Short-term plasticity of EI balance at single neurons can detect pattern transitions

Aditya Asopa, Upinder Singh Bhalla*,

National Centre for Biological Sciences, Tata Institute of Fundamental
Research, Bangalore, India

*bhalla@ncbs.res.in

*To whom correspondence should be addressed

Abstract

Sensory input and internal context converge onto the hippocampus as spatio-temporal patterns of activity. Transitions in these input patterns are frequently salient, yet the CA1 pyramidal neurons operate under conditions of divisive normalisation of summed patterned input by excitatory-inhibitory (EI) balance which suppresses most responses. We characterized the role of short-term potentiation (STP) in mediating change detection in the mouse hippocampus CA3-CA1 network using optogenetic patterned stimuli in CA3 while recording from CA1 pyramidal neurons. We parameterized STP and its effect on summation, and developed a multiscale model of network projections onto hundreds of E and I boutons each including stochastic signaling to mediate STP on the postsynaptic neuron. We show that STP modulates EI summation across patterns, and predicted and confirmed that single neurons can detect transitions in input patterns responses. Using the model we show that in feedforward networks, STP coupled with moderate sparsity of input due to input pattern sparsity and connections, strengthens rapid pattern change detection by single neurons.

Introduction

The balance between excitatory and inhibitory drives in neural networks has a wide range of computational functions. These include increased storage capacity (Kanamaru et al., 2023), optimal coding (Denève and Machens, 2016), fast temporal encoding and gating (Kremkow et al., 2010; van Vreeswijk and Sompolinsky, 1996; Vogels and Abbott, 2009), feature selectivity (Wilent et al., 2005), better noise robustness and stability (Rubin et al., 2017), higher dynamic range (Bhatia et al., 2019; Wehr and Zador, 2003), and stabilisation of ensembles (McKenzie, 2018; Roach et al., 2023). A recent study has shown the existence of tight temporal balance at the scale of arbitrary combinations of inputs at a single neuron level (termed precise balance) (Bhatia et al., 2019). Under a continuous but irregular firing regime (van Vreeswijk and Sompolinsky, 1996), balanced inputs help in efficient encoding by regulating firing rate irregularity (Miura et al., 2007). However, little is known about how EI balance itself evolves dynamically due to activity-driven plasticity and specifically, how single neurons can use STP changes in EI balance to implement functions of input gating, normalisation, and pattern change detection.

Continuously active neurons show history-dependence through short-term synaptic plasticity (STP) in which neurotransmitter release varies upon receiving successive action potentials at the boutons (Fortune and Rose, 2001; Regehr, 2012; Tsodyks and Markram, 1997; Zucker and Regehr, 2002). STP emerges from the competing time constants of pre-synaptic processes of vesicle depletion, calcium influx, calcium build-up, and neurotransmitter reuptake. Together, these phenomena create a temporal filter between activity in presynaptic and postsynaptic neurons or networks (Fortune and Rose, 2001; Galarreta and Hestrin, 1998; Klyachko and Stevens, 2006), resulting in a gain-control of external inputs (Abbott et al., 1997). Excitatory and inhibitory synapses show a range of short-term modifications in their synaptic weights in this manner (Bartley and Dobrunz, 2015; Klyachko and Stevens, 2006; Sun et al., 2018). Thus, the

resultant of E and I converging on a cell changes with time, suggesting that single neurons may extract interesting features of spatiotemporal inputs through STP (Asopa and Bhalla 2023).

In the current study, we show that short-term plasticity dynamics of E-I balance are tightly coupled for trains of inputs in the frequency range from 20 to 50 Hz. Using in vitro whole cell patch clamp experiments in the feedforward excitatory-inhibitory (FFEI) network of CA3-CA1 in the mouse hippocampus, we probed how the precise balance (Bhatia et al., 2019) between excitation and inhibition evolves over input trains of various frequencies optogenetically generated at CA3. We incorporated this STP data into mass-action stochastic models of presynaptic release for E and I synapses, and embedded a CA1 neuron receiving hundreds of such synapses into an abstract FFEI network model of CA3. This model predicted that the postsynaptic CA1 neuron could detect pattern changes, which we then confirmed experimentally.

Results

Optogenetic stimulation to create input combinations in the CA3-CA1 network

To test the dependence of the CA1 pyramidal cell's response on the short-term history of inputs in the feed-forward excitatory-inhibitory network, we carried out optogenetics experiments in the hippocampal CA3-CA1 FFEI network. We obtained whole-cell patch-clamp recordings from CA1 pyramidal cells (PCs) while stimulating channelrhodopsin (ChR2) expressing CA3 pyramidal cells using grid-based optical patterns (Fig 1A-B, see Methods). A pattern consisted of either 1, 5, or 15 spots (or squares) (spot size 13 x 7 μm, power 14.5 μW/spot) of blue light (470 nm) illuminated the CA3 layer in a series of eight pulses at a fixed frequency of 20, 30, 40, and 50 Hz (Fig 1D). A probe pulse preceded this train of 8 pulses to record the baseline postsynaptic response 300 ms before the onset of the train. We used this protocol to measure the dynamics of E-I balance in a three-dimensional stimulus space consisting of CA3 activation size (number of squares), activation frequency (stimulation frequency), and short-term plasticity dependence (pulse number) in the train (Fig1C).

To monitor the strength of the total resultant optogenetic activation of the CA3 layer, we used an extracellular field electrode in the stratum radiatum (Fig1A, methods). The field response correlated well with the intracellular optical depolarization (Fig 1E-G), allowing us to use the field response as a readout for stimulation strength in CA3 for further experiments. The field response scaled with the size of the pattern with larger patterns showing higher extracellular response (Fig 1F) and showed a correlation with single-cell optical depolarization recorded from a CA3 cell kept in the centre of the projection frame (Fig1G). The field response showed a slight decrease over the course of the pulse train of approximately 2% per pulse (regression fit slope=0.02, r^2=0.05). We attribute this to ChR2 desensitisation.

Fig 1: Experiment design and basic response properties. **A.** (Top) Network schematic of hippocampal CA3-IN-CA1 network. (Bottom) A transverse hippocampal section showing channelrhodopsin expression in orange (tdTomato), with the stimulation grid drawn to scale overlaid on the CA3 network. An extracellular (field) electrode was used to record the total optogenetic excitation of the CA3 layer (black arrowhead), and CA1 recordings were made from individual pyramidal cells with a whole-cell patch clamp electrode (white arrowhead). (scale bar = 500μm) **B.** (Top) A DIC image of the CA3 cell layer with a few spots of a stimulation grid overlaid, drawn to scale. (Black dotted lines mark the outlines of a few CA3 cells). (Bottom) A schematic of the CA3-IN-CA1 layer showing the recruitment of CA3 pyramidal cells with three different patterns. The downstream CA1 pyramidal cell receives direct monosynaptic excitation from the activated population and disynaptic inhibition via a heterogeneous, shared population of CA1 interneurons. **C.** The main stimulation protocol consists of a pattern containing 1, 5, or 15 squares illuminated in the form of a pulse train delivered at frequencies 20, 30, 40, and 50 Hz, creating a 3-dimensional probing of the FFEI circuit. **D.** An example recording in the current clamp mode showing the postsynaptic potential of a patched CA1 cell (black) in response to a triggered 15-square pattern (red) used for pulsed optical stimulation in CA3 (cyan), estimated using field potential (orange). **E.** Extracellular recording of CA3 layer activation for 1,5,15 square patterns. **F.** Distribution of CA3 field response to first pulse for number of squares per pattern (Kruskal-Wallis test). **G.** Correlation of single neuron optical current recorded from patched CA3 cells with the corresponding field response for different pattern size values. **H.** Post synaptic potentials recorded in a CA1 pyramidal cell in current clamp for different pattern sizes. **I.** Distribution of PSPs across all recorded CA1 cells (n=16, $p < 0.001$, Kruskal-Wallis test). **J.** A slight decrease in the strength of field response was seen along the pulse train. (linear regression fit with slope 2%, $r^2=0.05$, bars indicate 95% CI). **K.** Postsynaptic currents (PSC) recorded from CA1 cells in voltage clamp show the proportional relationship between excitatory and inhibitory PSCs for each pattern size. **L.** Similar to F and I, the distribution of PSC amplitude across the number of squares. **M.** A sample recording of excitatory (pink) and inhibitory (teal) currents recorded from a CA1 cell in the protocol described in panel D for a 15 square pattern. Blue traces in E,H,K represent the frame onset. Horizontal lines inside the violins represent quartiles.

To understand the short-term dynamics of excitation and inhibition, we recorded the CA1 responses across eight pulses along a train of four discrete frequencies (20, 30, 40, and 50Hz). Additionally, to test the effects of sparseness and strength of input, we recruited different fractions of CA3 using three pattern sizes (1, 5 or 15 spots/pattern). A pattern with a larger number of squares elicited a larger response in CA1 (Fig 1H-I). Next, we performed whole-cell patch clamp recordings in voltage-clamp (methods) to record the excitatory and inhibitory post-synaptic currents in CA1PCs. As reported previously (Bhatia et al., 2019), different activation patterns in CA3 produced proportional E and I post-synaptic currents in the recorded neuron (Fig 1K). Similar to post-synaptic depolarisations (Fig 1I), the postsynaptic currents showed a mutually proportional relationship with the pattern size (Fig 1L). Finally, we recorded excitatory and inhibitory post-synaptic currents for the discrete frequency pulse trains to measure the short-term plasticity effects on the strength of excitation and inhibition received by the CA1 pyramidal cells (Fig 1M). We observed that for each pattern and frequency of the train, the EPSC and IPSC followed a similar trend.

We concluded that grid-based optical stimulation could create input combinations at various stimulation frequencies to probe the monosynaptic excitation and disynaptic inhibition arriving at CA1 pyramidal cells.

Excitation and Inhibition both undergo depression

In the balanced feed-forward EI networks, inhibition follows excitation with a slight delay (Bhatia et al., 2019). However, the three synapses in the circuit, namely CA3-CA1 (excitation), CA3-CA1INs (excitation), and CA1INs-CA1PCs (inhibition), show significant history-dependent modulation of synaptic weights (Bartley and Dobrunz, 2015; Klyachko and Stevens, 2006). To understand how this precise EI balance at the level of CA1PCs dynamically evolves due to short-term plasticity processes, we first measured postsynaptic potentials and the probability of escape across a range of frequencies and input strengths. Next, we compared the strength of excitatory and inhibitory postsynaptic currents along those dimensions.

We varied the stimulus size and frequency while measuring the postsynaptic responses (PSP) elicited by each pulse (Fig 2A). All the PSPs of a train were normalised to the

probe pulse given 300 ms preceding the train onset. The PSP varied significantly on all the tested three axes of parameters: pattern size ($p<0.001$), stimulus frequency ($p<0.001$), and pulse index ($p<0.001$) (ANOVA). The response also showed a marked difference along the pulse train, where the summed EPSP in the first half was potentiated significantly larger than that in the second. We tested this by comparing the pulse responses divided between the two halves of the train against a shuffled pulse order ($p<0.001$, Wilcoxon rank test). Thus, STP adds significant history dependence to the input size and frequency.

While PSPs are a measure of depolarisation of the receiving CA1 pyramidal cells, neuronal firing is the readout of escape from EI balance. For the stimulus strengths in our experiments, the E-I balance made it unlikely for the receiving cell to cross the threshold and fire an action potential. Out of more than 2000 trial sweeps across different patched cells, pattern parameters, and pulse train frequencies, an escape led to spiking in only 7% of trials ($n=154/2201$). To estimate the effect of STP on escape, we measured the likelihood of escape from EI balance across the three dimensions (Fig 2B) by counting the fraction of total spiking trials that elicited for each combination of parameters. Similar to the PSP response amplitudes, spike likelihood was higher in the first half of the pulse train ($p < 0.001$, ANOVA), showing a tight control of escape dependent on STP. Similarly, across the frequency and number of squares, the spike likelihood varied significantly across the frequencies and number of squares (stimulus frequency $p<<0.001$, number of squares, pulse index $p << 0.001$, ANOVA).

Using voltage-clamp recordings, we separated the excitatory and inhibitory postsynaptic currents by holding the CA1 neuron at the reversal potentials for the GABAA receptor (-70 mV) and AMPA receptor (0 mV), respectively. We thus measured E and I currents across the three parameters. We used the postsynaptic response to the probe pulse as a kernel that was then fit to the responses for the rest of the train. This allowed us to estimate the postsynaptic currents received in every pulse, allowing for the buildup across multiple pulses. The EPSC showed a trend of early potentiation followed by depression (Fig2 Ei-Eii), while the inhibition underwent depression from the start (Fig2 Eiii-Eiv). This effect was significant across all the tested dimensions ($p < 10^{-6}$, ANOVA).

Thus, the balance briefly tilts in favour of E, after which both E and I undergo a proportional depression. As relative amplitudes and relative delays together form the components of the E-I balance (Bhatia et al., 2019), we next investigated how these two components trend across the stimulus dimensions.

Fig 2. (A) The trend of trial normalised excitatory postsynaptic potentials (across stimulation frequency and pulse index in the train for 5 square (Ai) and 15 square (Aii) patterns across all recorded cells (n=16). B. Spike probability as a function of the stimulus frequency, pulse index and number of squares in the current clamp cells showing higher likelihood in the first half of the train for five squares (top) and 15 square patterns (bottom). C. Post-synaptic responses were derived from kernel fits to valley to peak height for each response for both postsynaptic currents (voltage clamp, EPSC and IPSCs) and postsynaptic voltage (current clamp, not shown). D. Normalised EPSC and IPSC responses obtained for an example cell for a sample pattern showing consistent depression in inhibition and biphasic response in excitation. E. Both excitatory (Ei-Eii) and inhibitory (Eiii-Eiv) PSCs show short-term depression along the pulse train. However, this depression is much faster for inhibition than for excitation. Statistically, the short-term depression is significant ($p<0.001$) across the number of squares, pulse index, and excitation vs inhibition, and the input frequency (ANOVA).

Relative kinetics of excitation-inhibition and normalisation prevent escape from E-I balance

The E-I balance at the neuronal level emerges from the relative amplitude and time delays of the excitatory and inhibitory components. We next measured how the resultant E-I ratio and delays between onset as well as PSC peak depend on the STP. As before, we delivered pulse trains optically and monitored postsynaptic potentials. We observed short-term potentiation for the first two pulses (Fig 3A) in the train, followed by a steady depression due to the combined effects of STD in both excitation (Fig 2E) and inhibition (Fig 2E).

Interestingly, the stimulus onset timing showed an increase over the pulse train (Fig 3B), while the timing of the peak did not show a significant change (Fig 3C). The corresponding EPSC and IPSC responses both underwent depression as the pulse train progressed (Fig 3D). Additionally, the delay in onset of the PSCs increased over the pulse train for both EPSC and IPSCs (Fig 3E), although the timing of peaks did not change (Fig 3F).

Next, we calculated the E/I ratio across stimulus frequencies in the pulse trains. We found that the E/I ratio increased along the pulse train (Fig 3G). However, the delay between EPSC and IPSC responses for both the onset and peaks showed no significant change (Fig 3H-I).

As shown previously (Bhatia et al. 2019), an increase in input strength (by increasing the number of recruited cells) in CA3 produces increased excitation that is overcome both by increased inhibition and reduced delay between onsets of excitatory and inhibitory currents. This creates a strong normalising effect on the summation of responses at the receiving CA1 cell. In static summation, we observed sublinearity and divisive normalisation similar to the earlier work (Bhatia et al 2019). We next asked how these sublinearities evolve over the pulse train using the coefficient of sub-threshold divisive normalisation (γ)(Bhatia et al., 2019) as a metric. Gamma (γ) relates the

observed response of a CA1 cell to the expected response obtained by summing the inputs linearly in the following manner:

$$\text{Observed} = \frac{\gamma * \text{Expected}}{\gamma + \text{Expected}}$$

The smaller the gamma, the larger the deviation from the linear combination of inputs due to the divisive effects of inhibition. As a control, we also tested a linear fit, implying offsetting or subtractive effects of inhibition upon increasing the input strength (Fig. 3J). A larger slope, like a larger gamma, indicates a diminished influence of inhibition on the summation.

Using responses from the reference pulse, we replicated earlier observations showing divisive normalisation with a median gamma of 7.16 (95% CI = 4.76 - 10.2)(Fig 3Li) and a median slope of 0.57 (95% CI = 0.45 - 0.72). We found that both gamma and slope values show a rise and then a fall in a history-dependent manner (sample cell: Fig 3Ki-Kii, all cells: Fig 3Mi-Mii). In other words, the response becomes more linear around pulses 2 and 3 and subsequently becomes more sublinear. Across the measured frequencies, however, the normalisation does not significantly differ (Fig 3Mi, Mii, right bar).

Figure 3: A. Trend of normalised postsynaptic potentials in a pulse train received at CA1 pyramidal cells across stimulation frequencies. B. Response onset times measured from the start of light stimulation in CA3 to the start of postsynaptic response in CA1. C. Delay from light stimulation start to the peak of the response. D-F. Same as above for voltage clamp postsynaptic excitatory (red) and inhibitory (green) currents. G. E-I balance, measured as the ratio of E to I, showing a rise along the pulse train that is significantly high in the later half of the train (p<0.001, ANOVA) but not significantly different across stimulation frequencies. H-I. The differences between the onset delays and peak delays of excitation and inhibition do not show a significant trend across the frequencies or pulse train.

J. Scaling of observed responses in CA1 for patterns of different sizes (5, 7, 15 spots) as compared to the expected responses obtained by the sum of constituent single-spot responses, shown for a sample cell to illustrate the response normalisation that is fitted with a saturating (coefficient gamma, purple) and a linear function (coefficient slope, green). K. The trend of gamma (Ki) and slope (Kii) for the sample cell along the pulse train shows a biphasic response across all the frequencies with an inverse relation with pulse index. L. The distribution of gamma (Li) and slope (Lii) shows a significant left shift along the pulse train (only shown for the first and last pulse) (p<0.001, Kruskal-Wallis). M. A distribution of gamma and slope across the pulse index and stimulus frequencies, shown along their partial distributions on both axes (all cells).

Presynaptic release models are bounded by observations from pulse-train sequences.

Having characterised the dynamics of plasticity of excitatory and inhibitory synaptic currents over trains of pulses, we next used this data to develop models of presynaptic signalling. Our goal was two-fold: to better understand how STP mechanisms play out for patterned input in a network context; and to predict network mechanisms that might lead to escape from EI balance. Datasets with similar pulse trains have been used to obtain phenomenological models (e.g., (Barros-Zulaica et al., 2019)). Instead, we chose to build a chemical kinetics-based multi-step neurotransmitter vesicle release model similar to previously published models (Neher, 2015; Schneggenburger et al., 1999) (Figure 4A). This incorporated both calcium buffering, which is believed to underlie short-term potentiation, and depletion of docked vesicles, which is believed to underlie short-term depression. We implemented a reduced ball-and-stick cellular model having AMPA and GABA synapses on the dendrite. Like in the experiment, the model held the soma in a voltage clamp and delivered the various patterned stimuli at different frequencies to estimate EPSCs/IPSCs. The chemical-electrical model was implemented in MOOSE (Ray and Bhalla, 2008), and parameters were fit using the HOSS pipeline (Viswan et al., 2024).

We obtained fits across the range of frequencies for each individual cell, as illustrated in Fig 4. Each cell had somewhat distinct STP characteristics (Supplementary Figure 1). While the E fit applied directly to the CA3->CA1 pyramidal neuron glutamatergic synapse, the I fit was by necessity a composite of CA3->Interneuron and Interneuron->CA1. We made the further simplifying assumption based on the peak and onset times (Fig 3) that only one interneuron class was involved, most likely PV interneurons.

We used parameters for Cell 7492 for subsequent model-building. Other cells had qualitatively similar plasticity profiles (Supplementary Fig 2).

Figure 4: Fitting the presynaptic kinetic model to EPSC and IPSC data. A: Model of reaction system in each bouton. Synaptic input triggers a reaction causing extracellular Ca2+ to enter the bouton, and a pump removes the Ca. CaM buffers the Ca2+. Ca binds/unbinds from successive stages of the readily releasable (RR) pool of vesicles till they are docked, at which point the final Ca-binding step causes synaptic release. The released neurotransmitter (Glu for excitatory synapses, GABA for inhibitory) opens a ligand-gated receptor channel on the postsynaptic CA1 neuron, which is held voltage-clamped.(green arrow). The reaction scheme was identical for Glutamatergic and GABAergic synapses, but the rates were different to fit to the voltage-clamp recordings. B-E: EPSC fits for 20, 30, 40 and 50 Hz, respectively. Blue traces are experimental peaks and valleys averaged over all repeats and all 5-square patterns for neuron 7492. Red traces are simulated EPSCs. F: Cumulative release curve for glutamate vesicles. Blue:

expected transmitter release completion within 10 ms based on our observations of rapid postsynaptic responses from Figure 1 H,K. Red: simulated release. G: Distribution of model fitting scores for all cells and for glu and GABA receptors. Scores are normalised RMS differences between experimental and simulated points, as in panels B-F. Any score below 0.3 is a good fit.

Burst EPSP responses constrain single-cell parameters.

We next embedded our presynaptic signalling models as synaptic boutons on a simulated neuron and used burst-EPSP responses to constrain its properties. 100 glu synapses were placed on dendritic spines, and 200 GABA synapses on the dendrite (Figure 5B). In the CA1 the E:I synapse ratio is between 12:1 to 25:1 (Bezaire and Soltesz, 2013), but we implemented a subset of E synapses targeted by the CA3 pattern-responsive neurons in the interests of computational efficiency, as most of the inputs would be silent.

We implemented an abstract hippocampal network to drive activity on the synapses. The CA3 and interneuron layers were each implemented as 16x16 arrays of abstract neurons, with sparse randomised connectivity between them. Optical patterned input was mapped onto the CA3 neuron layer and included charge buildup from successive light pulses as well as ChR2 desensitisation (methods). The CA3 array projected sparsely to the 100 glu synapses on the CA1 model, and the interneuron array projected onto the 200 GABAergic synapses (Figure 5C). The initial projection statistics were derived from published work, but since many connections are cut in the slice preparation, we explored a range of probabilities of synaptic connections between the layers.

An example of CA3-triggered presynaptic signalling and postsynaptic currents is presented in Supplementary Figure 3. Note that the synaptic chemistry was modelled using the Gillespie Stochastic Systems Algorithm, hence the responses are noisy and synaptic release is probabilistic.

As a first-pass test of our model and comparison with experimental data, we modelled pulse train experiments performed in the current clamp (Figure 6). We ran the simulations deterministically to correspond to the averaged experimental traces. Since simulations permit access to internal state variables not available to the experiments, we also looked at how E and I currents, and E and I conductances balanced out during the pulse train (Figure 6F-O). Notably, the simulations predict that the E current builds

up, especially on the second pulse of the train, and is larger than the I current at this point. This corresponds well with the observations in Panels A-D and in Figures 2 and 3. Overall, the synapse and single neuron components of our model worked well together to replicate experiments.

Figure 5: Cell and network model. A: CA1 Pyramidal Neuron model. The cell has a compartment for soma, a 200-micron compartment for the dendrite, and 100 spines. Each spine is modelled as a head compartment, a neck compartment, and a presynaptic glutamatergic bouton. There are 200 GABAergic synapses positioned directly on the dendrite. B: Network model. The CA3 has 16x16 neurons, projecting randomly to the 16x16 interneurons. CA3 also projects directly to the Glu synapses. The interneurons project to GABA synapses.

Figure 6: Temporal summation and EI balance in burst stimuli. A-D: EPSP measured at soma for representative cell 3531 (maroon) and simulation (yellow) for each frequency. E: Distribution and mean of EPSP over all recorded cells for each frequency for the experiment (maroon) and for a simulated neuron (yellow). There is no significant dependence of EPSP on frequency for the recordings (Spearman's rank correlation p > 0.1). F-I: Simulated excitatory and inhibitory currents for each frequency. Note that the currents are measured close to resting potential, so the driving force (E_{GABA} - Vm) is small compared to the voltage-clamp experiments in Figure 1K, where Vm was held at 0 mV. J: Simulated excitatory and inhibitory current peaks as a function of frequency and their ratio. K-N: Simulated excitatory and inhibitory conductances for each frequency. O: Simulated excitatory and inhibitory conductance peaks. The ratio is almost independent of frequency. Note that simulated synaptic conductances are reduced by ~50x from nominal values to match the much smaller cell geometry and higher input impedance.

Poisson train stimuli constrain network parameters

Following the synapse and cell-level analysis of signal summation and short-term plasticity in the CA3-CA1 system, we turned to network-level readouts, including an analysis of scatter in the data, which we modelled as stochastic synaptic chemical kinetics and transmitter release. In the presence of stochasticity, one cannot fit individual trials to data and instead must examine statistics over many responses. We performed a series of Poisson-train experiments both to test the capacity of the model to replicate data and to constrain network-level parameters using the statistics of the responses (Figure 7, Panel A and left-hand panels below). As all trials of the Poisson-train experiments used 'frozen' pulse trains, we delivered the same frozen pulse sequence to the model (Figure 7, Panel B and right-hand panels below). Panels A and B show example EPSP traces for experiment and simulation, respectively.

In order to constrain network parameters, we looked at scatter plots of EPSP probability, of EPSPs over time, of EPSPs as a function of inter-stimulus interval, and distributions of EPSP amplitudes. We used simple parameter sweep runs to compare models to experiments and semiquantitatively compared the responses. The network parameters are listed below:

wtGlu	Weight (conductance) of glutamatergic synapses.
wtGABA	Weight (conductance) of GABAergic synapses
pCA3_CA1	probability of a given CA3 neuron synapsing onto a CA1 neuron
pCA3_Inter	probability of a given CA3 neuron synapsing onto an interneuron
pInter_CA1	probability of an interneuron synapsing onto a CA1 neuron
zeroIndices	Number of CA3 neurons to be excluded from optical input
ChR2_ampl	Scaling factor for optical stimulus to depolarize CA3 neuron.

We first compared the probability that each optical stimulus would elicit an EPSP (Panels C and D). As expected, 15-square patterns (Yellow dots) almost always gave an EPSP, while 5-square patterns frequently failed. In the simulation, we had perfect sensitivity to EPSPs, but we found that setting the detection threshold for an EPSP to 0.2 mV replicated the response probability. We then plotted scatter plots of EPSPs over time. Note that the initial high EPSP responses declined over a few seconds (Fig 7E).

We asked if ChR2 desensitisation of the stimulus might account for this. However, the field response does not show such a steep decline (Fig 1, Supplementary Figure 4). Instead, in the simulations, the initial decline was eliminated in the no-STP model (Supplementary Figure 5), suggesting that the effect is due to presynaptic depression that builds up over multiple input pulses.

We then plotted each EPSP against the inter-stimulus interval (ISI) preceding it (Figure 7 G, H). We observed that the EPSP was high for very short intervals and then declined to a lower steady value. The simulations matched the experimental scatter plots qualitatively. We initially assumed this was due to increased STP for closely succeeding input pulses. To our surprise, the no-STP model also showed the same profile of high EPSP at short ISIs (Supplementary Figure 6B). Removal of NMDA receptors also did not change the EPSP-vs-ISI profile (Supplementary Figure 6C). Thus, we conclude that the high EPSP was simply due to charge accumulation during rapid synaptic input.

Finally, to directly compare the statistics, we plotted the distributions of EPSP amplitudes (Figure 7 I, J). While the histograms were statistically different, the qualitative profiles matched. The 5-square EPSPs peaked close to zero and had a tail out to 5 mV, whereas the 15-square EPSPs had a peak shifted toward the right. The distributions of EPSP standard deviation for 5-square (blue) stimuli were similar, but the 15-square histograms (yellow) were smaller in the model than in the experiment, possibly reflecting an unaccounted source of variability (Figure 7 K, L).

We stress that our parameter runs revealed a continuum of responses in all the metrics we used for comparing experiments to models. All our experimental readings report a range of parameter values from EPSC and IPSC amplitudes to STP profiles. Hence we regard the current model as one of a family of models along the various axes of network, cellular and synaptic parameters which we explored.

Overall, our parameter searches resulted in network, cellular, and synaptic parameters in the multiscale network-presynaptic signalling-postsynaptic electrophysiology model which could replicate numerous qualitative features of experimental readouts and their distributions.

Figure 7: Experimental and simulated responses to Poisson spike train input, all patterns. A: Experimental EPSP aligned with light trigger pulses. Blue trace is recorded data; red trace is the fit used to compute peaks and subsequent statistics. B: Same for simulated data. C-L: Statistics for readouts of experiment (left column) and simulated (right column) data. Blue traces: 5-square stimulation, yellow traces: 15-square stimulation. C, D: probability of trigger to generate a peak in the EPSP trace. E, F: Scatter plot of EPSP peaks as a function of time over the Poisson train. Note the initial decline due to STP and ChR2 desensitisation G, H: Scatter plot of EPSP peaks vs. Inter-spike-interval. Note that the experimental data show a strong elevation of EPSP at short ISI. I, J: Distribution of EPSP peaks as estimated by peak-to-valley fitting. K, L: Distribution of standard deviations of EPSP peak amplitudes estimated over all recordings.

STP and sparse coding yield pattern transition detection

We next used our model to predict a single-cell computation reliant on patterned input and STP. As a heuristic argument, consider two patterns of activity, A and B, comprising different subsets of the population of CA3 pyramidal cells. Let these overlap to a small extent, say A∩B < 20% (Figure 8A). Since the CA3 projects directly to CA1 spines, the overlap onto excitatory synapses will also be around A∩B. CA3 also projects to a layer of inhibitory interneurons, and these then project to the CA1 neuron. If there is a high probability of connections to the much smaller interneuron population, then most interneurons will be activated for any given stimulus. Because of this, a given stimulus pattern A yields a broad inhibitory pattern I_A on the CA3 cell and similarly for I_B. Thus, the inhibitory overlap ($I_A \cap I_B$) >> A∩B for high inhibitory connectivity. In our default model, $I_A \cap I_B$ ~ 40 over 200 synapses, and A∩B ~ 15 over 100 synapses (Supplementary Figure 10).

Now, we consider STP effects. Repeated inputs of pattern A result in the initial strengthening of excitatory inputs, then desensitisation (Figure 8B upper cells). In parallel, the inhibitory pattern I_A undergoes desensitisation. At the transition from pattern A to B, most of the excitatory inputs have not been stimulated before, so the fresh synapses are stronger. This gives a larger excitatory drive (Figure 8B lower cell). Providing that the inhibitory overlap is moderate, the inhibitory drive is not much changed. From this heuristic argument, we predict two things: First, at transitions between patterns, there is a transient elevated response in the postsynaptic cell. Second, the size of the transient is larger if the input patterns are sparse because this lessens the overlap between them.

To test this in the model, we simulated a sequence of patterns of the form "AAAAAAAABBBBBBBBCCCCCCCCDDDDDDDD", that is, eight repeats of each of the four patterns. We ran these initial tests using deterministic calculations for the model of synaptic release kinetics. To calibrate the basal response, we ran the model without STP (Figure 8C). We observed that the successive input patterns gave output EPSP amplitudes that were similar but with small steps up and down between patterns. This was due to small differences in the summed input from each pattern, as mapped

through the randomised wiring of connections. When we introduced STP, each of the transitions generated a positive transient response (Figure 8D), even if the basal response led to a downward step.

Based on these predictions, we asked if we could detect such transitions in our slice preparation. We delivered the stimuli at repeat rates of 8, 20 and 50 Hz to span a range from theta to low gamma frequencies. In each case, we delivered the same 8x4 sequence of patterns. We did this for patterns with 5 and 15 squares, respectively. We found that in the 15-square patterns, there were significant (Wilcoxon test, $p < 0.05$) 'surprise' transients in 5 out of the 39 transitions in our dataset, mostly at 50 Hz. (Figure 8 E-H). The 5-square patterns showed much weaker transients, significant only in one case (Wilcoxon test, $p < 0.02$).

We then asked if our model with full stochastic synaptic kinetics could replicate these results. We obtained results close to the experimental data using patterns which led to the activation of ~75% of the synapses (Supplementary Movie 1). Like in the experiments, the 50 Hz pattern, in particular, showed strong 'surprise' detection at the transitions between patterns (Figure 8 I-L). The effect was weaker at 20 Hz and not visible at 8 Hz. Notably, the simulations also replicated the dip in EPSP following the onset of the pattern train, especially at 50 Hz. This is only partially explained by ChR2 desensitisation, as the field responses show a smaller decline (Figure 1J, Supplementary Figure 4). Instead, the dip appears to be substantially due to short-term depression of the excitatory synapses over prolonged pulse trains since it goes away in the noSTP model (Figure 8C, D).

We then explored a range of connectivity parameters to predict how 'surprise' computations might be affected. We varied the seven network parameters: pCA3_CA1, pCA3_Inter, pInter_CA1, wtGlu, wtGABA, the sparseness of input pattern, and channelRhodopsin weight onto CA3. One of the strongest factors was the sparsity of the stimulus pattern, such that sparse stimuli led to stronger transition responses than dense stimulus patterns (Compare Figure 8 I-L: 78% coverage with Figure 8 M-P: 35% coverage). In order to map this systematically, we ran simulations with variations for connectivity probabilities, synaptic weights, and sparsity (Supplementary Figure 7). We ran multiple stochastic simulations with different random seeds for connectivity and for

stochastic chemistry and compared the two pre-transition responses with the two post-transition responses. Surprisingly, the profiles for transition detection were dependent both on the network parameters and also on the identity of the patterns (e.g., Supplementary Figure 7). After averaging, the most consistent effect on the surprise response came from the two parameters related to the sparsity of excitatory input onto the CA1 neurons: pCA3_CA1 and input sparseness.

In summary, we predicted the occurrence of 'surprise' responses, in which the summed EPSPs are larger when one repeated pattern changes to another. We confirmed this experimentally, and used our model to explore network configurations which would support this. The broad trends were that the surprise response was stronger for faster repeats (50Hz), sparse patterns and sparse excitatory connectivity.

Figure 8: 'Surprise' responses to pattern changes. A. Schematic of connectivity from pattern A (red) and pattern B (blue) to four unique and one overlapping synapse on the CA1 pyramidal cell. B. Schematic of changes in synaptic strength before (upper cell) and after (middle cell) a series of repeats of the same pattern. Expanded red synapses in the top cell indicate STP-strengthened synapses after the first pulse. These synapses are weakened by depression on later pulses, as in the middle cell. The transition of patterns is indicated by the dashed line. The lower schematic indicates strong responses to new pattern B impinging on previously quiet synapses, excluding one which overlapped. C-P: Responses to pattern sequence of prepulse followed by four patterns, each with eight repeats, in the order AAAAAAAABBBBBBBBCCCCCCCCDDDDDDDD. In the scatterplots the four patterns are indicated by different colours. C. Deterministic model response where synapses lack STP. Due to small differences in the number of connecting synapses, the four patterns have slightly different EPSP amplitudes. D. Deterministic model response when synapses had STP. There is an initial decline in EPSP due to repeated inputs and presynaptic depression. Note that there is a positive transient every time the pattern changes, even if the new pattern elicited a weaker response in the no-STP case. E-H: Sample experimental data. E: Illustrative raw trace, 20Hz repeat rate. F, G, H: EPSP peaks (dots) for each trial and each pattern, mean shown as the blue line. The third transition at 50 Hz has a significant increase (Wilcoxon test, $p < 0.002$). I-L: Simulation output corresponding to experiment for dense stimulus activation. The second and third transitions at 50Hz are significant (Wilcoxon test, $p < 0.002$ CHECK). M-P. Simulation output for sparse stimulus activation shows much stronger transition detection.

Discussion

We used spatially and temporally patterned optical stimuli to drive synapses in the hippocampal CA3-CA1 circuit and analyse properties of short-term plasticity in this feedforward network, with a particular view to identifying when excitation might overcome inhibition in CA1 pyramidal neurons. We found that the second or third pulse in a train was most likely to have stronger excitation and further that summation of multiple synapses onto inputs also peaked and became more linear around the second pulse. We used our multidimensional dataset of different pulses, different patterns, and different frequencies to parameterise a multiscale model of the CA3-CA1 network, including multi-step neurotransmitter release on a conductance-based model for the CA1 neuron. We used the model to predict how STP dynamics of E and I inputs provide a mechanism for single cells to detect transitions in input pattern sequences, confirmed it experimentally, and used the model to explore network configurations which could robustly detect such transitions.

EI balance is retained despite distinct STP of E and I

EI balance is a well-studied phenomenon in the brain. Various processes like inhibitory plasticity, homeostatic plasticity, and developmental co-tuning have been proposed to be behind the origin and maintenance of the EI balance (Agnes et al., 2020; Agnes and Vogels, 2024; L. Chen et al., 2022; Jia et al., 2022; Lagzi and Fairhall, 2024; Mackwood et al., 2021; Miehl and Gjorgjieva, 2022; Trapp et al., 2018; Wu et al., 2022). At short time scales (< 0.5s), EI balance is dynamic in that it can detect and amplify faster temporal changes and efficiently gate continuous input streams (Cao and Van Hooser, 2022; Kremkow et al., 2010; Rotman et al., 2011; Vogels and Abbott, 2009; Wu and Zenke, 2021). Our results show that this balance is also tightly regulated at fast time scales (20Hz - 50Hz) by short-term plasticity, as reported previously (Bartley et al., 2015; Bartley and Dobrunz, 2015; Klyachko and Stevens, 2006; Sun et al., 2018), due to differential effects of STP on excitatory and inhibitory neurons. Thus, short-term

plasticity adds another layer of timing control on the gating by dynamic filtering of E-I inputs (Rotman et al., 2011).

Where can such temporally precise gating be useful? A putative example of that is place field formation in novel environments. In a study (Sheffield and Dombeck, 2019; Sheffield et al., 2017), researchers found that when a mouse is traversing in a new environment, the dendritic inhibition undergoes reduction, leading to dendritic spiking and possibly potentiation of the excitatory inputs through processes like BTSP (Bittner et al., 2017, 2015) and thus, formation of new place cells. Based on our results that show activity-dependent modulation in excitation and inhibition, we speculate that these short-timescale changes may lead to dendritic spiking driven by local transitory tilt in EI balance towards excitation.

STP modulates escape and normalisation

We have shown how STP supports transient amplification of responses (figure 3) and of subsets of synapses (Figure 8). Pyramidal cells in CA1 receive clusters of balanced excitatory and inhibitory inputs, forming co-tuned E-I assemblies that have been speculated or observed in vivo (Lee et al., 2012; Valero et al., 2022). Our experimental and modeling findings complement theoretical studies that have demonstrated that specific inputs can be gated through either gain modulation (Vogels and Abbott, 2009) or temporal delays (Kremkow et al., 2010). Our study reveals that short-term plasticity (STP) modulates the strength of subthreshold divisive normalisation (SDN) (Bhatia et al., 2019) in a frequency-dependent manner (Figure 3). As SDN emerges from both the relative amplitude and timing of excitatory and inhibitory inputs, its modulation by STP provides a novel mechanism for input gating in continuously firing balanced networks.

This modulation of subthreshold normalisation can be useful in input selection in networks where individual cells multiplex multiple input streams. Several studies have established that place-specific tuning is also seen in silent or non-place cells, which can be revealed by excitability changes, conjunctive inputs, or, at a low level, dendritic nonlinear events (Bittner et al., 2017, 2015; Lee et al., 2012; Sheffield et al., 2017). Thus, the CA3 cell network multiplexes a diverse set of inputs from which the downstream CA1 cells have to select (Payeur et al., 2019). Our results suggest how

subsets of synapses on a neuron may be gated through STP to implement this operation.

A resource to model the CA3-CA1 feedforward circuit

A major outcome of our study is the development and parameterisation of a multiscale, chemical+electrical+circuit level model of the CA3-CA1 feedforward circuit, which runs rapidly even on laptop hardware. We anticipate this will make it straightforward to deploy for studies on hippocampal single-neuron physiology and synaptic signalling, especially models of E and I input convergence onto plastic synapses. It takes 20 seconds of wall clock time to run 1 second of simulated time on a laptop with an AMD 6800HS processor (2022 vintage). Feature-wise, a few studies combine signalling (typically postsynaptic spine and dendrite signalling) with electrical signalling (Bhalla, 2017, 2011; Dainauskas et al., 2023; Dorman and Blackwell, 2022; Mattioni and Le Novère, 2013; Somashekar and Bhalla, 2024). Some of these adopt the same branched 1-dimensional diffusion approach and have similar speeds, but others include detailed 3-dimensional reaction-diffusion kinetics on spines (W. Chen et al., 2022) or boutons (Nadkarni et al., 2012, 2010) and these run much more slowly. However, none of the other models, to our knowledge, incorporates the network convergence onto a population of synapses on the CA1 PC dendrite.

As a network model, our simulation is more reduced and abstract than typical EI-balanced spiking networks and far more so than models with exhaustive cellular and network-level detail, such as the Blue Brain project (Markram et al., 2015). However, the simplicity of our model, with just seven network parameters, has enabled us to constrain it well to the recorded activity under a series of well-controlled stimulus patterns (Figures 7 and 8). To our knowledge not even the more detailed models have demonstrated fitting to experimental data for the input-output transforms of the CA3-CA1 network.

The utility of our model as a flexible resource is demonstrated by the facile manner in which reaction schemes can be switched in and out of the spine. In the current study, we employed two different spine kinetic models, one with STP and one without. This

was simply a matter of loading in a different SBML (or compatible) chemical model definition file. Thus the model is not only consistent with FAIR principles, it goes one step beyond the 'Interoperable' and 'Reusable' by lending itself to easy modification for other studies.

STP leads to single-neuron detection of pattern transitions

A key prediction of our study is that single CA1 neurons can perform 'surprise detection' to report changes in the spatial pattern of repeated inputs (Figure 8). At the single-synapse level, synaptic depression performs the operation of decorrelation, leading to more efficient coding (Goldman et al., 1999; Rotman et al., 2011). Our observations can be framed as decorrelation of the spatial pattern (alternatively, activity vector) of multi-synaptic input converging onto a neuron (Kremkow et al., 2010).

Our findings extend previous work on STP in regulating synaptic responses to temporally structured input (Klyachko and Stevens, 2006). The co-dependent modulation of excitatory and inhibitory synapses contributes to both circuit stability and plasticity, aligning with recent theories on how such dynamics are essential for encoding and maintaining memories (Agnes and Vogels, 2024).

More complex learning rules (Carvalho and Buonomano, 2011; Gütig and Sompolinsky, 2006) support learning of spatiotemporal activity sequences. We report a distinct form of spatiotemporal computation, that of discounting of familiar patterns, and the 'learning' here happens over a very few presentations and a few hundred milliseconds.

A similar observation of change detection has been reported for visual stimuli (Hu et al 2021). In this study, mice were trained on a visual change detection task while visual cortical activity was recorded using 2-photon calcium imaging. The authors trained a variety of ANN models to replicate the data and found that a rate model with short-term depression was able to do so after a few hundred epochs of training. Short-term depression has also been reported to support a similar computation of stimulus-specific adaptation in a recurrent network of the auditory cortex (Yarden and Nelken 2017). Our model achieves equivalent computations with simple feedforward connectivity at a single-neuron level, with much tighter linkages to the underlying physiological mechanisms.

In conclusion, we use experiments and models to explore mechanisms for how short-term synaptic plasticity modulates E-I balance in the CA3-CA1 hippocampal circuit. We describe an interplay between STP, EI balance and temporal dynamics of synaptic integration leading to filtering in the time domain and tight control over when cells may escape from EI balance. We predict and show the 'surprise' detection of changing spatial patterns of stimuli. We present a molecular-to-circuit scale model which is tightly constrained by and replicates a range of input-output transforms of the feed-forward + inhibition network of the CA3 to CA1, offering a mechanistic understanding of how hippocampal circuits process information relevant to learning and memory.

Methods

Model Repositories

All models and simulations to generate the figures are on github at: https://github.com/BhallaLab/STP_EI_paper_figs

Animals

Experiments were conducted on a mouse line with the expression of cre-recombinase specific to hippocampal CA3 (Grik4Cre) obtained from Jackson laboratory (ID: 006474). The animals were housed in a cage with up to 6 cage mates and under a 14/10 light/dark cycle. Food and water were provided ad libitum. The housing and health were managed as per the National Centre for Biological Sciences Institutional Animal Ethics Committee (IAEC) guidelines (Project: NCBS-IAE-2022/8 (R2ME). Only transgenic F1 mice, identified with genotyping, of both sexes were used for experiments.

Virus Injection

Transgenic mice between the ages 27-35 days were selected for virus injections. These animals were transferred to a separate holding facility where they were kept till the terminal experiment.

Virus: We used an adeno-associated virus AAV-LEX-rev-ChR2-tdtomato (Addgene 18917) to express channelrhodopsin2 in CA3 pyramidal cells in a cre-lox-dependent manner. The virus was diluted using sterilised PBS to a titre of 1-5 x 10^{12} GC/ml and stored at -80°C in aliquots. The animals were injected with a dose of 300-500 nl of virus suspension in 10-20ms long pulses of pressurised nitrogen using a Picospritzer (Parker-Hannifin, USA). The animals were then allowed to recover under supervision

and veterinary care. Terminal experiments were performed after an incubation period of 35 days or more, resulting in cre-lox-dependent expression of channelrhodopsin2 (ChR2) specific to the CA3 region of the hippocampus on the injected side. The virus AAV-FLEX-rev-ChR2-tdtomato was a gift from Scott Sternson (Addgene plasmid #18917; http://n2t.net/addgene:18917 ; RRID: Addgene_18917) (Atasoy et al., 2008).

Electrophysiology

Slice preparation: Animals aged between 56-120 days and at least 28 days after the virus injection were used for in vitro experiments. The animals were humanely euthanised using isoflurane (Forane) and brain was harvested into ice-cold slush of cutting solution (87mM NaCl, 2.5mM KCl, 7mM $MgCl_2.6H_2O$, 0.5mM $CaCl_2.2H_2O$, 25mM $NaHCO_3$, 1.25mM NaH_2PO_4, 75mM Sucrose) maintained over ice and aerated with carbogen (95% O_2 + 5% CO_2). After a rest period of 2 minutes, the hippocampus from the injected side was dissected out, and 350 μm thick sections were made using a vibrating microtome (Leica VT1200) submerged in the above-mentioned solution. The hippocampal sections were collected in a beaker filled with the recording solution (124mM NaCl, 2.7mM KCl, 1.3mM MgCl2.6H2O, 2mM CaCl2.2H2O, 26mM NaHCO3, 1.25mM NaH2PO4, 10mM D-(+)-glucose, pH 7.3-7.4 and osmolarity of 305-315 mOsm) at room temperature and allowed to recover for one hour in the dark before starting the experiments.

Recording Setup: Electrophysiological recording and optical stimulation was done on an upright DIC Olympus 53WI microscope. A slice hold-down (Warner Instruments WI 64-0246) was used to keep the slice from moving. The recording bath was perfused with the recording solution at a rate of 2 ml/min with a peristaltic pump. The recording solution was kept aerated with carbogen and the influx of the solution into the recording chamber was heated using an inline heater (Warner Instruments TC-324B) to 32-33°C.

Electrical activity was recorded using a headstage amplifier controlled by a patch clamp amplifier (Multiclamp 700B, Axon Instruments, Molecular Devices) and a data acquisition system (Digidata 1550B, Axon Instruments, Molecular Devices) using

Multiclamp commander (Molecular Devices) and custom written protocols in Clampex (Molecular Devices). The electrophysiology data was acquired at 20kHz and filtered in the 0-10kHz band.

In some experiments, another recording electrode was used to record extracellular responses in the CA3 cell layer. Whole cell patch clamp recordings from CA3 and CA1 cells were made using borosilicate glass micropipettes pulled using tungsten filament puller (P1000, Sutter Instruments) and filled with an internal solution containing 130mM K-gluconate, 5mM NaCl, 10mM HEPES, 1mM Na_4-EGTA, 2mM $MgCl_2$, 2mM Mg-ATP, 0.5mM Na-GTP and 10mM Phosphocreatine, pH ~7.3, and osmolarity ~290 mOsm. For voltage-clamp recordings, K-gluconate was replaced with Cs-Gluconate. Field response was often biphasic with multiple cycles, hence, a peak-to-peak value was used.

The cells were either recorded in the current clamp or the voltage clamp. In current clamp recordings, the cells were maintained at a membrane potential of -70 mV. Cells were rejected if the membrane potential changed by 5mV over the course of the recording, holding current either was above 100 pA or changed by more than 25 pA to maintain the mentioned membrane potential. In voltage-clamp recordings, the cells were rejected if the series resistance crossed 25 MΩ or changed by 30% of the initial value. Liquid junction potential was calculated to be 16.926 mV according to the stationary Nernst–Planck equation (Marino et al., 2014) using LJPcalc (RRID: SCR_025044) and was not compensated.

In optogenetic preparations, the additional stimulation that comes in due to the excitation of axons and dendrites can not be ruled out. To assess that, we patched a CA3 cell and recorded its optical depolarization caused by a single spot of light in a grid using a whole-cell voltage clamp. The CA3 cells have a large receptive field that arises due to the expression of ChR2 on their dendritic arbour. This receptive field is stronger closer to the soma but spans on both sides of the dendritic tree (Supplementary Figure 8A). A patched CA3 cell fired action potential with 100% success (Supplementary Figure 8B) when the cell layer was stimulated with patterns with multiple spots. A single spot, on the other hand, was only sufficient to induce action potentials when directly incident on soma (Supplementary Figure 8B).

Optical Stimulation

Pyramidal cells in the CA3 cell layer expressing channelrhodopsin were stimulated using a DMD-based patterned projector system (Polygon 400G, Mightex Systems, Canada) using a spot size of approximately 14µm x 7.8µm under a 40x water immersion objective. The spots (or squares) made a hexagonal grid with an inter-spot spacing equal to 32 µm. The grid occupied the middle 2/3rd of the projector frame (336µm x 187.2µm), and the microscope was translated on a stage to make the grid part overlay on the CA3 cell layer under a 40x objective. In 1-square stimulation experiments, the CA3 cell layer was illuminated one spot at a time with an interstimulus time of 3 seconds to avoid desensitisation of channelrhodopsin. In other experimental protocols, 5-spot and 15-spot patterns were used at various stimulation frequencies. Each presentation of a pattern was for 2-5 milliseconds. The light source had a total estimated maximum intensity of 6.7 mW, and for the grid size used, each spot had an absolute power of approximately 14.47 µW and a power density of 100 mW/mm^2. For all experiments, unless mentioned, the intensity of light was kept at 100%. Optical stimulation patterns were made and loaded onto Polygon using Polyscan2 (Mightex, Canada) and triggered using time series generated from custom protocols written in Clampex (Molecular Devices) and communicated to Polygon using Digidata 1440 DAC. A phototransistor (OPT101) was placed behind the dichroic to acquire the light stimulus waveform using the DAq and recorded along with the patch-clamp time series data.

Presynaptic Plasticity Model

Voltage clamp data from the pulse-train experiments (Figure 4) was used to train a model of presynaptic short-term plasticity (Figure 5A) implemented in MOOSE (Ray and Bhalla, 2008). The model included presynaptic signalling, activation of postsynaptic glutamate and GABA receptors, and electrophysiology of a ball-and-stick model of the postsynaptic neuron held in the voltage-clamp. The glutamate receptors were placed on dendritic spines implemented as separate electrical compartments, and the GABA receptors were on the dendritic shaft. Model fitting was performed in two stages. First, the raw voltage-clamp experimental data files in hdf5 format were analysed using a

Python script and the peaks and valleys for each trial were extracted and stored in an intermediate hdf5 file. Second, the HOSS pipeline (Hierarchical Optimization of Systems Simulations) (Viswan et al., 2024) was used to fit the presynaptic signalling model to this data, uniquely for each cell and separately for excitatory and inhibitory synapses. There were ten to eighteen 'experiments' run on the model to perform the fitting for each cell: Two experiments for each frequency, two to four pulse frequencies, and an additional experiment for time-course of neurotransmitter release. All this was done for 5-square and 15-square stimulation, respectively. For each frequency (typically 20, 30, 40, or 50 Hz pulse rate), we performed peak-to-valley comparisons with experiments for eight pulses in a trial and paired-pulse response experiments focusing on the first couple of pulses. The simulations for the experiments were performed using MOOSE to compute the somatic voltage-clamp response to synaptic input in the dendrites in a ball-and-stick cellular model.

The synaptic input was obtained using mass-action simulations of presynaptic signalling for the model in Figure 5A. This, too, was part of the multiscale MOOSE model. Excitatory synapses were implemented as boutons connecting to a dendritic spine with a compartment for the spine shaft and spine head, respectively. The simulated neurotransmitter release triggered the opening of ligand-gated receptor channels for AMPA and NMDA receptors placed on the spine head. Inhibitory synapses were implemented as presynaptic boutons with the same chemical topology but different kinetics (Figure 5A) and coupled now to GABA receptors placed on the dendritic shaft.

This entire process was repeated for excitatory and inhibitory synapses, respectively. All calculations for the optimisation were run deterministically since, otherwise, massive averaging would be needed for each estimate of voltage-clamp current. Equations, kinetic parameters, and concentrations for the STP and no-STP models are in Supplementary Tables 1 and 2.

Postsynaptic neuron model

The postsynaptic neuron model (Figure 5B) was a ball-and-stick model with a soma compartment of 10 um length and 10 um diameter and a single dendritic compartment of 200 um length and 2 um diameter. The dendritic compartment was ornamented with

100 evenly spaced spines, each containing a spine shaft and a spine head compartment. The spine head had a model of the AMPA receptor, the NMDA receptor, and calcium dynamics resulting from influx through the NMDA receptor. However, calcium dynamics were not utilised for any further calculations. A separate presynaptic bouton with the glutamate release kinetic model was apposed to each spine head.

The dendritic compartment also had 200 evenly spaced inhibitory synapses, likewise apposed to presynaptic boutons with the GABA release kinetic model.

The soma had voltage-gated Na and delayed rectifier K_DR channels in the Hodkin-Huxley formalism but with kinetics drawn from previous models. For most models, the Na and K_DR conductances were held to small values (6 and 3.5 Siemens/m^2, respectively) as we were interested in linear somatic integration, but for a few test simulations, we increased Na and K to 400 and 450 Siemens/m^2 to obtain stimulus-driven spiking. Equations and kinetics for the voltage-gated and ligand-gated channels in the model are presented in Supplementary Table 3.

Network model

We implemented a simple thresholding matrix connectivity network model. This had a 16x16 array of CA3 neurons and another 16x16 array of inhibitory interneurons (Figure 5C). Excitatory projections from CA3 to the 100 synapses on the CA1 neuron were implemented as a 256x100 matrix. Default entries were zeros, and ones were filled in at random with a probability pCA3_CA1 = 0.02 and weight wtGlu = 2 in our default model. Similarly, projections from CA3 to interneurons were defined as a 256x256 matrix with connection probability pCA3_Inter = 0.005. Projections from the interneurons to the 200 GABA synapses on the CA1 neuron were a 256x200 matrix with default connection probability pInter_CA1 = 0.004 and weight wtGABA = 50. All told, there were just five parameters in this model.

Each optically stimulated spatial pattern of activity in CA3 was converted into a vector of ones and zeros by thresholding the CA3 membrane potential, calculated as per the next section. Activity on the vector of glutamatergic synapses was:

Glu synapse vector:

$$\text{Glu synapse vector} = \text{matrix of connections} \ \times \text{CA3 activity vector}$$

Synaptic activity was thresholded at 1 to drive the presynaptic signalling. Similarly, Interneuron vector :

$$\text{interneuron vector} = thresh((\text{CA3\_InterneuronConnections} \times \text{CA3\_vector}))$$

GABA synapse vector:

$$\text{GABA synapse vector} = thresh((\text{Interneuron\_CA1 Connections} \times \text{interneuron vector}))$$

Stimulus model

To deliver the stimulus, we defined a set of eight input patterns. These were hand-coded as 8x8 grids of ones and zeros, set up to resemble the letters A through H. In keeping with the experimental protocol, the first five patterns (A through E) were sparse, with 22 ones. These represented the 5-square patterns. The last three patterns (F through H) were dense, with 38 ones.

We expanded the 8x8 patterns into 16x16 patterns by replicating the pattern four times over the array. For the first five patterns we zeroed out a specified number of entries on the 16x16 pattern using the parameter 'zeroIndices' to adjust sparseness. This zeroing out was done irrespective of the original state of the entry.

The product term:

$$\text{Glu synapse vector} = \text{matrix of connections} * \text{CA3 activity vector}$$

sometimes led to very different numbers of active glutamate synapses, differing by as much as 50% between patterns. Hence we carried out a preliminary calculation to fill in entries in the connection matrix to bring the final synaptic activation for all the 5-square patterns to roughly the same level (within 15%).

During the simulation runs we provided a time-series of patterns according to frozen Poisson timing series with a mean of 20 Hz (Figure 7) or a time-series of four patterns each repeated 8 times, at 8, 20 and 50 Hz (Figure 8). The time-series was generated with a resolution of 0.5 ms, and the stimulus duration (activation of Ca influx into the presynaptic boutons) was 2 ms. We used a fixed delay of 5 ms between the Glu and GABA synapse activation.

We implemented ChR2 desensitisation and membrane potential buildup over multiple pulses on the CA3 neurons using the following equations:

$$\frac{dG}{dt} = \frac{1 - G}{\tau_R}$$

$$\frac{dV_m}{dt} = \frac{E_{\text{rest}} - V_m}{\tau_M}$$

For each light pulse, we updated G and Vm as follows:

$$G = G \cdot \exp\left(-\left(D_0 + \frac{D_1}{\Delta t}\right)\right)$$

$$V_m = V_m + \tau_O \cdot \frac{E_{\text{ChR2}} - V_m}{\tau_C} \cdot G$$

Where

G = conductance of ChR2

V_m = CA3 pyramidal neuron membrane potential

Δt = time since last stimulus

D_0 = basal desensitization of ChR2 = 0.01

D_1 = light pulse induced desensitization of ChR2 = 0.0004

τ_R = Recovery time course of ChR2 from desensitization = 1.5 s

τ_M = Membrane time course of CA3 neuron = 10 ms

τ_C = Time course to charge CA3 neuron when ChR2 is open = 5 ms

τ_O = ChR2 channel open time = 1 ms

Example pulse trains with these parameters are illustrated in Supplementary Figure 9.

Simulation parallelization and Visualization

Model fitting (Figure 4) and data generation for multi-trial stochastic runs (Figures 7 and 8) involved parallel calculations. The stochastic runs, in particular, were embarrassingly parallel in that they required multiple repeats of the same model with different random

seeds. This was orchestrated through Python scripts using the multiprocessing library. The calculations were carried out on a 128-core AMD Epyc 7763 server. Data from these runs was organised into pandas data frames and dumped into hdf5 files in the same format as the experimental data. Figures 4, 7 and 8 were generated from these data files using pandas, Matplotlib and Python scripts.

3-D model visualisation (Supplementary movie) was performed using the MOOSE built-in interface to vpython.

References

Abbott, L.F., Varela, J.A., Sen, K., Nelson, S.B., 1997. Synaptic Depression and Cortical Gain Control. Science 275, 221–224. https://doi.org/10.1126/science.275.5297.221

Agnes, E.J., Luppi, A.I., Vogels, T.P., 2020. Complementary Inhibitory Weight Profiles Emerge from Plasticity and Allow Flexible Switching of Receptive Fields. The Journal of Neuroscience 40, 9634–9649. https://doi.org/10.1523/JNEUROSCI.0276-20.2020

Agnes, E.J., Vogels, T.P., 2024. Co-dependent excitatory and inhibitory plasticity accounts for quick, stable and long-lasting memories in biological networks. Nat Neurosci 27, 964–974. https://doi.org/10.1038/s41593-024-01597-4

Barros-Zulaica, N., Rahmon, J., Chindemi, G., Perin, R., Markram, H., Muller, E., Ramaswamy, S., 2019. Estimating the Readily-Releasable Vesicle Pool Size at Synaptic Connections in the Neocortex. Front. Synaptic Neurosci. 11, 29. https://doi.org/10.3389/fnsyn.2019.00029

Bartley, A.F., Dobrunz, L.E., 2015. Short-term plasticity regulates the excitation/inhibition ratio and the temporal window for spike integration in CA 1 pyramidal cells. Eur J of Neuroscience 41, 1402–1415. https://doi.org/10.1111/ejn.12898

Bartley, A.F., Lucas, E.K., Brady, L.J., Li, Q., Hablitz, J.J., Cowell, R.M., Dobrunz, L.E., 2015. Interneuron Transcriptional Dysregulation Causes Frequency-Dependent Alterations in the Balance of Inhibition and Excitation in Hippocampus. The Journal of Neuroscience. https://doi.org/10.1523/jneurosci.1834-15.2015

Bezaire, M.J., Soltesz, I., 2013. Quantitative assessment of CA1 local circuits: Knowledge base for interneuron-pyramidal cell connectivity: Quantitative Assessment Of Ca1 Local Circuits. Hippocampus 23, 751–785. https://doi.org/10.1002/hipo.22141

Bhalla, U.S., 2017. Synaptic input sequence discrimination on behavioral timescales mediated by reaction-diffusion chemistry in dendrites. eLife 6, e25827. https://doi.org/10.7554/eLife.25827

Bhalla, U.S., 2011. Multiscale interactions between chemical and electric signaling in LTP induction, LTP reversal and dendritic excitability. Neural Networks 24, 943–949. https://doi.org/10.1016/j.neunet.2011.05.001

Bhatia, A., Moza, S., Bhalla, U.S., 2019. Precise excitation-inhibition balance controls gain and timing in the hippocampus. eLife. https://doi.org/10.7554/elife.43415

Bittner, K.C., Grienberger, C., Vaidya, S.P., Milstein, A.D., Macklin, J.J., Suh, J., Tonegawa, S., Magee, J.C., 2015. Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. Nat Neurosci 18, 1133–1142. https://doi.org/10.1038/nn.4062

Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., Magee, J.C., 2017. Behavioral time scale synaptic plasticity underlies CA1 place fields. Science 357, 1033–1036. https://doi.org/10.1126/science.aan3846

Cao, A.S., Van Hooser, S.D., 2022. Paired Feed-Forward Excitation With Delayed

Inhibition Allows High Frequency Computations Across Brain Regions. Frontiers in Neural Circuits 15. https://doi.org/10.3389/fncir.2021.803065

Carvalho, T.P., Buonomano, D.V., 2011. A novel learning rule for long-term plasticity of short-term synaptic plasticity enhances temporal processing. Front Integr Neurosci 5, 20. https://doi.org/10.3389/fnint.2011.00020

Chen, L., Li, X., Tjia, M., Thapliyal, S., 2022. Homeostatic plasticity and excitation-inhibition balance: The good, the bad, and the ugly. Current Opinion in Neurobiology 75, 102553. https://doi.org/10.1016/j.conb.2022.102553

Chen, W., Carel, T., Awile, O., Cantarutti, N., Castiglioni, G., Cattabiani, A., Del Marmol, B., Hepburn, I., King, J.G., Kotsalos, C., Kumbhar, P., Lallouette, J., Melchior, S., Schürmann, F., De Schutter, E., 2022. STEPS 4.0: Fast and memory-efficient molecular simulations of neurons at the nanoscale. Front. Neuroinform. 16, 883742. https://doi.org/10.3389/fninf.2022.883742

Dainauskas, J.J., Marie, H., Migliore, M., Saudargiene, A., 2023. GluN2B-NMDAR subunit contribution on synaptic plasticity: A phenomenological model for CA3-CA1 synapses. Front. Synaptic Neurosci. 15, 1113957. https://doi.org/10.3389/fnsyn.2023.1113957

Denève, S., Machens, C.K., 2016. Efficient codes and balanced networks. Nat Neurosci 19, 375–382. https://doi.org/10.1038/nn.4243

Dorman, D.B., Blackwell, K.T., 2022. Synaptic Plasticity Is Predicted by Spatiotemporal Firing Rate Patterns and Robust to In Vivo-like Variability. Biomolecules 12, 1402. https://doi.org/10.3390/biom12101402

Fortune, E.S., Rose, G.J., 2001. Short-term synaptic plasticity as a temporal filter. Trends in Neurosciences 24, 381–385. https://doi.org/10.1016/S0166-2236(00)01835-X

Galarreta, M., Hestrin, S., 1998. Frequency-dependent synaptic depression and the balance of excitation and inhibition in the neocortex. Nat Neurosci 1, 587–594. https://doi.org/10.1038/2822

Goldman, M.S., Nelson, S.B., Abbott, L.F., 1999. Decorrelation of spike trains by synaptic depression. Neurocomputing 26–27, 147–153. https://doi.org/10.1016/S0925-2312(99)00068-5

Gütig, R., Sompolinsky, H., 2006. The tempotron: a neuron that learns spike timing–based decisions. Nat Neurosci 9, 420–428. https://doi.org/10.1038/nn1643

Jia, D.W., Vogels, T.P., Costa, R.P., 2022. Developmental depression-to-facilitation shift controls excitation-inhibition balance. Commun Biol 5, 873. https://doi.org/10.1038/s42003-022-03801-2

Kanamaru, T., Hensch, T.K., Aihara, K., 2023. Maximal Memory Capacity Near the Edge of Chaos in Balanced Cortical E-I Networks. Neural Computation 35, 1430–1462. https://doi.org/10.1162/neco_a_01596

Klyachko, V.A., Stevens, C.F., 2006. Excitatory and feed-forward inhibitory hippocampal synapses work synergistically as an adaptive filter of natural spike trains. PLoS Biology 4, 1187–1200. https://doi.org/10.1371/journal.pbio.0040207

Kremkow, J., Aertsen, A., Kumar, A., 2010. Gating of Signal Propagation in Spiking Neural Networks by Balanced and Correlated Excitation and Inhibition. J. Neurosci. 30, 15760–15768. https://doi.org/10.1523/JNEUROSCI.3874-10.2010

Lagzi, F., Fairhall, A.L., 2024. Emergence of co-tuning in inhibitory neurons as a network phenomenon mediated by randomness, correlations, and homeostatic plasticity. Sci. Adv. 10, eadi4350. https://doi.org/10.1126/sciadv.adi4350

Lee, D., Lin, B.J., Lee, A.K., 2012. Hippocampal place fields emerge upon single-cell manipulation of excitability during behavior. Science 337, 849–853. https://doi.org/10.1126/science.1221489

Mackwood, O., Naumann, L.B., Sprekeler, H., 2021. Learning excitatory-inhibitory neuronal assemblies in recurrent networks. eLife 10. https://doi.org/10.7554/ELIFE.59715

Marino, M., Misuri, L., Brogioli, D., 2014. A new open source software for the calculation of the liquid junction potential between two solutions according to the stationary Nernst-Planck equation. https://doi.org/10.48550/arXiv.1403.3640

Markram, H., Muller, E., Ramaswamy, S., Reimann, M.W., Abdellah, M., Sanchez, C.A., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., Kahou, G.A.A., Berger, T.K., Bilgili, A., Buncic, N., Chalimourda, A., Chindemi, G., Courcol, J.-D., Delalondre, F., Delattre, V., Druckmann, S., Dumusc, R., Dynes, J., Eilemann, S., Gal, E., Gevaert, M.E., Ghobril, J.-P., Gidon, A., Graham, J.W., Gupta, A., Haenel, V., Hay, E., Heinis, T., Hernando, J.B., Hines, M., Kanari, L., Keller, D., Kenyon, J., Khazen, G., Kim, Y., King, J.G., Kisvarday, Z., Kumbhar, P., Lasserre, S., Le Bé, J.-V., Magalhães, B.R.C., Merchán-Pérez, A., Meystre, J., Morrice, B.R., Muller, J., Muñoz-Céspedes, A., Muralidhar, S., Muthurasa, K., Nachbaur, D., Newton, T.H., Nolte, M., Ovcharenko, A., Palacios, J., Pastor, L., Perin, R., Ranjan, R., Riachi, I., Rodríguez, J.-R., Riquelme, J.L., Rössert, C., Sfyrakis, K., Shi, Y., Shillcock, J.C., Silberberg, G., Silva, R., Tauheed, F., Telefont, M., Toledo-Rodriguez, M., Tränkler, T., Van Geit, W., Díaz, J.V., Walker, R., Wang, Y., Zaninetta, S.M., DeFelipe, J., Hill, S.L., Segev, I., Schürmann, F., 2015. Reconstruction and Simulation of Neocortical Microcircuitry. Cell 163, 456–492. https://doi.org/10.1016/j.cell.2015.09.029

Mattioni, M., Le Novère, N., 2013. Integration of Biochemical and Electrical Signaling-Multiscale Model of the Medium Spiny Neuron of the Striatum. PLoS ONE 8, e66811. https://doi.org/10.1371/journal.pone.0066811

McKenzie, S., 2018. Inhibition shapes the organization of hippocampal representations. Hippocampus 28, 659–671. https://doi.org/10.1002/hipo.22803

Miehl, C., Gjorgjieva, J., 2022. Stability and learning in excitatory synapses by nonlinear inhibitory plasticity. PLoS Computational Biology 18. https://doi.org/10.1371/journal.pcbi.1010682

Nadkarni, S., Bartol, T.M., Sejnowski, T.J., Levine, H., 2010. Modelling Vesicular Release at Hippocampal Synapses. PLoS Comput Biol 6, e1000983. https://doi.org/10.1371/journal.pcbi.1000983

Nadkarni, S., Bartol, T.M., Stevens, C.F., Sejnowski, T.J., Levine, H., 2012. Short-term plasticity constrains spatial organization of a hippocampal presynaptic terminal. Proc. Natl. Acad. Sci. U.S.A. 109, 14657–14662. https://doi.org/10.1073/pnas.1211971109

Neher, E., 2015. Merits and Limitations of Vesicle Pool Models in View of Heterogeneous Populations of Synaptic Vesicles. Neuron 87, 1131–1142. https://doi.org/10.1016/j.neuron.2015.08.038

Payeur, A., Béïque, J.-C., Naud, R., 2019. Classes of dendritic information processing. Curr Opin Neurobiol 58, 78–85. https://doi.org/10.1016/j.conb.2019.07.006

Ray, S., Bhalla, U.S., 2008. PyMOOSE: Interoperable Scripting in Python for MOOSE. Front Neuroinform 2, 6. https://doi.org/10.3389/neuro.11.006.2008

Regehr, W.G., 2012. Short-Term Presynaptic Plasticity. Cold Spring Harbor Perspectives in Biology 4, a005702–a005702. https://doi.org/10.1101/cshperspect.a005702

Roach, J.P., Churchland, A.K., Engel, T.A., 2023. Choice selective inhibition drives stability and competition in decision circuits. Nat Commun 14, 147. https://doi.org/10.1038/s41467-023-35822-8

Rotman, Z., Deng, P.-Y., Klyachko, V.A., 2011. Short-Term Plasticity Optimizes Synaptic Information Transmission. Journal of Neuroscience 31, 14800–14809. https://doi.org/10.1523/JNEUROSCI.3231-11.2011

Rubin, R., Abbott, L.F., Sompolinsky, H., 2017. Balanced excitation and inhibition are required for high-capacity, noise-robust neuronal selectivity. Proceedings of the National Academy of Sciences 114, E9366–E9375. https://doi.org/10.1073/pnas.1705841114

Schneggenburger, R., Meyer, A.C., Neher, E., 1999. Released Fraction and Total Size of a Pool of Immediately Available Transmitter Quanta at a Calyx Synapse. Neuron 23, 399–409. https://doi.org/10.1016/S0896-6273(00)80789-8

Sheffield, M.E., Dombeck, D.A., 2019. Dendritic mechanisms of hippocampal place field formation. Current Opinion in Neurobiology 54, 1–11. https://doi.org/10.1016/j.conb.2018.07.004

Sheffield, M.E.J., Adoff, M.D., Dombeck, D.A., 2017. Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation. Neuron 96, 490-504.e5. https://doi.org/10.1016/j.neuron.2017.09.029

Somashekar, B.P., Bhalla, U.S., 2024. Discriminating neural ensemble patterns through dendritic computations in randomly connected feedforward networks. https://doi.org/10.7554/eLife.100664.1

Sun, H.Y., Li, Q., Bartley, A.F., Dobrunz, L.E., 2018. Target-cell-specific Short-term Plasticity Reduces the Excitatory Drive onto CA1 Interneurons Relative to Pyramidal Cells During Physiologically-derived Spike Trains. Neuroscience 388, 430–447. https://doi.org/10.1016/j.neuroscience.2018.07.051

Trapp, P., Echeveste, R., Gros, C., 2018. E-I balance emerges naturally from continuous Hebbian learning in autonomous neural networks. Scientific Reports 8. https://doi.org/10.1038/s41598-018-27099-5

Tsodyks, M.V., Markram, H., 1997. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. Proc. Natl. Acad. Sci. U.S.A. 94, 719–723. https://doi.org/10.1073/pnas.94.2.719

Valero, M., Zutshi, I., Yoon, E., Buzsáki, G., 2022. Probing subthreshold dynamics of hippocampal neurons by pulsed optogenetics. Science 375. https://doi.org/10.1126/science.abm1891

van Vreeswijk, C., Sompolinsky, H., 1996. Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity. Science 274, 1724–1726. https://doi.org/10.1126/science.274.5293.1724

Viswan, N.A., Tribut, A., Gasparyan, M., Radulescu, O., Bhalla, U.S., 2024. Hierarchical

    Optimization of Biochemical Networks.

Vogels, T.P., Abbott, L.F., 2009. Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. Nat Neurosci 12, 483–491. https://doi.org/10.1038/nn.2276

Wehr, M., Zador, A.M., 2003. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. Nature 426, 442–446. https://doi.org/10.1038/nature02116

Wilent, W.B., Contreras, D., Contreras, D., 2005. Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex. Nature Neuroscience. https://doi.org/10.1038/nn1545

Wu, Y.K., Miehl, C., Gjorgjieva, J., 2022. Regulation of circuit organization and function through inhibitory synaptic plasticity. Trends in Neurosciences 45, 884–898. https://doi.org/10.1016/j.tins.2022.10.006

Wu, Y.K., Zenke, F., 2021. Nonlinear transient amplification in recurrent neural networks with short-term plasticity. eLife 10, e71263. https://doi.org/10.7554/eLife.71263

Zucker, R.S., Regehr, W.G., 2002. Short-Term Synaptic Plasticity. Annual Review of Physiology 64, 355–405. https://doi.org/10.1146/annurev.physiol.64.092501.114547

Supplementary figures.

Supplementary Figure 1: STP profiles for each of the 12 voltage-clamp recordings with the pulse-train stimulus.

Supplementary Figure 2: Example model fits for cell 6201. Here, the 40 Hz data was not recorded; hence, panel C is blank. A, B, D: EPSC fits for 20, 30, and 50 Hz, respectively. Blue traces are experimental peaks and valleys averaged over all repeats and all 15-square patterns for neuron 6201. Red traces are simulated EPSCs. Note that the experimental traces are noisy hence the fits miss some of the peaks. E: Cumulative release curve for glutamate vesicles. Blue: target release complete within 10 ms. Red: simulated release. G: Distribution of model fitting scores for all cells and for glu and GABA receptors. Scores are normalized RMS differences between experimental and simulated points, as in panels B-F. Any score below 0.3 is a good fit.

Supplementary Figure 3: Presynaptic signalling model dynamics. All quantities are units of numbers of molecules per bouton unless otherwise stated. Signalling calculations are stochastic. Note that these are only two representative active synapses from among the 100 excitatory and 200 inhibitory synapses on the cell, and there are other synapses with distinct activity. Currents are reported as total currents for all synapses of the specified type. A: Optical stimulus train and corresponding somatic potential. B: Glutamaterigic bouton signalling. C. GABAergic bouton signalling.

Supplementary Figure 4 .   Response distributions for Poisson pulse train: field and EPSP data.

A: Sample data. Blue: Raw EPSP data, maroon: field potentials, green: stimulus pulse trigger. B-M: 5 square pattern responses in blue, 15-square pattern responses in orange. B, C: Scatter plots of EPSP vs fEPSP for 5 and 15 square input, respectively. Linear regression fit is plotted in black. P-value < 0.001 in both cases. D: Field EPSP as a function of time. E: EPSP as a function of time. F. Probability that a pulse will elicit an fEPSP peak. This is almost 100% reliable. G. Probability that a pulse will elicit an EPSP. Note that the initial ~2 seconds of the train have fewer events of low probability. H: Field EPSP as a function of inter-stimulus interval (ISI) for the pair of stimuli immediately preceding the fEPSP. I: EPSP as a function of ISI. Note that, unlike panel H, short ISIs elicit larger EPSPs.   J: Distribution of field EPSP responses. K: Distribution of EPSP responses. This differs qualitatively from panel J, except that in both cases, the 15-square distributions are shifted over to the right. L: distribution of standard deviation in field EPSP for successive trials of each given pattern. M: Distribution of standard deviation in EPSP for successive trials for each pattern.

Supplementary figure 5. Contribution of desensitization and STP to the initial dip in EPSPs. A: Experimental data. Note the decline in EPSP from t = 0 to t ~2 seconds. B. Reference simulated mode has a similar dip in EPSPl. C. Simulated model with STP disabled. In this case there is no initial dip in EPSP,.

Supplementary figure 6. Ruling out possible contributions to the dependence of EPSP on ISI. A: Reference model. B: Model without STP. The ISI dependence remains. Note that the response amplitudes are larger in the absence of STP, suggesting that presynaptic depression reduces EPSP amplitudes in the reference model. C: Model without NMDA receptors. The response is almost identical to the reference response, suggesting that most of the current is carried by the Glu receptors, and the NMDA receptor opening does not lead to amplification of EPSP at short ISIs.

Supplementary Figure 7. Dependence of 'surprise' response on network parameters. The three columns are the transitions between the first and second, second and third, and third and fourth patterns in the sequence. Note that 'surprise' is quite sensitive to pattern. Overall, the most consistent effects are that surprise is largest for small values of the connection probability from CA3 to CA1 (pCA3_CA1) and for large values of zeroIndices. zeroIndices controls the number of squares in the 16x16 input array, which are set to zero, in other words, it is a measure of sparseness. Responses averaged over four stochastic runs.

8A

8B

Supplementary Figure 8. A) Heatmap of optogenetic depolarisation of a single CA3 pyramidal cell obtained by single spots of light in a 24x24 grid while keeping the patched PC in the centre of the frame. B) Distribution of optical depolarization in CA3 PC against patterns of size 1 square (single spots), 5 squares, or 15 squares. Filled circles denote the instances when the patched cell fired an action potential. For patterns of size 5 and 15, all the trials resulted in spikes, while for 1-square patterns, 4/576 spots induced a spike.

Supplementary Figure 9. Time course of charging and desensitisation of optically stimulated CA3 neuron model. At 0.2 seconds we deliver a burst with intervals of 5, 10, 15, 20, 25, 30 ms. At t=0.7 seconds we deliver a single pulse. At t=1 seconds we deliver a train of 32 pulses at the frequencies of 8, 20 and 50 Hz as indicated above each panel.

Supplementary Figure 10. Overlap among synapses stimulated by each of 8 stimulus patterns. Patterns 0 through 4 are 5-square patterns, and patterns 5,6,7 are 15-square patterns. Diagonals indicate self-overlap, and are therefore the count of number of stimulated synapses. A: Glutamatergic synapses. B: GABAergic synapses.