their neurons across the five repetitions were first averaged and then binned (using the same 20 ms bins as above). Last, the Pearson correlation of all pairs of the preprocessed spike trains were calculated, and averaged across the population.

4.7 Calculation of 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 (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.

4.8 Synaptic clustering coefficient and likelihood of plastic changes in synapse clusters

Synaptic clustering coefficient (SCC), quantify the co-localization of synapses on the dendrites of a neuron from its presynaptic assembly with a single number, is fully described in (Ecker et al., 2023), but a minimal description of it can be found below. First, the nearest neighbor distance (along the dendrites) between all pairs of synapses from the presynaptic assembly were

computed and averaged (*mean nnd*). Second, 20 controls were generated by always selecting the same number of random presynaptic E cells from the circuit and *mean nnd*s of the control populations were calculated. Last, *SCC* was defined as the negative z-score of assembly *mean nnd* with respect to the distribution of control *mean nnd*s. *SCC* is thus a parameter-free metric, centered at zero, and is positive for intersynaptic distances that are lower than expected (indicating clustering) and negative otherwise (indicating dispersion). Additionally, the significance of the clustering or dispersion of the synapse locations was determined with a two-tailed t-test of assembly *mean nnd* against the 20 random samples with an alpha level of 0.05. *SCC* was implemented using `NeuroM` and `ConnectomeUtilities`.

Synapse clusters were also detected based on synapse neighbour distances. In order to be part of a spatial 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 similar to Druckmann et al. (2014). The distribution of synapse neighbour distances of the ten selected synapses were compared with a Poisson model (assuming exponentially distributed inter-synapse distances) based on all (same branch) synapse neighbour distances on the given neuron. Clusters were merged in a subsequent step, thus synapse clusters with more than ten synapses, spanning more than $20\,\mu$ms were also feasible. As plastic changes in synapse clusters were only analyzed for a small subpopulation of assemblies (ten L5 PCs per assembly), *SCC* was used to select subpopulations with high probability of finding synapse clusters. To this end, assembly neurons with positive, significant *SSC* values with respect to an assembly (either the same assembly for within-assembly analysis, or other ones for analysing cross-assembly interactions) were selected, and the ones with the ten highest assembly indegree (with respect to the same assembly) selected (see Ecker et al., 2023 for the same selection method). 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:

$$\frac{P(changed \,|\, category) - P(changed)}{P(changed \,|\, category) + P(changed)} \tag{9}$$

where *changed* was split to be either potentiated or depressed.

4.9 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 using the `connalysis` package, based on a fast C++ implementation from `flagsercount` (Lütgehetmann et al., 2020).

4.10 Conditional mutual information

To address the independence of the probability of change of a connection given its pre- and postsynaptic pattern-indegree three random variables over the set of all connections in the network were considered. The discrete random variables X and Y, denoting pattern-indegree of the pre- and postsynaptic cells respectively, and the binary random variable Ch, which is 1 if the connection changed and 0 otherwise. The conditional mutual information $I(X, Y|\mathrm{Ch})$, which gives the expected value of the mutual information between X and Y given that the connection changes or not (Gray, 2011) was computed as follows:

$$I(X, Y|\mathrm{Ch}) = \sum_{\substack{x,y\in\mathcal{P}_{in}\\ c\in\{0,1\}}} p(x, y|c)\log\frac{p(x, y|c)}{p(x|c)p(y|c)} \tag{10}$$

where $\mathcal{P}_{in}$ is the set of all possible values of pattern-indegree, $p(x, y)$ denotes joint probability and $p(x|z)$ denotes conditional probability. To determine the significance of this value, it was compared with the conditional mutual information of a random control, in which the values of the pattern-indegrees were independently shuffled.

4.11 MICrONS dataset

For the comparison of one of our findings with a rodent electron microscopy dataset, the 117 version of MICrONS (2021) was used. Synapses from sources other than one of the 60,048 classified neurons inside the bigger (*minnie65*) volume were discarded. The dataset is 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.

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 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., 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.

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., 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).

Crick, F. (1984). Memory and molecular turnover. *Nature*, 312:101.

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.

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.

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. (2023). Cortical cell assemblies and their underlying connectivity: an in silico study. *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.

Fusi, S., Drew, P. J., and Abbott, L. F. (2005). Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611.

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.

Gray, R. M. (2011). *Entropy and information theory*. Springer Science & Business Media.

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.

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.

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., Bolanos-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.

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*.

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).

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*.

Miller, J. E. K., Ayzenshtat, I., Carrillo-Reid, L., and Yuste, R. (2014). Visual stimuli recruit intrinsically generated cortical ensembles. *PNAS*, 111(38):E4053–E4061.

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.

Naud, R. and Sprekeler, H. (2018). Sparse bursts optimize information transmission in a multiplexed neural code. *PNAS*, 115(27):E6329–E6338.

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., 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., Marie, H., O'Donnell, C., and Veltz, R. (2022). A stochastic model of hippocampal synaptic plasticity with geometrical readout of enzyme dynamics. *bioRxiv*.

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.

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).

Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., and Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437).

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.

Trägenap, S., Whitney, D. E., Fitzpatrick, D., and Kaschube, M. (2022). Experience drives the development of novel, reliable cortical sensory representations from endogenously structured networks. *bioRxiv*.

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.

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

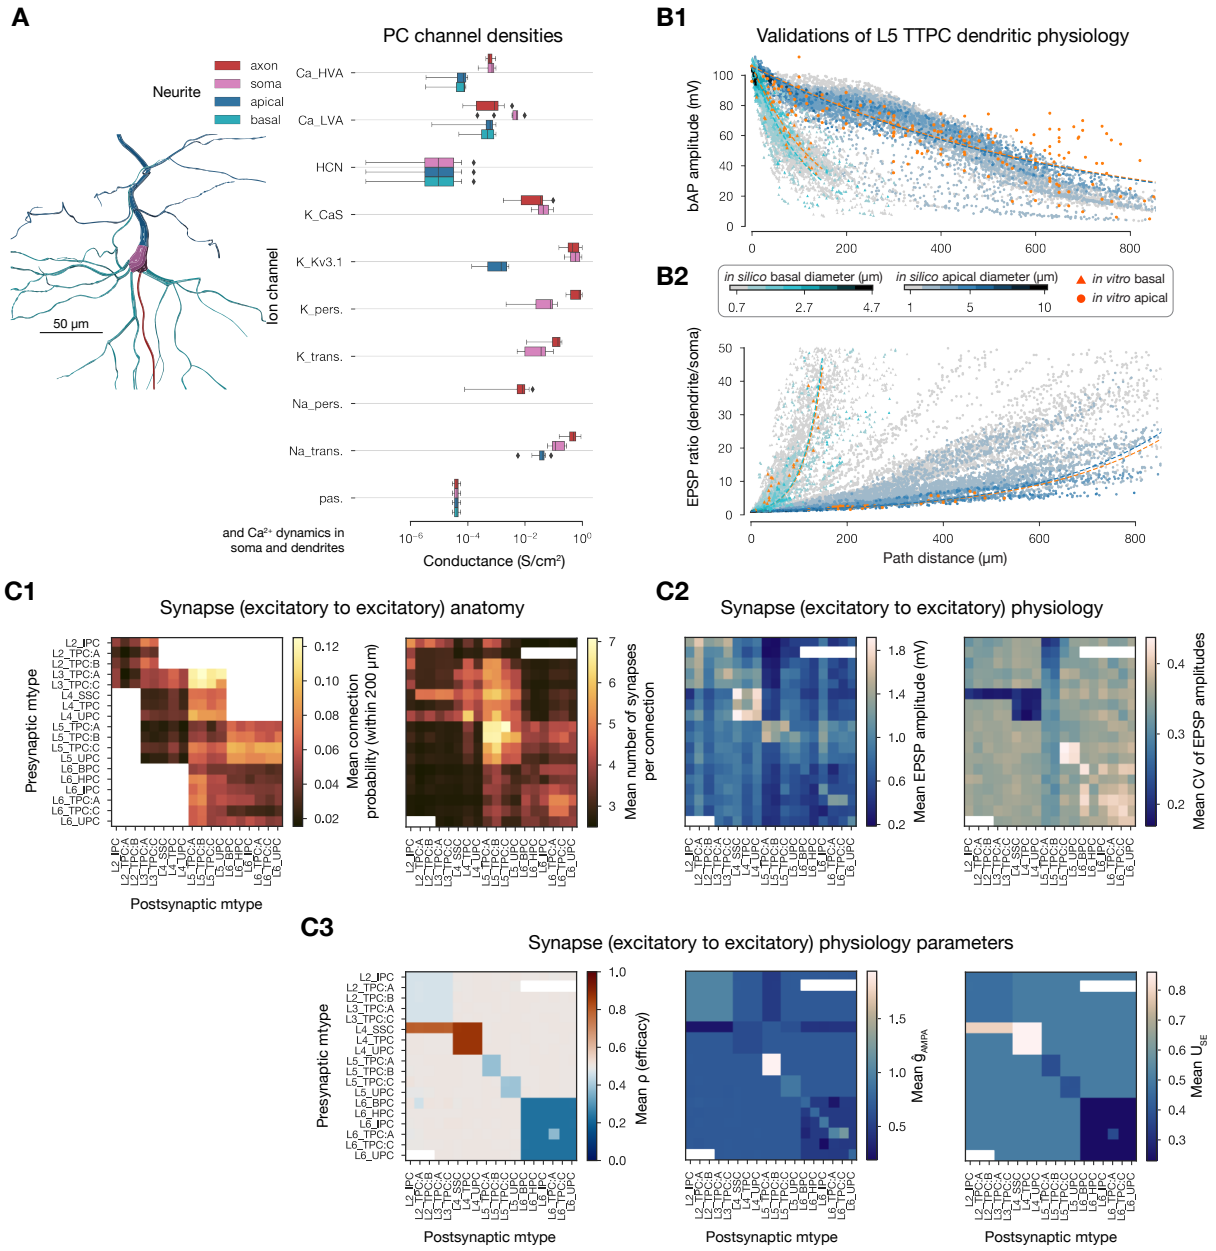

Figure S1: Physiology of excitatory cells and E to E connections. A: Distribution of ion-channel densities in the excitatory (cADpyr) electrical type (etype). B: Validation of dendritic physiology of the cADpyr e-type on L5 TTPC mtypes. **B1**: Validation of back-propagating action potential (bAP) amplitude 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). (A similar panel has also been shown Isbister et al., 2023). **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}.

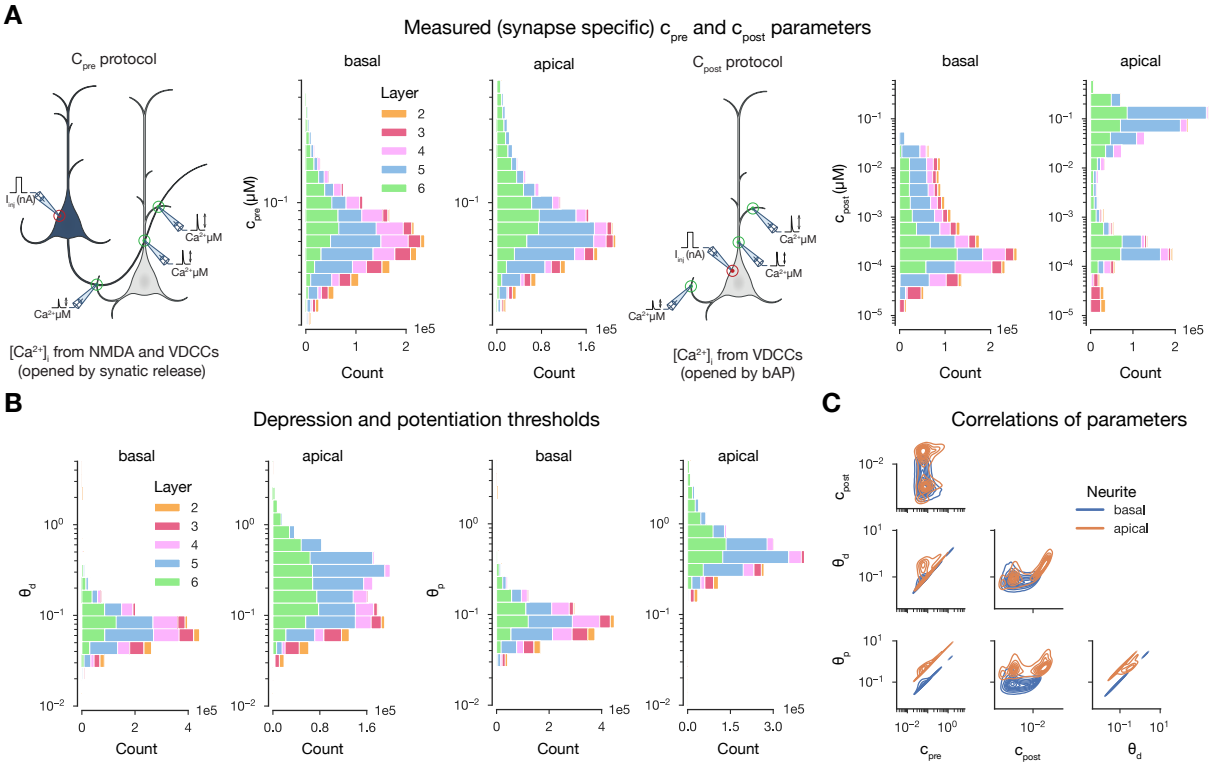

Figure S2: Synapse-specific parameters of the plasticity model. A: Layer- and neurite type-wise distribution of measured $[Ca^{2+}]_i$ peaks (used to derive parameters of the plasticity model shown in B). Synapses are grouped based on the soma location of the postsynaptic cell. 10% of all synapses are shown. Schematics on their lefts illustrate the measurement protocols. **B:** 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 A. **C**: Correlations of the parameters shown in A and B.

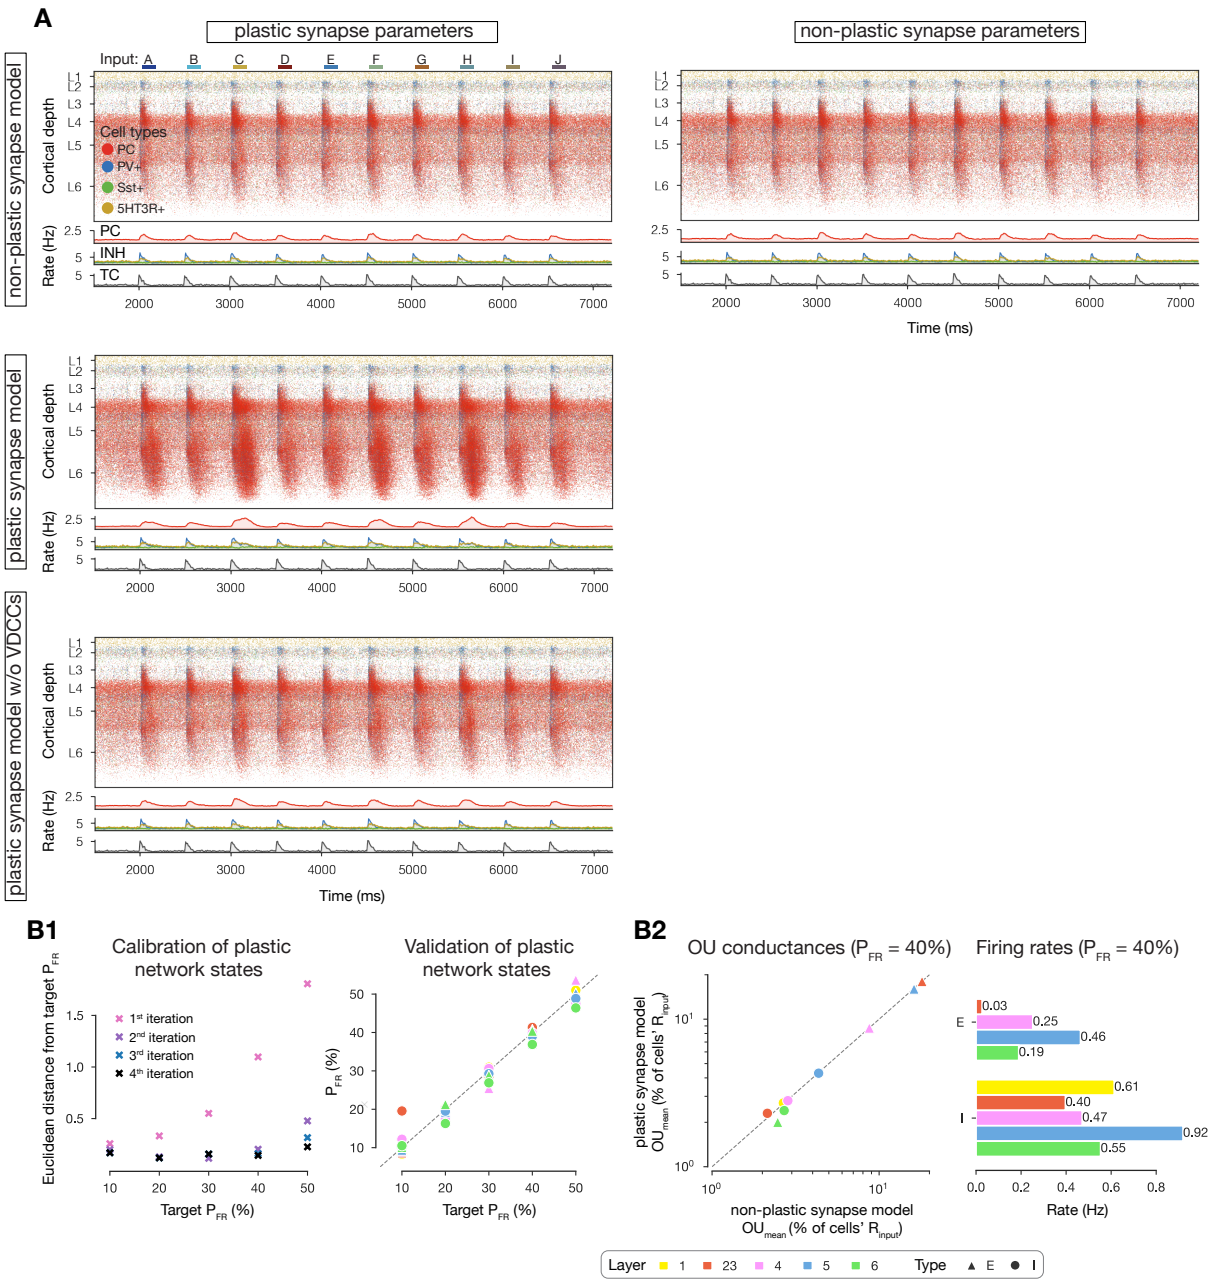

Figure S3: Calibration of the *in vivo*-like network state. A: Same as Figure 2B (i.e., raster plots of the microcircuit's activity) under different synapse setups. The microcircuit equipped with the plasticity model of Chindemi et al. (2022) only resembles that of the non-plastic network's of Isbister 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 (4^{th}) 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.

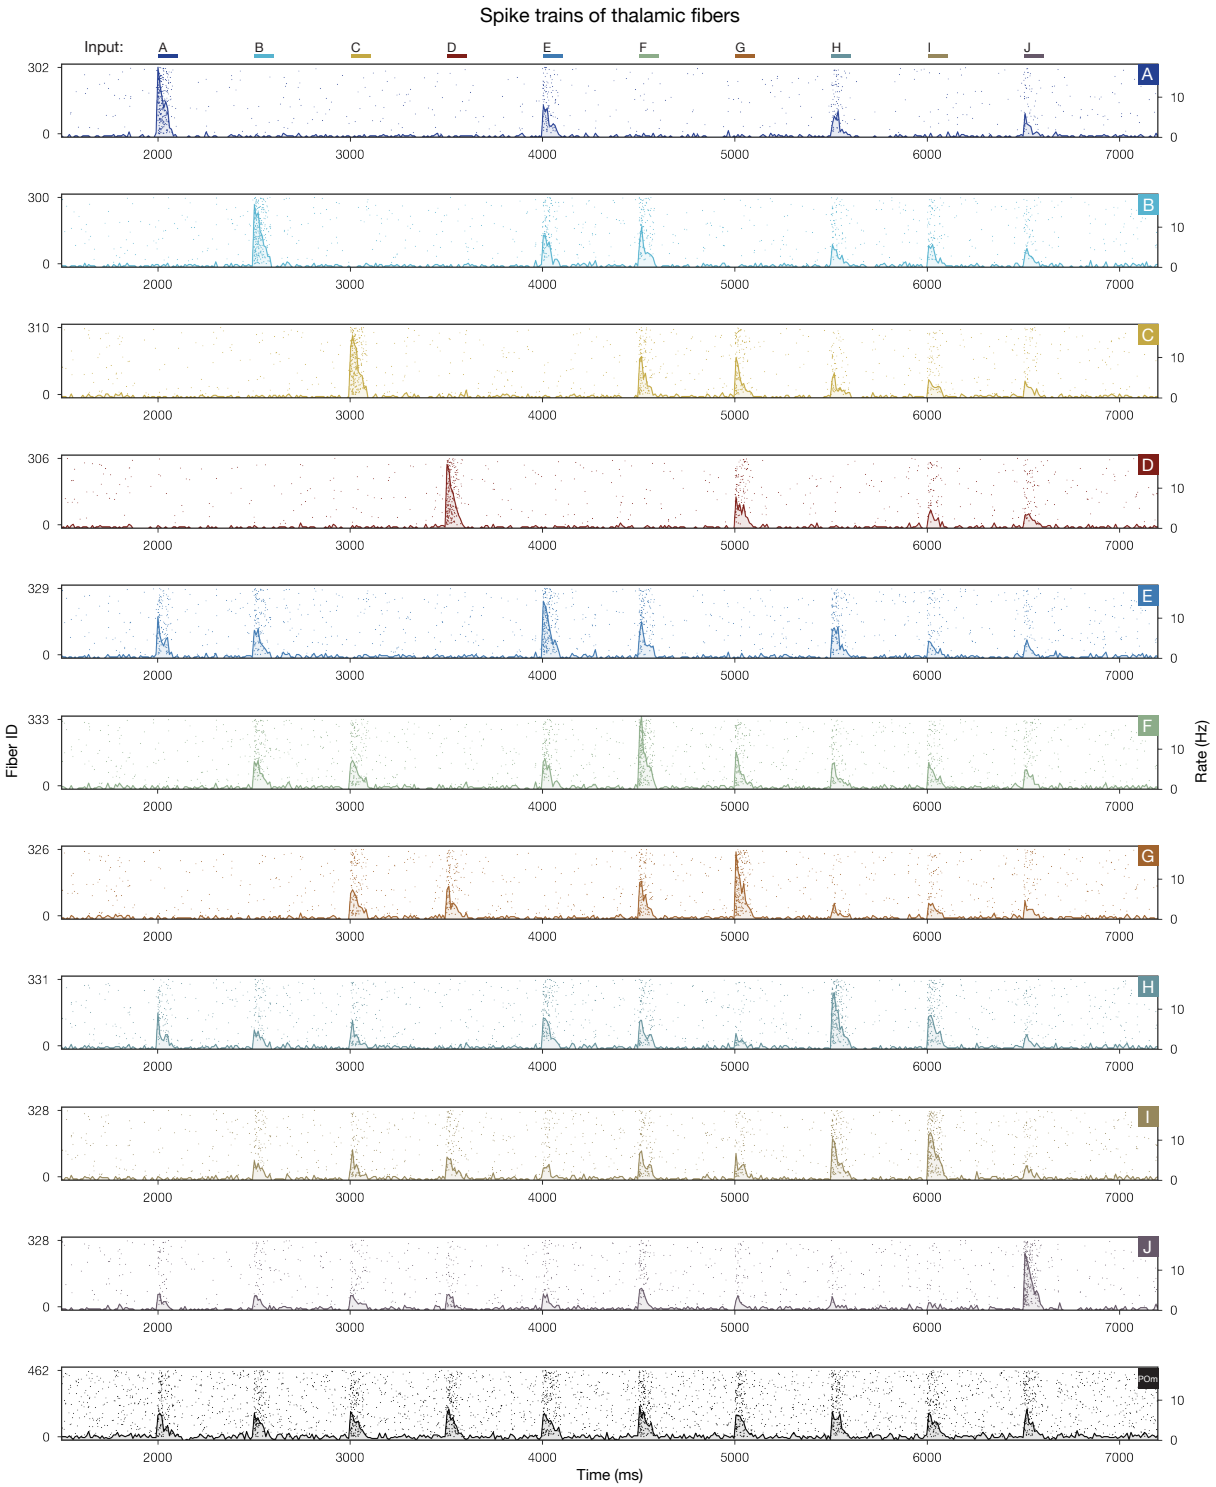

Figure S4: Activity of the thalamic fibers. Raster plots of VPM fibers forming each of the 10 input patterns (Figure 2A) for the stimulus stream (i.e., from pattern A at 2000 ms to pattern J at 6500 ms). Bottom row shows the same for non-specific POm fibers. (A similar panel has been shown in Ecker et al., 2023.)

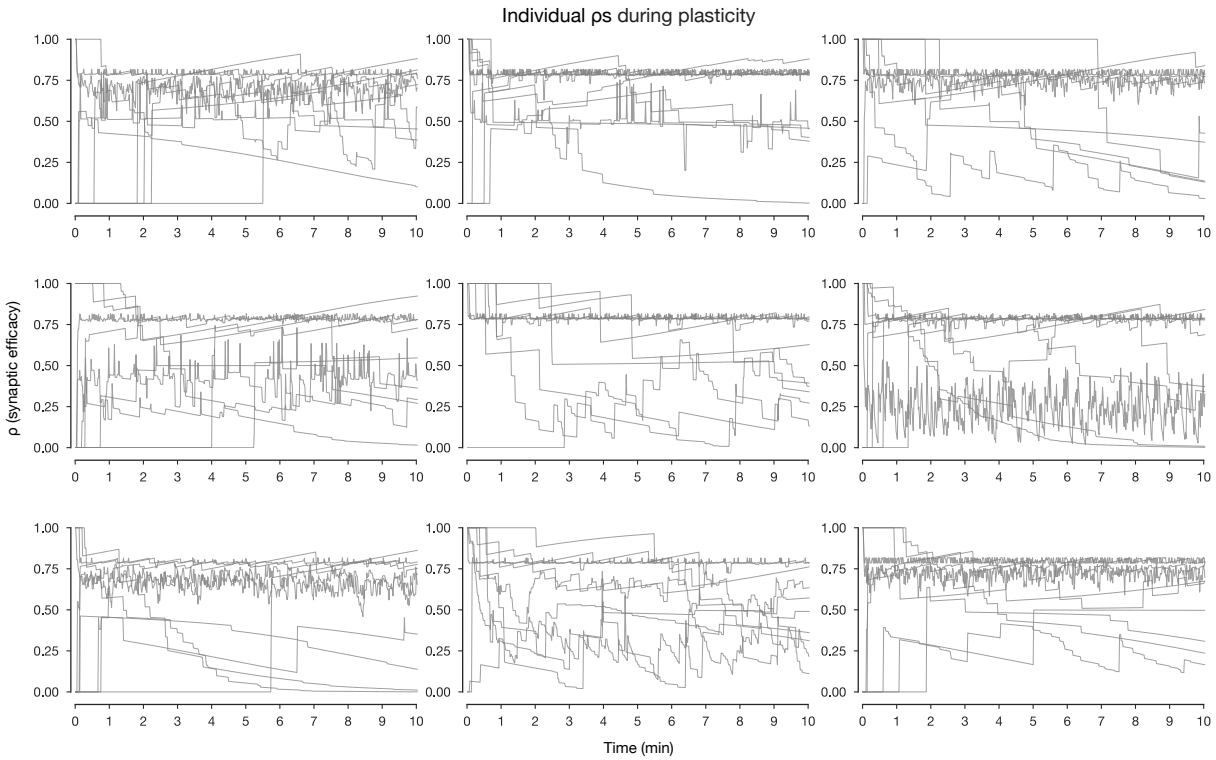

Figure S5: Changes in synaptic efficacy during plasticity. Individual ρ traces (10 per panel) during plasticity.

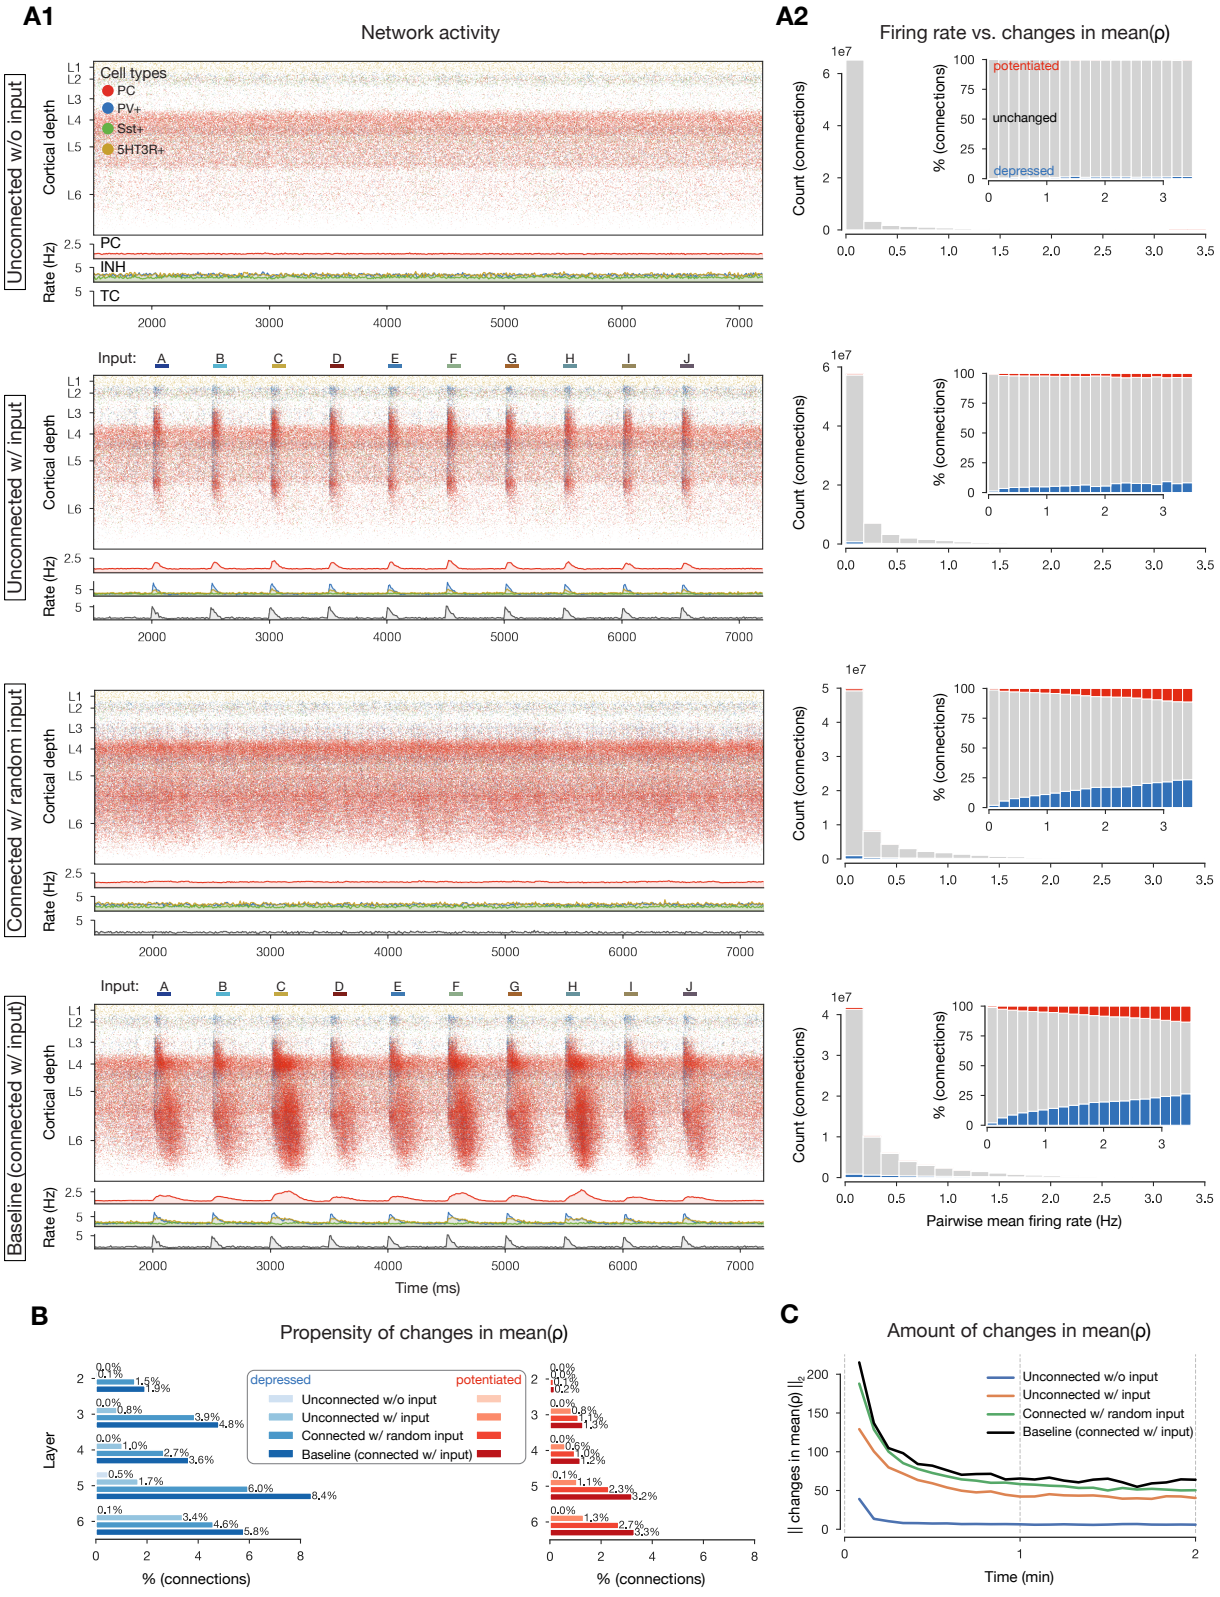

Figure S6: Changing connections in plastic control simulations. **A:** Same as Figure 2B and 3D2 (i.e., raster plots of the microcircuit's activity and plastic changes in mean ρ vs. firing rates under different conditions. The last row of A2 is not an exact replica of Figure 3E2 as these simulations were run for 2 minutes. **B:** Similar to Figure 3D1 (i.e., layer-wise propensity of changes in mean ρ) but split across conditions. **C:** Similar to Figure 2E (i.e., L2 norm of changes in mean ρ values for all conditions.)

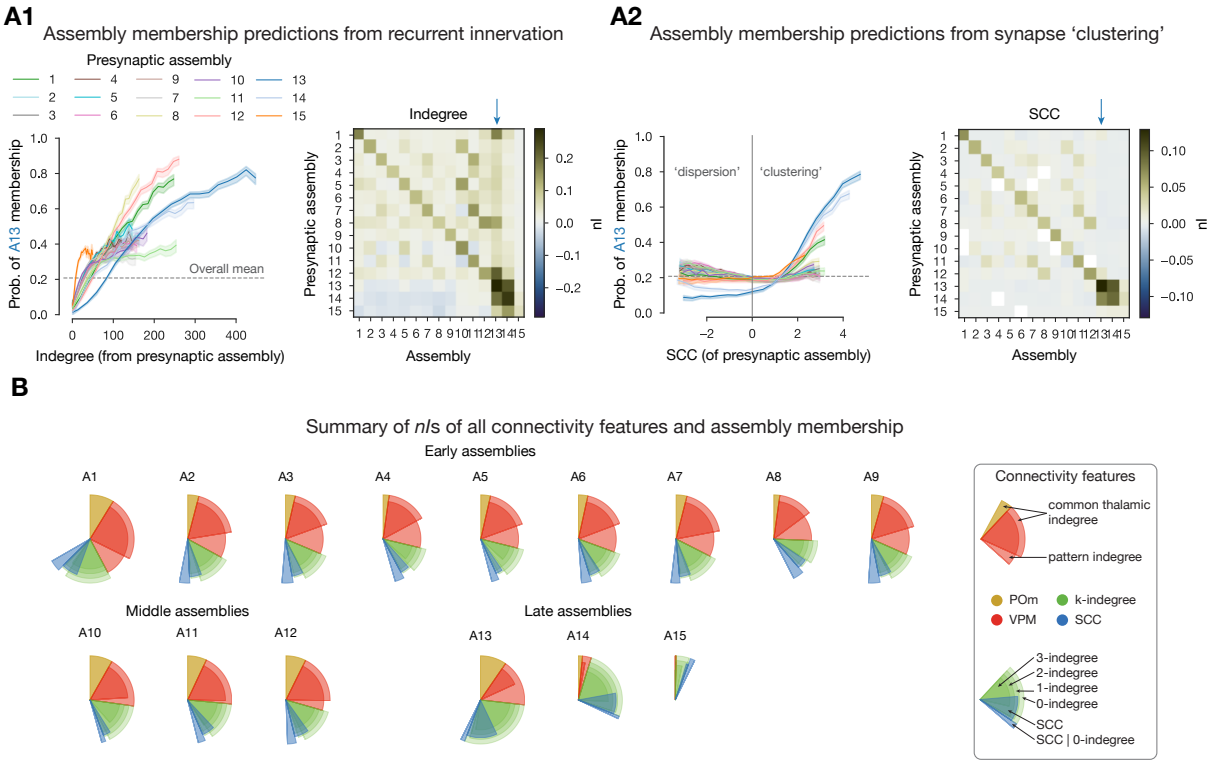

Figure S7: Connectivity features underlying cell assembly membership. A: Selected connectivity features of assemblies. **A1:** Left: Probability of membership in an exemplary assembly (A13) against assembly-indegree with respect to all assemblies. Solid lines indicate the mean and the shaded areas indicate 95% confidence interval. Right: nI (normalized mutual information, see Ecker et al., 2023) of assembly-indegree and assembly membership (blue arrow indicates postsynaptic assembly A13, shown in detail on its left). **A2:** Probability of membership in the same exemplary assembly against synapse clustering coefficient (SCC, see Methods and Ecker et al., 2023) with respect to all assemblies; nI of SCC and assembly membership. White boxes indicate non-significant nI. **B:** Summary of within-assembly interactions (diagonals of nI matrices) for all connectivity features considered in Ecker et al. (2023). (Similar panels have been shown in Ecker et al., 2023.)

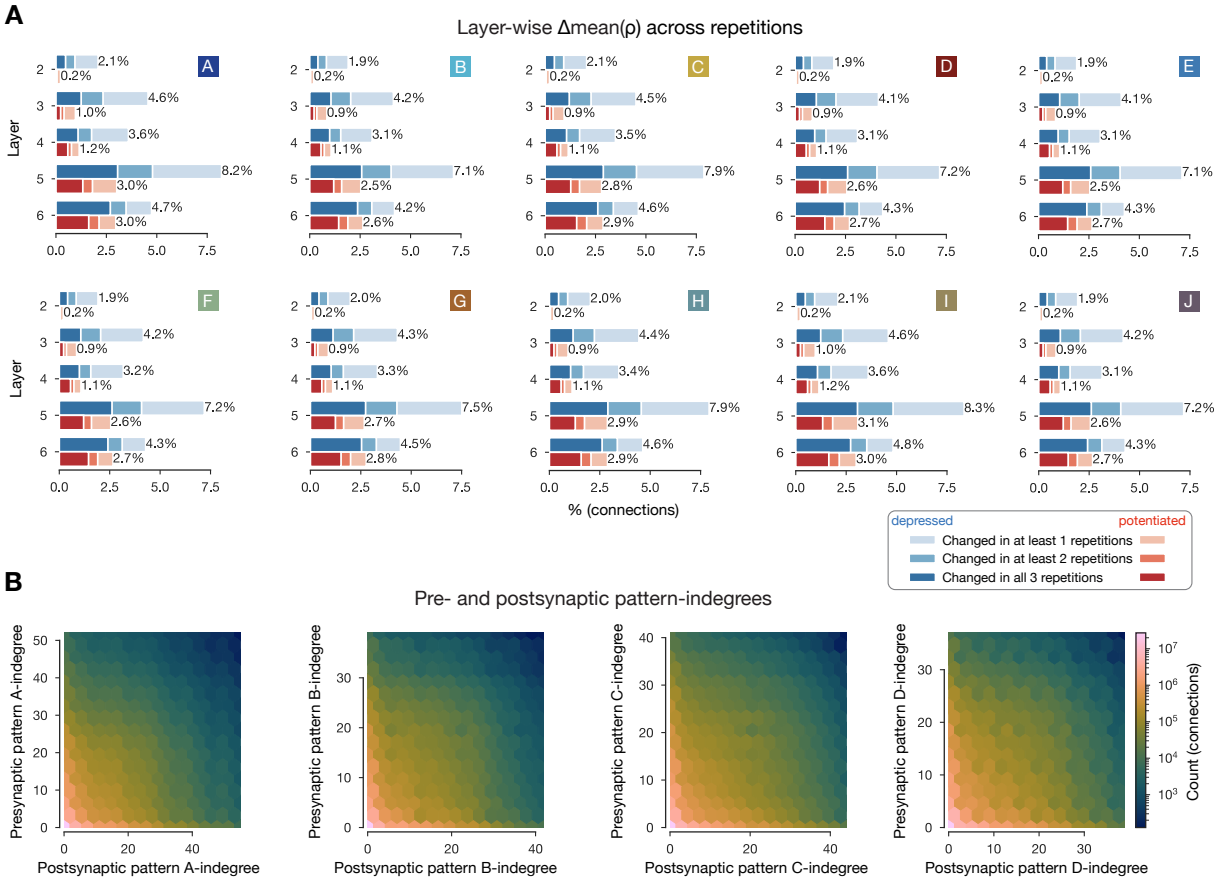

Figure S8: **A:** Layer-wise propensity of changes per single pattern. As on Figure 3D1, layer corresponds to the soma location of the postsynaptic cells. **B:** Pattern-indegree of pre- and postsynaptic cells (for the four base patterns only).

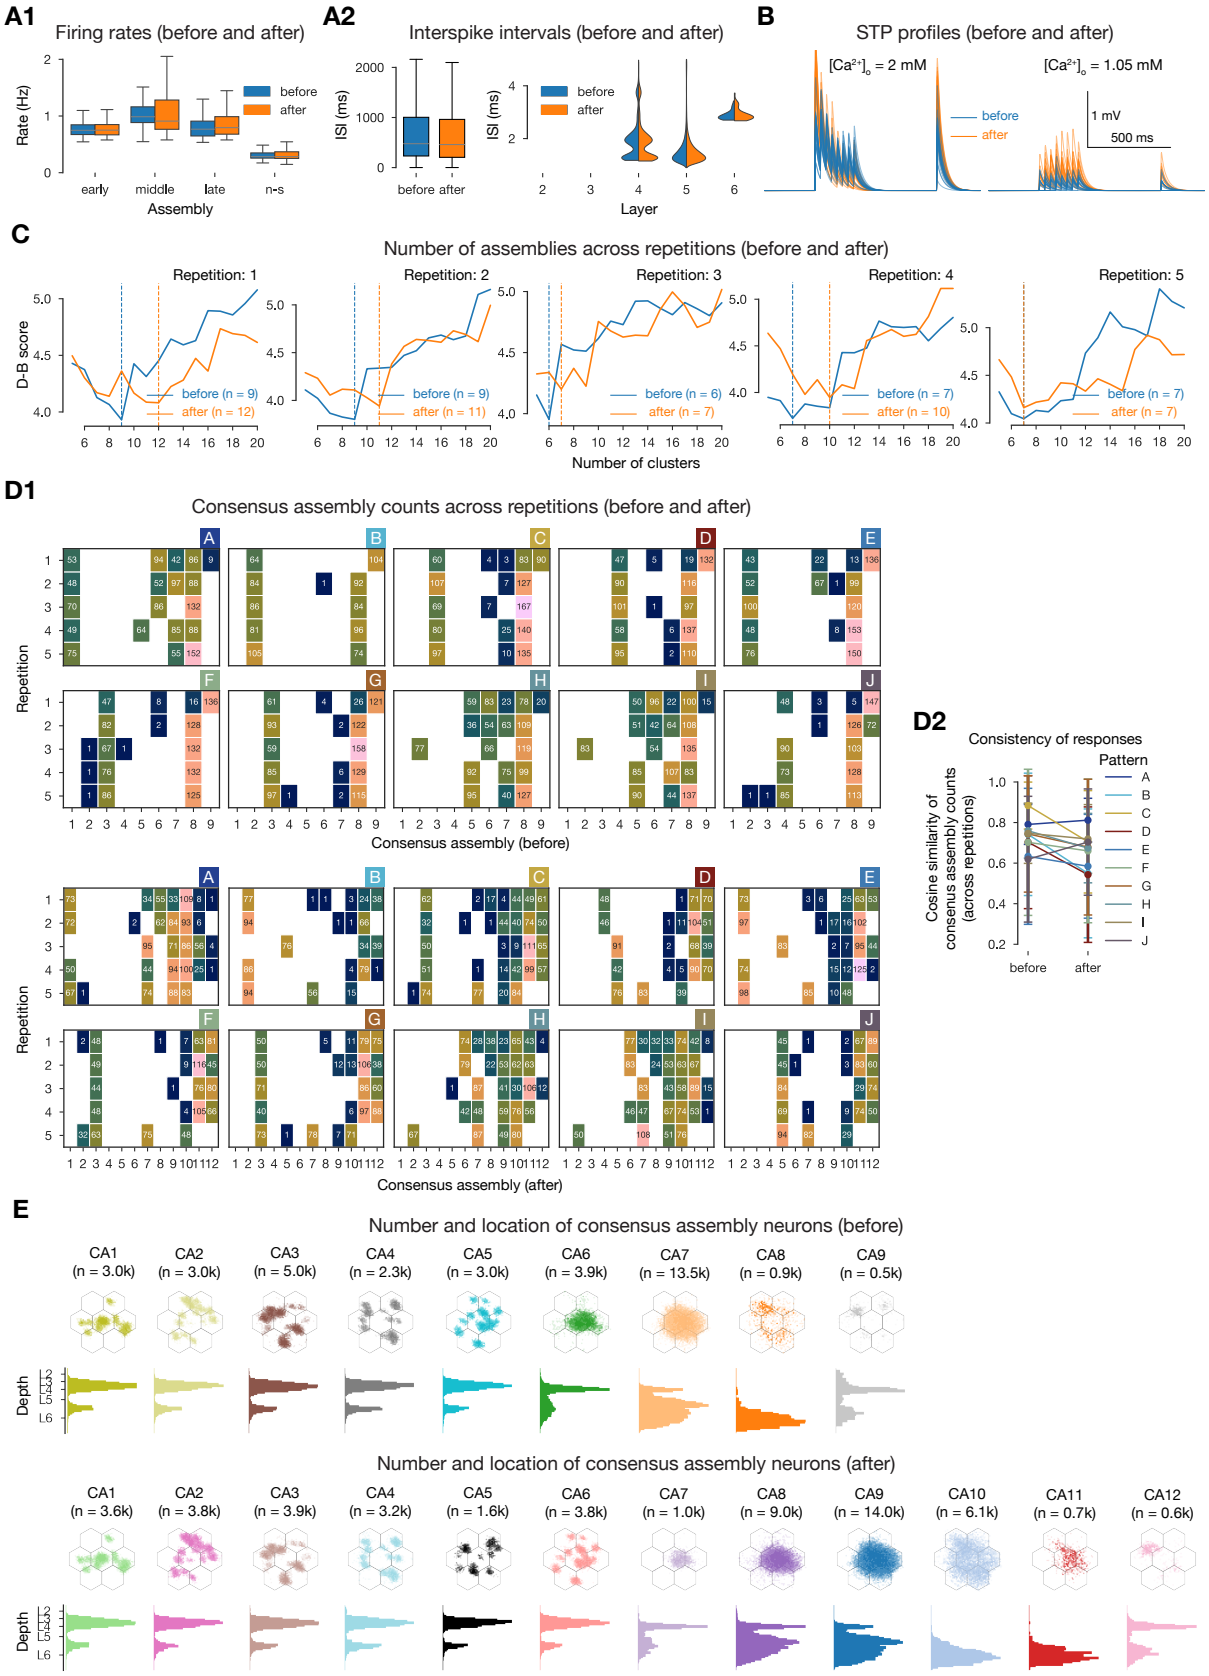

Figure S9: Comparison of cell assemblies before and after plasticity. A1: Firing rates before and after plasticity grouped by assemblies (n-s stands for non-significant time bins). (Caption continues on the next page.)

A2 Interspike interval (ISI) distribution (of all excitatory spikes) before and after plasticity on the left. On the right: Zoom in to low ISIs (≤ 5) ms split by layer. **B:** STP profiles before and after plasticity. At *in vitro* $[Ca^{2+}]_o$ on the left, and *in vivo* on the right. Thin lines represent the 20 individual repetitions, while the thicker ones their means. **C:** Davis-Bouldin index (see Ecker et al., 2023 and Methods) of different number of assemblies before and after plasticity across repetitions. (The index is to be minimized to achieve optimal number of clusters.) **D1:** Number of times a consensus assembly is active over repetitions before and after plasticity. E.g. the first rows per patterns are the counts of colored boxes from Figure 7C2 and D2. This representation can be used to judge the grouping of assemblies (see D2), and also for calculating their normalized Euclidean distance (see Figure 7G2). **D2:** Cosine similarity of rows of consensus assembly matrices (split by patterns before vs. after plasticity). **E:** Similar to Figure 4B (i.e., number and location of consensus assembly neurons before and after plasticity.)