Single neurons detect spatiotemporal activity transitions through STP and EI imbalance

Aditya Asopa and 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 activity patterns. Transitions in these input patterns are frequently salient. We demonstrate that short-term potentiation (STP) mediates escape from EI balance to implement mismatch detection in spatiotemporally patterned activity sequences. We characterized STP in the mouse hippocampus CA3-CA1 network using optogenetic patterned stimuli in CA3 while recording from CA1 pyramidal neurons. STP modulates EI summation across patterns, first amplifying, then reducing responses. We parameterized a multiscale model of network projections onto hundreds of E and I boutons on a CA1 neuron, each including stochastic signaling to mediate STP. The model detected mismatches in trains of input patterns, which we experimentally confirmed. Mismatch selectivity depends on pattern sparsity, network weights, and connectivity, and neuronal spiking makes it sharper. Finally, we predict strong mismatch detection for gamma-frequency bursts between theta cycles, consistent with theta-tuned snapshots of novel input.

Introduction

The balance between excitatory and inhibitory drives (EI balance) in neural networks has a wide range of computational functions. These include increased storage capacity [1], optimal coding [2], fast temporal encoding and gating [3–5], and higher dynamic range [6,7]. EI balance applies over tight time scales and even combinations of inputs at a single neuron level, termed precise balance[6]. However, little is known about how EI balance itself evolves dynamically due to activity-driven plasticity in sparsely active networks.

Active synapses undergo STP, during which neurotransmitter release varies upon receiving successive action potentials at the boutons [8–11]. STP emerges from the interaction and competition between pre-synaptic processes of vesicle depletion, calcium influx, calcium build-up, and neurotransmitter reuptake. Together, these processes create a temporal filter between activity in presynaptic and postsynaptic neurons and networks [8,12,13]. Excitatory and inhibitory synapses show distinct short-term modifications in their synaptic weights [13–15].

Activity in the CA3 is believed to converge to ensembles of cells representing distinct stimuli and contexts [16,17]. In abstract terms, these ensembles form instantaneous patterns, or vectors of activity, which are decoded in a heteroassociative manner by the CA1[18]. At a cellular physiology level, this static view implies differential responses by individual CA1 neurons to different patterns. However, precise EI balance in the feedforward excitatory-inhibitory (FFEI) network of CA3-CA1 normalizes these differences[6]. Given that activity is not static, it may be more salient for neurons to detect transitions between input patterns. Differential changes on E and I synapses may provide a mechanism for single neurons to extract interesting features of spatiotemporal inputs through STP[19].

In the current study, we used in vitro whole cell patch clamp recordings and patterned optical stimulation of the CA3-CA1 in the mouse hippocampus, to probe how EI balance evolves over spatiotemporal input trains 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 model of the hippocampal FFEI network. This model predicted that the postsynaptic CA1 neuron could detect pattern changes, which we confirmed experimentally. Finally we showed that change detection is enhanced when gamma frequency inputs trigger CA1 neuron spiking.

Results

Optogenetic patterned stimulation in the CA3-CA1 network

We obtained whole-cell patch-clamp recordings from CA1 pyramidal cells (PCs) while stimulating channelrhodopsin (ChR2) expressing CA3 pyramidal cells using optical patterns generated by a digital micromirror device (Fig 1A-B, see Methods). A pattern consisted of either 1, 5, or 15 spots (spot size 13 x 7 μm, power 14.5 μW/spot) of blue light (470 nm). We 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 by 300 ms to provide a reference baseline response. We measured the dynamics of E-I balance in a three-dimensional stimulus space consisting of CA3 activation size (number of spots), activation frequency (stimulation frequency), and short-term plasticity dependence (pulse number) in the train (Fig 1C).

To monitor the strength of the total resultant optogenetic activation of the CA3 layer (Extended Data Fig 1), we used an extracellular field electrode in the CA3 stratum radiatum (Fig 1A, methods). The field response correlated well with optically-driven CA3 PC depolarization (Fig 1E-G), and scaled with the size of the pattern (Fig 1F). 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 desensitization.

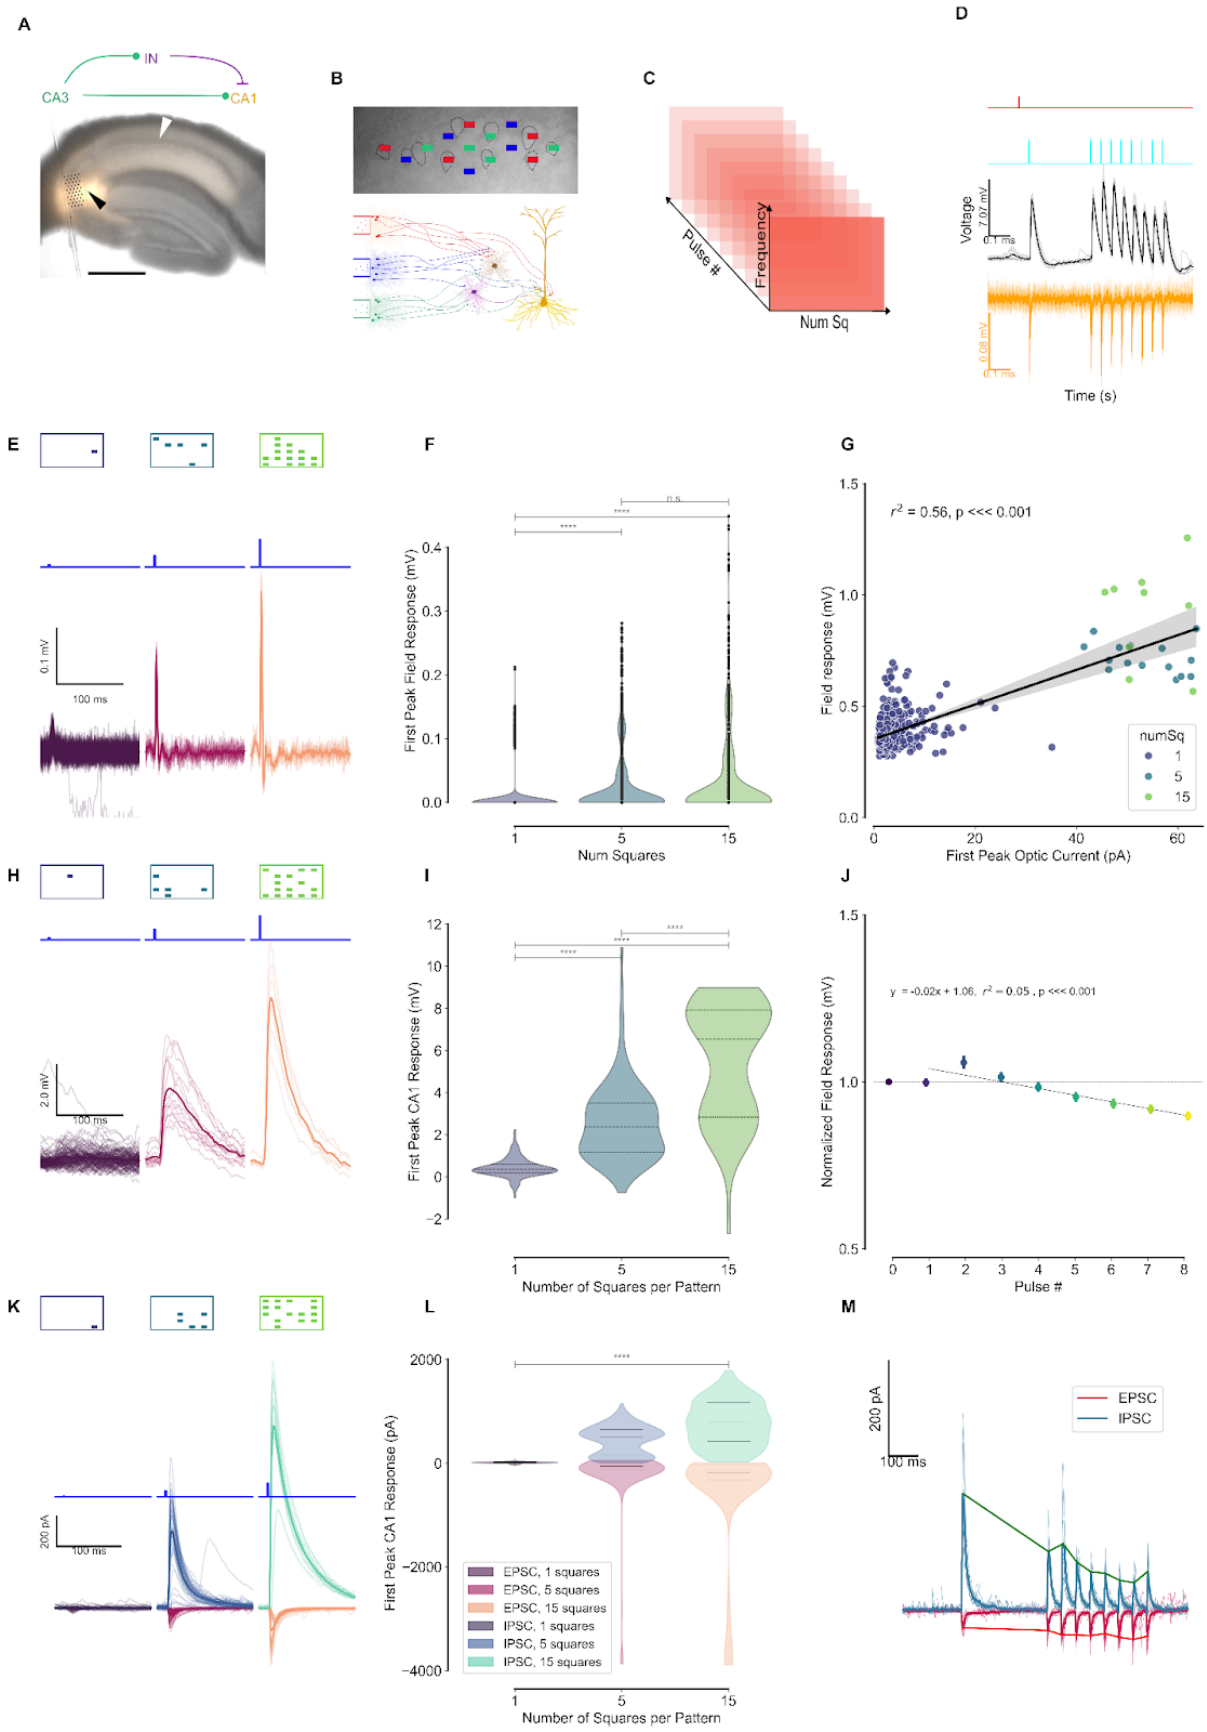
A
IN
CA3
CA1
B
C
Pulse #
Frequency
Num Sq
D
Voltage
7.07 mV
0.1 ms
0.08 mV
0.1 ms
Time (s)
E
0.1 mV
100 ms
F
First Peak Field Response (mV)
n.s.
****
****
Num Squares
G
r^2 = 0.56, p <<< 0.001
Field response (mV)
First Peak Optic Current (pA)
numSq
1
5
15
H
2.0 mV
100 ms
I
First Peak CA1 Response (mV)
****
****
****
Number of Squares per Pattern
J
y = -0.02x + 1.06, r^2 = 0.05 , p <<< 0.001
Normalized Field Response (mV)
Pulse #
K
200 pA
100 ms
L
First Peak CA1 Response (pA)
****
EPSC, 1 squares
EPSC, 5 squares
EPSC, 15 squares
IPSC, 1 squares
IPSC, 5 squares
IPSC, 15 squares
Number of Squares per Pattern
M
200 pA
100 ms
EPSC
IPSC

80

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. All violin plots show responses to probe pulse, thus without STP.

As expected, patterns with more spots elicited larger responses in CA1 in current clamp (Fig 1 H-I) and voltage clamp (Fig 1 K-L). We separated E and I components by holding the cell at -70 mV (GABAR1 reversal potential) and 0 mV (gluR reversal potential) respectively. As reported previously[6], different activation patterns in CA3 produced proportional E and I post-synaptic currents in the recorded neuron (Fig 1K). However, over the pulse train, E and I underwent distinct STP profiles (Fig 1 M).

Thus our patterned optical stimulation provided multiple input combinations to probe the dynamics of monosynaptic excitation and disynaptic inhibition arriving at CA1 pyramidal cells.

EI balance evolves over the pulse train

We next asked how precise EI balance at the level of CA1 PCs dynamically evolves due to short-term plasticity processes across a range of frequencies and input strengths (Fig 2A). All the PSPs of an 8-pulse train were normalised to the probe pulse. 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). Notably, the PSP of the first four pulses was significantly larger than that in the last four. 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).

While PSPs are a measure of subthreshold EI balance dynamics, neuronal firing is the network readout of escape from EI balance. Spiking was rare for the stimulus strengths in our experiments. We observed spiking in only 7% of trials (n=154/2201). We measured the likelihood of escape from EI balance by counting the fraction of total spiking trials that elicited for each combination of parameters (Fig 2B). Similar to the PSP response amplitudes, spike likelihood was higher in the first half of the pulse train ($p < 0.001$, ANOVA) Similarly, the spike likelihood varied significantly across the frequencies and number of spots (stimulus frequency $p<<0.001$, number of spots, pulse index $p << 0.001$, ANOVA).

We next measured E and I currents under voltage clamp across pattern size, pulse frequency, and pulse number. The EPSC showed a trend of early potentiation followed

by depression (Fig 2 Ei-Eii), while the inhibition underwent depression from the start (Fig 2 Fi-Fii). This effect was significant at $p < 10^{-6}$ across all the tested dimensions (ANOVA).

Thus EI balance over a burst briefly tilts in favour of E, after which both E and I undergo depression. This may result in a brief opportunity of escape from EI balance two or three pulses into a burst.

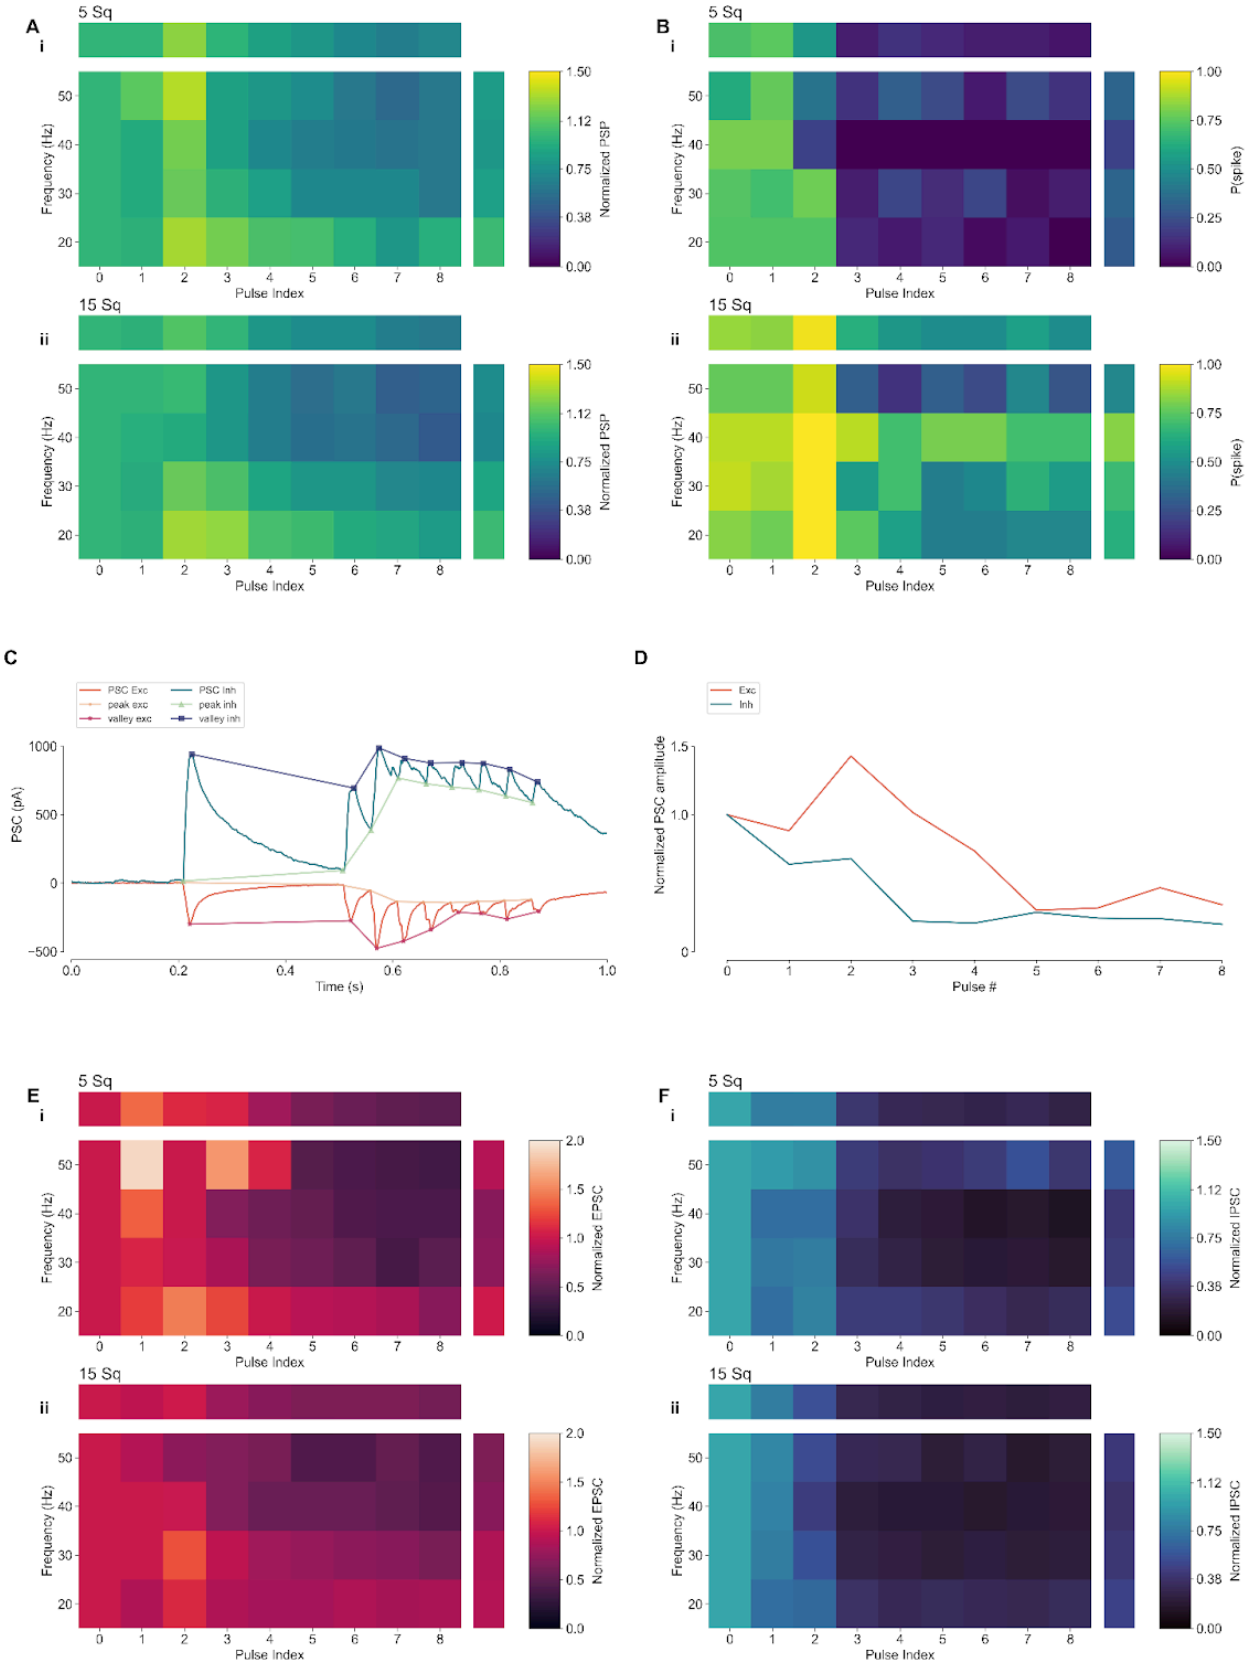
A
i
5 Sq
Frequency (Hz)
Pulse Index
Normalized PSP
ii
15 Sq
B
P(spike)
C
PSC Exc
peak exc
valley exc
PSC Inh
peak inh
valley inh
PSC (pA)
Time (s)
D
Exc
Inh
Normalized PSC amplitude
Pulse #
E
Normalized EPSC
F
Normalized IPSC

142

Fig 2. (A) 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 (Fi-Fii) 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 across the number of squares, pulse index, and excitation vs inhibition ($p<0.001$, ANOVA)., and the input frequency ($p<0.05$, ANOVA).

Multi-synapse summation and divisive normalization evolve over a pulse train

How might a neuron sum distinct input ensembles if they have distinct histories? We used the pattern-size axis of our stimulus protocol to probe history dependence of summation and subthreshold divisive normalization. We have previously shown that summed PSPs are not direct outcomes of EI ratios, because the delay between E and I onset is also a function of stimulus strength[6]. We therefore measured each of these terms: postsynaptic potentials, postsynaptic excitatory and inhibitory currents, and their onset and peak delays (Fig 3 A-D, extended data figure 2) over the pulse series, and obtained gamma (γ) as a measure of linearity [6] (Fig 3 E-G). Gamma (γ) relates the observed response of a CA1 cell to the expected response obtained by summing the inputs linearly in the following manner:

$$\gamma = \frac{expected * observed}{expected - observed}$$

The smaller the gamma, the larger the deviation from the linear combination of inputs due to the divisive effects of inhibition.

Using responses from the reference pulse, we replicated earlier observations[6,7] showing divisive normalisation, and obtained a median gamma of 7.16 (95% CI = 4.76 - 10.2)(Fig 3G). For comparison, we also performed a linear regression and obtained 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 (Fig 3G). Thus the differential STP profiles for E and I synapses yield a history-dependence of pattern-specific summation, with greater linearity around pulse two or three in a train.

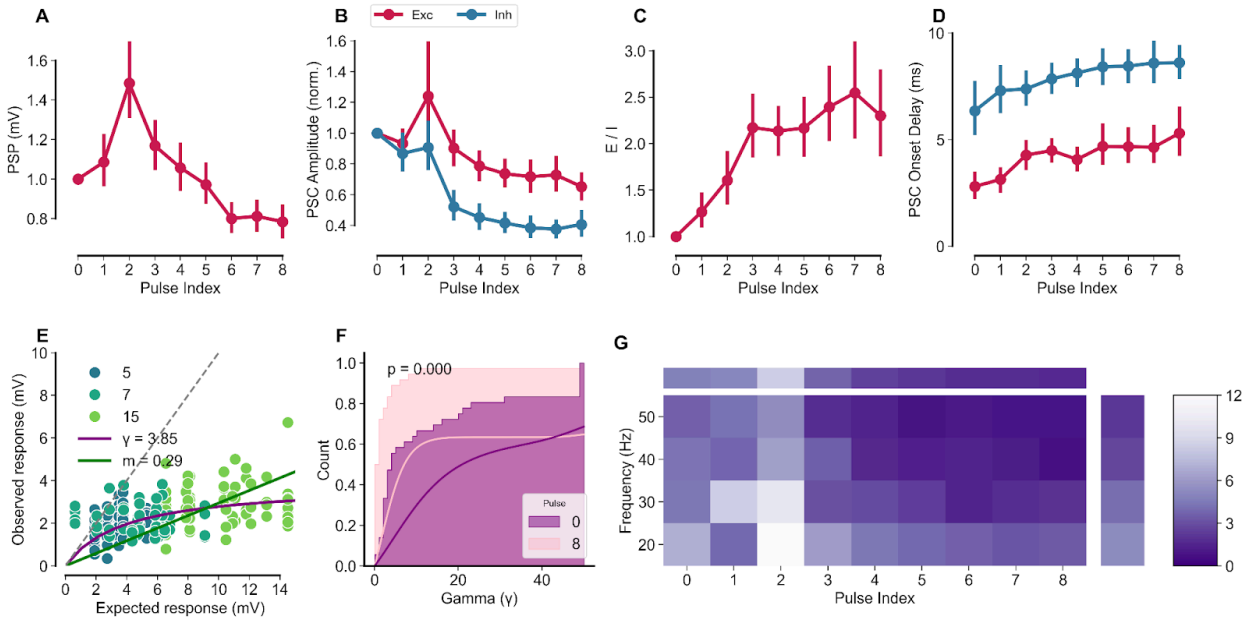

Figure 3: A. Trend of normalised postsynaptic potentials in a 20 Hz pulse train received at CA1 pyramidal cells against patterns of size 5 and 15. B. Same as A for voltage clamp postsynaptic excitatory (red) and inhibitory (teal) currents. C. E-I balance, measured as the ratio of E to I, showing a rise along the pulse train that is significantly high in the later half of the train ($p<0.001$, ANOVA). D. The difference between the onset delays of excitation and inhibition.

E. 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). F. Gamma shows a leftward shift along the pulse train, here shown by comparing cumulative distribution of gamma for probe pulse vs pulse 8 of the train. ($p<1e-3$, Mann-Whitney test) G. Gamma depends both on pulse index and frequency. Gamma averaged over the first four pulses is larger than in the later four ($p<0.001$, ANOVA).

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. Datasets with similar pulse trains have been used to obtain phenomenological models (e.g.,[20]). Instead, we chose to build a chemical kinetics-based multi-step neurotransmitter vesicle release model similar to previously published models [21,22] (Figure 4A). This incorporated both calcium buffering, which contributes to short-term potentiation, and depletion of docked vesicles, which underlies short-term depression. We implemented a ball-and-stick cellular model having AMPA and GABA synapses on the dendrite, with current or voltage-clamp recordings from the soma to compare to experiments. This chemical-electrical model was implemented in MOOSE [23], and parameters were fit using a multi-objective optimization pipeline to match waveforms and dynamics over all frequencies and the entire pulse train (Methods, Extended Fig 3) [24]. Each cell had somewhat distinct STP characteristics (Supplementary Fig 1). We independently fit presynaptic signaling models for E and I to their respective voltage-clamp waveforms (methods, Supplementary Fig 2,3).

While the E fit applied directly to the CA3->CA1 pyramidal neuron glutamatergic synapse, the I fit was a composite of CA3->Interneuron and Interneuron->CA1 projections (Fig 4B, C). Based on the peak and onset times, we made the further simplifying assumption that only one interneuron class was involved, most likely PV interneurons. An example of CA3-triggered presynaptic signalling and postsynaptic currents is presented in Extended Fig 4. Note that the synaptic chemistry was modelled using the Gillespie Stochastic Systems Algorithm [25], hence the responses are noisy and synaptic release is probabilistic.

We used parameters for Cell 7492 for subsequent model-building (Extended Fig 3).

Burst EPSP responses constrain single-cell parameters.

We next embedded our presynaptic signalling models as synaptic boutons in a conductance-based electrophysiology model of a 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 4B). In the CA1 the E:I synapse ratio is between 12:1 to 25:1 [26], 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 E inputs would be silent.

We implemented an abstract hippocampal network to drive activity on the synapses. CA3 was implemented as a 16x16 array of integrate-and-fire neurons. Optical patterned input was mapped onto this array (Extended Fig 5) and included charge buildup from successive light pulses as well as ChR2 desensitisation (methods, Extended Fig 6). The interneuron layer was a 16x16 array of binary neurons. The CA3 array projected sparsely to the 100 glutamatergic synapses on the CA1 model and to the 256 interneurons. The interneuron array projected onto the 200 GABAergic synapses (Fig 4C).

As a first-pass test of our model and comparison with experimental data, we modelled pulse train experiments performed using current clamp. We ran the simulations deterministically to correspond to the averaged experimental traces (Fig 5 A-D). Since simulations permit access to internal state variables not available to the experiments, we also looked at how E and I currents and conductances balanced out during the pulse train (Figure 5F-O). Notably, the simulations predict that the E current builds up, especially on the second pulse of the train. This corresponds well with the observations in Figure 5 A-D and in Figures 2 and 3.

Overall, the synapse and single neuron components of our model worked well together to replicate burst inputs in current clamp.

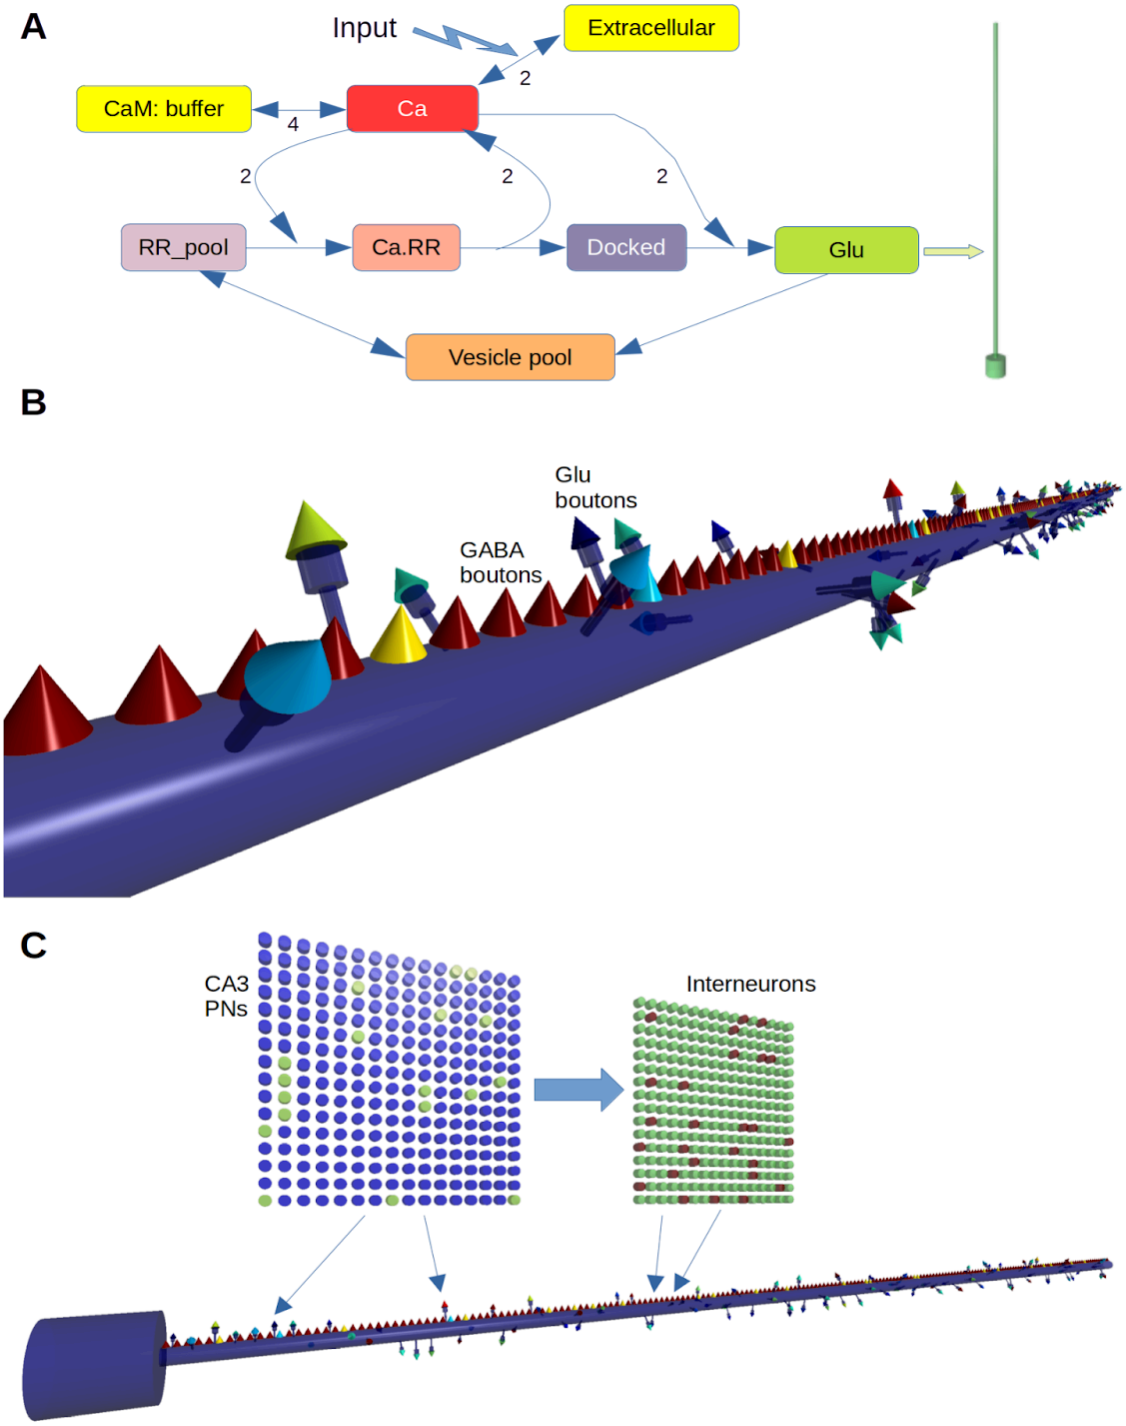

Figure 4: Multiscale model. A: Model of reaction system in each bouton. Synaptic input triggers entry of extracellular Ca2+ into the bouton, and a pump removes the Ca. CaM buffers the Ca2+. Ca binds/unbinds from successive stages of the readily releasable (RR) pool of vesicles till they are docked, at 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 the respective voltage-clamp recordings. B: Close-up of dendrite of 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 presynaptic boutons positioned directly on the dendrite. C: Network model. The CA3 has 16x16 neurons, projecting randomly to 16x16 interneurons. CA3 also projects directly to the Glu synapses on the spines. The interneurons project to GABA synapses.

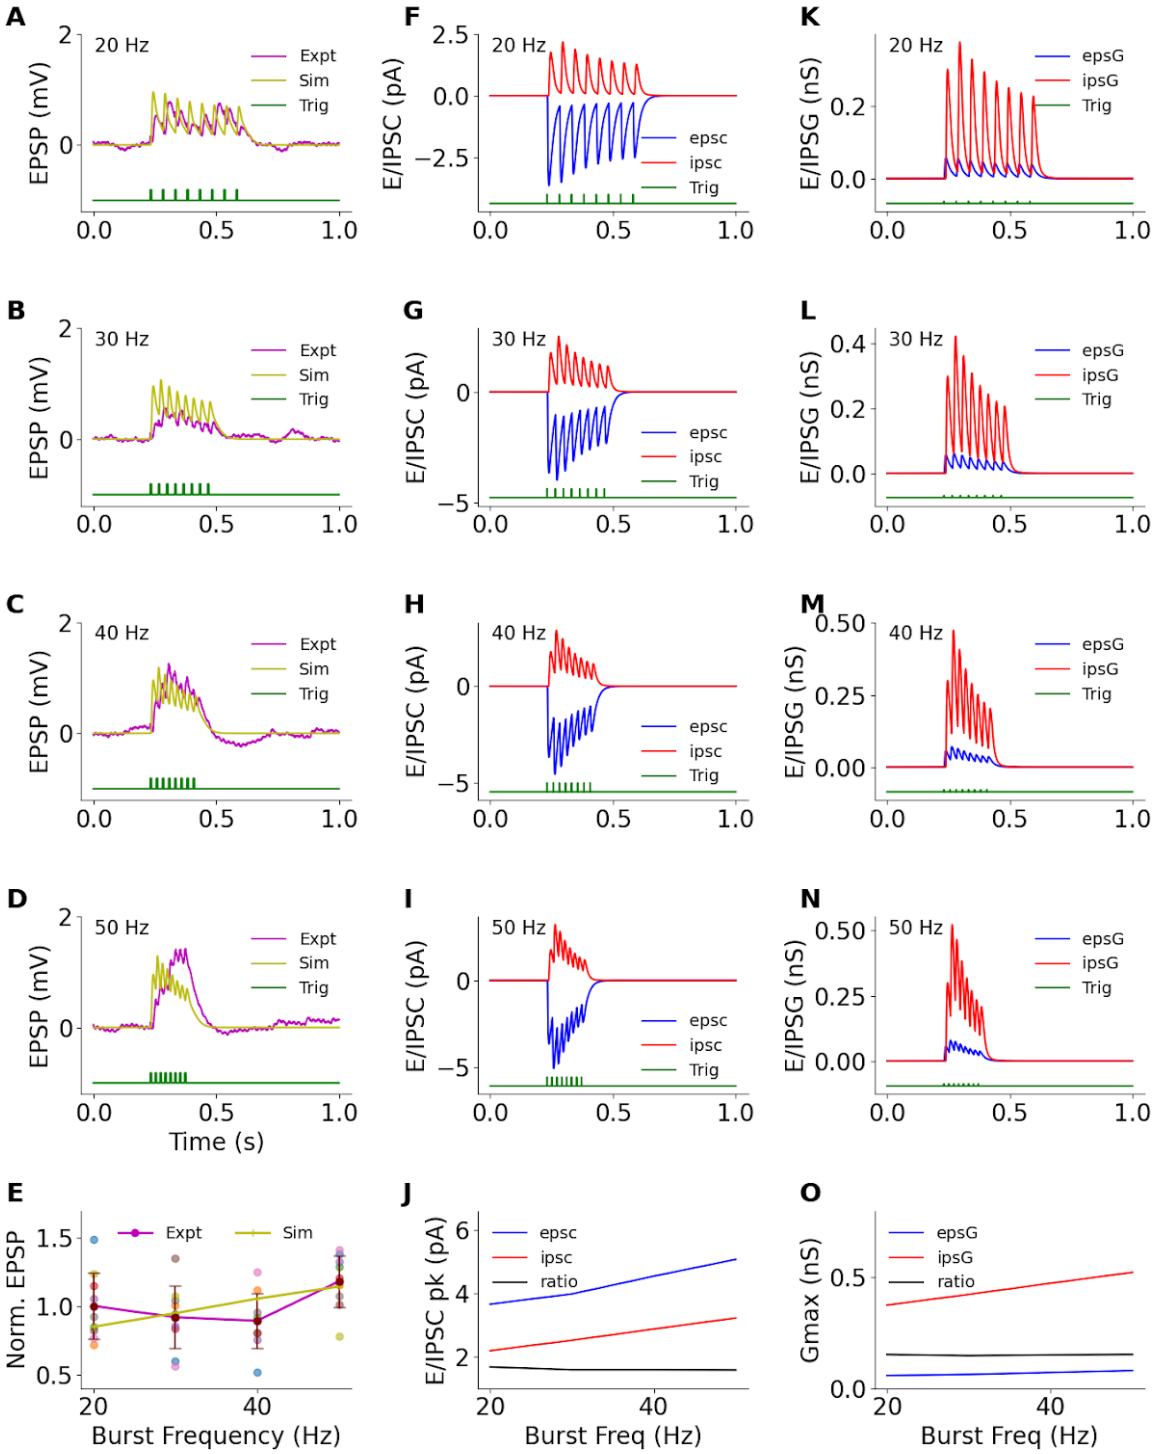

Figure 5: Temporal summation and EI balance in burst stimuli. A-D: EPSP measured at soma for representative cell 3531 (maroon) and simulation (yellow) for each frequency. E: Distribution and mean of EPSP over all recorded cells for each frequency for the experiment (maroon) and for a simulated neuron (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

265 currents are measured close to resting potential, so the driving force (E_{GABA} - Vm) is small compared to the voltage-clamp
266 experiments in Figure 1K, where Vm was held at 0 mV. J: Simulated excitatory and inhibitory current peaks as a function of
267 frequency and their ratio. K-N: Simulated excitatory and inhibitory conductances for each frequency. O: Simulated excitatory and
268 inhibitory conductance peaks. The ratio is almost independent of frequency. Note that simulated synaptic conductances are reduced
269 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 next tuned our model to network-level readouts. Here we explicitly included an analysis of scatter in the experimental data. We modeled the origin of this scatter as stochastic synaptic chemical kinetics and transmitter release (methods). We delivered a 'frozen' Poisson-train sequence to experiment and model, to test the capacity of the model to replicate data and to further constrain network-level parameters using the statistics of the responses. Figure 6 A and B show example EPSP traces for a single trial of an experiment and simulation, respectively. We used simple parameter sweep runs to tune the model network parameters.

We first compared the probability that each optical stimulus would elicit an EPSP (Fig 6 C and D). As expected, 15-square patterns (yellow dots) almost always gave an EPSP, while 5-square patterns frequently failed. In the simulation there was no instrumentation noise, but we found that setting the detection threshold for an EPSP to 0.5 mV yielded similar scatter plots to experiment. In Fig 6 E and F we compare scatter plots of EPSP amplitudes over time. Note that the initial high EPSP responses declined over the first two seconds (Fig 6E). We asked if ChR2 desensitisation of the stimulus might account for this. We found that the field response does indeed show a decline over multiple pulses (Fig 1, Extended Fig 7). We also checked for STD as a mechanism by comparing the decline between the reference and the no-STP model (Extended Fig 8). This suggests that there is a small additional contribution due to presynaptic depression.

We then plotted each EPSP against the inter-stimulus interval (ISI) preceding it (Fig 6 G, H). We observed that the EPSP was high for very short intervals and then declined to a lower steady value, both in experiments and simulations. We initially assumed this was due to increased STP for closely succeeding input pulses. Unexpectedly, the no-STP model also showed the same profile of high EPSP at short ISIs (Extended Fig 9). Removal of NMDA receptors in the model also did not change the EPSP-vs-ISI profile (Extended Figure 9C). Thus, we conclude that the high EPSP was simply due to charge accumulation during rapid synaptic input.

We then compared the distributions of EPSP amplitudes (Figure 6 I,J). While the large sample size meant that the histograms were statistically different at $p<1e-50$, the K-S test statistics were small (<0.18) indicating that the qualitative profiles were similar. As a further comparison between experiment and model, we estimated the power spectral density of the EPSP waveforms (Figure 7 K, L). These correlated very closely between experiment and simulation (Spearman's coeff. = 0.92 and 0.94 for 5 and 15 square, $p<1e-80$ for both).

Overall, we were able to replicate numerous features of the experimental dataset in a multiscale model incorporating presynaptic signalling, postsynaptic electrophysiology, and abstracted network connectivity and responses. While there was considerable scatter in both datasets, we obtained semi-quantitative matches between experiment and model for several readouts.

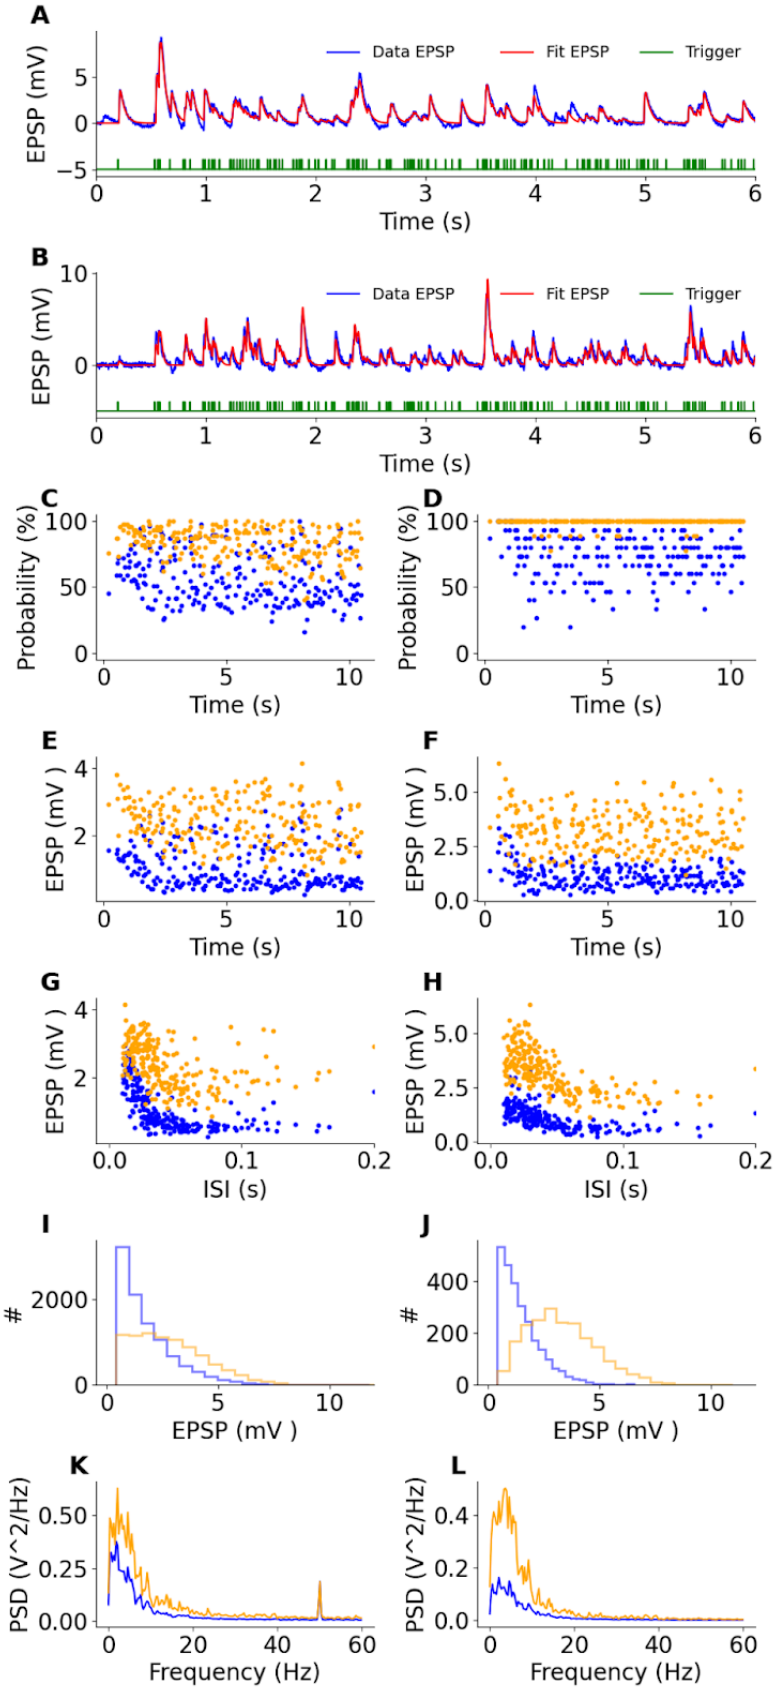
A
Data EPSP
Fit EPSP
Trigger
EPSP (mV)
Time (s)
B
Data EPSP
Fit EPSP
Trigger
EPSP (mV)
Time (s)
C
Probability (%)
Time (s)
D
Probability (%)
Time (s)
E
EPSP (mV)
Time (s)
F
EPSP (mV)
Time (s)
G
EPSP (mV)
ISI (s)
H
EPSP (mV)
ISI (s)
I
#
EPSP (mV)
J
#
EPSP (mV)
K
PSD (V^2/Hz)
Frequency (Hz)
L
PSD (V^2/Hz)
Frequency (Hz)

312

Figure 6: 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: Comparisons for readouts of experiment (left column) and simulated (right column) data. Blue traces: 5-square stimulation, orange 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. Slope for 5 square: Experiment: -0.338±0.0727 mV/s, R^2=0.089; simulation: -1.09±0.14 mV/s, R^2=0.11. Slope for 15 square: Experiment 0.679±0.12 mV/s, R^2=0.022; simulation -1.11 ± 0.243 mV/s, R^2=0.072. Note that the R^2 values are small, due to large scatter in the data. G, H: Scatter plot of EPSP peaks vs. Inter-spike-interval. Note that the experimental data show an elevation of EPSP at short ISI. We fit this using a simple exponential decay function y=y0.exp(-t/tau)+y1. Tau for 5 square: Experiment: 13.5±0.54ms, R^2=0.152; Sim: 30.3±4.19ms, R^2=0.102. Tau for 15 square: Experiment: 18.3±1.93ms, R^2=0.038; Simulation:47.1±7.25ms, R^2=0.169. Note that the R^2 values are small, due to large scatter in the data. I, J: Distribution of EPSP peak amplitudes. The Kolmogorov-Smirnoff statistic comparing experiment and simulation was 0.176 for 5-square distributions, and 0.183 for 15-square, indicating that the histograms were similar, though the large number of observations meant that they could each be distinguished with p < 1e-50. K, L: Power spectral density of EPSP response over the entire dataset. Other than the 50Hz line noise spike in the experimental dataset in panel K, the distributions are very similar (Spearman's coeff. = 0.92 and 0.94 for 5 and 15 square, p<1e-80 for both).

STP and sparse coding yield pattern mismatch detection

We next used our model to predict a single-cell computation emerging from patterned input and STP. As a heuristic argument, consider two patterns of activity, A and B, comprising different ensembles of CA3 pyramidal cells. Let these overlap to a small extent, say A∩B < 20% (Figure 7A). Since the CA3 projects directly to CA1 spines, the overlap onto excitatory synapses will also be around A∩B. We ignore inhibitory effects for simplicity, but these are considered in the detailed simulations below. Now, we consider STP. Repeated inputs of pattern A result in the initial strengthening of excitatory inputs, then depression (Figure 7B). 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 7B rightmost cell). From this heuristic argument, we predict two things: First, at transitions between patterns (mismatches), we expect a transient elevated response in the postsynaptic cell. Second, the size of the transient should be larger if the input patterns are sparse in E, 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. To calibrate the basal response, we ran the model without STP and without stochasticity (Figure 7C). We observed that the output EPSP amplitudes were similar for the successive input patterns. When we introduced STP, each of the transitions generated a positive transient (Figure 7D), even if the basal response led to a small downward step. With denser connectivity, the transients were smaller (Figure 7E).

Based on these predictions, we asked if we could experimentally detect pattern transitions in our slice preparation. We delivered the same 8x4 stimulus sequence at repeat rates of 8, 20 and 50 Hz to span a range from theta to gamma frequencies (Figure 7 F). We did this for patterns with 5 and 15 squares, respectively. Figure 7 Fiv shows a strong response to the third transition at 50 Hz for the 15-square case. We then asked if our model with full stochastic synaptic kinetics could replicate these results. Our sparse connectivity model produced multiple strong mismatch responses,

consistent with sparse overlap of synapses (extended Fig 9), and these were reduced in a denser model (Fig 7 G, H, Supplementary Movie 1). This agrees with our heuristic analysis above. Like in the experiments, the 50 Hz pattern, in particular, showed strong mismatch responses (Figure 7 G, H). We used the Wilcoxon test to compare experimental EPSPs immediately before and after pattern transitions (Fig 7I-K). More mismatch responses were seen for 15 square patterns and gamma-frequency pulses (Fig 7 K).

We highlight the observation that mismatch detection was selective for specific transitions. This occurred both in experiments (e.g, Fig 7 Fiv for the third transition) and in simulations (e.g., Fig 7 Giv for transitions one and three).

We then explored a range of connectivity parameters in our model to predict how mismatch computations might be affected (Table 1, Figure 7L-Q).

Parameter	Description	Default value
wtGlu	Weight conductance (nS) of glutamatergic synapses. Defined in model as relative value with reference = 3.	0.474 ± 0.135 nS
wtGABA	Weight (conductance) of GABAergic synapses. Relative reference value in model = 10	12.57 nS
pCA3_CA1	CA3 to CA1 connection probability, i.e., probability that a given CA3 neuron synapses onto a CA1 neuron	0.02
pCA3_Inter	CA3 to Interneuron connection probability	0.01
pInter_CA1	Interneuron to CA1 connection probability	0.01
zeroIndices	Number of silent CA3 neurons, ‘masking’ optical input.	192
ChR2_ampl	Scaling factor for optical stimulus to depolarize CA3 neuron.	0.1

Table 1: Network parameters and their definitions.

We first confirmed that sparse patterns had stronger mismatch responses (Figure 7L). There was a sharp peak in selectivity for low values of wtGlu, and a plateau for higher values (Fig 7M). Mismatch was small for low values of wtGABA, supporting the role of inhibition in mismatch detection (Fig 7N). Greater values of excitatory synaptic projection probability, pCA3_CA1 led to lower selectivity (Fig 7O). This is complementary to the trend for sparseness (Fig 7L), which is expected since they reciprocally affect the number of activated synapses for each stimulus pattern. Both the inhibitory projection terms pCA3_Inter and pInter_CA1 showed increasing selectivity with increasing connectivity (Fig 7P,Q). We interpret this as due to increasing overlap between input patterns, so that depression of inhibitory inputs carries over between patterns.

Thus, at this stage we had predicted STP-dependent single-neuron detection of mismatches in sequences of inputs, confirmed it experimentally, and then explored network conditions which supported it. Notably, sparse connectivity was a prerequisite.

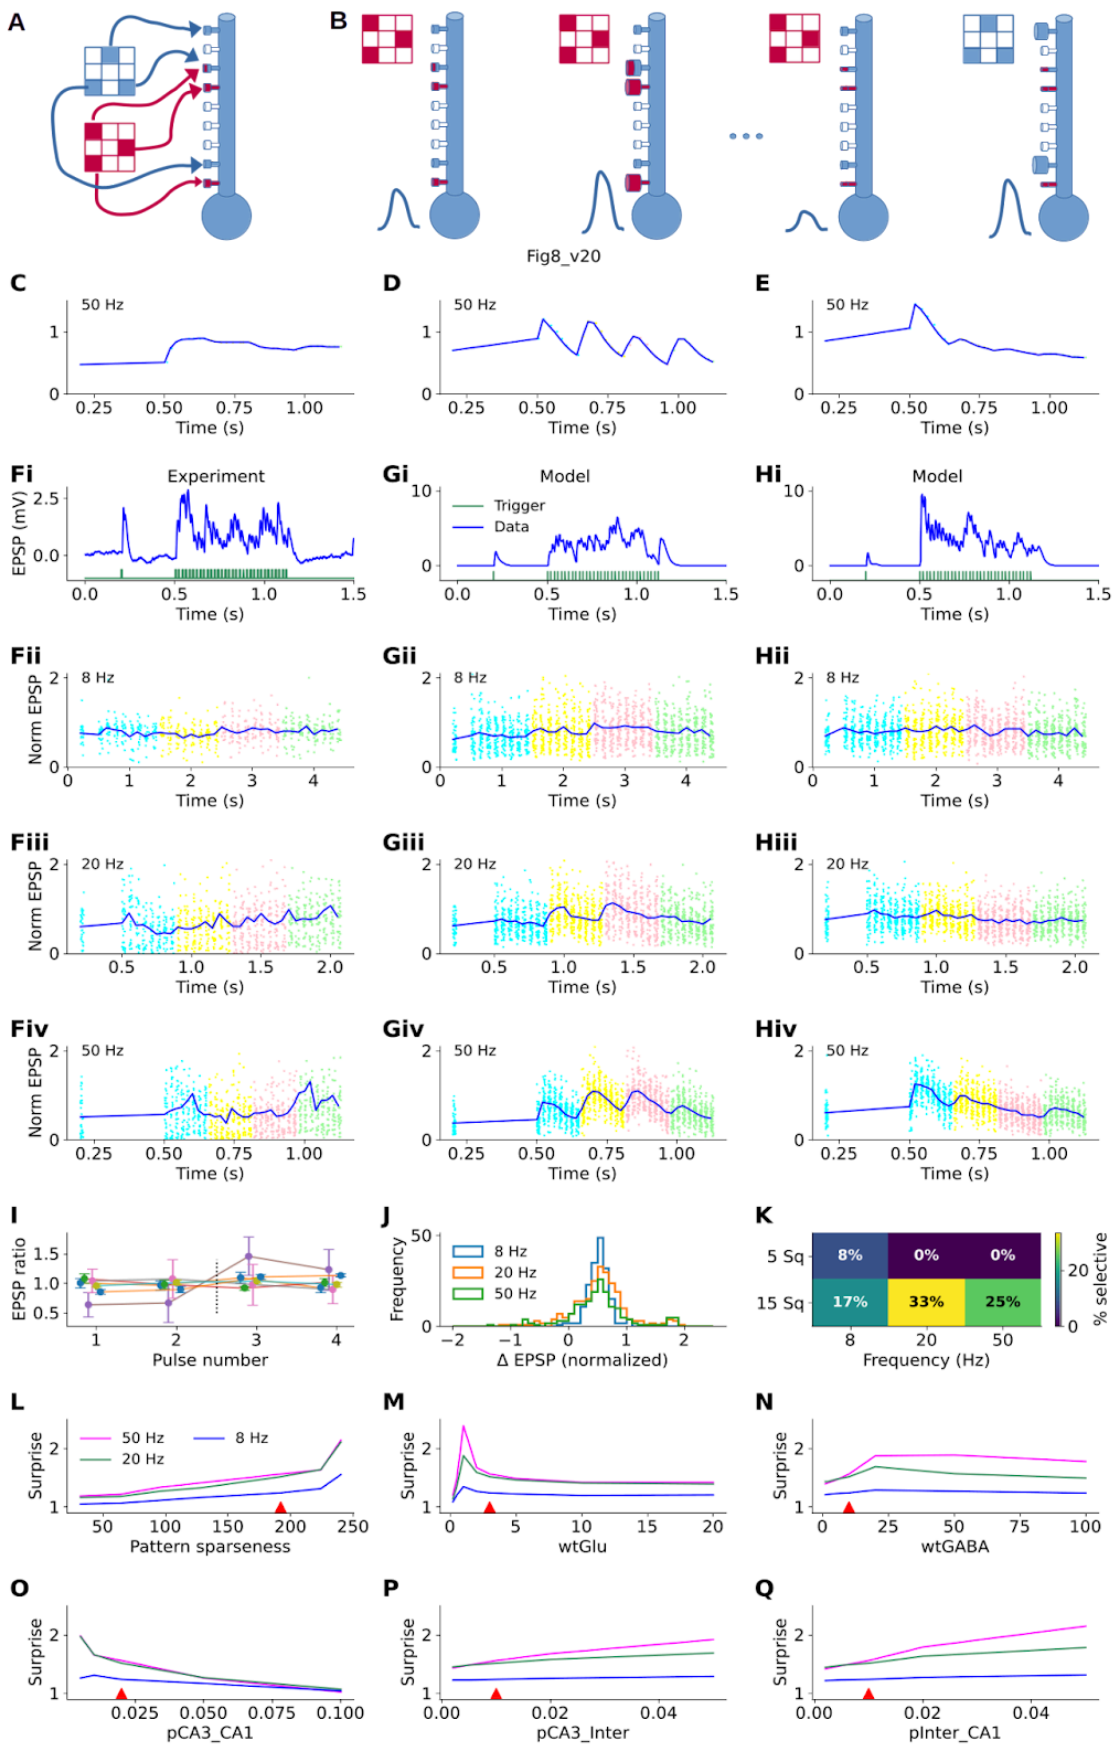
A
B
Fig8_v20
C
D
E
50 Hz
Time (s)
Fi
Experiment
EPSP (mV)
Gi
Model
Trigger
Data
Hi
Fii
Gii
Hii
8 Hz
Norm EPSP
Fiii
Giii
Hiii
20 Hz
Fiv
Giv
Hiv
I
EPSP ratio
Pulse number
J
Frequency
8 Hz
20 Hz
50 Hz
Δ EPSP (normalized)
K
5 Sq
15 Sq
8%
0%
0%
17%
33%
25%
Frequency (Hz)
% selective
L
Surprise
Pattern sparseness
M
wtGlu
N
wtGABA
O
pCA3_CA1
P
pCA3_Inter
Q
pInter_CA1

393

Figure 7: Mismatch responses to pattern changes. A. Schematic of connectivity from pattern 1 (red) and pattern 2 (blue) to four unique and one overlapping synapse on the CA1 pyramidal cell. B. Schematic of changes in synaptic strength over a series of repeats of pattern 1, followed by pattern 2. Response amplitude is indicated by the height of EPSP in blue, near the soma, and synaptic strength by size of the synapse. The first cell is unstimulated, the second undergoes STP on the red synapses, and the third cell experiences depression of the synapses and hence lower amplitude response. The last cell receives pattern 2 (blue) which activates and strengthens two previously quiet synapses, and one overlapping one. This results in a larger response. C, D, E: Deterministic normalized EPSP responses to pattern sequence of prepulse followed by four patterns, each with eight repeats, in the order AAAAAAAABBBBBBBBCCCCCCCCDDDDDDDD. 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, but no transients . D. Deterministic model response when synapses have STP. There is a transient elevation in response every time the pattern changes, even if the new pattern elicited a weaker response in the no-STP case. E. Deterministic model response loses mismatch detection when pCA3_CA1 is raised to 0.1, that is, there is 10% connection probability between every presynaptic CA3 cell and the postsynaptic CA1 neuron. F, G, H: example traces and scatterplots for responses to the same pattern sequence. F: Experimental neuron 2821. G: reference model. H: Model with dense connections. Fi, Gi, Hi: sample EPSP traces. F,G,H remaining panels: Scatterplots for each EPSP peak over multiple repeats for the series of patterns. Cyan, yellow, pink and green indicate repeats for patterns A, B, C and D respectively. Blue line is mean. The strongest response occurs for the 50Hz repeats. There is a significant mismatch response going from pattern C to D (pink to green) for experimental data in Fiv (Wilcoxon test, $p < 0.0001$). Giv: In reference model (75% of spots set to zero), transitions one and two but not three show significance ($p<0.0001$, $p<0.0001$, $p=0.674$)p < 0.001). In Hiv (only 12.5% of spots set to zero) the first and third transitions are significant (Wilcoxon test, $p=0.0003$, $p=0.0143$). I: Methodology for estimation of mismatch response significance. Pulses 1 and 2 are before the transition, and 3 and 4 are after it. For clarity, data are offset to the mean of the four illustrated pulses. Dashed vertical line indicates transition. Peak EPSP is found for pulses 1 and 2, and compared with pulses 3 and 4 using the Wilcoxon test. In this example there is an elevation for only one of the transitions out of six. J: Distribution of mismatch responses in normalized EPSP. Note that the 8Hz case (blue) is narrow, so has fewer large transients. K: Heatmap of percentage of mismatch responses which were significant (Wilcoxon test, $p < 0.05$). Gamma frequency (20 and 50 Hz) and 15 square cases have more significant transitions. L-Q: Parameter dependence of mismatch response, computed as the ratio of the mean of the two samples before the pattern change, to the mean of the two samples after. Red triangles indicate the value of parameters in the reference model. L: Sparseness is defined as the number of stimulus points out of the total 16x16 stimulus which are set to zero. There is a steady rise in mismatch. M: mismatch vs. synaptic weight of glutamatergic synapses has a sharp peak at around 1 (corresponding to ~0.158 nS per synapse). N. Mismatch against GABAergic synaptic weight peaks at around 20 (~26 nS per synapse), and then almost plateaus. L, M, N: Mismatch vs. connectivity probability. O: Excitatory synapses from CA3 to the target CA1 pyramidal neuron. Note that this has a complementary slope to panel I. P: Excitatory synapses from CA3 to interneurons. Q: Inhibitory synapses from interneurons to postsynaptic CA1 pyramidal neuron. Both P and Q have a steady positive slope, which we interpret as due to increased saturation of connectivity of the inhibitory input to the CA1 cell, hence ensuring that depression in the GABAergic input is not pattern selective.

Mismatch detection is effective in spiking neurons.

Finally, we utilized the model to predict how mismatch responses might manifest in spiking neurons. To do this we introduced voltage-dependent Na and K_DR channels in the soma of the CA1 neuronal model and increased the wtGlu and wtGABA each by 67%. We obtained spiking responses for 50 stochastic runs with patterned input at 50 Hz (Figure 8 A,B), and averaged the resulting firing rates with a 10ms window, as a proxy for population activity (Figure 8 C-K). To our surprise, the spiking model yielded strong mismatch detection (Figure 8C). Furthermore, spiking could track each optical pulse. As a consistency test, we replicated the finding from Figure 7 that high sparsity of connections impaired selectivity (Figure 8D, E).

We next explored a range of stimulus conditions for functional implications. First, we considered the frequently used behavioural paradigm of 'oddball' detection, which produces a rapid ~100-250ms evoked potential in several modalities [27,28]. We delivered a constant, repeating pattern interspersed every 8 stimuli with the oddball (deviant) stimulus. Under these conditions the mean spiking approached zero, but the oddballs stood out with brief, brisk spiking (Figure 8 F).

Gap stimuli, where the deviant stimulus is an absence of stimulus, are also known to be strong behavioural triggers for oddball responses (refs). However, in our model the gap stimulus sequence led to nearly zero spiking (Fig 8 G), suggesting that a more complete CA1 connectivity network model may be required to detect the gap.

We then asked how differential STP in E and I might impact mismatch detection. We first removed STP from glutamatergic synapses (Fig 8 H), then from GABA synapses (Fig 8 I), then from both (Fig 8 J). Remarkably, mismatch detection was impaired in all three cases. Lack of Glu STP, either alone or with GABA STP, eliminated the primary mechanism of mismatch detection through depression of repeated excitatory inputs. Lack of just GABA STP kept the inhibitory tone high, while the Glu synapses became weaker. Hence the response almost shut down though some selectivity remained (Fig 8 I). When we removed GABA input entirely we observed elevated spiking but selectivity was retained for the first transition (Figure 8 K).

We next performed two controls to confirm the role of transitions in eliciting mismatch responses. In the first control we delivered the same spatial pattern for all 32 pulses

(Figure 8J). As expected, this led to a rapid decline in spiking. Next, we randomly selected one pattern among our set of five for each pulse (Fig 8 K). Surprisingly, this also gave a decline in spiking, though it was not as steep as the uniform pattern. Upon further analysis, it turned out that random sequences had sufficient overlap to trigger a subset of the same synapses (Extended Fig 10). The net effect was that excitatory synapses were stimulated at ~21 Hz for random patterns, as compared to the 50 Hz for uniform patterns. This led to slower synaptic depression.

How might the mismatch response manifest in more natural stimulus contexts than a train of 32 pulses? To investigate this, we considered the well-known phenomenon of gamma bursts modulated by a theta background[29]. To set the gamma burst parameters, we confirmed that mismatch selectivity persisted broadly over the gamma frequency range (50-120 Hz, Figure 8 N). We then provided theta-modulated gamma bursts, where each theta cycle had five pulses of input at 100Hz, with the start of each burst separated by 130 ms to give a 7.69 Hz theta rhythm (Fig 8 O). We compared the case where each burst was a different pattern with the control where each burst was the same pattern (Fig8 P, Q). The CA1 neuronal response was much larger for the case with changing input patterns between theta cycles. Thus mismatch detection also works when the ensemble of neurons comprising a gamma burst changes between theta cycles. We also checked the case where each cycle of the gamma burst had a different ensemble, which could precess, as in place cells[29]. Both non-precessing and precessing bursts were efficiently transmitted (Fig 8, R,S).

Thus, in this final set of tests of our model, we predicted that mismatch detection remains sharp in spiking neurons. The mismatch response also occurs for single oddball patterns, but not gaps. We confirmed that STP on both E and I synapses are required, and showed that mismatch detection also works for 100 Hz gamma-frequency bursts riding on the theta rhythm. Together, these support possible functional relevance of these mechanisms for spatio-temporal pattern selectivity in in-vivo-like network conditions.

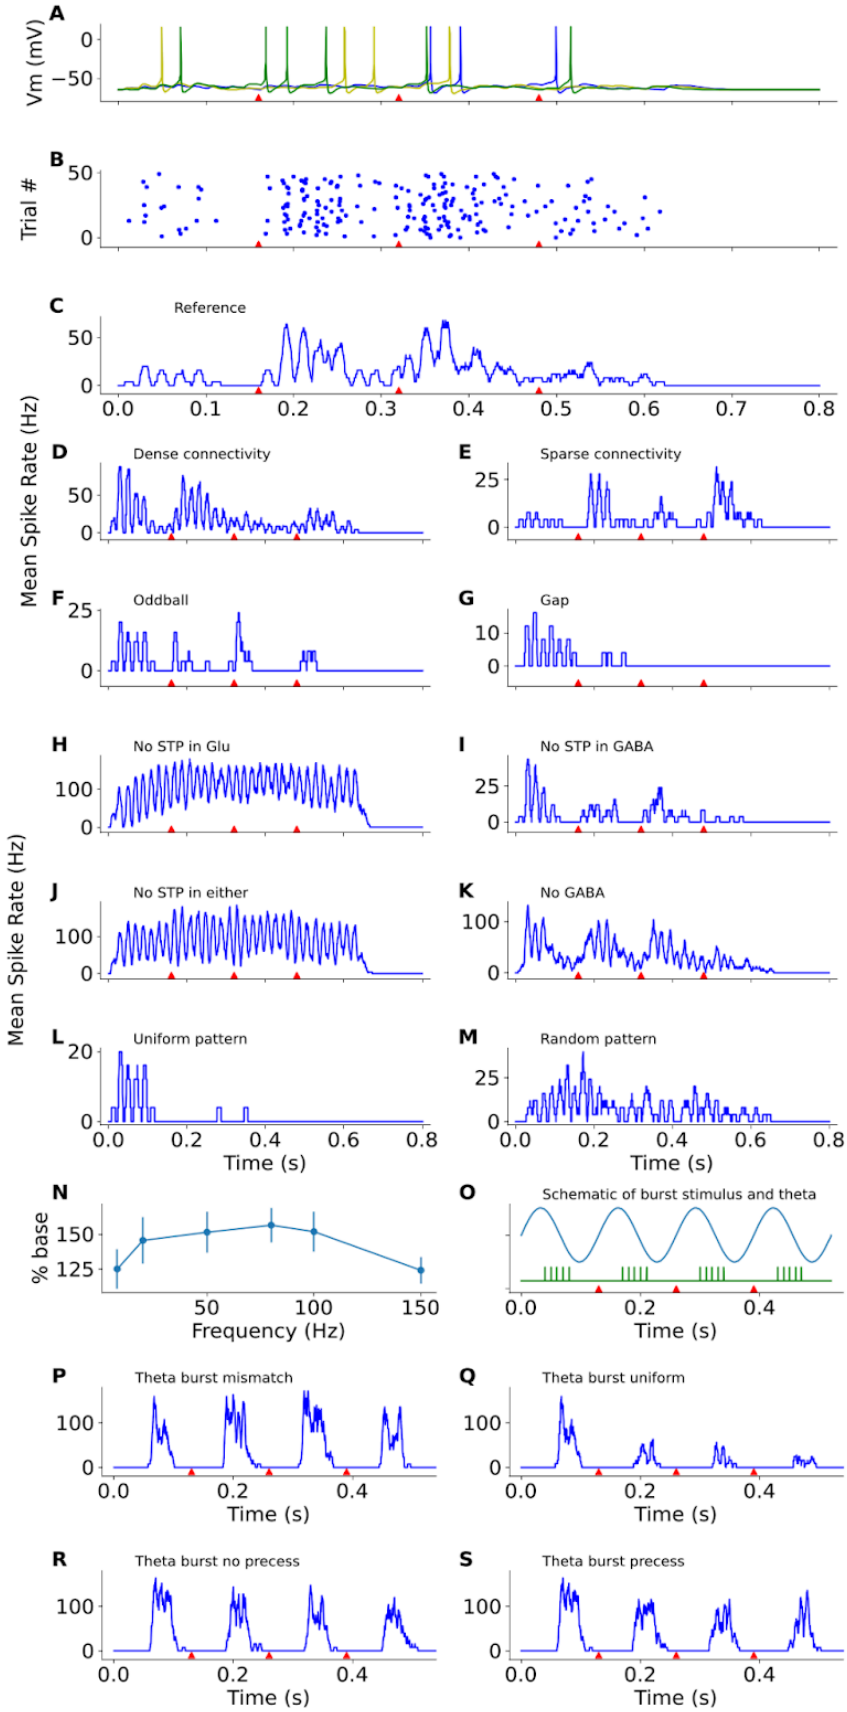

Figure 8. Spiking responses are selective for mismatch in patterned sequences. All runs were performed using a

spiking neuron model with stochastic synaptic chemical kinetics and each run consisted of 50 trials. Red triangles indicate transition times between patterns. For panels A to E and H to K, the stimulus consists of eight repeats of each four patterns as in AAAAAAAABBBBBBBBCCCCCCCCDDDDDDDD. Traces from C to M and P to S are spike rates with a moving window of 10 ms and 5 ms respectively. Mismatch statistics from C to M use the Mann-Whitney U test to compare the three pulses immediately before vs. the three immediately after a transition. P to us use the Mann-Whitney U test to compare 40 ms of response in the first theta cycle with the corresponding 40ms in the next three cycles. A: Spiking responses on three illustrative trials. B. Raster plot of spiking for reference stimulus. C: Reference model exhibits mismatch responses for the first two pattern transitions (p=0.002, 0.032, 0.434). D: Dense stimuli (25% spots nulled) gives a mismatch signal (p=0.014) only for the first transition. E: Sparse patterns (87.5% spots nulled) are mismatch selective but with less spiking (p=0.037, 0.407, 0.008). F: Oddball (deviant) stimulus presented every 8th pulse, against a uniform background: AAAAAAAABAAAAAAACAAAAAAADAAAAAAA. Though there is less spiking, all deviant stimuli trigger a response (p=0.044, 0.006, 0.035). G: Gap stimulus presented every 8th pulse, where instead of a deviant stimulus, no stimulus pattern was delivered. There is no tuning and activity rapidly drops to zero. H. Lack of STP in Glu removes mismatch detection, and spike rates are high. I. Lack of STP in GABA results in sustained strong inhibition and sparse firing, but some mismatch selectivity remains for the first two transitions (p=0.016, 0.003, 0.924). J. Complete removal of STP results in sustained brisk firing and no mismatch detection. K: Control with GABA synapses inactivated. There is elevated firing but the first pattern transition is still resolved (p = 0.006, 0.129, 0.851). L: Control with uniform pattern (A repeated 32 times). The spiking response rapidly drops to close to zero. M. Control with randomized pattern (one of pattern A to pattern E) on each pulse. Spiking response continues without any tuning, but with slow decay over the 32 pulses. N: Mismatch response (methods) is tuned to frequency of pattern repeats. The strongest responses are in the gamma range between 20 and 100 Hz. O: Schematic of gamma burst stimulus (green) repeated at theta frequency of 7.69 Hz (blue). Each burst consisted of a pattern repeated 5 times at 100 Hz. P: Theta modulated gamma burst responses showed strong spiking responses when each burst had a different pattern. Q: Spiking in second, third and fourth theta cycles was lower than the first when each burst had the same pattern (p=0.057, 0.029, 0.029). R, S: Continued strong firing when different patterns are presented in each of the five pulses of theta-modulated gamma. There is a small decline over successive theta cycles. R: No precession. The last burst is smaller than the first (p=0.042). S: Precession present. The last burst is again smaller than the first (p=0.029).

Discussion

We probed CA3-CA1 input-output properties for functional outcomes of pattern-specific short-term plasticity using spatially and temporally patterned optical stimuli. EI balance and summation properties tilted briefly to excitation on the second or third pulse of a burst due to differential STP profiles for excitation and inhibition. We parameterized a molecule-to-network multiscale model of CA3-CA1 network and plasticity using a series of burst, summation, and Poisson train inputs of optically defined input patterns. The model predicted that STP dynamics of E and I inputs provide a mechanism for single cells to detect transitions in input pattern sequences, which we then confirmed experimentally. We used the model to explore network configurations which could detect such transitions, and showed that spiking is strongly mismatch-tuned in continuous as well as theta-modulated gamma bursts.

EI balance is dynamically tuned through differential STP of E and I inputs

EI balance is a ubiquitous mechanism for controlling neural excitability. Numerous processes, including inhibitory plasticity, homeostatic plasticity, and developmental co-tuning have been proposed to underlie the origin and maintenance of EI balance [30–38]. EI balance is dynamic, due in part to differential effects of STP on excitatory and inhibitory inputs to the postsynaptic neuron [13–15,39,40]. Using field electrode stimulation of the Shaffer collaterals, these studies report a sustained shift to excitation during burst input[13]. In contrast, our sparse optical patterned stimuli results in a small window of escape from EI balance shortly after the start of a burst (Figure 3), following which both E and I undergo depression to restore balance (Figure 3, 8). Thus, short-term plasticity adds another layer of timing control through gating of E-I balance [41].

Adding further detail to the changes in EI balance, we find that short-term plasticity (STP) modulates neuronal summation over the course of a burst. Specifically, the strength of subthreshold divisive normalisation (SDN)[6] evolves in a frequency-dependent manner, and is most linear around pulse 3 of a burst (Figure 3).

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 (~20x slower than real time on an AMD 6800HS). Feature-wise, several studies combine signalling (typically postsynaptic spine and dendrite signalling) with electrical signalling[42–47]. Some of these combine compartmentalized signalling with standard branched neuronal electrophysiology calculations, as we do. Others include detailed 3-dimensional reaction-diffusion kinetics on spines[48] or boutons[49,50] and these run much more slowly. Our current model is distinct in that it incorporates the network, as well as modeling a population of stochastic chemical synapses on the dendrite of a spiking CA1 pyramidal neuron.

Our network model is much reduced compared to models with exhaustive cellular and network-level detail[51]. Its simplicity enables extensive exploration of the network parameters and comparison with recorded activity under a series of well-controlled stimulus patterns (Figures 4-7).

Our model is modular. One can swap out different cellular geometries, provide distinct synaptic signaling models using SBML or compatible standards, and alter the detail of the network components. In the current study, we employed four different spine kinetic models, with different combinations of STP in E and I, simply by loading in different kinetic definitions. Thus the model is consistent with the 'Interoperable' and 'Reusable' aspects of FAIR principles [52].

STP leads to single-neuron mismatch detection in pattern sequences

Where can temporally precise gating be useful? At the single-synapse level, synaptic depression performs the operation of decorrelation, leading to more efficient coding [41,53]. Our observations can be framed as decorrelation of the spatial pattern (alternatively, activity vector) of multi-synaptic input converging onto a neuron[3]. A putative example of this is place field formation in novel environments. Dombeck and co-workers [54,55], found that dendritic inhibition is reduced when a mouse traverses a new environment. Consistent with this, our simulations predict elevated spiking when novel patterns occur, in part due to presynaptic depression of inhibition (Figure 7,8).

There is a vigorous debate on the mechanisms of auditory mismatch negativity (e.g., [56–58]) of which the fresh-afferent model maps to part of our proposed mechanism (Fig 7A,B). Leaving aside the obvious differences between auditory cortex and hippocampus, we frame our model as a transient differential tilt in EI balance (Fig 3, Fig 8 H-K), rather than just afferent adaptation. We note that the time-scale of our mismatch response (~125 ms) is similar to the 100 to 200 ms of the classic auditory mismatch negativity [59]. A similar computation of stimulus-specific adaptation in the auditory cortex has been modeled using a recurrent network with short-term depression[60]. Our findings of pattern mismatch detection may also provide a mechanistic interpretation of published results showing behavioural and physiological selective responses to changes in visually delivered patterns[54]. The authors found that a rate model with short-term depression was able to replicate the experiments. In a recent study, signatures of prediction error (implemented by absence of conditioned response) were observed in recurrent cortical networks in organotypic preparation[61]. These prediction errors were visible after a 24-hour long training period that achieved the associative conditioning. Thus the study showed the role of inhibition stabilized attractor networks of cortical circuits to produce prediction-error like responses in single neurons. Our model achieves equivalent computations with simple feedforward connectivity at a single-neuron level, without training, and with direct linkages to the underlying physiological mechanisms.

Both our experiments and models show selectivity for specific transitions. This suggests a further computation of sequence disambiguation. Consider sequence S1 AAAABBBBCCCC and sequence S2 FFFFBBBBGGGG, overlapping in ensemble B. We stipulate the presence of neuron $S1_{AB}$ which detects the A->B transition, and so on. Then S1 is uniquely identified by the activity of $S1_{AB}$,$S1_{BC}$, and S2 by $S2_{FB}$,$S2_{BG}$. This is a rapid and parsimonious mechanism for distinguishing two intersecting sequences such as time-cell sequences [62], on sub-second time-scales.

Mismatch detection is pronounced in the gamma band

Gamma oscillations have been implicated in memory formation and communication between brain regions [63,64]. Gamma bursts are typically modulated by the theta rhythm [29] , which is itself closely coupled to sensory intake such as sniffing[65,66] and whisking[67], and also to movement [68]. Our proposed mechanism for mismatch detection is enhanced precisely in the gamma range (Figure 8). Further, we find that theta-modulated gamma bursts perform mismatch detection between theta cycles (Figure 8P,Q). We propose that sensory snapshots obtained through sniffs or whisks provide the CA1 with theta-modulated gamma bursts of stimulus-specific activity vectors. When inputs change from one theta cycle to the next, individual cells in CA1 light up to detect specific changes between stimuli, such as odor 1 changing to odor 2, or texture 1 to texture 2 (Figure 7, 8). This implements a dimensionality reduction of a sequence of sensory inputs into a trajectory through a space of stimulus transitions.

Methods

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)[69].

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

672 puller (P1000, Sutter Instruments) and filled with an internal solution containing 130mM
673 K-gluconate, 5mM NaCl, 10mM HEPES, 1mM Na_4-EGTA, 2mM $MgCl_2$, 2mM Mg-ATP,
674 0.5mM Na-GTP and 10mM Phosphocreatine, pH ~7.3, and osmolarity ~290 mOsm. For
675 voltage-clamp recordings, K-gluconate was replaced with Cs-Gluconate. Field response
676 was often biphasic with multiple cycles, hence, a peak-to-peak value was used.
677 The cells were either recorded in the current clamp or the voltage clamp. In current
678 clamp recordings, the cells were maintained at a membrane potential of -70 mV. Cells
679 were rejected if the membrane potential changed by 5mV over the course of the
680 recording, holding current either was above 100 pA or changed by more than 25 pA to
681 maintain the mentioned membrane potential. In voltage-clamp recordings, the cells
682 were rejected if the series resistance crossed 25 MΩ or changed by 30% of the initial
683 value. Liquid junction potential was calculated to be 16.926 mV according to the
684 stationary Nernst–Planck equation [70] using LJPcalc (RRID: SCR_025044) and was not
685 compensated.
686 In optogenetic preparations, the additional stimulation that comes in due to the
687 excitation of axons and dendrites can not be ruled out. To assess that, we patched a
688 CA3 cell and recorded its optical depolarization caused by a single spot of light in a grid
689 using a whole-cell voltage clamp. The CA3 cells have a large receptive field that arises
690 due to the expression of ChR2 on their dendritic arbour. This receptive field is stronger
691 closer to the soma but spans on both sides of the dendritic tree (Extended Figure 1A). A
692 patched CA3 cell fired action potential with 100% success (Extended Figure 1B) when
693 the cell layer was stimulated with patterns with multiple spots. A single spot, on the
694 other hand, was only sufficient to induce action potentials when directly incident on
695 soma (Extended Figure 1B).

696 Optical Stimulation

697 Pyramidal cells in the CA3 cell layer expressing channelrhodopsin were stimulated
698 using a DMD-based patterned projector system (Polygon 400G, Mightex Systems,
699 Canada) using a spot size of approximately 14μm x 7.8μm under a 40x water
700 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. There were five different non-overlapping patterns of 5-spots, and three non-overlapping patterns of 15-spots. 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.

Code and Model accessibility

All scripts and models for performing simulations for figures are presented in the GitHub site https://github.com/BhallaLab/STP_EI_paper_figs. Simulations were run using MOOSE 4.0 compiled using gcc version 11.4.0 and Python version 3.10. Short calculations were run on Ubuntu 22.04 on an 8-core laptop, and large runs were done on a 128-core server also running Ubuntu 22.04. All software used is open-source and hosted on github. MOOSE is available at https://moose.ncbs.res.in and https://github.com/BhallaLab/moose. Optimization of presynaptic models was performed using HOSS 1.0[71] (https://github.com/BhallaLab/HOSS) and FindSim[72] (https://github.com/BhallaLab/FindSim).

Presynaptic Plasticity Model

Voltage clamp data from the pulse-train experiments (Figure 3) was used to train a model of presynaptic short-term plasticity (Figure 4A, Extended Figure 3, Supplementary fig 2, 3) implemented in MOOSE [23]. 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) [24] 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 (depending on how many were feasible in each patch recording), 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, and 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 on the dendrite of a ball-and-stick cellular model.

The synaptic input was obtained using mass-action simulations of presynaptic signalling for the model in Figure 4A. 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 4B) was a ball-and-stick model with a soma compartment of 10 μm length and 10 μm diameter and a single dendritic compartment of 200 μm length and 2 μm 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.
Under the passive conditions of our model (RM =1.0 $\Omega.m^2$, RA = 1.0 $\Omega.m$) the length constant of this model was ~700 μm, which is much longer than the geometrical length. Hence space clamp problems are unlikely. Furthermore, the short electrotonic length means it does not matter if the Glu and GABA receptors are interspersed as opposed to GABA being more proximal.
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 spiking simulations, we increased Na and K to 400 and 450 Siemens/m^2. 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 as indicated in the table. Similarly, projections from CA3 to interneurons were defined as a 256x256 matrix with connection probability pCA3_Inter. Projections from the interneurons to the 200 GABA synapses on the CA1 neuron were a 256x200 matrix with default connection probability pInter_CA1. All told, there were just seven parameters in this model (Table 1).

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. We expanded the 8x8 patterns into 16x16 patterns by replicating the pattern twice over the array (Extended Fig 5). 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. For the default level of sparseness (75%) the first five patterns (A through E) had 18 to 26 nonzero

squares. The last three patterns (F through H) were dense, with 144, 152 and 152 squares.

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%). For the default 75% sparseness this came to between 35 and 38 active glutamatergic synapses (Extended Figure 10).

During the simulation runs we provided a time-series of patterns according to frozen Poisson timing series with a mean of 20 Hz (Figure 6) or a time-series of four patterns each repeated 8 times, at 8, 20 and 50 Hz (Figure 7). 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:

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

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

Where

$G = \text{conductance of ChR2}$

$V_m = \text{CA3 pyramidal neuron membrane potential}$

$\Delta t = \text{time since last stimulus}$

$D_0 = \text{basal desensitization of ChR2} = 0.01$

$D_1 = \text{light pulse induced desensitization of ChR2} = 0.0004$

$\tau_R = \text{Recovery time course of ChR2 from desensitization} = 1.5 \text{ s}$

$\tau_M = \text{Membrane time course of CA3 neuron} = 10 \text{ ms}$

$\tau_C = \text{Time course to charge CA3 neuron when ChR2 is open} = 5 \text{ ms}$

$\tau_O = \text{ChR2 channel open time} = 1 \text{ ms}$

Example pulse trains with these parameters are illustrated in Extended Figure 6.

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, 6, 7 and 8 were generated from these data files using pandas, Matplotlib and Python scripts.

3-D model visualisation (Fig 4 and Supplementary movie) was performed using the built-in interface within MOOSE, to the 3-D graphics library *vpython*.

Funding Information

AA and USB are at NCBS-TIFR which receives the support of the Department of Atomic Energy, Government of India, under Project Identification No. RTI 4006. The study received funding from SERB Grant CRG/2022/003135-G.

Declaration of competing interests:

The authors declare no competing interests.

Author Contribution:

Aditya Asopa: Conceptualization, Data curation, Analysis, Investigation, Visualization, Methodology, Writing original draft, Writing review and editing;
Upinder Singh Bhalla: Conceptualization, Resources, Simulation, Visualization, Analysis, Software, Supervision, Funding acquisition, Writing original draft, Project administration, Writing review and editing

References

1. Kanamaru, T., Hensch, T. K. & Aihara, K. Maximal Memory Capacity Near the Edge of Chaos in Balanced Cortical E-I Networks. *Neural Computation* **35**, 1430–1462 (2023).
2. Denève, S. & Machens, C. K. Efficient codes and balanced networks. *Nat Neurosci* **19**, 375–382 (2016).
3. Kremkow, J., Aertsen, A. & Kumar, A. Gating of Signal Propagation in Spiking Neural Networks by Balanced and Correlated Excitation and Inhibition. *J. Neurosci.* **30**, 15760–15768 (2010).
4. van Vreeswijk, C. & Sompolinsky, H. Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity. *Science* **274**, 1724–1726 (1996).
5. Vogels, T. P. & Abbott, L. F. Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nat Neurosci* **12**, 483–491 (2009).
6. Bhatia, A., Moza, S. & Bhalla, U. S. Precise excitation-inhibition balance controls gain and timing in the hippocampus. *eLife* (2019) doi:10.7554/elife.43415.
7. Wehr, M. & Zador, A. M. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* **426**, 442–446 (2003).
8. Fortune, E. S. & Rose, G. J. Short-term synaptic plasticity as a temporal filter. *Trends in Neurosciences* **24**, 381–385 (2001).
9. Regehr, W. G. Short-Term Presynaptic Plasticity. *Cold Spring Harbor Perspectives in Biology* **4**, a005702–a005702 (2012).
10. Tsodyks, M. V. & Markram, H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci.*

U.S.A. **94**, 719–723 (1997).

11. Zucker, R. S. & Regehr, W. G. Short-Term Synaptic Plasticity. *Annual Review of Physiology* **64**, 355–405 (2002).

12. Galarreta, M. & Hestrin, S. Frequency-dependent synaptic depression and the balance of excitation and inhibition in the neocortex. *Nat Neurosci* **1**, 587–594 (1998).

13. Klyachko, V. A. & Stevens, C. F. Excitatory and feed-forward inhibitory hippocampal synapses work synergistically as an adaptive filter of natural spike trains. *PLoS Biology* **4**, 1187–1200 (2006).

14. Bartley, A. F. & Dobrunz, L. E. 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 (2015).

15. Sun, H. Y., Li, Q., Bartley, A. F. & Dobrunz, L. E. 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 (2018).

16. Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M. & Tanila, H. The Hippocampus, Memory, and Place Cells. *Neuron* **23**, 209–226 (1999).

17. Yuste, R., Cossart, R. & Yaksi, E. Neuronal ensembles: Building blocks of neural circuits. *Neuron* **112**, 875–892 (2024).

18. Rolls, E. T. & Kesner, R. P. A computational theory of hippocampal function, and empirical tests of the theory. *Progress in Neurobiology* **79**, 1–48 (2006).

19. Asopa, A. & Bhalla, U. S. A computational view of short-term plasticity and its

implications for E-I balance. *Current Opinion in Neurobiology* **81**, 102729 (2023).

20. Barros-Zulaica, N. *et al.* Estimating the Readily-Releasable Vesicle Pool Size at Synaptic Connections in the Neocortex. *Front. Synaptic Neurosci.* **11**, 29 (2019).

21. Neher, E. Merits and Limitations of Vesicle Pool Models in View of Heterogeneous Populations of Synaptic Vesicles. *Neuron* **87**, 1131–1142 (2015).

22. Schneggenburger, R., Meyer, A. C. & Neher, E. Released Fraction and Total Size of a Pool of Immediately Available Transmitter Quanta at a Calyx Synapse. *Neuron* **23**, 399–409 (1999).

23. Ray, S. & Bhalla, U. S. PyMOOSE: Interoperable Scripting in Python for MOOSE. *Front Neuroinform* **2**, 6 (2008).

24. Viswan, N. A., Tribut, A., Gasparyan, M., Radulescu, O. & Bhalla, U. S. Hierarchical Optimization of Biochemical Networks. Preprint at https://hal.science/hal-04593669 (2024).

25. Gillespie, D. T. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**, 2340–2361 (1977).

26. Bezaire, M. J. & Soltesz, I. Quantitative assessment of CA1 local circuits: Knowledge base for interneuron-pyramidal cell connectivity: Quantitative Assessment Of Ca1 Local Circuits. *Hippocampus* **23**, 751–785 (2013).

27. Butler, R. A. Effect of Changes in Stimulus Frequency and Intensity on Habituation of the Human Vertex Potential. *The Journal of the Acoustical Society of America* **44**, 945–950 (1968).

28. Garrido, M. I., Kilner, J. M., Stephan, K. E. & Friston, K. J. The mismatch negativity: A review of underlying mechanisms. *Clinical Neurophysiology* **120**, 453–463 (2009).

29. Lisman, J. E. & Jensen, O. The Theta-Gamma Neural Code. *Neuron* **77**, 1002–1016 (2013).

30. Agnes, E. J., Luppi, A. I. & Vogels, T. P. Complementary Inhibitory Weight Profiles Emerge from Plasticity and Allow Flexible Switching of Receptive Fields. *The Journal of Neuroscience* **40**, 9634–9649 (2020).

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

32. Chen, L., Li, X., Tjia, M. & Thapliyal, S. Homeostatic plasticity and excitation-inhibition balance: The good, the bad, and the ugly. *Current Opinion in Neurobiology* **75**, 102553 (2022).

33. Jia, D. W., Vogels, T. P. & Costa, R. P. Developmental depression-to-facilitation shift controls excitation-inhibition balance. *Commun Biol* **5**, 873 (2022).

34. Lagzi, F. & Fairhall, A. L. Emergence of co-tuning in inhibitory neurons as a network phenomenon mediated by randomness, correlations, and homeostatic plasticity. *Sci. Adv.* **10**, eadi4350 (2024).

35. Mackwood, O., Naumann, L. B. & Sprekeler, H. Learning excitatory-inhibitory neuronal assemblies in recurrent networks. *eLife* **10**, (2021).

36. Miehl, C. & Gjorgjieva, J. Stability and learning in excitatory synapses by nonlinear inhibitory plasticity. *PLoS Computational Biology* **18**, (2022).

37. Trapp, P., Echeveste, R. & Gros, C. E-I balance emerges naturally from continuous Hebbian learning in autonomous neural networks. *Scientific Reports* **8**, (2018).

38. Wu, Y. K., Miehl, C. & Gjorgjieva, J. Regulation of circuit organization and function

through inhibitory synaptic plasticity. *Trends in Neurosciences* **45**, 884–898 (2022).

39. Bartley, A. F. *et al.* Interneuron Transcriptional Dysregulation Causes Frequency-Dependent Alterations in the Balance of Inhibition and Excitation in Hippocampus. *The Journal of Neuroscience* (2015) doi:https://doi.org/10.1523/jneurosci.1834-15.2015.

40. Wu, Y. K. & Zenke, F. Nonlinear transient amplification in recurrent neural networks with short-term plasticity. *eLife* **10**, e71263 (2021).

41. Rotman, Z., Deng, P.-Y. & Klyachko, V. A. Short-Term Plasticity Optimizes Synaptic Information Transmission. *Journal of Neuroscience* **31**, 14800–14809 (2011).

42. Bhalla, U. S. Synaptic input sequence discrimination on behavioral timescales mediated by reaction-diffusion chemistry in dendrites. *eLife* **6**, e25827 (2017).

43. Bhalla, U. S. Multiscale interactions between chemical and electric signaling in LTP induction, LTP reversal and dendritic excitability. *Neural Networks* **24**, 943–949 (2011).

44. Dainauskas, J. J., Marie, H., Migliore, M. & Saudargiene, A. GluN2B-NMDAR subunit contribution on synaptic plasticity: A phenomenological model for CA3-CA1 synapses. *Front. Synaptic Neurosci.* **15**, 1113957 (2023).

45. Dorman, D. B. & Blackwell, K. T. Synaptic Plasticity Is Predicted by Spatiotemporal Firing Rate Patterns and Robust to In Vivo-like Variability. *Biomolecules* **12**, 1402 (2022).

46. Mattioni, M. & Le Novère, N. Integration of Biochemical and Electrical Signaling-Multiscale Model of the Medium Spiny Neuron of the Striatum. *PLoS ONE* **8**, e66811 (2013).

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

48. Chen, W. *et al.* STEPS 4.0: Fast and memory-efficient molecular simulations of neurons at the nanoscale. *Front. Neuroinform.* **16**, 883742 (2022).

49. Nadkarni, S., Bartol, T. M., Stevens, C. F., Sejnowski, T. J. & Levine, H. Short-term plasticity constrains spatial organization of a hippocampal presynaptic terminal. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 14657–14662 (2012).

50. Nadkarni, S., Bartol, T. M., Sejnowski, T. J. & Levine, H. Modelling Vesicular Release at Hippocampal Synapses. *PLoS Comput Biol* **6**, e1000983 (2010).

51. Markram, H. *et al.* Reconstruction and Simulation of Neocortical Microcircuitry. *Cell* **163**, 456–492 (2015).

52. Eriksson, O. *et al.* Combining hypothesis- and data-driven neuroscience modeling in FAIR workflows. *Elife* **11**, e69013 (2022).

53. Goldman, M. S., Nelson, S. B. & Abbott, L. F. Decorrelation of spike trains by synaptic depression. *Neurocomputing* **26–27**, 147–153 (1999).

54. Sheffield, M. E. & Dombeck, D. A. Dendritic mechanisms of hippocampal place field formation. *Current Opinion in Neurobiology* **54**, 1–11 (2019).

55. Sheffield, M. E. J., Adoff, M. D. & Dombeck, D. A. Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation. *Neuron* **96**, 490-504.e5 (2017).

56. Garrido, M. I., Kilner, J. M., Stephan, K. E. & Friston, K. J. The mismatch negativity: A review of underlying mechanisms. *Clinical Neurophysiology* **120**, 453–463 (2009).

57. May, P. J. C. & Tiitinen, H. Mismatch negativity (MMN), the deviance-elicited auditory deflection, explained. *Psychophysiology* **47**, 66–122 (2010).

58. Näätänen, R., Jacobsen, T. & Winkler, I. Memory-based or afferent processes in mismatch negativity (MMN): A review of the evidence. *Psychophysiology* **42**, 25–32 (2005).

59. Butler, R. A. Effect of Changes in Stimulus Frequency and Intensity on Habituation of the Human Vertex Potential. *The Journal of the Acoustical Society of America* **44**, 945–950 (1968).

60. Yarden, T. S. & Nelken, I. Stimulus-specific adaptation in a recurrent network model of primary auditory cortex. *PLoS Comput Biol* **13**, e1005437 (2017).

61. Liu, B. & Buonomano, D. V. Ex Vivo Cortical Circuits Learn to Predict and Spontaneously Replay Temporal Patterns. Preprint at https://doi.org/10.1101/2024.05.30.596702 (2024).

62. Modi, M. N., Dhawale, A. K. & Bhalla, U. S. CA1 cell activity sequences emerge after reorganization of network correlation structure during associative learning. *eLife* **3**, e01982 (2014).

63. Van Vugt, M. K., Schulze-Bonhage, A., Litt, B., Brandt, A. & Kahana, M. J. Hippocampal Gamma Oscillations Increase with Memory Load. *J. Neurosci.* **30**, 2694–2699 (2010).

64. Misselhorn, J., Schwab, B. C., Schneider, T. R. & Engel, A. K. Synchronization of Sensory Gamma Oscillations Promotes Multisensory Communication. *eNeuro* **6**, ENEURO.0101-19.2019 (2019).

65. Eichenbaum, H., Kuperstein, M., Fagan, A. & Nagode, J. Cue-sampling and

goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *J. Neurosci.* **7**, 716–732 (1987).

66. Tsanov, M., Chah, E., Reilly, R. & O'Mara, S. M. Respiratory cycle entrainment of septal neurons mediates the fast coupling of sniffing rate and hippocampal theta rhythm. *Eur J of Neuroscience* **39**, 957–974 (2014).

67. Grion, N., Akrami, A., Zuo, Y., Stella, F. & Diamond, M. E. Coherence between Rat Sensorimotor System and Hippocampus Is Enhanced during Tactile Discrimination. *PLoS Biol* **14**, e1002384 (2016).

68. Joshi, A. *et al.* Dynamic synchronization between hippocampal representations and stepping. *Nature* **617**, 125–131 (2023).

69. Atasoy, D., Aponte, Y., Su, H. H. & Sternson, S. M. A FLEX Switch Targets Channelrhodopsin-2 to Multiple Cell Types for Imaging and Long-Range Circuit Mapping. *The Journal of Neuroscience* **28**, 7025–7030 (2008).

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

71. Viswan, N. A., Tribut, A., Gasparyan, M., Radulescu, O. & Bhalla, U. S. Hierarchical Optimization of Biochemical Networks. (2024).

72. Viswan, N. A., HarshaRani, G. V., Stefan, M. I. & Bhalla, U. S. FindSim: A Framework for Integrating Neuronal Data and Signaling Models. *Front. Neuroinform.* **12**, 38 (2018).