Bidirectional synaptic plasticity rapidly modifies hippocampal representations independent of correlated activity

Aaron D. Milstein¹, Yiding Li^{2*}, Katie C. Bittner^{3*}, Christine Grienberger², Ivan Soltesz¹, Jeffrey C. Magee^{2†} and Sandro Romani^{3†}

¹Dept. of Neurosurgery and Stanford Neurosciences Institute, Stanford University School of Medicine, Stanford, CA

²Howard Hughes Medical Institute, Baylor College of Medicine, Houston, TX

³Howard Hughes Medical Institute, Janelia Research Campus, Ashburn, VA

*These authors contributed equally to this work.

†Co-corresponding authors. Please send correspondence to:

Jeffrey C. Magee: Jeffrey.Magee@bcm.edu

Sandro Romani: romanis@janelia.hhmi.org

Abstract

According to standard models of synaptic plasticity, correlated activity between connected neurons drives changes in synaptic strengths to store associative memories. Here we tested this hypothesis *in vivo* by manipulating the activity of hippocampal place cells and measuring the resulting changes in spatial selectivity. We found that the spatial tuning of place cells was rapidly reshaped via bidirectional synaptic plasticity. To account for the magnitude and direction of plasticity, we evaluated two models – a standard model that depended on synchronous pre- and post-synaptic activity, and an alternative model that

depended instead on whether active synaptic inputs had previously been potentiated. While both models accounted equally well for the data, they predicted opposite outcomes of a perturbation experiment, which ruled out the standard correlation-dependent model. Finally, network modeling suggested that this form of bidirectional synaptic plasticity enables population activity, rather than pairwise neuronal correlations, to drive plasticity in response to changes in the environment.

Main Text

Activity-dependent changes in synaptic strength can flexibly alter the selectivity of neuronal firing for particular features of the environment, providing a cellular substrate for learning and memory. Various forms of Hebbian synaptic plasticity have been considered for decades to be the main or even only synaptic plasticity mechanisms present within most brain regions of a number of species. The core feature of such plasticity mechanisms is that they are autonomously driven by the repeated presence of correlated presynaptic and postsynaptic activity that leads to either increases or decreases in synaptic strength depending on the exact temporal coincidence (1-4).

The hippocampus plays an important role in many forms of learning and memory, and the spatial firing rates of hippocampal place cells have been shown to change with alterations in environmental context or the locations of salient features, like reward (5-11). Furthermore in CA1 neurons, place cell activity can emerge in a single trial following a dendritic calcium spike (also called a plateau potential) (12-14). The form of synaptic plasticity responsible for this rapid change in selectivity, termed behavioral timescale

synaptic plasticity (BTSP), modifies synaptic inputs active within a multi-second time window around the plateau potential. That BTSP strengthens many synaptic inputs whose activation did not cause or even coincide with postsynaptic activity suggests that it might be a fundamentally different form of plasticity than classical correlation-driven Hebbian plasticity (1-3). Such a plasticity rule could enable representation learning in cortical brain regions like the hippocampus to be guided by delayed behavioral outcomes, rather than by short timescale associations of neuronal input and output. However, it was not clear from previous experiments if short timescale correlations would modulate changes in synaptic strength induced by BTSP, which may reveal similarities with other correlative forms of plasticity.

In the current study, we sought to directly determine the dependence of BTSP on the correlation of presynaptic activity and postsynaptic depolarization in individual place cells. Intracellular voltage recordings from CA1 place cells were established in head-fixed mice trained to run for a water reward on a circular treadmill decorated with visual and tactile cues to distinguish spatial positions (187 cm in length). We began by examining how the induction of BTSP changes the membrane potential (V_m) dynamics in neurons already exhibiting location specific firing (i.e. place cells). To do so we injected brief step currents (300 ms duration) through the intracellular electrode for a small number of consecutive laps (5-8 laps; Fig. 1, A to C) to evoke plateau potentials at a second location along the track some distance from the initial place field (labeled “Induction 2” in Fig. 1, A to H). We observed that the plasticity induced by experimentally-evoked dendritic plateaus both increased V_m ramp amplitude near the plateau induction position, and also decreased ramp amplitude at the peak location of the original place field (Fig. 1, B and

Fig 1.

Fig. 1. Dendritic calcium spikes translocate hippocampal place fields via bidirectional synaptic plasticity. (A) Spatial firing of a CA1 pyramidal cell recorded intracellularly from a mouse running laps on a circular treadmill. Dendritic plateau potentials evoked by intracellular current injection first induce a place field at ~120 cm

(Induction 1), then induce a second place field at ~10 cm and suppress the first (Induction 2). **(B)** Black: intracellular V_m traces from individual laps in (A); blue: example low-pass filtered V_m ramp superimposed on unfiltered trace, and duplicated with expanded scale (inset). **(C)** Animal position vs. time during laps in which place fields were induced by evoked plateaus (locations marked with colored circles). t refers to the inter-event intervals between traversals of the initial place field peak location and evoked plateaus. **(D)** Spatially binned subthreshold V_m ramp depolarizations averaged across laps after plasticity induction. Colored dashes mark the average locations of evoked plateaus. **(E)** For each position, induced changes in V_m ramp amplitude are plotted against the time interval between plateau onset and traversal of that position during plasticity induction laps. **(F – H)** Same as (C – E) for a different example cell in which the amplitude of the original place field was not reduced following the second plasticity induction. **(I)** For all recorded neurons with a pre-existing place field in which plasticity was induced at a second location ($n=13$), changes in ramp amplitude are compared to initial ramp amplitude for each spatial bin (1.87 cm). Induced changes in ramp amplitude are negatively correlated with initial ramp amplitude. Explained variance (R^2) and statistical significance ($p < 0.05$) reflect Pearson's correlation and a 2-tailed null hypothesis test. **(J)** For the same neurons as in (I), changes in ramp amplitude are compared to time to plateau onset. Trace color indicates initial ramp amplitude before plasticity induction. **(K)** Two-dimensional Gaussian regression and interpolation of data from all recorded plasticity inductions (20 inductions from 13 cells) was used to estimate the plasticity rule that relates initial ramp amplitude and time to plateau onset to induced changes in ramp amplitude (trace color).

D). The time course of the V_m changes showed that decreases in ramp amplitude occurred at positions in space that were traversed multiple seconds before or after induced plateaus (Fig. 1, E and J). Furthermore, the exact magnitude of decreases in V_m ramp amplitude was greatest at spatial positions where initial ramp amplitude was largest (Fig. 1I). Interestingly, in a subset of recordings the ramp amplitude at the original place field location was not reduced (Fig. 1G). Inspection of the animals' run trajectories during such instances revealed that long pauses in running just before the plateau induction position on multiple laps "protected" the original place field from depression by excluding the underlying location-selective inputs from the plasticity time window (Fig. 1, F to H). When both initial ramp amplitude and relative input timing are considered, it is apparent that the preferred conditions for large synaptic depression are that spatial inputs 1) have already been strengthened by previous plasticity, resulting in elevated postsynaptic depolarization at the time of presynaptic spikes, and 2) are activated within a time window ~2 – 4 seconds away from a plateau (Fig. 1J, trace color indicates initial ramp amplitude; Fig 1K; two-dimensional interpolation from data, trace color indicates change in ramp amplitude, see Materials and Methods). In summary, BTSP can either strengthen or weaken synapses in a small number of trials, providing a bidirectional learning mechanism capable of both rapid memory storage and erasure.

The above analysis revealed a relationship between initial V_m ramp amplitude and bidirectional changes in V_m depolarization induced by BTSP (Fig. 1, I to K), though it was inverted compared to most common formulations of correlative plasticity in which small depolarizations induce synaptic depression and large depolarizations induce synaptic potentiation (3, 15-17). We next sought to investigate this possible causal

relationship between postsynaptic depolarization and plasticity induced by BTSP. First, we formulated a set of mathematical models of the underlying synaptic learning rule to generate testable hypotheses and to predict changes in synaptic strength given the following quantities: presynaptic spike times, postsynaptic depolarization, postsynaptic plateaus, and the strengths of each synaptic input before each plateau.

We compared two classes of plasticity models - a standard model dependent on coincident presynaptic spiking and postsynaptic depolarization (correlative), and an alternative model dependent instead on the timing of presynaptic spikes and the strengths of each synaptic input at the time of activation (non-correlative) (see Materials and Methods for details). To account for the long time course of BTSP, both models required temporal filters of synaptic activity to generate slow biochemical intermediate signals marking synapses as eligible for either synaptic potentiation or depression (local synaptic eligibility; Fig. 2, A to C; fig. S1, C and D; fig. S2, C and D) (1, 3, 15, 16, 18). Biologically, these traces could correspond to the enzymatic activity of calcium-dependent kinases and phosphatases, and post-translational modification and synaptic localization of proteins that regulate synaptic function (18-23). While in the voltage-dependent model, the amplitudes of these eligibility signals were modulated by the value of postsynaptic depolarization at the time of presynaptic activation (Fig. 2A and fig. S1, A, C and D), in the weight-dependent model (Fig. 2C and fig. S2, C and D), eligibility for plasticity depended only on presynaptic firing rate. Both models also required a temporal filter of the plateau potential to generate a second intermediate signal that extended in time long enough to interact with synaptic activity occurring up to seconds after a plateau potential (global dendritic instructive signal; Fig. 2, A to C; fig.

Fig 2.

Fig. 2. Two alternative models of bidirectional synaptic plasticity account for

measured changes in V_m . (A) Diagram depicts a “voltage-dependent” model of

bidirectional BTSP. Three factors influence changes in synaptic strength at each input:

1) presynaptic firing rate and timing, 2) postsynaptic V_m depolarization at the time of presynaptic spiking, and 3) postsynaptic plateau timing and duration. The product

(degree of correlation) of presynaptic firing rate and postsynaptic depolarization determines the amplitude of long duration “synaptic eligibility signals” that mark each

synapse as eligible for later synaptic potentiation or depression. Synaptic eligibility

signals, following an additional nonlinear transformation, are later converted into

changes in synaptic strength when in the presence of a second required “instructive

signal” generated downstream of postsynaptic plateaus. (B) Example traces depict the

signals described in (A) for a single presynaptic input onto a neuron that exhibited a pre-

existing place field before plateau induction. Shown is a single lap on the circular treadmill for a trial in which a plateau was evoked by intracellular current injection. Top: while presynaptic firing at this input (red) does not overlap in time at all with the postsynaptic plateau (black), it does coincide with the spatially-tuned depolarization underlying the cell's initial place field (grey). Middle: this generates long duration eligibility signals (blue: potentiation eligibility; red: depression eligibility) that overlap in time with the delayed instructive signal (black) (shading marks area of signal overlap). Bottom: at this input a large rate of synaptic depression and a small rate of synaptic potentiation result in a net decrease in synaptic strength. **(C)** Diagram depicts an alternative “weight-dependent” model of bidirectional BTSP. Three factors influence changes in synaptic strength at each input: 1) presynaptic firing rate and timing, 2) postsynaptic plateau timing and duration, and 3) the current synaptic weight of each input before each evoked plateau. In this model, synaptic eligibility signals depend only on presynaptic firing. Like in (A), a plateau-related instructive signal is required to convert synaptic eligibility signals into changes in synaptic weight. However, in this model, the current weight of each input influences the magnitude and direction of synaptic plasticity such that weak synapses favor potentiation, and strong synapses favor depression. **(D)** The voltage-dependent model was optimized to generate predicted ramp depolarizations for each neuron in the experimental dataset (20 inductions from 13 neurons). Changes in ramp amplitude are compared to time to plateau onset, with color indicating initial ramp amplitude (compare to Fig. 1J). **(E)** Estimate of plasticity rule obtained by regression and interpolation of model simulation data (compare to Fig. 1K). **(F – G)** Same as (D – E) for the weight-dependent model.

S1, C and D; fig. S2, C and D). This plateau-related “instructive signal” was broadcast globally to all synapses, and was required for plasticity. Changes in synaptic weight in the models occurred only during periods of temporal overlap between localized eligibility signals and the global instructive signal (Fig. 2B; fig. S1, C and D; fig. S2, C and D).

Model parameters (fig. S3, A to K and fig. S4, A to K, see Materials and Methods) were optimized by minimizing the difference between measured and predicted V_m ramp depolarizations. Both model variants generated predictions in good agreement with experimental data (fig. S3, L to O; fig. S4, L to O; and fig. S5), which underscores the importance of the components common to both models – long timescale intermediate signals downstream of synaptic activation that are transformed into bidirectional changes in synaptic weight by dendritic plateau potentials. However, the two models made qualitatively different predictions about the causal role of the activation state of the postsynaptic neuron in controlling the magnitude and direction of plasticity. While in the voltage-dependent model, correlation between presynaptic activity and postsynaptic depolarization influences the sign of plasticity, in the alternative model, the sign of plasticity is independent of postsynaptic voltage, and is modulated instead by current synaptic weight such that weak synapses tend to potentiate and strong synapses tend to depress (3, 24-29).

We next sought to distinguish between the two model classes with an *in vivo* perturbation experiment where neuronal V_m was experimentally depolarized by intracellular current injection during plasticity induction trials. In otherwise silent CA1 neurons exhibiting no spatial tuning or spiking during treadmill running, we injected current through the intracellular recording pipette to depolarize the neuron at steady-state

by ~10 mV, which often exceeded threshold for spiking (Fig. 3A; see fig. S6, A to C for simulation results supporting significant dendritic depolarization during this manipulation). Then, for 4 – 5 consecutive laps, plateau potentials were experimentally induced by an additional large, brief step current (~ 300 ms) at a fixed location along the track. In all neurons tested, this procedure resulted in the emergence of a large place field near the plateau induction site, as evidenced by spiking and a large amplitude V_m ramp depolarization (Fig. 3, A and C). Consistent with previous control experiments converting silent cells to place cells without V_m depolarization (Fig. 3B) (13), we observed only increases, and no decreases in ramp amplitude at spatial positions surrounding the plateau location (Fig. 3C and fig. S6D). This absence of synaptic depression is inconsistent with the prediction of the standard voltage-dependent model (Fig. 3D), but instead favors the alternative model, which predicted that only previously potentiated inputs would be eligible for synaptic depression, independent of postsynaptic voltage.

These results strongly indicate that bidirectional BTSP is fundamentally different from other previously characterized forms of associative synaptic plasticity that depend on three factors – presynaptic spiking, postsynaptic voltage, and a delayed reinforcement signal (1, 18, 30, 31). A particular advantage of a voltage-independent plasticity rule (Fig. 2C and fig. S2) is that changes in strength at each synapse are determined independently by signals generated locally, whereas plasticity rules that depend on the global activation state of the postsynaptic neuron may not allow independent credit to be assigned to the subset of synapses that contributed to a desired outcome (32).

We next aimed to explore how this form of plasticity could impact memory storage at the network level. During goal-directed navigation, hippocampal neurons have been

Fig 3.

A

B

C

D

Fig. 3. Bidirectional BTSP is insensitive to postsynaptic depolarization. (A)

Intracellular V_m traces from individual laps in which plasticity was induced by experimentally-evoked plateau potentials in an otherwise silent CA1 cell. During plasticity induction laps, the neuron was experimentally depolarized by ~10 mV at steady-state with an intracellular current injection. On the background of this elevated depolarization at every spatial position, step current injections (300 ms) evoked plateau potentials at the same spatial position for five consecutive laps and induced a place field. (B – C) Place field ramp depolarizations induced by experimentally-evoked

plateaus (individual cells in grey). Similar to control neurons that were converted from silent cells to place cells without steady-state depolarization (n=25, average in black) (B), neurons that underwent plasticity induction during steady-state depolarization (n=5, average in purple) (C) exhibited only synaptic potentiation, and no synaptic depression, at all spatial positions. (D) The data regression and interpolation in Fig. 1K was used to predict the changes in ramp amplitude that would result if BTSP was either dependent on (red) or independent of (blue) postsynaptic voltage.

shown to preferentially acquire new place fields near behaviorally-relevant locations, and to translocate existing place fields towards those locations (8-10, 33). Based on previous evidence that plateau probability in CA1 is facilitated by long-range feedback inputs onto distal CA1 dendrites from entorhinal cortex (12, 34, 35), and diminished by dendrite-targeting inhibition (35-39), we constructed a network model of the CA1 microcircuit where the probability of plateau initiation and thus BTSP induction was regulated by feedback inhibition and an instructive input from entorhinal cortex (Fig. 4, A and B) (10, 40-44).

We simulated a virtual animal running at a constant velocity on a circular treadmill for three separate phases of exploration (Fig. 4C). The first phase (ten laps) simulated exploration of a novel environment. The next phase (ten laps) simulated a goal-directed search for a target placed at a single fixed location (90 cm). Finally, the stability of acquired spatial representations were assessed by five additional laps with the goal removed (Fig. 4C). At each time step (10 ms), instantaneous plateau probabilities were computed for each cell (Fig. 4B), determining which neurons would initiate a dendritic plateau and undergo plasticity, following the experimentally validated bidirectional synaptic learning rule from Fig. 2C.

During the first few laps of simulated exploration, CA1 pyramidal neurons rapidly acquired place fields that, as a population, uniformly tiled the track (Fig. 4; C, D, and F). As neurons increased their activity over time, feedback inhibition increased proportionally and prevented further plasticity (Fig. 4, A to C). During laps with a goal presented at a fixed location, an additional population of silent neurons acquired place fields nearby the goal location (Fig. 4, E and F), while a separate population shifted their place field

Fig. 4.

Fig. 4. Bidirectional BTSP enables rapid adaptation of population representations

in a network model. (A) Diagram depicts components of a hippocampal network model. A population of CA1 pyramidal neurons receives spatially tuned excitatory input from a population of CA3 place cells and an instructive input from entorhinal cortex (EC) that signals the presence of a behavioral goal. The output of CA1 pyramidal neurons recruits feedback inhibition from a population of interneurons. (B) The probability that model CA1 neurons emit plateau potentials and induce bidirectional plasticity is

negatively modulated by feedback inhibition. As the total number of active CA1 neurons increases (labeled “normalized population activity”), feedback inhibition increases, and plateau probability decreases until a target level of population activity is reached, after which no further plasticity can be induced (black). An instructive input signaling the presence of a goal increases plateau probability, resulting in a higher target level of population activity inside the goal region (red). **(C)** Each row depicts the summed activity of the population of model CA1 pyramidal neurons across spatial positions during a lap of simulated running. Laps 1-10 reflect exploration of a previously unexplored circular track. During laps 11-20, a goal is added to the environment at a fixed location (90 cm). During laps 21-25, the goal is removed for additional exploration of the now familiar environment. **(D – E)** Activity of individual model CA1 pyramidal neurons during simulated exploration as described in (C). **(D)** The activity of neurons are sorted by the peak location of their spatial activity following 10 laps of novel exploration. A fraction of the population remains inactive and untuned. **(E)** The activity of neurons after 10 laps of goal-directed search are first sorted by their original peak locations (left), and then re-sorted by their peak locations following exposure to the fixed goal (right). An increased fraction of neurons express place fields near the goal position. **(F)** Histogram depicts neurons recruited to express new place fields in each spatial bin (epochs: novel explore, blue; fixed goal, red; familiar explore, grey). **(G)** Histogram depicts absolute distance of translocated place fields. **(H)** Histogram depicts relative distance of translocated place fields to the goal location.

positions towards the goal (Fig. 4; E, G, and H). Overall this resulted in an increased proportion of place cells with place fields nearby the goal position (Fig. 4E). Simulated place cell activity remained stable in the final phase of exploration of the now familiar environment (Fig. 4, C and F to H). These network modeling results recapitulate experimentally observed statistics of CA1 place cell translocation during goal-directed behavior (9).

This model provides a proof of principle that bidirectional BTSP can enable populations of place cells to rapidly adapt their spatial representations to changes in the environment without any compromise to spatial selectivity. An important feature of this model is that plasticity is regulated by the global activity of populations of neurons, rather than by pairwise correlations between single neurons and their inputs. This allows the network to rapidly acquire population-level representations of previously unencountered stimuli, as well as to modify outdated representations to better reflect changes to behaviorally relevant stimuli like goal location. Interestingly, bidirectional BTSP is asymmetric, tending to potentiate inputs more strongly when they are active before a plateau rather than during or after a plateau. In the network model, this caused the population representation of the goal to actually peak before the goal location itself (see also (6)), producing a predictive memory representation that could potentially be used by an animal to recall the path leading to the goal (45).

In summary, we found that dendritic plateaus could induce both potentiation and depression of subsets of synaptic inputs, resulting in translocation of a cell's place field towards the position where the plateaus were evoked. Quantitative inference of the underlying learning rule from the experimental data revealed that the direction and

magnitude of changes in synaptic strength depended on the current strength of each input at the time of a plateau, but not the degree of correlation between presynaptic and postsynaptic activity. In addition, bidirectional BTSP exhibited saturability (fig. S7) and state-dependence, two important features for stability of learned neuronal representations (3, 17, 27, 46, 47).

Together our experimental and modeling results establish bidirectional BTSP as a non-correlative mechanism for rapid and reversible learning. Rather than acting to autonomously reinforce pre-existing short timescale correlations between pre- and postsynaptic activity like standard Hebbian learning, bidirectional BTSP is capable of completely reshaping pairwise neuronal correlations in response to instructive input signals that promote dendritic plateau potentials. While it could be argued that BTSP is still a correlative form of plasticity due to its requirement that presynaptic spikes occur within a time window surrounding a plateau, this long timescale correlation is not between presynaptic spiking and the activation or output state of the postsynaptic neuron, but rather between presynaptic spiking and activation of a separate instructive input pathway that drives the postsynaptic plateau (12, 34, 35). Furthermore, since this long timescale coincidence between two inputs is only permissive for BTSP, but does not determine the sign of the plasticity, it cannot be classified as correlative in the same sense as classical Hebbian or even anti-Hebbian forms of plasticity. As suggested by our network model, if plateau potentials are generated by mismatch between a target instructive input and the output of the local circuit, as reflected by dendritically-targeted feedback inhibition, bidirectional BTSP can implement objective-based learning (48, 49). In addition to providing insight into the fundamental mechanisms of spatial memory formation in the

hippocampus, these findings suggest new directions for general theories of biological learning and the development of artificial learning systems (44, 50).

Materials and Methods

In vivo intracellular electrophysiology

All experimental methods were approved by the Janelia Institutional Animal Care and Use Committee (Protocol 12-84 & 15-126). All experimental procedures in this study, including animal surgeries, behavioral training, treadmill and rig configuration, and intracellular recordings, were performed identically to a previous detailed report (13) in an overlapping set of experiments, and are briefly summarized here.

In vivo experiments were performed in 6-12 week-old mice of either sex. Craniotomies above the dorsal hippocampus for simultaneous whole-cell patch clamp and local field potential (LFP) recordings, as well as affixation of head bar implants were performed under deep anesthesia. Following a week of recovery, animals were prepared for behavioral training with water restriction, handling by the experimenter, and addition of running wheels to their home cages. Mice were trained to run on the cue-enriched linear treadmill for a dilute sucrose reward delivered through a licking port once per lap (~187 cm). A MATLAB GUI interfaced with a custom microprocessor-controlled system for behavioral tracking and control. Position-dependent reward delivery and intracellular current injection were triggered by photoelectric sensors, and animal run velocity was measured by an encoder attached to one of the wheel axles. In a subset of experiments (Fig. 3), in addition to position-dependent step current to evoke plateau potentials, steady-state current was injected to depolarize neurons beyond

threshold for axosomatic action potentials during plasticity induction laps. While steady-state depolarization of the soma is expected to attenuate along the path to distal dendrites (51), the pairing of back-propagating action potentials with synaptic inputs has been shown to significantly amplify dendritic depolarization by inactivating A-type potassium channels and activating voltage-gated sodium channels and NMDA-type glutamate receptors (52-55). Simulations of a biophysically-detailed CA1 place cell model with realistic morphology and distributions of dendritic ion channels (39) suggest that steady-state somatic depolarization of a silent CA1 pyramidal cell *in vivo* results in levels of distal dendritic depolarization comparable to place cells at the peak of their place field (fig. S6, A to C).

To establish whole-cell recordings from CA1 pyramidal neurons, an extracellular LFP electrode was lowered into the dorsal hippocampus using a micromanipulator until prominent theta-modulated spiking and increased ripple amplitude was detected. Then a glass intracellular recording pipette was lowered to the same depth while applying positive pressure. The intracellular solution contained (in mM): 134 K-Gluconate, 6 KCl, 10 HEPES, 4 NaCl, 0.3 MgGTP, 4 MgATP, 14 Tris-phosphocreatine, and in some recordings, 0.2% biocytin. Current-clamp recordings of intracellular membrane potential (V_m) were amplified and digitized at 20 kHz, without correction for liquid junction potential.

Place field analysis

To analyze subthreshold V_m ramps, action potentials were first removed from raw V_m traces and linearly interpolated, then the resulting traces were low-pass filtered (<3 Hz).

For each of 100 equally sized spatial bins (1.87 cm), V_m ramp amplitudes were averaged across periods of 5-10 minutes of running laps on the treadmill. The spatially-binned ramp traces were then smoothed with a Savitzky-Golay filter with wrap-around. Ramp amplitude was quantified as the difference between the peak and the baseline (average of the 10% most hyperpolarized bins). For cells with a second place field induced, the same baseline V_m value determined from the period before the second induction was also used to quantify ramp amplitude after the second induction. Ramp width was quantified as the peak-normalized area under the curve. Plateau duration was estimated as the duration of intracellular step current injections, or as the full width at half maximum V_m in the case of spontaneous naturally-occurring plateaus. For each spatial bin, the elapsed time between traversal of that position and onset of a plateau was variable across induction laps, depending on lap-by-lap differences in run velocity (e.g. Fig. 1, C and F). To analyze the relationship between this time interval and changes in ramp amplitude, time relative to plateau onset (Fig. 1, E, H, J and K) was conservatively estimated as the minimum time delay across all induction laps. Since not all possible pairs of initial ramp amplitude and time delay relative to plateau onset were sampled in the experimental dataset, expected changes in ramp amplitude (Fig. 1K; Fig. 2; E and G; Fig. 3D; fig. S5; and fig. S8) were predicted from the sampled experimental or model data points by a two-dimensional Gaussian process regression and interpolation procedure using a rational quadratic covariance function, implemented in the open-source python package sklearn (56, 57).

Computational modeling

Two classes of mathematical models of the synaptic learning rule underlying bidirectional BTSP were built and optimized to predict the spatially tuned V_m ramp depolarizations of experimentally recorded CA1 place cells. All code (python) necessary to reproduce the modeling results is open-source and publicly available (58, 59). The following components and notation were shared across all model variants. CA1 place cell ramp depolarization ΔV as a function of position x was modeled as a weighted sum of the spatial firing rates of a population of 200 CA3 place cell inputs with place fields spaced uniformly across a 187 cm circular track, with a background level of depolarization V_b subtracted. This background is equivalent to the level of activation that would be produced if all inputs had a uniform weight of 1.

$$(1) \quad \Delta V(x) = \alpha \cdot \sum_i W_i \cdot R_i(x) - V_b$$

The firing rate R_i of CA3 place cell i with place field peak location y_i was modeled as a Gaussian function of position x , accounting for wraparound of the circular track with length ℓ :

$$(2) \quad d_i(x) = \begin{cases} -\ell + (x - y_i), & x - y_i \geq \frac{\ell}{2} \\ x - y_i, & -\frac{\ell}{2} < x - y_i < \frac{\ell}{2} \\ \ell + (x - y_i), & x - y_i \leq -\frac{\ell}{2} \end{cases}$$

$$(3) \quad R_i(x) = e^{-\left(\frac{d_i(x)}{\sigma}\right)^2}$$

The place field widths of CA3 inputs, controlled by σ , were set to have a full floor width ($3 \cdot \sqrt{2} \cdot \sigma$) of 90 cm (half-width of ~34 cm) throughout the study (60), though models tuned with alternative values of input field widths generated quantitatively similar predictions (fig. S8). Initial synaptic weights in silent cells before the first place field

induction were assumed to have a value of 1. Initial synaptic weights before the second place field induction in neurons already expressing a place field were estimated from the recorded ramp depolarization by least squares approximation. The scaling factor α was calibrated such that if the synaptic weights of CA3 place cell inputs varied between 1 and 2.5 as a Gaussian function of their place field locations, the postsynaptic CA1 cell would express a V_m ramp with 90 cm width and 6 mV amplitude, consistent with previous measurements of place field properties and the degree of synaptic potentiation by BTSP (13).

For each experimental recording, the position of an animal as a function of time during a plasticity induction lap j determined the firing rates of model CA3 inputs as a function of time $R_{j,i}(t) = R_{j,i}(x_j(t))$ (Fig. 2B, fig. S1B, and fig. S2B). In accordance with experimental data (12, 39), the firing rates of model place cell inputs decreased to zero during periods when the animal stopped running. Postsynaptic dendritic plateau potentials during each induction lap were modeled as binary functions of time $P_j(t)$. To generate long duration plasticity eligibility signals $\tilde{L}_{j,i}(t)$ specific to each synaptic input i and instructive plasticity signals $\tilde{P}_j(t)$ shared by all synapses, input firing rates and plateaus were convolved with causal temporal filters with exponential rise and decay (Fig. 2; fig. S1, C and D; fig. S2, C and D), denoted by a raised tilde.

$W_{j,i}$ corresponds to the synaptic weight of each input i prior to each plasticity induction lap j . Changes in synaptic weight $\Delta W_{j,i}$ were calculated once per lap by integrating a net rate of change of synaptic weight $\frac{dW_{j,i}(t)}{dt}$ (defined separately for each model below) over the duration of lap j . For all models, additional nonlinear gain functions $q_{\pm}$ transformed synaptic eligibility signals $\tilde{L}_{j,i,\pm}(t)$ and contributed to the rate of

1 change $\frac{dW_{j,i}(t)}{dt}$. These scaled and rectified sigmoidal functions $q(\tilde{L})$ were parameterized
2 and expressed as follows:

$$3 \quad (4) \quad gain_{slope} = \frac{2}{gain_{peak} - gain_{th}}$$

$$4 \quad (5) \quad \hat{q}(\tilde{L}) = \frac{1}{1 + e^{(-gain_{slope} \cdot (\tilde{L} - gain_{th}))}}$$

$$5 \quad (6) \quad gain_{scale} = \frac{1}{\hat{q}(1) - \hat{q}(0)}$$

$$6 \quad (7) \quad q(\tilde{L}) = \begin{cases} 0, & \hat{q}(\tilde{L}) \leq \hat{q}(0) \\ gain_{scale} \cdot (\hat{q}(\tilde{L}) - \hat{q}(0)), & \hat{q}(0) < \hat{q}(\tilde{L}) < \hat{q}(1) \\ 1, & \hat{q}(\tilde{L}) \geq \hat{q}(1) \end{cases}$$

7

8 *Voltage-dependent model:*

9 The voltage-dependent model (Fig. 2A, fig. S1, and fig. S3) contained the following 11
10 free parameters:

11 1) *signal_{eligibility}* τ_{rise} , 2) *signal_{eligibility}* τ_{decay} , 3) *signal_{instructive}* τ_{rise} , 4) *signal_{instructive}* τ_{decay} ,
12 5) ΔV_{max} , 6) k_+ , 7) k_- , 8) $gain_{th,+}$, 9) $gain_{peak,+}$, 10) $gain_{th,-}$, 11) $gain_{peak,-}$.

13 In this model, distinct eligibility signals for synaptic potentiation and synaptic depression
14 were oppositely sensitive to postsynaptic voltage. The instantaneous postsynaptic
15 depolarization amplitude $\Delta V_j(t)$ during plasticity induction lap j , including the large brief
16 depolarization produced by the dendritic plateau potential, was normalized to a
17 saturating amplitude ΔV_{max} and rectified:

$$18 \quad (8) \quad \Delta \hat{V}_j(t) = \begin{cases} 0 & \Delta V_j(t) \leq 0 \\ \frac{\Delta V_j(t)}{\Delta V_{max}}, & 0 < \Delta V_j(t) < \Delta V_{max} \\ 1 & \Delta V_j(t) \geq \Delta V_{max} \end{cases}$$

19 $\Delta \hat{V}_j(t)$ influenced synaptic eligibility signals according to:

$$(9) \quad \tilde{L}_{j,i,+}(t) = (1 - \Delta\hat{V}_j(t)) \cdot R_{j,i}(t); \text{ filtered by } \text{signal}_{\text{eligibility}} \tau_{\text{rise}} \text{ and } \tau_{\text{decay}}$$

$$(10) \quad \tilde{L}_{j,i,-}(t) = \Delta\hat{V}_j(t) \cdot R_{j,i}(t); \text{ filtered by } \text{signal}_{\text{eligibility}} \tau_{\text{rise}} \text{ and } \tau_{\text{decay}}$$

The sigmoidal gain functions $q_+(\tilde{L}_{j,i,+}(t), \text{gain}_{th,+}, \text{gain}_{peak,+})$ and

$q_-(\tilde{L}_{j,i,-}(t), \text{gain}_{th,-}, \text{gain}_{peak,-})$, as well as the plateau-related instructive signal $\tilde{P}_j(t)$

contributed to the net rate of change $\frac{dW_{j,i}(t)}{dt}$ according to:

$$(11) \quad \frac{dW_{j,i}(t)}{dt} = k_+ \cdot q_+(\tilde{L}_{j,i,+}(t)) \cdot \tilde{P}_j(t) - k_- \cdot q_-(\tilde{L}_{j,i,-}(t)) \cdot \tilde{P}_j(t)$$

where k_+ and k_- are scalar learning rate constants. Synaptic weights were bounded such that $W_{j,i} \geq 0$.

Weight-dependent model:

The weight-dependent model (Fig. 2C, fig. S2, and fig. S4) contained the following 11 free parameters:

1) $\text{signal}_{\text{eligibility}} \tau_{\text{rise}}$, 2) $\text{signal}_{\text{eligibility}} \tau_{\text{decay}}$, 3) $\text{signal}_{\text{instructive}} \tau_{\text{rise}}$, 4) $\text{signal}_{\text{instructive}} \tau_{\text{decay}}$,
5) ΔW_{max} , 6) k_+ , 7) k_- , 8) $\text{gain}_{th,+}$, 9) $\text{gain}_{peak,+}$, 10) $\text{gain}_{th,-}$, 11) $\text{gain}_{peak,-}$.

This model was formulated with the aim of obtaining a first order dependence of changes in synaptic weights on the current value of synaptic weight just before each plasticity-inducing plateau potential. We chose a two-state non-stationary kinetic model of the form shown in fig. S2A as a concrete example of a model that satisfies this dependency. Independent and finite synaptic resources at each synapse could occupy either an inactive or an active state, and the synaptic weight of each input $W_{j,i}$ was defined as proportional to the occupancy of the active state. Since the occupancy of each state in a kinetic model constrains the flow of finite resources between states, the

net change in synaptic weight at each input $\frac{dW_{j,i}(t)}{dt}$ naturally depended on the current value of synaptic weight $W_{j,i}$. This occupancy could be equivalently interpreted as a proportion of synapses that have been potentiated among a subpopulation of inputs with shared place field locations and binary weights. Synaptic weights were constrained such that

$$(12) \quad W_{j,i} = w_{j,i} \cdot W_{max}$$

where $0 \leq w_{j,i} \leq 1$, and W_{max} depends on a parameter ΔW_{max} that specifies a maximum change in weight above a baseline of 1:

$$(13) \quad W_{max} = 1 + \Delta W_{max}$$

In the absence of any voltage sensitivity of synaptic eligibility signals, both eligibility for potentiation and eligibility for depression depended only on a presynaptic firing rate and were therefore equivalent in the formulation of this model variant (fig. S3, C and D):

$$(14) \quad \tilde{L}_{j,i,+}(t) = \tilde{L}_{j,i,-}(t) = \tilde{L}_{j,i}(t) = R_{j,i}(t); \text{ filtered by } signal_{eligibility} \tau_{rise} \text{ and } \tau_{decay}$$

Sigmoidal gain functions $q_+(\tilde{L}_{j,i}(t), gain_{th,+}, gain_{peak,+})$ and $q_-(\tilde{L}_{j,i}(t), gain_{th,-}, gain_{peak,-})$, as well as the plateau-related instructive signal $\tilde{P}_j(t)$ and the current normalized synaptic weight $w_{j,i}$ contributed to the net rate of change $\frac{dW_{j,i}(t)}{dt}$ according to:

$$(15) \quad \frac{dW_{j,i}(t)}{dt} = \frac{W_{max} \cdot (1 - w_{j,i}) \cdot k_+ \cdot q_+(\tilde{L}_{j,i}(t)) \cdot \tilde{P}_j(t) - W_{max} \cdot w_{j,i} \cdot k_- \cdot q_-(\tilde{L}_{j,i}(t)) \cdot \tilde{P}_j(t)}{1}$$

where k_+ and k_- are scalar learning rate constants.

For both of the above model variants, the values of the bounded free parameters were automatically explored using a population-based version of the simulated annealing algorithm (59) to minimize an objective error function based on the difference between target and predicted ramp waveforms for each cell in the experimental dataset.

Additional model variants with either fewer (7) or more (13) free parameters were also tested (fig. S5). In the simpler model variants, the nonlinear gain functions $q_{\pm}$ were not applied to the eligibility signals and the change in synaptic weights instead depended linearly on the amplitude of the synaptic eligibility signals $\tilde{L}_{j,i,\pm}(t)$. These models were unable to account for the depression component of bidirectional BTSP (fig. S5). In the more complex model variants, eligibility signals for potentiation $\tilde{L}_{j,i,+}(t)$ and depression $\tilde{L}_{j,i,-}(t)$ were filtered with distinct time constants ($signal_{eligibility,+} \tau_{rise}$, $signal_{eligibility,+} \tau_{decay}$, $signal_{eligibility,-} \tau_{rise}$, $signal_{eligibility,-} \tau_{decay}$). The additional complexity of these model variants did not result in improved predictions relative to the original set of models (fig. S5).

Goal-directed spatial learning model:

To investigate the implications of bidirectional BTSP for reward learning by a population of CA1 place cells (Fig. 4), we constructed a network model comprised of 1000 CA1 pyramidal cells each receiving input from a population of 200 CA3 place cells with place fields spaced at regular intervals spanning the 187 cm circular track. Lap running was simulated at a constant run velocity of 30 cm/s. The synaptic weights at inputs from model CA3 place cells to model CA1 cells were controlled by the weight-dependent model of bidirectional BTSP described above. For this purpose, the 11 free parameters of that model were tuned to match synthetic data under the following constraints: 1) 5 consecutive induction laps with one 300 ms duration plateau per lap at a fixed location resulted in a place field ramp depolarization that peaked 10 cm before the location of plateau onset, had an asymmetric shape (80 cm rise, 40 cm decay), and had a peak

amplitude of 8 mV; 2) 5 subsequent plasticity induction at a location 90 cm away from the initial peak for 5 consecutive laps resulted in a 5 mV decrease in ramp amplitude at the initial peak location, and an 8 mV peak ramp amplitude at the new translocated peak position.

Before simulated exploration, all synaptic weights were initialized to a value of 1, which resulted in zero ramp depolarization in all model CA1 cells. Under these baseline conditions, each model CA1 neuron k had a probability $p_k(t) = p_{basal} = 0.0075$ of emitting a single dendritic plateau potential in 1 second of running. During each 10 ms time step, this instantaneous probability $p_k(t)$ was used to weight biased coin flips to determine which cells would emit a plateau. If a cell emitted a plateau, it persisted for a fixed duration of 300 ms, and was followed by a 500 ms refractory period during which $p_k(t)$ was transiently set to zero.

After the first lap, CA1 neurons that had emitted at least one plateau and had induced synaptic potentiation produced nonzero ramp depolarizations (Fig. 4C). The output firing rates $R_k(t)$ of each CA1 neuron k were considered to be proportional to their ramp depolarizations $V_k(t)$. The activity $R_{inh}(t)$ of a single inhibitory feedback element was set to be a normalized sum of the activity of the entire population of CA1 pyramidal neurons:

$$(16) \quad R_{inh}(t) = \beta \cdot \sum_k R_k(t)$$

where the normalization constant β was chosen such that the activity of the inhibitory feedback neuron would be 1 if every CA1 pyramidal neuron expressed a single place field and as a population their place field peak locations uniformly tiled the track. Then,

the probability that any CA1 neuron k would emit a plateau $p_k(t)$ was negatively regulated by the total population activity via the inhibitory feedback term $R_{inh}(t)$:

$$(17) \quad p_k(t) = \begin{cases} f(R_{inh}(t), \gamma_{basal}, p_{basal}), & R_{inh}(t) < \gamma_{basal} \\ 0, & R_{inh}(t) \geq \gamma_{basal} \end{cases}$$

where γ_{basal} defined a target normalized population activity (set to 0.3) and f is a descending sigmoid (Fig. 4B).

In some laps, a specific location was assigned as the target of a goal-directed search. To mimic the activation of an instructive input from entorhinal cortex signaling the presence of the goal, for a period of 500 ms starting at the goal location, the probability that a CA1 neuron would emit a plateau potential $p_k(t)$ was transiently increased. Within the goal region, the relationship between $p_k(t)$ and $R_{inh}(t)$ was instead:

$$(18) \quad p_k(t) = \begin{cases} f(R_{inh}(t), \gamma_{goal}, p_{goal}), & R_{inh}(t) < \gamma_{goal} \\ 0, & R_{inh}(t) \geq \gamma_{goal} \end{cases}$$

where p_{goal} is an elevated peak plateau probability of 0.03 per second, γ_{goal} is an elevated target normalized population activity (set to 0.5) and f is a descending sigmoid (Fig. 4B).

Acknowledgements

We are grateful to Nicolas Brunel and Karel Svoboda for discussions and comments on the manuscript, Grace Ng for contributing to software development, Ivan Raikov for technical assistance with high-performance computing, and Kristopher Bouchard at LBNL for sharing large-scale computing resources provided by the National Energy Research Scientific Computing Center, a Department of Energy Office of Science User

- 1 Facility (DE-AC02-05CH11231). This work was also made possible by computing
- 2 allotments from NSF (XSEDE Comet and NCSA Blue Waters) and supported by HHMI
- 3 and NIH (BRAIN U19 award NS104590).

References

1. W. Gerstner, M. Lehmann, V. Liakoni, D. Corneil, J. Brea, Eligibility Traces and Plasticity on Behavioral Time Scales: Experimental Support of NeoHebbian Three-Factor Learning Rules. *Front Neural Circuits* **12**, 53 (2018).
2. T. Keck *et al.*, Integrating Hebbian and homeostatic plasticity: the current state of the field and future research directions. *Philos Trans R Soc Lond B Biol Sci* **372**, (2017).
3. H. Z. Shouval, S. S. Wang, G. M. Wittenberg, Spike timing dependent plasticity: a consequence of more fundamental learning rules. *Front Comput Neurosci* **4**, (2010).
4. S. Song, K. D. Miller, L. F. Abbott, Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* **3**, 919-926 (2000).
5. J. O'Keefe, D. H. Conway, Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp Brain Res* **31**, 573-590 (1978).
6. M. R. Mehta, C. A. Barnes, B. L. McNaughton, Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci U S A* **94**, 8918-8921 (1997).
7. C. Lever, T. Wills, F. Cacucci, N. Burgess, J. O'Keefe, Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* **416**, 90-94 (2002).
8. D. Dupret, J. O'Neill, B. Pleydell-Bouverie, J. Csicsvari, The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat Neurosci* **13**, 995-1002 (2010).
9. J. D. Zaremba *et al.*, Impaired hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion. *Nat Neurosci* **20**, 1612-1623 (2017).
10. G. F. Turi *et al.*, Vasoactive Intestinal Polypeptide-Expressing Interneurons in the Hippocampus Support Goal-Oriented Spatial Learning. *Neuron*, (2019).
11. Y. Ziv *et al.*, Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* **16**, 264-266 (2013).

12. K. C. Bittner *et al.*, Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat Neurosci* **18**, 1133-1142 (2015).
13. K. C. Bittner, A. D. Milstein, C. Grienberger, S. Romani, J. C. Magee, Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* **357**, 1033-1036 (2017).
14. M. Diamantaki *et al.*, Manipulating Hippocampal Place Cell Activity by Single-Cell Stimulation in Freely Moving Mice. *Cell Rep* **23**, 32-38 (2018).
15. C. Clopath, L. Busing, E. Vasilaki, W. Gerstner, Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci* **13**, 344-352 (2010).
16. C. Clopath, W. Gerstner, Voltage and Spike Timing Interact in STDP - A Unified Model. *Front Synaptic Neurosci* **2**, 25 (2010).
17. P. Jedlicka, L. Benuskova, W. C. Abraham, A Voltage-Based STDP Rule Combined with Fast BCM-Like Metaplasticity Accounts for LTP and Concurrent "Heterosynaptic" LTD in the Dentate Gyrus In Vivo. *PLoS Comput Biol* **11**, e1004588 (2015).
18. K. He *et al.*, Distinct Eligibility Traces for LTP and LTD in Cortical Synapses. *Neuron* **88**, 528-538 (2015).
19. M. Graupner, N. Brunel, STDP in a bistable synapse model based on CaMKII and associated signaling pathways. *PLoS Comput Biol* **3**, e221 (2007).
20. B. E. Herring, R. A. Nicoll, Long-Term Potentiation: From CaMKII to AMPA Receptor Trafficking. *Annu Rev Physiol* **78**, 351-365 (2016).
21. T. Rossetti *et al.*, Memory Erasure Experiments Indicate a Critical Role of CaMKII in Memory Storage. *Neuron* **96**, 207-216 e202 (2017).
22. S. Incontro *et al.*, The CaMKII/NMDA receptor complex controls hippocampal synaptic transmission by kinase-dependent and independent mechanisms. *Nat Commun* **9**, 2069 (2018).
23. I. M. Mansuy, Calcineurin in memory and bidirectional plasticity. *Biochem Biophys Res Commun* **311**, 1195-1208 (2003).
24. C. Bellone, R. A. Nicoll, Rapid bidirectional switching of synaptic NMDA receptors. *Neuron* **55**, 779-785 (2007).
25. M. C. Lee, R. Yasuda, M. D. Ehlers, Metaplasticity at single glutamatergic synapses. *Neuron* **66**, 859-870 (2010).

26. N. Rebola, M. Carta, F. Lanore, C. Blanchet, C. Mulle, NMDA receptor-dependent metaplasticity at hippocampal mossy fiber synapses. *Nat Neurosci* **14**, 691-693 (2011).
27. W. C. Abraham, Metaplasticity: tuning synapses and networks for plasticity. *Nat Rev Neurosci* **9**, 387 (2008).
28. J. Cichon, W. B. Gan, Branch-specific dendritic Ca(2+) spikes cause persistent synaptic plasticity. *Nature* **520**, 180-185 (2015).
29. C. O'Donnell, M. F. Nolan, M. C. van Rossum, Dendritic spine dynamics regulate the long-term stability of synaptic plasticity. *J Neurosci* **31**, 16142-16156 (2011).
30. Z. Brzosko, W. Schultz, O. Paulsen, Retroactive modulation of spike timing-dependent plasticity by dopamine. *Elife* **4**, (2015).
31. Z. Brzosko, S. Zannone, W. Schultz, C. Clopath, O. Paulsen, Sequential neuromodulation of Hebbian plasticity offers mechanism for effective reward-based navigation. *Elife* **6**, (2017).
32. B. A. Richards, T. P. Lillicrap, Dendritic solutions to the credit assignment problem. *Curr Opin Neurobiol* **54**, 28-36 (2019).
33. S. A. Hollup, S. Molden, J. G. Donnett, M. B. Moser, E. I. Moser, Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J Neurosci* **21**, 1635-1644 (2001).
34. H. Takahashi, J. C. Magee, Pathway interactions and synaptic plasticity in the dendritic tuft regions of CA1 pyramidal neurons. *Neuron* **62**, 102-111 (2009).
35. A. D. Milstein *et al.*, Inhibitory Gating of Input Comparison in the CA1 Microcircuit. *Neuron* **87**, 1274-1289 (2015).
36. L. Palmer, M. Murayama, M. Larkum, Inhibitory Regulation of Dendritic Activity in vivo. *Front Neural Circuits* **6**, 26 (2012).
37. S. Royer *et al.*, Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat Neurosci* **15**, 769-775 (2012).
38. M. Lovett-Barron *et al.*, Regulation of neuronal input transformations by tunable dendritic inhibition. *Nat Neurosci* **15**, 423-430, S421-423 (2012).
39. C. Grienberger, A. D. Milstein, K. C. Bittner, S. Romani, J. C. Magee, Inhibitory suppression of heterogeneously tuned excitation enhances spatial coding in CA1 place cells. *Nat Neurosci* **20**, 417-426 (2017).

40. T. Stefanelli, C. Bertollini, C. Luscher, D. Muller, P. Mendez, Hippocampal Somatostatin Interneurons Control the Size of Neuronal Memory Ensembles. *Neuron* **89**, 1074-1085 (2016).
41. P. J. Sjostrom, E. A. Rancz, A. Roth, M. Hausser, Dendritic excitability and synaptic plasticity. *Physiol Rev* **88**, 769-840 (2008).
42. H. J. Pi *et al.*, Cortical interneurons that specialize in disinhibitory control. *Nature* **503**, 521-524 (2013).
43. L. Tyan *et al.*, Dendritic inhibition provided by interneuron-specific cells controls the firing rate and timing of the hippocampal feedback inhibitory circuitry. *J Neurosci* **34**, 4534-4547 (2014).
44. J. Guerguiev, T. P. Lillicrap, B. A. Richards, Towards deep learning with segregated dendrites. *Elife* **6**, (2017).
45. K. L. Stachenfeld, M. M. Botvinick, S. J. Gershman, The hippocampus as a predictive map. *Nat Neurosci* **20**, 1643-1653 (2017).
46. E. L. Bienenstock, L. N. Cooper, P. W. Munro, Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci* **2**, 32-48 (1982).
47. L. N. Cooper, M. F. Bear, The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nat Rev Neurosci* **13**, 798-810 (2012).
48. B. A. Richards *et al.*, A deep learning framework for neuroscience. *Nat Neurosci* **22**, 1761-1770 (2019).
49. J. Sacramento, R. P. Costa, Y. Bengio, W. Senn, Dendritic error backpropagation in deep cortical microcircuits. *arXiv preprint arXiv:1801.00062*, (2017).
50. J. Bono, C. Clopath, Modeling somatic and dendritic spike mediated plasticity at the single neuron and network level. *Nat Commun* **8**, 706 (2017).
51. N. L. Golding, T. J. Mickus, Y. Katz, W. L. Kath, N. Spruston, Factors mediating powerful voltage attenuation along CA1 pyramidal neuron dendrites. *J Physiol* **568**, 69-82 (2005).
52. T. Jarsky, A. Roxin, W. L. Kath, N. Spruston, Conditional dendritic spike propagation following distal synaptic activation of hippocampal CA1 pyramidal neurons. *Nat Neurosci* **8**, 1667-1676 (2005).
53. G. J. Stuart, M. Hausser, Dendritic coincidence detection of EPSPs and action potentials. *Nat Neurosci* **4**, 63-71 (2001).

54. M. Migliore, D. A. Hoffman, J. C. Magee, D. Johnston, Role of an A-type K⁺ conductance in the back-propagation of action potentials in the dendrites of hippocampal pyramidal neurons. *J Comput Neurosci* **7**, 5-15 (1999).
55. J. Schiller, Y. Schiller, NMDA receptor-mediated dendritic spikes and coincident signal amplification. *Curr Opin Neurobiol* **11**, 343-348 (2001).
56. A. Abraham *et al.*, Machine learning for neuroimaging with scikit-learn. *Front Neuroinform* **8**, 14 (2014).
57. C. E. Rasmussen, C. K. I. Williams, *Gaussian processes for machine learning*. Adaptive computation and machine learning (MIT Press, Cambridge, Mass., 2006), pp. xviii, 248 p.
58. A. D. Milstein, Code repository for computational model of bidirectional behavioral timescale plasticity in hippocampal CA1 place cells. <http://github.com/neurosutras/BTSP> (2020).
59. A. D. Milstein, Code repository for nested: parallel multi-objective optimization software. <https://github.com/neurosutras/nested> (2020).
60. K. Mizuseki, S. Royer, K. Diba, G. Buzsaki, Activity dynamics and behavioral correlates of CA3 and CA1 hippocampal pyramidal neurons. *Hippocampus* **22**, 1659-1680 (2012).

Fig S1.

Fig. S1. Voltage-dependent model: example simulation results. (A) Voltage-sensitivity of synaptic eligibility signals for potentiation and depression. **(B)** Model presynaptic firing rates and postsynaptic V_m during simulation of one lap of treadmill

running during plasticity induction. The timing of a dendritic plateau potential is indicated with a black dash. Example traces depict two presynaptic inputs, one that potentiates (left, blue) and one that depresses (right, red). **(C)** Temporal filters with exponential rise and decay (synaptic eligibility signal filter, grey; dendritic instructive signal filter, black). **(D)** Presynaptic input firing rates (B) were convolved with the eligibility signal filter (C) and multiplied by a voltage modulation factor (A) to generate long duration synaptic eligibility signal traces for potentiation (solid blue and red lines) and depression (dashed blue and red lines) (potentiating input example from (B), left; depressing input example from (B), right). The dendritic plateau potential was convolved with the instructive signal filter (C) to generate a long duration dendritic instructive signal (black). Periods of temporal overlap of eligibility and instructive signals that drive plasticity are shaded. **(E)** Sigmoidal gain functions sensitive to the amplitude of synaptic eligibility signals regulate the rates of synaptic potentiation (blue) and depression (red). **(F)** Net rates of change in synaptic weight (potentiating input example from (B) and (D), blue, left; depressing input example from (B) and (D), red, right). **(G)** Spatially-binned V_m ramp depolarization predicted by the model (before Induction 2, grey; after Induction 2; black). **(H)** Weighted contributions of individual presynaptic inputs to the postsynaptic ramp (before Induction 2, left; after Induction 2, right). The potentiating (blue) and depressing (red) input examples shown in (B), (D) and (F) are highlighted in color.

Fig. S2.

Fig. S2. Weight-dependent model: example simulation results. (A) State diagram for two-state kinetic model describing the flow of finite synaptic resources. **(B – H)** Same as fig. S1, B to H for simulation of the weight-dependent model of bidirectional BTSP. **(D)** Same as fig. S1D, except only a single synaptic eligibility signal is shown (red and

blue lines). In the weight-dependent model, postsynaptic voltage does not modulate the amplitude of synaptic eligibility traces, so eligibility for potentiation and depression are both marked by a single trace generated by convolving presynaptic firing rates (B) with the synaptic eligibility signal filter (C).

Fig S3.

Fig. S3. Voltage-dependent model: model parameter distributions and comparison

of model predictions to experimental data. (A – K) The free parameters of the voltage-dependent model of bidirectional BTSP were optimized for each cell in the experimental dataset. Shown here are the distributions of parameter values across cells. **(L – O)** Model predictions and experimental data are compared for features measured from spatially binned V_m ramp depolarizations. In 7/13 neurons with a second place field induced (Induction 2, red), the first place field was also experimentally-

induced (Induction 1, blue), so the model was fit to predict both place field inductions with the same set of parameters. Explained variance (R^2) and statistical significance ($p < 0.05$) reflect Pearson's correlation and 2-tailed null hypothesis tests. (**L**) Peak ramp amplitude. (**M**) Ramp width. (**N**) Shift of ramp peak location relative to mean location of plateau onset. (**O**) Minimum ramp amplitude across spatial bins.

Fig. S4.

Fig. S4. Weight-dependent model: model parameter distributions and comparison of model predictions to experimental data. (A – O) Same as fig. S3, A to O for parameters and predictions of the weight-dependent model of bidirectional BTSP.

Fig. S5.

Fig. S5. Evaluation of additional model variants. (A – G) Additional model variants of varying complexity were also evaluated for their capability to predict experimentally measured bidirectional changes in V_m ramp amplitude by BTSP. The plasticity rule was estimated by two-dimensional interpolation from model data (see Materials and Methods). (A – C) Model predictions from variants of the voltage-dependent model with 7 (A), 11 (B), or 13 (C) free parameters (see Materials and Methods). (D) Plasticity rule estimated from the experimental data. (E – G) Model predictions from variants of the weight-dependent model with 7 (E), 11 (F), or 13 (G) free parameters (see Materials and Methods). (H) Residual error of ramp depolarizations predicted by each model is averaged across spatial bins and cells.

Fig. S6.

Fig. S6. Postsynaptic voltage perturbation experiment: simulation of dendritic depolarization and quantification of experimental data. (A – C) Simulation of a

biophysically-detailed CA1 pyramidal cell model with realistic morphology and dendritic ion channel distributions (39) to estimate the effect of steady-state somatic depolarization on distal dendritic V_m . Three conditions are compared: a silent cell with uniform input weights (black), a silent cell with ~10 mV of steady-state depolarization induced by somatic current injection (purple), and a place cell receiving potentiated inputs at the peak of its place field (blue). Under conditions of somatic current injection, a combination of attenuated propagating depolarization and back-propagating action potentials amplifies local synaptic input by activating dendritic voltage-gated ion channels, resulting in a level of dendritic depolarization comparable to the place field condition. **(A)** Simulated V_m traces recorded from soma (left), distal apical oblique dendrite (center), and a distal apical dendritic spine (right). **(B)** Mean low-pass filtered V_m at simulated dendritic recording sites at varying distances from the soma. **(C)** Same as (B) for simulated recordings from dendritic spines. **(D – F)** Quantification of experimental ramp depolarizations induced by BTSP under conditions of somatic depolarization (Fig. 3). **(D)** Peak ramp amplitude ($p < 0.015$). **(E)** Ramp width ($p > 0.143$). **(F)** Shift of ramp peak location relative to mean location of plateau onset ($p > 0.255$). p-values reflect two-sided Mann-Whitney U tests.

Fig. S7.

Fig. S7. Evidence for saturability of synaptic potentiation by BTSP. Peak V_m ramp amplitudes are not correlated with the total accumulated duration of plateau potentials across laps during plasticity induction, indicating a saturating nonlinearity. Shown are data from both silent cells in which *de novo* place fields were induced (light grey, $n=25$), and place cells in which a second plasticity induction translocated a pre-existing place field (dark grey, $n=13$). Explained variance (R^2) and statistical significance ($p < 0.05$) reflect Pearson's correlation and a 2-tailed null hypothesis test.

Fig. S8.

Fig. S8. Evaluation of bidirectional BTSP models with alternative CA3 input place

field widths. (A – D) Additional model variants with alternative values of place field width for presynaptic CA3 place cell inputs were also evaluated for their capability to predict experimentally measured bidirectional changes in V_m ramp amplitude by BTSP. Plasticity rules were estimated by two-dimensional interpolation from model data (see Materials and Methods). **(A – C)** Model predictions from variants of the weight-dependent model with 75 cm (A), 90 cm (B), or 105 cm (C) CA3 input place field widths. **(D)** Residual error of ramp depolarizations predicted by each model is averaged across spatial bins and cells.