

Signatures of rapid plasticity in hippocampal CA1 representations during novel experiences

Highlights

- Population recordings in mouse hippocampus during familiar and novel exploration
- CA1 place field dynamics predicted by behavioral timescale synaptic plasticity (BTSP)
- BTSP-like features were highly experience dependent and enriched during novelty

Authors

James B. Priestley, John C. Bowler,
Sebi V. Rolotti, Stefano Fusi,
Attila Losonczy

Correspondence

jbp2150@columbia.edu (J.B.P.),
al2856@columbia.edu (A.L.)

In brief

By combining virtual reality and two-photon calcium imaging, Priestley et. al. find pervasive signatures of rapid neural plasticity mechanisms in the mouse hippocampus during new experiences. Their study highlights a strong regulation of neural feature tuning in CA1 by novelty, which could contribute to quickly encoding new memories.



Article

Signatures of rapid plasticity in hippocampal CA1 representations during novel experiences

James B. Priestley,^{1,2,3,*} John C. Bowler,^{1,2} Sebi V. Rolotti,^{1,2} Stefano Fusi,^{1,3,4} and Attila Losonczy^{1,4,5,*}

¹Department of Neuroscience, Columbia University, New York, NY 10027, USA

²Doctoral Program in Neurobiology and Behavior, Columbia University, New York, NY 10027, USA

³Center for Theoretical Neuroscience, Columbia University, New York, NY 10027, USA

⁴Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10027, USA

⁵Lead contact

*Correspondence: jbp2150@columbia.edu (J.B.P.), al2856@columbia.edu (A.L.)

<https://doi.org/10.1016/j.neuron.2022.03.026>

SUMMARY

Neurons in the hippocampus exhibit a striking selectivity for specific combinations of sensory features, forming representations that are thought to subserve episodic memory. Even during completely novel experiences, hippocampal “place cells” are rapidly configured such that the population sparsely encodes visited locations, stabilizing within minutes of the first exposure to a new environment. What mechanisms enable this fast encoding of experience? Using virtual reality and neural population recordings in mice, we dissected the effects of novelty and experience on the dynamics of place field formation. During place field formation, many CA1 neurons immediately modulated the amplitude of their activity and shifted the location of their field, rapid changes in tuning predicted by behavioral timescale synaptic plasticity (BTSP). Signatures of BTSP were particularly enriched during the exploration of a novel context and decayed with experience. Our data suggest that novelty modulates the effective learning rate in CA1, favoring rapid mechanisms of field formation to encode a new experience.

INTRODUCTION

Learning in neuronal systems is complicated by a fundamental tension between stability and plasticity (Carpenter and Grossberg, 1991). Networks with fast learning rates encode new information with high fidelity at the expense of overwriting older patterns, whereas a slow learning rate can preserve existing structure yet stymie the encoding of novel information. Theoretical studies of neuronal memory capacity suggest that optimal solutions involve concerted processes operating on a spectrum of timescales (Roxin and Fusi, 2013; Benna and Fusi, 2016). These ideas harmonize with models of multistage memory systems in the brain (McClelland et al., 1995): fast-learning circuits can quickly capture detailed memories of new episodes, which are progressively transferred and integrated into slower systems downstream.

The mammalian hippocampus is intimately involved in the formation of episodic memories and likely mediates an intermediate stage of processing and storage of experiential information prior to long-term storage in the cortex (McClelland et al., 1995). Although generally viewed as a short-term memory system, hippocampal dynamics exhibit a diversity of time constants, both at the level of its subnetworks (Mankin et al., 2015; Ziv et al., 2013) and cellular plasticity mechanisms (Bittner et al., 2015; Mehta

et al., 1997; Magee and Grienberger, 2020). These results are most often derived from the study of “place cells,” excitatory neurons in the hippocampus that are active in specific locations in an environment during exploration (O’Keefe and Dostrovsky, 1971; Moser et al., 2008). Spatial behaviors provide a convenient model for studying memory, as the various sensory settings that animals encounter in the environment are organized into relational neural representations in the hippocampus (Eichenbaum, 2017). Novel population codes develop with remarkable speed, requiring only a few exposures to an environment before a new set of place fields is learned that spans the available space (Wilson and McNaughton, 1993; Frank et al., 2004).

A unique synaptic learning rule was recently discovered in the CA1 subregion of the hippocampus, by which pyramidal neurons formed stable place fields within just a few trials after burst firing was recorded from the neuron at a particular location in the environment (Bittner et al., 2015, 2017). The sudden emergence of tuning in previously silent neurons marks a rapid reconfiguration of the weights of synapses that were active around the time of the burst event (Bittner et al., 2017; Milstein et al., 2021), with an asymmetric envelope that extends to inputs active several seconds before the event occurred. This behavioral timescale synaptic plasticity (BTSP) is a notable departure from conventional plasticity schemes. Its expression depends on the



presence of somatic burst firing driven by plateau potentials, which reflect nonlinear input integration in the dendritic arbor of pyramidal cells (Epsztein et al., 2011). These events could be gated by the presence of other factors such as inhibition, neuromodulation, or “instructive” inputs signaling reinforcement or novelty (Gerstner et al., 2018; Milstein et al., 2021; Grienberger and Magee, 2021; Rolotti et al., 2022), which could enable the circuit to rapidly construct new representations during salient experiences. It remains unknown what connection this plasticity mechanism has to the reorganization of hippocampal responses that occurs when animals are exposed to different environments (“global remapping,” Muller and Kubie, 1987) or in response to salient cues or reinforcement (Hollup et al., 2001; Zaremba et al., 2017; Dupret et al., 2010). The presence of plateau potentials or associated somatic burst spiking reliably leads to the formation of place fields *de novo* (Bittner et al., 2015; Diamantaki et al., 2018), but other experiments have reported that new place fields can appear in the absence of these signatures in both familiar and novel environments (Cohen et al., 2017).

Clearly, the network’s synaptic matrix does not start from a blank slate, when learning from each new episode; prior experience in other environments may already provide a weight distribution that produces location-specific, suprathreshold spiking for some neurons or subthreshold tuning that could be amplified to unmask new receptive fields through other plasticity mechanisms (Lee et al., 2012; McKenzie et al., 2021). There is also evidence that CA1 ensembles may arise to some degree from pre-configured network structures (Druckmann et al., 2014; Groszmark and Buzsáki, 2016; Bocchio et al., 2020; Geiller et al., 2022), compatible with observations of hippocampal “pre-play” of future experiences (Farooq and Dragoi, 2019), although these findings remain controversial (Silva et al., 2015; Muessig et al., 2019). Other work has implicated local dendritic spikes in plasticity-driven place field formation (Sheffield et al., 2017), which could involve anatomical clustering of similarly tuned input (Sheffield and Dombeck, 2019; Adoff et al., 2021). However, more global mechanisms such as BTSP could be valuable for rapidly modifying the synaptic landscape to accommodate novel learning (Milstein et al., 2021), and its skewed, seconds-long timescale can endow place fields with predictive information (Zhao et al., 2022). We lack clarity on the extent to which BTSP contributes to learning hippocampal representations in these different scenarios, which could provide insight on how the network learning rate may change in response to factors such as novelty or salience.

In this work, we conducted a longitudinal analysis of place field formation during familiar and novel experiences in order to search for correlates of BTSP and how they may change as a function of experience. Using 2-photon functional calcium imaging, we surveyed thousands of place fields and identified an enrichment of BTSP-like dynamics during the initial exposures to a new environment, which then decayed over the course of several days. Our findings are compatible with widespread BTSP in CA1 and illustrate an experience-dependent regulation of plasticity that could be controlled by internal or external factors to dynamically tune the learning rate of hippocampal representations.

RESULTS

Novel CA1 representations develop rapidly in virtual reality

We constructed a virtual reality (VR) system for head-restrained mice, comprising 5 liquid crystal displays (LCD) surrounding a running wheel; movement through the virtual environments was yoked to a rotary encoder on the wheel axle. We combined this apparatus with 2-photon functional calcium imaging in order to record CA1 neural populations as mice explored the virtual contexts (Figures 1A and S1; see STAR Methods for details). Mice were stereotactically injected with an rAAV vector encoding the calcium sensor GCaMP6f under the control of the synapsin promoter, targeted to the CA1 pyramidal layer. We then implanted a chronic window above the hippocampus to provide optical access for imaging experiments (Lovett-Barron et al., 2014). All imaging data were postprocessed in Suite2p (Pachitariu et al., 2017) for motion correction, cell detection, and extraction of raw fluorescence traces (Figure 1B). Signals were neuropil corrected and detrended for baseline drift and were deconvolved (Friedrich et al., 2017) to reduce the impact of calcium autocorrelation on our analysis. As described previously (Ahmed et al., 2020), we referred to the resulting signals as events, given the limited ability to resolve precise spike times from calcium dynamics. Imaging data were collected from approximately the same area of CA1 for each 3-day experiment sequence, although we did not attempt to track the identity of individual neurons across sessions. The resulting dataset consisted of an average of 391 ± 39 CA1 neurons per session (35 sessions from 12 experiments, comprising two 3-day experiment sequences each from 6 mice).

We first trained mice to run for sucrose rewards in a 3 m virtual environment. All virtual environments began and ended in a darkened tunnel; entering the exit tunnel triggered a 2 s intertrial interval (ITI) during which the screens would remain dark, and after which, the animal was teleported back to the start of the track. After 1–2 weeks of training, most mice reliably ran over 100 trials in under an hour. To study the dynamics of place field formation during novel experience, we used the VR system to rapidly alternate between familiar and novel environments multiple times during a single recording (Figure 1C). In each recording session, mice ran trials through alternating blocks in either a familiar (the training environment) or novel context. Context switches were uncued; the mouse was simply teleported to the other environment at the end of the last ITI of a given block. We repeated the context alternation procedure across 3 consecutive days using the same familiar/novel contexts to examine the effects of increasing familiarity on CA1 coding. The complete 3-day experiment was repeated twice per mouse using two different novel contexts (the familiar context remained the same).

Pyramidal cells’ activity tiled the virtual track (Sheffield et al., 2017; Zhao et al., 2020), forming a reliable sequence of spatial responses on each trial (Figure 1D). Switching the virtual context during the session recruited an approximately orthogonal ensemble of place cells, indicating our apparatus could trigger robust global remapping across VR scenes (Figures S1G and S1H). Correspondingly, we could reliably decode the position

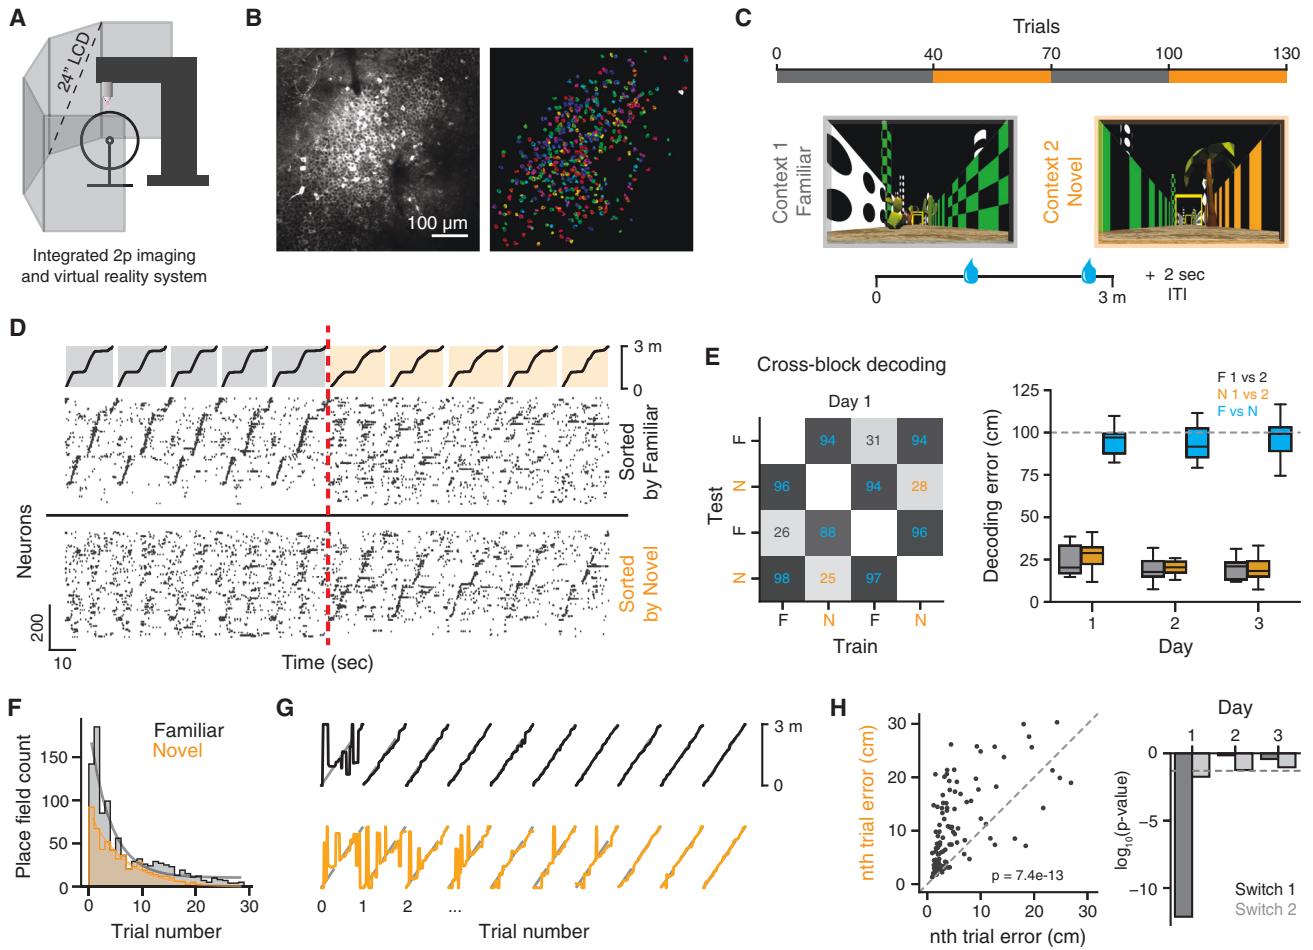


Figure 1. Formation of CA1 representations in novel virtual environments

(A) Schematic of integrated 2-photon microscope and virtual reality system.

(B) Left: example 2-photon field of view (FOV) of CA1 pyramidal neurons expressing GCaMP6f. Right: spatial masks of ROIs from the same FOV.

(C) Each session consisted of 130 trials, alternating in blocks between a familiar (gray) or novel (orange) context (40 trials for the first block, then 30 trials per block). Both environments were 3 m in length. Sucrose rewards were delivered at ~1.2 and 2.7 m on each trial; the distance to rewards was the same in both contexts. Each experiment comprised 3 consecutive days of recordings during the context switch paradigm, using the same familiar/novel contexts.

(D) Position of the animal in the environment (top) and ensemble neural activity from an example day 1 session (bottom), shown for ± 5 trials around the very first exposure to the novel context (transition marked in red). Neurons are sorted by their peak firing location on the track in either familiar (top) or novel (bottom) trials. A new sequence of place fields rapidly organizes within the first few trials in the novel context.

(E) On each day, we trained a decoder on the neural data from each context trial block separately and used it to predict position in all other trial blocks. Left: average absolute position error for each pair of training/testing blocks on day 1. Right: summary of decoding performance for different comparison types ($n = 12$ experiments).

(F) Place field accumulation in the first trial blocks of the familiar and novel contexts on day 1, shown as a histogram of the formation trials (the first trial that a cell fired within its place field).

(G) Performance of spatial decoding tested on the initial 10 trials in a block, after training on the remaining trials from the same block, shown for the first presentation of each context for an example session on day 1.

(H) Left: trial number-matched comparison of initial decoding performance between the two contexts for the first context switch on day 1. Data pooled from all experiments. Initial errors are consistently greater in the novel context (Wilcoxon signed-rank test). Right: p values for this comparison computed for each context switch for all days. The difference in decoding error is eliminated with experience.

of the animal with a classifier trained and tested on trial blocks sharing the same VR context but achieved only chance-level performance when tested on the opposite context (Figure 1E). Only a very small fraction of neurons exhibited similar place tuning in at least one trial block of both environments (<4% of all place cells). Similar to prior work, we measured an increased

density of place fields near the reward locations (Hollup et al., 2001; Dupret et al., 2010; Zaremba et al., 2017) and also the end of the track (Figures S1D–S1F). Some of these cells showed evidence of context-invariant coding of the reward zones (Gauthier and Tank, 2018) and the VR track end tunnel (which is visually similar between contexts), but they represented a

minute fraction of the total place cells across the dataset (Figure S1H, ~1%).

Examining the very first exposure to the novel context on day 1, a sequence of place fields was already visible in the first few trials following the context switch; new place fields in the novel environment continued to appear over time, with field accumulation decaying roughly exponentially over trials (Figure 1F) similar to prior reports (Sheffield et al., 2017; Grienberger and Magee, 2021). Gross place field accumulation in the familiar context was 2-fold greater in the initial trials and decayed faster compared with the novel context ($\tau = 3.9$ in familiar and $\tau = 6.4$ in novel for exponential decay fit to Figure 1E). Trial-by-trial decoding similarly showed that it took more time to reduce position errors from the start of the novel context block compared with the familiar context, but only on day 1, mainly the very first context switch in the experiment (Figures 1G and 1H). We noticed an overall smaller place coding fraction in the novel context, particularly on the first exposures to the environment (Figure S1B), which could reflect a slower accumulation of fields during novel experience (Sheffield et al., 2017; Dong et al., 2021; Grienberger and Magee, 2021). Given that place field firing rate, precision, and stability also increase with experience (Frank et al., 2004; Chen et al., 2013), it is possible that we undercount some nascent fields due to the sensitivity of calcium imaging. Overall though, these findings agree with our intuition that many place fields in the familiar environment are due to prior learning (and so they immediately and robustly appear within the first trials in the session), although the novel representation may continue to grow through ongoing plasticity over a longer time period.

CA1 population activity shifts transiently in the novel context

We analyzed the dynamics of place coding across trials in each context block in order to identify activity signatures that may be consistent with different underlying plasticity mechanisms. One hallmark of plateau-induced field formation in CA1 is the backward shift of spatial tuning, relative to the location of burst firing during the formation trial (i.e., the trial when the cell first fires near its place field). This transient shift is a consequence of the asymmetric BTSP kernel that produces a large ramping membrane depolarization leading up to the plateau location (Bittner et al., 2017). This has been shown to drive somatic activity prior to the plateau location (Zhao et al., 2020, 2022; Rolotti et al., 2022). If new place fields in the novel context form predominantly through BTSP, then this should induce a transient backward drift that is measurable in the population tuning of space during the first few trials in the environment. Since this is the time period during which the majority of place fields form (Figure 1F), the population drift could arise from the cumulative effect of many fields shifting on the same trial, an effect that can notably be measured without first identifying specific place fields or their formation trial.

To quantify any spatial shift between population representations on different trials, we estimated the population spatial cross-correlation between pairs of trials (± 75 cm, Figure 2A). In Figure 2B, we summarized the resulting spatial shifts between all trial pairs as a matrix, shown for day 1 sessions

(when the animal is exposed to the novel context for the very first time) separately for each context trial block. Since the sign of the spatial shift between two trials is reversed when the order of comparison is reversed (i.e., trial *a* is after trial *b*, so *b* is before *a*), these matrices are antisymmetric. In general, the upper triangle of the shift matrices was slightly positive, indicating that the population tuning tended to drift backward in space relative to earlier trials in each block. However, the shift pattern during the first exposure to the novel context (switch 1) was markedly different from other trial blocks: early trials in the context showed a far more exaggerated shift forward in space relative to later trials. This pattern was mirrored by changes in the raw value of the peak cross-correlation over those trial pairs (Figure S2), suggesting overall greater coding stability between later trials. We summarized the shift trend by computing the average pairwise shift for each trial per block (Figures 2C and S2), which showed a large but transient forward shift of population spatial tuning during the early trials of the first block in the novel context. This shift decayed rapidly within the first 10 trials in the new context, notably overlapping with the period during which the majority of new place fields appeared in the novel context (Figure 1E).

The representation in the novel context consistently showed a transient, backward drift during the first exposure across mice on day 1 (Figures 2D and S2), but this effect was highly experience dependent. We repeated the context switch protocol over a 3-day period and found that the greatest drift was reliably observed during day 1, when the animals were exposed to the novel context for the very first time (Figures 2D and 2E). It is possible that this transient drift is due to the rapid acquisition of new place fields during the initial trials in the novel context. If many neurons acquire their field through BTSP, the population shift can arise as a consequence of averaging over many place fields that shifted acutely after field formation. We simulated this condition and found that it produced a qualitatively similar pattern in the population shift matrices (Figure S3). Other recent reports have suggested that many individual CA1 neurons drift continuously over trials (Dong et al., 2021), but our simulations of this alternative scenario produced population shift patterns that were incongruous with the observed data (Figure S3). These effects were also not explained by changes in animals' velocity (Figures S7A–S7C). Given these comparisons, we reasoned that the population drift was most compatible with frequent BTSP-mediated place field formation in the first exposure to the novel context.

New place fields appear with characteristics of BTSP

We sought to connect this population-level observation with the behavior of individual place fields (Figure 3A). Due to the asymmetric plasticity kernel of BTSP, individual place fields should acutely shift backward (relative to the direction of animals' motion) between the formation trial and subsequent trials (Bittner et al., 2017; Zhao et al., 2020, 2022; Rolotti et al., 2022). Simulations of BTSP-mediated place field formation confirmed that this shift in place field location could be robustly measured across a wide range of animal velocities and noise levels using calcium imaging (Figure S4). To examine this in the data, we detected place fields in each context block separately for each recording

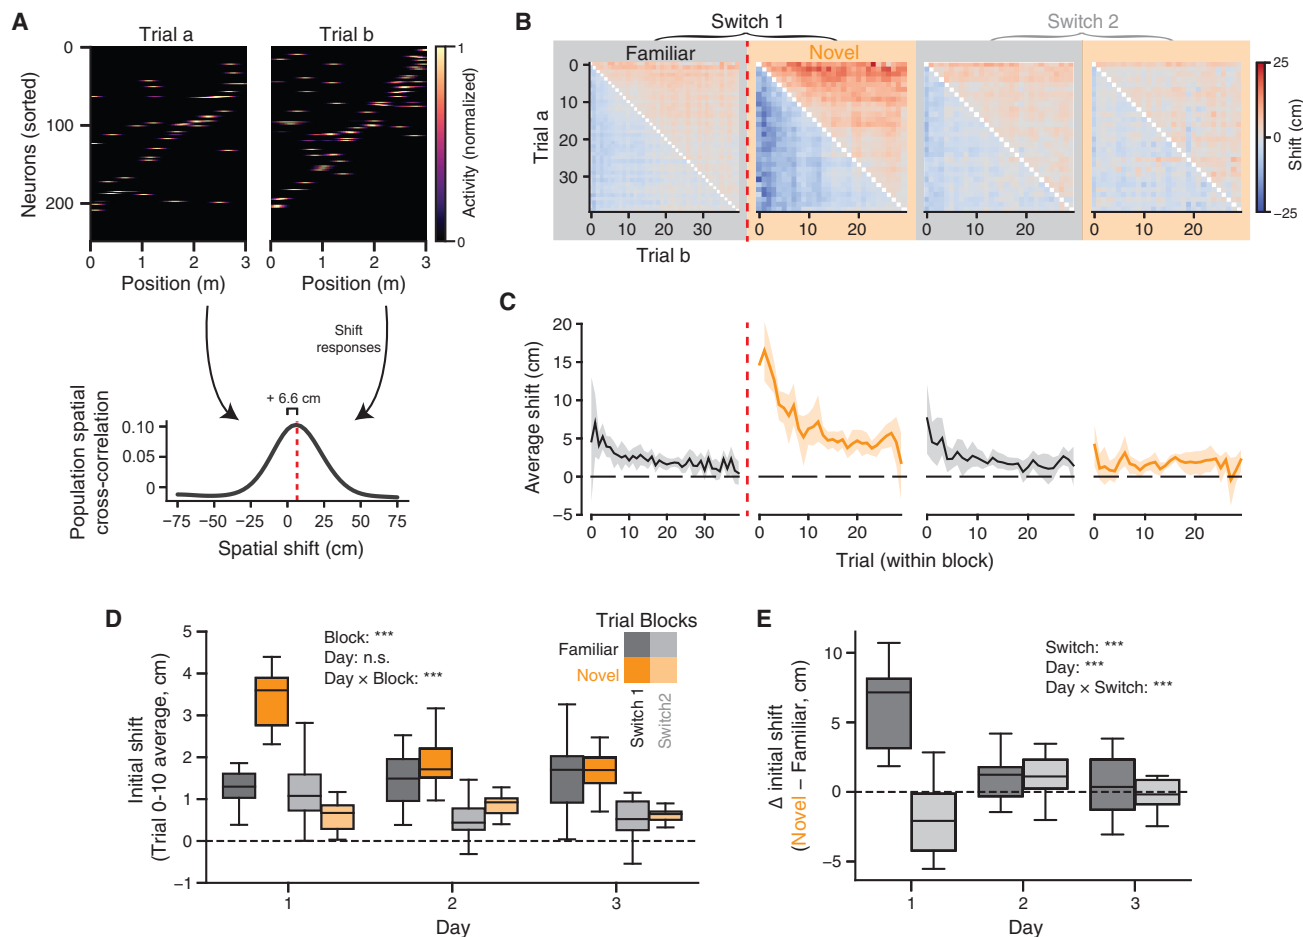


Figure 2. Transient, experience-dependent drift of CA1 ensemble coding

(A) Top: spatial tuning of the neural population during two example trials in the same context. Bottom: population spatial cross-correlation computed between the two trials. The correlation at different spatial lags is obtained by shifting each neuron's tuning in trial *b* before forming the population vector. Red-dashed line indicates the center of mass (COM) of the spatial cross-correlation curve. A shift at positive lags indicates the responses in trial *a* are at positions ahead of the responses in trial *b*.

(B) Spatial shifts (COM) between all pairs of trials during each trial block on day 1 (average of $n = 11$ recordings).

(C) Average shift for each trial in the session, mean and 95% confidence interval across day 1 recordings.

(D) Average representation shifts in the first 10 trials of each block, summarized over all 3 days of the experiment; linear mixed-effects model with main effects of trial block (categorical) and day (continuous), significance shown in inset.

(E) Difference in initial representation drift between the novel versus familiar contexts, for each context switch and day; linear mixed-effects model with main effects of switch (categorical) and day (continuous), significance inset. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

session and measured how displaced spatial activity on the formation trial was from activity on remaining trials (Figures 3B and 3C). Note that we use “formation trial” to denote the trial where we first detected stable firing within a place field in a given trial block, but this is not meant to imply that all fields form through ongoing plasticity (i.e., many likely appear simply due to prior learning). On comparing formation trial shifts across different trial blocks on day 1, we found a clear increase in the displacement of the formation trial's tuning curve specifically during the first exposure to the novel context, relative to the other trial blocks (Figure 3D). Similar to the population drift, the greatest field shifts were observed during day 1 on the first context switch (Figure 3E). Examining trial-by-trial displacements of place field activity, we also found that the shifting was most pronounced on

the formation trial and that place fields did not generally continue to drift after the first few trials following place field formation (Figures S3E–S3G). These results align with the experience-dependent drift in population tuning described in Figure 2, suggesting that the latter arises due to the cumulative effect of many individual cells undergoing acute tuning shifts as they formed their place fields.

The long timescale of BTSP also induces a correlation between the width of place fields and the speed of the animal during the plasticity event: if the animal runs faster, the potentiated inputs will span a larger region on the track (Bittner et al., 2017). For every identified place field, we measured the velocity of the animal as it traversed the place field on the trial of field formation. Examining the joint distribution of formation trial velocities and

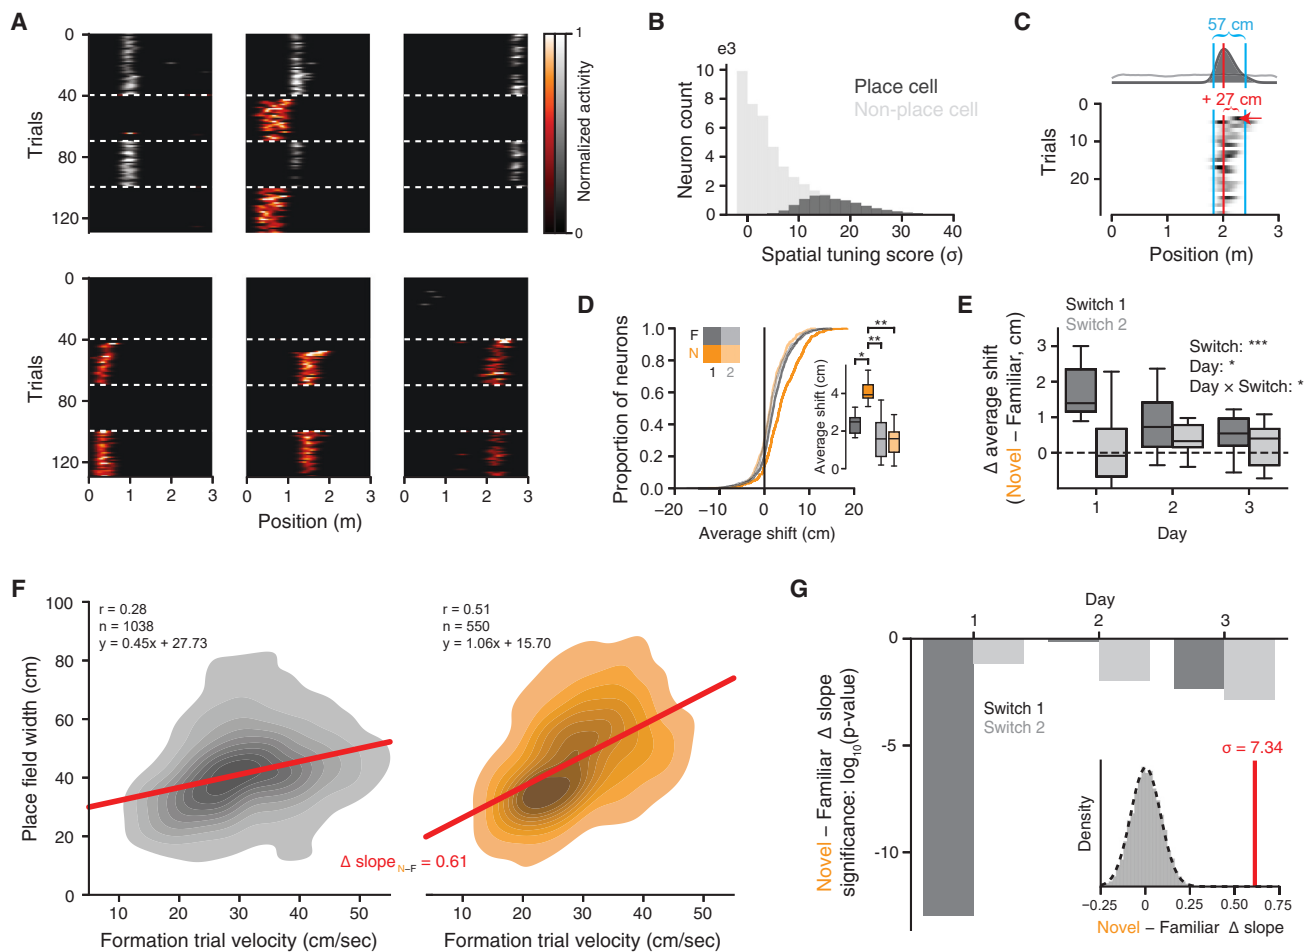


Figure 3. Place field dynamics are consistent with BTSP-mediated field formation in the novel context

(A) Example simultaneously recorded CA1 neurons with place fields during day 1 of the experiment (gray: familiar trials; orange: novel trials). Stable place fields appeared in both the familiar and novel contexts.

(B) Distribution of spatial tuning scores for place and non-place cells.

(C) Characterizations of place field formation for an example neuron: field detection, field width, and formation trial shift.

(D) Distribution of place field shifts (distance between formation trial activity and the remaining in-field activity) during day 1 of the experiment. Place fields in the first exposure to the novel context exhibit exaggerated spatial shifts on their formation trial. Inset: session-averaged place field shifts on day 1 ($n = 11$ day 1 recordings, Wilcoxon signed-rank test with Bonferroni correction).

(E) Difference in place field shifts between novel and familiar contexts, for each context switch and day. The greatest change is seen during the first context switch on day 1 ($n = 12$ experiments; linear mixed-effects model with main effects of switch [categorical] and day [continuous], significance inset).

(F) Distribution of place field widths and the velocity of the animal as it traversed the place field during the formation trial on day 1, switch 1. The linear fit is shown in red. The correlation between velocity and field width is stronger in the novel context.

(G) Significance of the difference in regression slopes between familiar and novel as shown in (F), for all days and switches. The Δ slope was compared with a null distribution created by randomly permuting the context labels of place fields for each context switch before recomputing the within-context regressions and between-context Δ slope (depicted in the inset for the day 1, switch 1 results shown in F). Plotted is the \log_{10} p value derived from a Gaussian fitted to the null distribution. The most significant difference is seen for day 1, switch 1. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

the width of the associated place fields, we found that these variables were correlated in both the familiar and novel contexts (Figure 3F, data shown for day 1, switch 1). However, this correlation was stronger in the novel context, and the linear fit produced a significantly greater slope compared with the familiar context (Figure 3G). The difference between novel and familiar slopes was also experience dependent; the first context switch on day 1 exhibited the most significant slope difference by far across the entire experiment sequence. These results are again

compatible with a greater fraction of place cells forming via BTSP during the first exposure to the novel context.

Shared variability across place field formation events

Ideally, we would like to segregate individual place fields according to their pattern of spatial drift over trials to potentially identify a subgroup of place fields exhibiting BTSP-like characteristics or other dynamics and study additional properties of these classes. Toward this aim, we first took an unsupervised

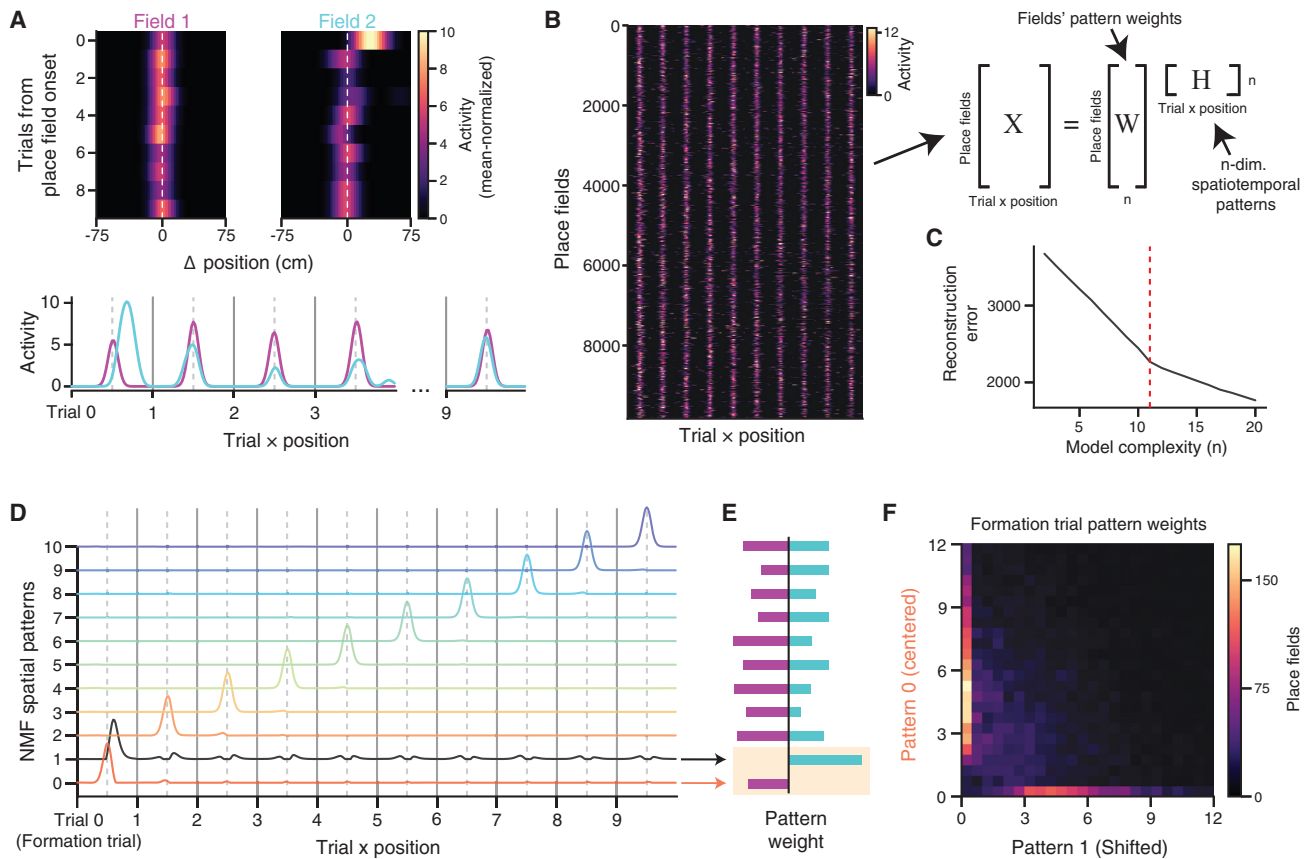


Figure 4. Factorizing place field responses across trials reveals distinct dynamics during the trial of place field appearance

(A) Top: two example place fields, with data shown in a 150 cm window aligned to the place field center and truncated to the first 10 trials from the onset of place field formation. Bottom: the same fields, “flattened” so that each trial’s spatial tuning curve is sequentially concatenated.

(B) Left: all identified place fields in the dataset with at least 10 trials of activity from onset, cropped and aligned as in (A). Right: schematic of non-negative matrix factorization (NMF). Each place field’s activity over trials is modeled as a weighted sum (**W**) of a small set of shared spatiotemporal patterns, **H**.

(C) Reconstruction error of **X** for different choices of the number of shared patterns *n* in **H**. There is a clear “elbow” at *n* = 11, where model improvement slows. In Figure S5A, we also show that this choice of *n* gives the most interpretable patterns.

(D) Spatiotemporal patterns learned by NMF in the *n* = 11 model. The model learns a separate pattern for each trial that produces the centered place field on that trial, plus an additional pattern (in black) on the formation trial that is shifted forward in space relative to the place field center.

(E) Weights of the spatiotemporal patterns in (D) for the two example place fields shown in (A) (Field 1: magenta; Field 2: cyan). Trial-to-trial fluctuations in place field amplitude are captured by modulating the weight of the corresponding spatiotemporal pattern. Note the weight differences for the two “formation trial” patterns (yellow shading), which reflect the forwarded shifted activity for the cyan field.

(F) Joint distribution of pattern weights across all place fields for the two “formation trial” components. Field weights for the two patterns are anticorrelated.

approach to identify a set of spatial activity patterns across trials that were shared between place fields (Figures 4A–4C). For every place field in the dataset, we examined the first 10 trials of spatial activity, starting from the formation trial. On each of these trials, 150 cm of the spatial tuning curve was extracted, centered around the place field (Figure 4A). We then concatenated all trials from the truncated spatial tuning matrices into a vector for every place field. Consider the example place fields in Figure 4A: since we have centered each field’s activity relative to its field center, the two “flattened” place fields appear as a relatively well-aligned series of bumps. However, there is a large deviation in the location of firing during the formation trial: Field 2 exhibits forward-shifted activity relative to its place field center, whereas Field 1 remains centered. We hypothesized

that many place fields will exhibit similar or other stereotyped patterns of spatial shifts on certain trials and used dimensionality reduction to attempt to discover these dynamics.

All place fields in the dataset were gathered into a matrix **X**, where each row is the spatiotemporal pattern of a single field during the first 10 trials from its formation trial (Figure 4B). Each row was normalized by its mean to encourage the model to focus on shared variability between place fields. We sought a matrix decomposition $\mathbf{X} = \mathbf{WH}$, where each row of **H** would describe a pattern of spatial activity over trials, and each column of **W** would describe how that pattern contributed to individual place fields. Here, we use non-negative matrix factorization (NMF), since its strict non-negativity aids in interpreting the extracted components (activity described by different

components cannot be “cancelled out” due to negative values). In this setting, it is straightforward to interpret the rows of \mathbf{H} as a pattern of spatial activity over trials, where each individual place field in \mathbf{X} is modeled as a weighted sum of those different patterns, with the weights given by the rows of \mathbf{W} . As with any dimensionality reduction, it is necessary to choose the number of components n learned by the model. We considered the quality of the reconstruction of \mathbf{X} across a range of different model complexities (Figure 4C) and identified a clear “elbow” (Milligan and Cooper, 1985) between two linear regimes of the reconstruction error at $n = 11$, where adding additional patterns to the model yielded smaller improvements in the fit quality.

Strikingly, the 11-component model results in a very clean and interpretable partitioning of variance in the place field dataset (Figure 4D). Ten of the 11 components encoded the presence of a centered place field on each of the 10 trials included in the data, whereas the 11th component encoded a forward-shifted place field on the first trial (i.e., the trial of field formation). This is a sensible way to decompose place fields: any individual field can now be reconstructed by appropriately weighting each per-trial component according to the place field’s amplitude on that trial, and the shifted formation trial component can be used to account for variability due to BTSP-like field formation (Figure 4E). Notably, it is not immediately obvious that NMF should find this per-trial representation; the model could instead identify components that are active on several or all trials, representing longer trends of place field amplitude and shifts that are shared between many neurons. Instead, our model’s representation suggests that the variability in place field activity in this dataset is best captured independently trial-by-trial for each place field. Other choices for n invariably harmed the interpretability of the model components (Figure S5A) and in light of the clear inflection point in the loss function at Figure 4C, we focused our remaining analysis on this decomposition.

The representation of formation trial activity in the model illustrates the additional variability present during the trial of place field formation across the dataset. Inspecting the joint distribution of fields’ weights on these two components, we found a high concentration of place fields that exclusively weighted one pattern or the other (Figure 4F, note the concentration about the axes). In fact, the two formation trial components had the most anticorrelated weight vectors out of any pair in the model (Figure S5B). The shifted formation trial component was also weakly anticorrelated with all remaining trial patterns, which could reflect the tendency for plateau potentials to evoke higher, burst firing rates on the formation trial compared with later trials. Given these features, we clustered place fields according to their weights on the 11 NMF components to attempt to isolate a group of fields with BTSP-like characteristics. We found that using two clusters was sufficient to reliably segregate these place fields (Figures S5C–S5F). In Figure 5A, we plotted the cluster centers for the two groups in the NMF component space, i.e., the average weight placed on each spatial pattern for the place fields within a given clusters. The two groups are mainly distinguished by their strong, opposing weights on the two formation trial components. Notably, the cluster that highly weighted the shifted “plateau”

component also showed lower weights on all subsequent trial components compared with the other cluster. We labeled the clusters as “BTSP-like” (cyan) and “Other” (magenta) based on these features.

Place fields show experience-dependent expression of BTSP characteristics

How do these two groups of place fields differ? We found that the BTSP-like fields exhibited strongly forward shifted activity on the formation trial and relatively reduced amplitude firing on later trials (Figures 5B–5D). We quantified the amplitude change by computing a formation trial gain for each place field and found that the BTSP-like fields consistently exhibited higher gain (Figure 5C). The two clusters were distinguished by several other features: BTSP-like fields generally exhibited greater place field width on later trials and reduced trial-to-trial stability compared with Other fields (Figures S5H and S5I). These post hoc characteristics (formation trial gain and shift, field width, and stability) could be used to decode the BTSP label of place fields with high accuracy using a linear classifier (85% on average, Figures S5J–S5K). One barrier to perfect classification is that BTSP characteristics like first trial shift were present to some extent on a continuum throughout the dataset, likely due in part to the dependence of this effect on velocity and noise levels (Figure S4), but our analysis identified a consistent group of BTSP-like place fields that did not change with the total number of clusters used (Figures S5D–F). We note that although we have labeled the clusters by visual inspection of their properties, the difference in first trial shift and other measures are not a pre-configured requirement of the model but rather one of several distinguishing feature that are discovered in an unsupervised manner directly from the data. Our model indicates that BTSP-like characteristics are intrinsically strong sources of variance across place fields.

Since we could reliably identify a subset of BTSP-like place fields in the dataset, we asked how the fraction of place fields forming with these dynamics changes as a function of experience (Figure 5E). On most days and trial blocks, the fraction of BTSP-like fields was below 40% of all fields, but on the first exposure to the novel context, this increased to nearly 60% on average (Figure 5E). The enrichment of BTSP-like place field formation was also experience dependent, decaying over the course of the 3-day experiment sequence. The fraction of fields classified as BTSP-like in each condition also correlated with the strength of the correlation between velocity and field width (Figure 5F) and the distance of population tuning drift during the initial trials in each context block (Figure 5G), suggesting that these population-level measures are good proxies for the prevalence of BTSP. As noted earlier, we repeated all of these 3-day experiment sequences in a second novel context for each mouse, and we found similar experience-dependent trends for each novel context for all main analyses (Figure S6). Overall, our field classification analysis is in good agreement with the population-level measures of BTSP prevalence and lends further support to an increased rate of BTSP-mediated acquisition of feature tuning during the initial encounters with novel experiences.

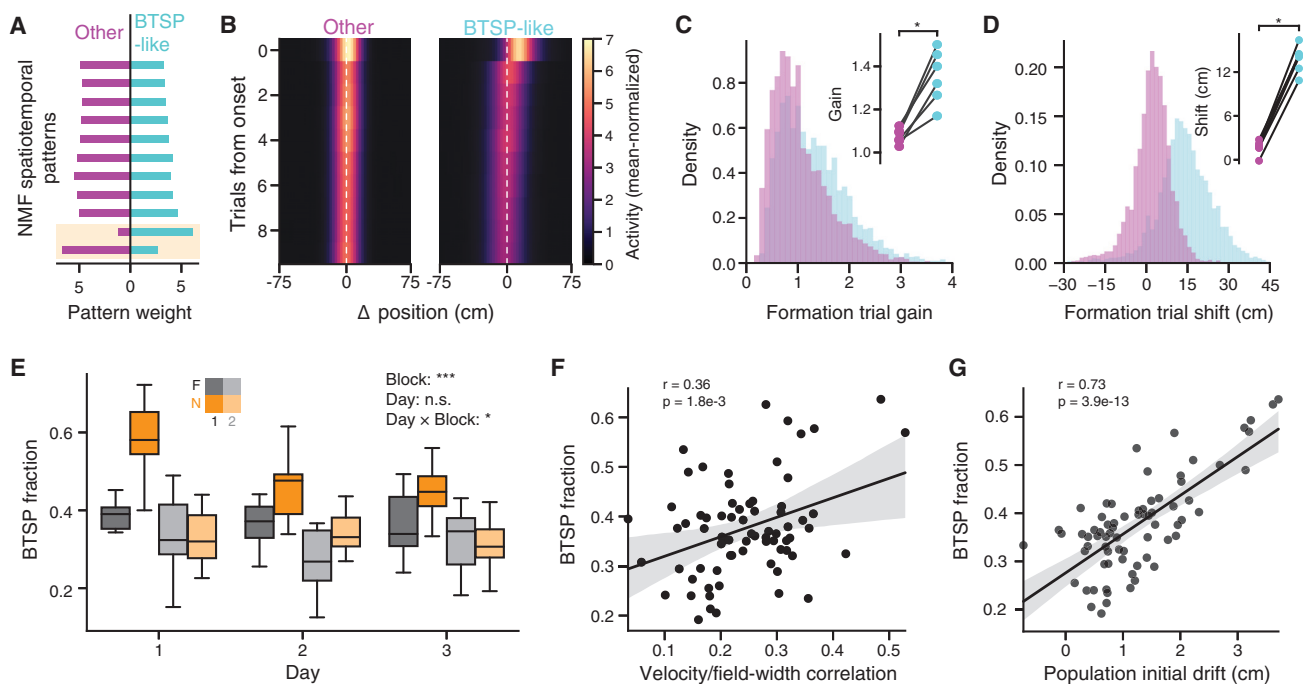


Figure 5. BTSP-like field formation is enriched in the novel context and decays with experience

(A) Place fields were divided into two groups via K-means clustering in NMF space. Plotted is the average NMF pattern weights for place fields in the two groups, labeled as BTSP-like and Other. BTSP-like fields have far greater weight on the shifted, plateau-like component on the formation trial, and generally less weight on later trial components compared with the Other cluster. Components are sorted as in Figures 4D and 4E.

(B) Average spatial tuning over trials for place fields in each group. As suggested by the pattern weights in (A), the BTSP-like group exhibits forward-shifted activity on the formation trial, and the formation trial amplitude is greater than later trials.

(C) Formation trial activity gain for each place field, calculated as the peak activity on the formation trial divided by the peak activity from the average of the remaining trials. BTSP-like cells show higher formation trial gain (Kolmogorov-Smirnov test, $p = 1.86 \times 10^{-77}$). Inset: average gain for place fields in each group, for each mouse (Wilcoxon signed-rank test).

(D) Formation trial shift for each place field, plotted as in (C) (Kolmogorov-Smirnov test, $p = 0$, inset: Wilcoxon signed-rank test).

(E) Fraction of place cells classified as BTSP-like throughout the experiment, shown separately for each trial block on each day. BTSP-like place field formation is enriched during the first exposure to the novel context and decays with experience (linear mixed effect model with main effects of trial block [categorical] and day [continuous], significance inset).

(F) Correlation between BTSP fractions and the slope of the velocity/field-width regression.

(G) Correlation between BTSP fraction and the population-level initial drift score. In (F) and (G), each point is a trial block from a single day and single mouse. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Spatiotemporal dynamics of place field formation

Our data delineated a transient period of enriched, BTSP-like place field formation during novel learning, but it is possible that these dynamics do not affect all components of new experiences homogeneously. Numerous prior works have detailed how hippocampal representations are biased by the presence of salient cues (Bourboulou et al., 2019) and reinforcement (Hollup et al., 2001; Zaremba et al., 2017; Dupret et al., 2010); these effects may be driven in part by the engagement of different plasticity mechanisms. On examining the spatial distribution of place field formation events, we found that BTSP-like and Other fields tended to concentrate in opposing regions of the virtual environments: BTSP-like fields accumulated in the regions between reward zones, whereas Other fields were particularly enriched near the reward zones (Figure 6A). Notably both of these distributions were correlated with the spatial distribution of velocities (Figure 6B): animals tended to run quickly between reward zones and reliably decelerated as they approached each

reward. It is important to account for this correlation, as our classification of BTSP-like fields is determined principally from the spatial shift in formation trial firing (Figures 4, 5, S5J, and S5K), which will be more difficult to detect when the animal is running at slower velocities (Figure S4). We addressed this confound using a linear model that related velocity to field formation density across the environment.

Our model learned a spatial velocity filter for each animal and each field type (Figure 6C; see STAR Methods). For Other cells, the velocity at the current position consistently predicted a strong negative effect on field density, whereas for BTSP-like cells, field density was positively predicted by velocity across a range of spatial lags. Overall, the velocity models explained a majority of variance in the distribution of field formation events for both field types (Figure 6D; see also Figure S7). In Figure 6E, we plotted the field distributions for each mouse and field type. Subtracting the model-predicted field distributions from these curves, we found that the residual field counts were largely

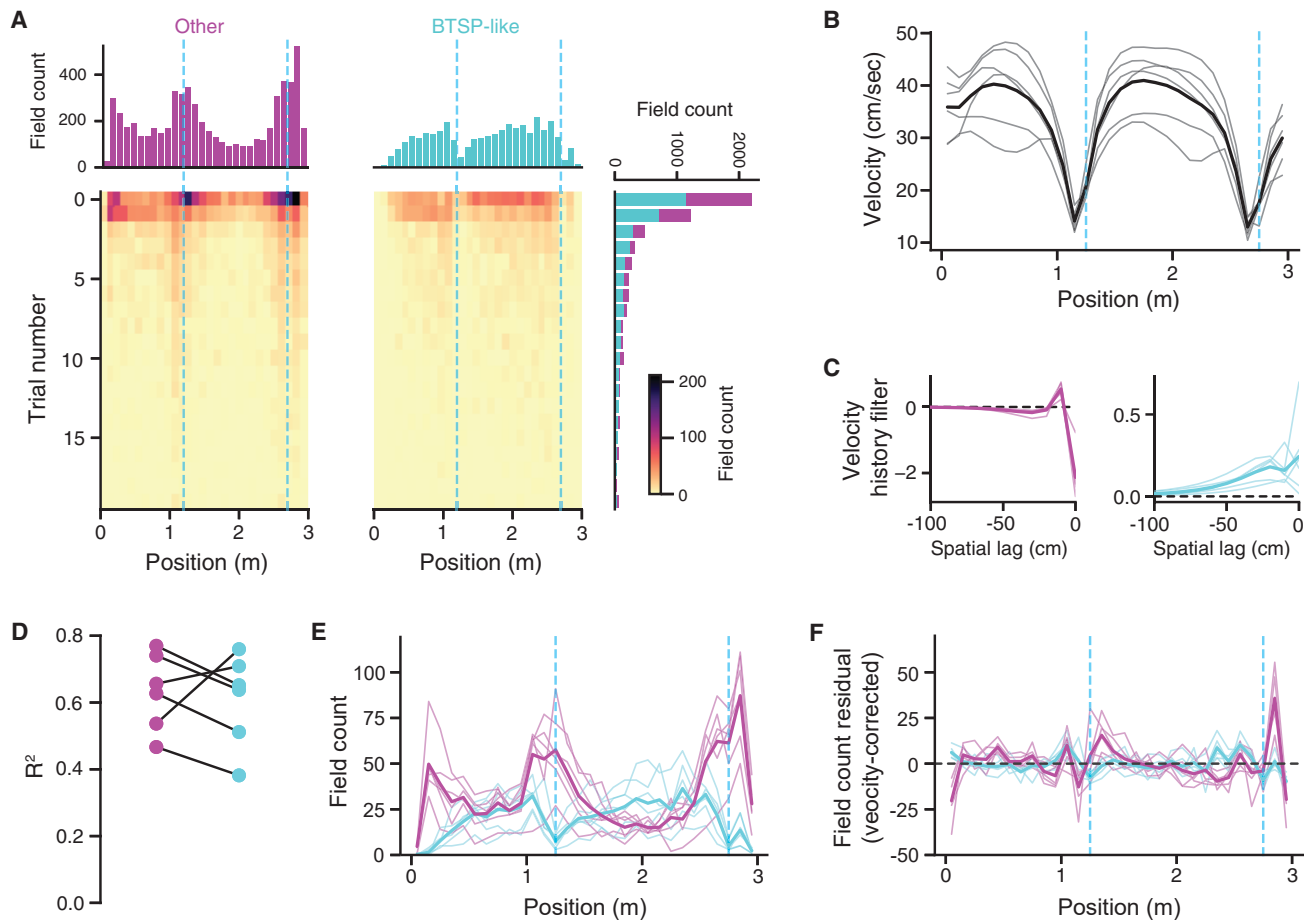


Figure 6. The spatial distribution of place field formation events

(A) Histogram of place field formation events over positions and trial number (from the start of the context block), with marginal distributions. Data are pooled from all conditions. Dashed-blue lines indicate reward zones.
 (B) Average velocity profile across the environment, for each mouse. The mice slow down considerably on approach to the reward zones.
 (C) Spatial kernels fit to predict place field density from the average velocity profile from each mouse. These describe the weighted average of velocity over past and present positions that best predicts the place field density at the present position. Velocity has a large negative effect on “Other” field density on a short spatial scale. Conversely, BTSP-like field density is positively affected by velocity, on a longer spatial scale.
 (D) Prediction quality of the velocity-to-field density model learned in (C).
 (E) As in (A), the distribution of place field formation events for Other and BTSP-like fields, shown separately for each mouse.
 (F) As in (E), but showing the residual field counts obtained after subtracting off the predictions of the model in (C and D). The residuals are largely spatially homogeneous and centered at zero, reflecting that the majority of variance is explained by velocity.

spatially homogenized and centered at zero, effectively removing the effect of the reward zones (Figure 6F). The residual increased frequency of Other cells near the very end of the track is compatible with the slight enrichment of neurons that exhibit place fields in both contexts at this location (Figure S1J). Since this region of the track is highly visually similar between both contexts, this may correspond to a reduction in plasticity-driven field formation here in favor of recruiting the same familiar representation in both contexts. In sum, given the current data and our field classification methods, we were unable to determine conclusively whether BTSP differentially contributed to the spatial distribution of place fields.

The prior analysis focused on the spatial distribution of place field formation events pooled over all trials. On any given trial

though, we hypothesized that the locations of new place fields might be constrained by the representation assembled over prior learning. This could arise from the recruitment of lateral inhibition at locations with existing place fields (Rolotti et al., 2022; Milstein et al., 2021; Robinson et al., 2020), promoting competitive interactions that repel accumulating fields away from one another. To test this idea, we examined the formation trial for every place field and measured the average spatial distance between that field and all fields of the same class that formed on the next trial (Figures 7A and 7B). Since we were interested in the next-trial field distance that was not explained by the overall distribution of place field locations, we normalized this distance by a null distribution that was constructed by randomly permuting the formation trial between fields within a trial block (Figure 7B). In

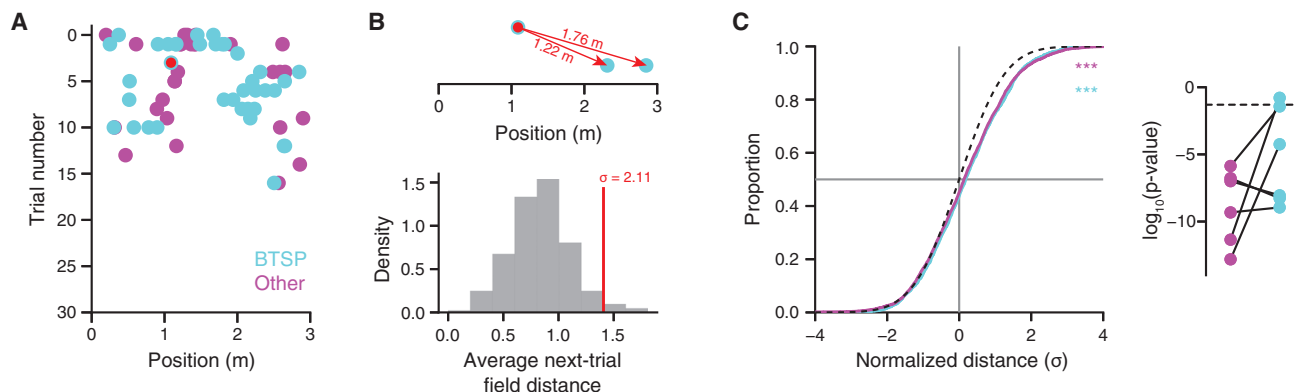


Figure 7. Place field accumulation follows a nonstationary distribution over positions

(A) Scatter plot of place field formation event over trials and positions, for an example session (day 1, first novel context exposure).

(B) Top: for every place field, the distance was calculated between its location and the location of all place fields that appeared on the next trial (of the same BTSP classification). Bottom: the average next-trial field distance was compared with a null distribution built by randomly permuting the formation trials of all place fields in the trial block. A normalized distance (σ) was derived by standardizing the true distance relative to the null distribution.

(C) Cumulative distribution of normalized distances across the dataset for each place field formation type. The standard normal distribution (chance) is plotted in black. Both BTSP and Other fields show greater next-trial field distances than expected by chance (inset: one-sample Kolmogorov-Smirnov [KS] test against the standard normal distribution). Right: significance of KS test computed for each mouse separately. Place fields accumulated at farther distances compared with prior trial fields than expected by random sampling in nearly all mice and conditions. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure 7C, we plotted the distribution of normalized distances for BTSP and Other fields. For both groups of fields, place fields tended to form farther away from fields that formed on the previous trial than would be expected from random sampling of the environment. This history dependence indicates the presence of competitive interactions that act to disperse new place fields, which could help to drive pattern separation of nearby locations.

DISCUSSION

Hippocampal circuits are remarkably plastic, with the ability to construct precise representations of novel experiences with very few exposures (Wilson and McNaughton, 1993). This feature is likely critical to the role of the hippocampus in memory storage: it can serve as a fast learning module that rapidly encodes new information, where it is retained and refined for a short period prior to its transfer to and long-term storage in the cortex. (McClelland et al., 1995; Roxin and Fusi, 2013). The recent discovery of BTSP is a striking demonstration of the speed of synaptic learning in the hippocampus (Bittner et al., 2015, 2017). In our work here, we have quantified the potential impact of this novel plasticity rule at the scale of large neuronal populations and tested a specific role in encoding novel experiences. Considering multiple lines of analysis, our results are consistent with the hypothesis that a sizable fraction of *de novo* place fields form via BTSP and that the probability of BTSP events is regulated by the novelty of ongoing experience.

Other reports have questioned the ubiquity of plateau-dependent plasticity events during place field formation (Cohen et al., 2017; Dong et al., 2021). New place fields can appear in a familiar or novel environment in the absence of plateau-associated complex spiking (Cohen et al., 2017), and CA1 neurons can exhibit

latent, subthreshold spatial tuning (Lee et al., 2012), which could possibly be amplified to suprathreshold place fields through other synaptic learning processes (McKenzie et al., 2021). In all conditions during our experiments, we detected features of place field formation that are correlated with BTSP, both at the level of neural ensembles (Figures 2 and 3) and single neurons (Figures 4 and 5). Although prior work on BTSP mainly used whole-cell patch recordings to unambiguously identify plateau potentials (Bittner et al., 2015, 2017; Zhao et al., 2020, 2022), our population-scale approach necessitates that we identify these events indirectly through changes in response amplitude and spatial tuning shifts, derived, respectively, from the prolonged somatic depolarization induced by the putative plateau potential (Epsztein et al., 2011; Bittner et al., 2015) and the asymmetric plasticity kernel of BTSP (Bittner et al., 2017). Although we cannot detect the presence of burst firing directly, this firing mode is ubiquitous among CA1 pyramidal neurons (Harvey et al., 2009; Grienberger et al., 2014) and strongly implicated in plateau-mediated plasticity (Bittner et al., 2015). Fluorescent calcium reporters are biased toward detecting spike bursts at the expense of isolated action potentials (Schoenfeld et al., 2021), as fluorescence generally increases with the number of underlying action potentials (Chen et al., 2013a). It is likely that location shifts we calculate are derived mainly from high-frequency spiking events near the center of the place field (Harvey et al., 2009). Although indirect, our method allowed us to study place field dynamics at scale and compare the relative frequency of field formation events as animals became increasingly familiar with new environments over days, questions that are impossible to tackle with intracellular recordings limited to short recordings in single neuron preparations (Bittner et al., 2015, 2017). Ultimately though, our approach can only detect features correlated with BTSP, with direct measurement of the membrane potential required to unambiguously identify these plasticity events.

All correlates of BTSP in the dataset exhibited strong experience-dependent effects, with an enrichment of BTSP-compatible dynamics specifically during the first exposures to the novel environment that decayed over subsequent days. Of course, it is obvious that over multiple days of experience in the same environment, we should expect progressively fewer place fields to appear via ongoing plasticity, with much of the representation simply recalled due to prior learning (although representations continue to drift in familiar environments, see [Mankin et al., 2012](#); [Ziv et al., 2013](#)). However, our data indicate that correlates of burst-dependent plasticity are particularly present during early learning in a novel experience, where our classification estimates that the majority of place fields appear *de novo* with BTSP-like characteristics ([Figure 5](#)). We also persistently identified a sizable fraction of BTSP-like fields in other conditions, suggesting a continuous baseline level of place field turnover due to BTSP that may contribute to representational drift ([Mankin et al., 2012](#); [Ziv et al., 2013](#)). Plateau potentials are generated through the convergence of coincident presynaptic inputs that triggers active dendritic conductances in the apical tuft ([Takahashi and Magee, 2009](#)), but the dendritic arbor of CA1 pyramidal neurons is also extensively regulated by local feedback inhibition ([Royer et al., 2012](#); [Lovett-Barron et al., 2012](#)). These circuits actively limit dendritic electrogenesis and tightly restrict the number of place fields that can be simultaneously induced via BTSP ([Rolotti et al., 2022](#)). This result implies that other factors must be present that transiently free the circuit of these constraints, in order to rapidly build new representations during novel experience.

There is already evidence for a temporary reduction in dendritic inhibition during novel experience in the hippocampus ([Sheffield et al., 2017](#); [Geiller et al., 2020](#)). We hypothesize that these inhibitory dynamics are at least partially regulated by factors originating outside of the hippocampus, such as neuromodulatory inputs from the basal forebrain or brainstem ([Palacios-Filardo and Mellor, 2019](#)) that could be involved in broader, brain-wide signaling of novelty detection. In particular, the locus coeruleus is highly sensitive to novelty and provides dense dopaminergic and noradrenergic input to CA1 ([Takeuchi et al., 2016](#)), and these projections have been optogenetically manipulated to increase place field density during a reward learning behavior in mice ([Kaufman et al., 2020](#)). Cholinergic neurons are also known to respond to reinforcement and are sensitive to stimulus uncertainty ([Hangya et al., 2015](#)), and cholinergic projections from the medial septum ramify extensively in CA1, where they are thought to modulate network function and plasticity in response to arousal ([Teles-Grilo Ruivo and Mellor, 2013](#)). Characterization of these inputs *in vivo* remains a relatively nascent endeavor and, given the diversity of receptor expression in the hippocampal circuit ([Teles-Grilo Ruivo and Mellor, 2013](#)), it is likely that they affect both the dendritic excitability of pyramidal neurons and also the recruitment of distinct inhibitory microcircuits. Overall, we hypothesize that the interplay of external novelty signals and regulation of local inhibition may provide a temporary window of hyper-excitability during new experience, which may permit rapid assembly of representations through BTSP. Abstractly, the temporary shift to frequent burst-mediated place field formation may contribute to an adaptive learning rate in the hippocampus, where the circuit responds

to a large change in its input statistics by scaling up the speed of synaptic updates in order to encode new information.

We found that place fields putatively forming via BTSP were generally less precise than the remaining fields ([Figure S5](#)). Although many neurons may first reach threshold via BTSP, it seems likely that this initial tuning is refined through secondary mechanisms, such as conventional Hebbian rules and local competitive interactions with other pyramidal neurons mediated through lateral inhibition ([Mehta et al., 2000](#); [Cohen et al., 2017](#); [McKenzie et al., 2021](#); [Robinson et al., 2020](#)). Considering that many of the non-BTSP fields are pre-existing place fields learned during previous experience, this could explain why BTSP fields were consistently wider, having not yet undergone further refinement. This is also compatible with the theory that the hippocampus is actively learning a compressed, decorrelated representation of the environment ([Gluck and Myers, 1993](#); [McClelland et al., 1995](#); [Schapiro et al., 2017](#); [Benna and Fusi, 2021](#)), as narrower place fields will decrease the correlations between nearby, similar locations. In line with this argument, we also found that the accumulation of place fields was history dependent ([Figure 7](#)). New fields tended to appear at locations farther away from those forming on the previous trial, suggesting that local CA1 networks may facilitate competitive interactions between pyramidal neurons ([Rolotti et al., 2022](#)) in order to disperse the representation and minimize the neural correlation between nearby locations.

Earlier studies reported that many CA1 place fields undergo a transient period of backward, asymmetric expansion during initial traversals of an environment ([Mehta et al., 1997](#)), an effect that encoded the direction of travel and could be modeled by plasticity at feedforward CA3-to-CA1 synapses ([Mehta et al., 2000](#)). In our experiments, animals were constrained to run in a single direction through the track, and hence, it is possible that some component of field width and drift is due to these effects as well, especially residual drift on the trials immediately after place field formation ([Figure S3](#)). However, this mechanism cannot explain our main field shift effect, as the finding in [Mehta et al. \(1997\)](#) is produced by an increasing negative skew in the place field that translates into a low firing rate tail along the region of the track leading into the place field, although the location of peak firing is mostly unchanged ([Mehta et al., 2000](#)). Due to the sensitivity of fluorescence calcium indicators, we cannot reliably detect very low firing rate activity, and hence, our ability to resolve this asymmetric expansion is likely limited. Further, the shift observed in the study of [Mehta et al. \(1997\)](#) is small (only ~2.5 cm even in a novel context) and occurs gradually over many trials, whereas our effect is abrupt and much larger (~15 cm for BTSP-like cells on average). However, this discussion highlights a general caveat to the use of calcium indicators, in that we underappreciate more subtle fluctuations in activity or the dynamics of low firing rate neurons, due to our bias toward detecting higher amplitude, burst-driven activity. Given that low firing rate neurons have been linked to novelty-related plasticity in CA1 ([Grosmark and Buzsáki, 2016](#)), it will be critical to examine these changes as well using this longitudinal approach, which may be facilitated by the development of next-generation calcium and voltage imaging tools ([Villette et al., 2019](#); [Adam et al., 2019](#)).

In a similar vein, there is a large body of research that highlights the role of theta oscillations and sequential spike sequences for orchestrating synaptic plasticity in the hippocampus (Buzsáki and Moser, 2013; Dragoi and Buzsáki, 2006), which could contribute to backward expansion of spatial tuning via Hebbian learning (Mehta et al., 2000). Notably, theta sequences are organized within the first few trials in a novel environment (Feng et al., 2015). Theta also appears to exert control over the probability of dendritic plateau potentials (Bittner et al., 2015). Future work will be required to study the interplay between these mechanisms during theta states and their differential contribution to plasticity across contexts, but it is unlikely that they operate in mutually exclusive regimes. In a more recent study of place field plasticity, Dong et al. (2021) reported pervasive and continual backward drift of CA1 place fields in VR that they suggested could arise from STDP (although they reported forward-shifting in CA1 during initial trials). However, the shifts in field locations in our data were overwhelmingly limited to the first few trials following place field formation (Figures 3, 4, and S4), which agrees with prior work in freely moving animals where individual place fields do not continuously drift (Frank et al., 2004; Mehta et al., 2000).

Although here we focused on overall environmental novelty, the hippocampal place code is also heavily influenced by the density and salience of environmental cues (Manns and Eichenbaum, 2009; Bourboulou et al., 2019) and the presence of reinforcement (Hollup et al., 2001; Zarella et al., 2017; Dupret et al., 2010), representations of which may actively shape behavior (Robinson et al., 2020). We searched for evidence of differential concentration of BTSP-like field formation along the virtual track but found that the distribution was mainly uniform over space after regressing out the component correlated with velocity (Figure 6). In our experiment, the location of reward is highly confounded with the velocity profile of the animal, and hence, we cannot determine concretely which is the causal factor for the variability in field distributions. Since the field shift associated with BTSP is most apparent at higher running speeds and this is a critical component of our field classification method, it is impossible for us to distinguish a reduction in BTSP events from a failure to detect them at the lower running speeds typically seen near the reward zones, and hence, our work does not preclude a connection between burst-dependent plasticity mechanisms and reinforcement (Grienberger and Magee, 2021; Milstein et al., 2021).

Our results are congruent with burst-dependent plasticity as an important contributor to representation learning in hippocampal area CA1. It remains unclear whether BTSP is also present in pyramidal neurons in other hippocampal or neocortical networks and what functional consequences this would have, given the differing circuit architecture (Van Strien et al., 2009). In CA1, pyramidal neurons receive a convergence of inputs from intrahippocampal recurrent networks in CA3, which are believed to store memories through attractor dynamics (Rolls, 2007), and from the superficial entorhinal cortex, which can directly relay sensory information about ongoing experience. Novelty detection may be central to the function of CA1 within the broader hippocampal circuit (McClelland et al., 1995; Lisman and Otmakhova, 2001). We suggest that the regulation of burst-dependent plasticity in

CA1 may selectively permit the integration of novel information into the hippocampus, instigating a cascade of plasticity throughout the hippocampal-cortical loop that could optimize internal representations to adapt to large changes in the statistics of ongoing experience.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Behavior and imaging
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Image preprocessing
 - Neural data analysis

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2022.03.026>.

ACKNOWLEDGMENTS

We thank Rick Warren for the design of the lightweight running wheel and for technical advice. We thank Tanya Tabachnik and the Advanced Instrumentation platform for design and machining of head-fixation hardware. We thank Zhenrui Liao, Pamela Rivière, and Nick Robinson for helpful discussions, and for comments on an earlier version of this manuscript. J.B.P. is supported by National Institute of Mental Health (NIMH) F31MH121058. J.C.B. is supported by NIMH F31NS110316. S.V.R. was supported by NIMH F31MH117892. S.F. is supported by NSF's NeuroNex program award DBI-1707398, the Gatsby Charitable Foundation, the Simons Foundation, and the Swartz Foundation. A.L. is supported by NIMH 1R01MH124047 and 1R01MH124867; and National Institute of Neurological Disorders and Stroke (NINDS) 1R01NS121106, 1U19NS104590, and 1U01NS115530.

AUTHOR CONTRIBUTIONS

J.B.P., S.F., and A.L. conceived the project. J.B.P. designed experiments, collected data, and performed analysis and modeling. J.C.B. built the virtual reality and behavioral control systems, with assistance from J.B.P. S.V.R. provided critical input on analysis. J.B.P., S.F., and A.L. wrote the paper with input from all authors.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: July 14, 2021
 Revised: January 19, 2022
 Accepted: March 16, 2022
 Published: April 20, 2022

REFERENCES

Adam, Y., Kim, J.J., Lou, S., Zhao, Y., Xie, M.E., Brinks, D., Wu, H., Mostajir-Radji, M.A., Kheifets, S., Parot, V., et al. (2019). Voltage imaging and

- optogenetics reveal behaviour-dependent changes in hippocampal dynamics. *Nature* 569, 413–417.
- Adoff, M.D., Climer, J.R., Davoudi, H., Marvin, J.S., Looger, L.L., and Dombeck, D.A. (2021). The functional organization of excitatory synaptic input to place cells. *Nat. Commun.* 12, 3558.
- Ahmed, M.S., Priestley, J.B., Castro, A., Stefanini, F., Solis Canales, A.S.S., Balough, E.M., Lavoie, E., Mazzucato, L., Fusi, S., and Losonczy, A. (2020). Hippocampal network reorganization underlies the formation of a temporal association memory. *Neuron* 107, 283, e6–291.e6.
- Benna, M.K., and Fusi, S. (2016). Computational principles of synaptic memory consolidation. *Nat. Neurosci.* 19, 1697–1706.
- Benna, M.K., and Fusi, S. (2021). Place cells may simply be memory cells: memory compression leads to spatial tuning and history dependence. *Proc. Natl. Acad. Sci. USA* 118, e2018422118.
- Bittner, K.C., Grienberger, C., Vaidya, S.P., Milstein, A.D., Macklin, J.J., Suh, J., Tonegawa, S., and Magee, J.C. (2015). Conjunctive input processing drives feature selectivity in hippocampal ca1 neurons. *Nat. Neurosci.* 18, 1133–1142.
- Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* 357, 1033–1036.
- Bocchio, M., Gouny, C., Angulo-Garcia, D., Toulat, T., Tressard, T., Quiroli, E., Baude, A., and Cossart, R. (2020). Hippocampal hub neurons maintain distinct connectivity throughout their lifetime. *Nat. Commun.* 11, 4559.
- Bourboulou, R., Marti, G., Michon, F.-X., El Feghaly, E., Nougier, M., Robbe, D., Koenig, J., and Epsztein, J. (2019). Dynamic control of hippocampal spatial coding resolution by local visual cues. *eLife* 8, e44487.
- Buzsáki, G., and Moser, E.I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat. Neurosci.* 16, 130–138.
- Carpenter, G.A., and Grossberg, S. (1991). Pattern Recognition by Self-Organizing Neural Networks (MIT Press).
- Chen, G., King, J.A., Burgess, N., and O’Keefe, J. (2013). How vision and movement combine in the hippocampal place code. *Proc. Natl. Acad. Sci. USA* 110, 378–383.
- Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013a). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300.
- Cohen, J.D., Bolstad, M., and Lee, A.K. (2017). Experience-dependent shaping of hippocampal CA1 intracellular activity in novel and familiar environments. *eLife* 6, e23040.
- Diamantaki, M., Coletta, S., Nasr, K., Zeraati, R., Laturnus, S., Berens, P., Preston-Ferrer, P., and Burgalossi, A. (2018). Manipulating hippocampal place cell activity by single-cell stimulation in freely moving mice. *Cell Rep.* 23, 32–38.
- Dong, C., Madar, A.D., and Sheffield, M.E.J. (2021). Distinct place cell dynamics in ca1 and ca3 encode experience in new environments. *Nat. Commun.* 12, 2977.
- Dragoi, G., and Buzsáki, G. (2006). Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* 50, 145–157.
- Druckmann, S., Feng, L., Lee, B., Yook, C., Zhao, T., Magee, J.C., and Kim, J. (2014). Structured synaptic connectivity between hippocampal regions. *Neuron* 81, 629–640.
- Dupret, D., O’Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* 13, 995–1002.
- Eichenbaum, H. (2017). The role of the hippocampus in navigation is memory. *J. Neurophysiol.* 117, 1785–1796.
- Epsztein, J., Brecht, M., and Lee, A.K. (2011). Intracellular determinants of hippocampal ca1 place and silent cell activity in a novel environment. *Neuron* 70, 109–120.
- Farooq, U., and Dragoi, G. (2019). Emergence of preconfigured and plastic time-compressed sequences in early postnatal development. *Science* 363, 168–173.
- Feng, T., Silva, D., and Foster, D.J. (2015). Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *J. Neurosci.* 35, 4890–4902.
- Frank, L.M., Stanley, G.B., and Brown, E.N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *J. Neurosci.* 24, 7681–7689.
- Friedrich, J., Zhou, P., and Paninski, L. (2017). Fast online deconvolution of calcium imaging data. *PLoS Comput. Biol.* 13, e1005423.
- Gauthier, J.L., and Tank, D.W. (2018). A dedicated population for reward coding in the hippocampus. *Neuron* 99, 179.e7, 193.e7.
- Geiller, T., Sadeh, S., Rolotti, S.V., Blockus, H., Vancura, B., Negrean, A., Murray, A.J., Rózsa, B., Polleux, F., Clopath, C., and Losonczy, A. (2022). Local circuit amplification of spatial selectivity in the hippocampus. *Nature* 601, 105–109.
- Geiller, T., Vancura, B., Terada, S., Troullinou, E., Chavlis, S., Tsagkatakis, G., Tsakalides, P., Ócsai, K., Poirazi, P., Rózsa, B.J., and Losonczy, A. (2020). Large-scale 3d two-photon imaging of molecularly identified ca1 interneuron dynamics in behaving mice. *Neuron* 108, 968.e9, 983.e9.
- Gerstner, W., Kistler, W.M., Naud, R., and Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition* (Cambridge University Press).
- Gerstner, W., Lehmann, M., Liakoni, V., Corneil, D., and Brea, J. (2018). Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules. *Front. Neural Circuits* 12, 53.
- Gluck, M.A., and Myers, C.E. (1993). Hippocampal mediation of stimulus representation: a computational theory. *Hippocampus* 3, 491–516.
- Grienberger, C., Chen, X., and Konnerth, A. (2014). NMDA receptor-dependent multidendrite Ca(2+) spikes required for hippocampal burst firing in vivo. *Neuron* 81, 1274–1281.
- Grienberger, C., and Magee, J.C. (2021). Entorhinal cortex directs learning-related changes in ca1 representations. Preprint at bioRxiv <https://www.biorxiv.org/content/early/2021/12/12/2021.12.10.472158>.
- Grosmark, A.D., and Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science* 351, 1440–1443.
- Hangya, B., Ranade, S.P., Lorenc, M., and Kepecs, A. (2015). Central cholinergic neurons are rapidly recruited by reinforcement feedback. *Cell* 162, 1155–1168.
- Harvey, C.D., Collman, F., Dombeck, D.A., and Tank, D.W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* 461, 941–946.
- Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.-B., and Moser, E.I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J. Neurosci.* 21, 1635–1644.
- Kaifosh, P., Zaremba, J.D., Danielson, N.B., and Losonczy, A. (2014). Sima: python software for analysis of dynamic fluorescence imaging data. *Front. Neuroinform.* 8, 80.
- Kaufman, A.M., Geiller, T., and Losonczy, A. (2020). A role for the locus coeruleus in hippocampal ca1 place cell reorganization during spatial reward learning. *Neuron* 105, 1018.e4, 1026.e4.
- Lee, D., Lin, B.-J., and Lee, A.K. (2012). Hippocampal place fields emerge upon single-cell manipulation of excitability during behavior. *Science* 337, 849–853.
- Lisman, J.E., and Otmakhova, N.A. (2001). Storage, recall, and novelty detection of sequences by the hippocampus: elaborating on the Socratic model to account for normal and aberrant effects of dopamine. *Hippocampus* 11, 551–568.
- Lovett-Barron, M., Kaifosh, P., Kheirbek, M.A., Danielson, N., Zaremba, J.D., Reardon, T.R., Turi, G.F., Hen, R., Zemelman, B.V., and Losonczy, A. (2014).

- Dendritic inhibition in the hippocampus supports fear learning. *Science* 343, 857–863.
- Lovett-Barron, M., Turi, G.F., Kaifosh, P., Lee, P.H., Bolze, F., Sun, X.-H., Nicoud, J.-F., Zemelman, B.V., Sternson, S.M., and Losonczy, A. (2012). Regulation of neuronal input transformations by tunable dendritic inhibition. *Nat. Neurosci.* 15, 423–430, S1.
- Magee, J.C., and Grienberger, C. (2020). Synaptic plasticity forms and functions. *Annu. Rev. Neurosci.* 43, 95–117.
- Mankin, E.A., Diehl, G.W., Sparks, F.T., Leutgeb, S., and Leutgeb, J.K. (2015). Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts. *Neuron* 85, 190–201.
- Mankin, E.A., Sparks, F.T., Slayeh, B., Sutherland, R.J., Leutgeb, S., and Leutgeb, J.K. (2012). Neuronal code for extended time in the hippocampus. *Proc. Natl. Acad. Sci. USA* 109, 19462–19467.
- Manns, J.R., and Eichenbaum, H. (2009). A cognitive map for object memory in the hippocampus. *Learn. Mem.* 16, 616–624.
- McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* 102, 419–457.
- McKenzie, S., Huszár, R., English, D.F., Kim, K., Christensen, F., Yoon, E., and Buzsáki, G. (2021). Preexisting hippocampal network dynamics constrain optogenetically induced place fields. *Neuron* 109, 1040.e7, 1054.e7.
- Mehta, M.R., Barnes, C.A., and McNaughton, B.L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA* 94, 8918–8921.
- Mehta, M.R., Quirk, M.C., and Wilson, M.A. (2000). Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25, 707–715.
- Milligan, G.W., and Cooper, M.C. (1985). An examination of procedures for determining the number of clusters in a data set. *Psychometrika* 50, 159–179.
- Milstein, A.D., Li, Y., Bittner, K.C., Grienberger, C., Soltesz, I., Magee, J.C., and Romani, S. (2021). Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *eLife* 10, e73046.
- Moser, E.I., Kropff, E., and Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.* 31, 69–89.
- Muessig, L., Lasek, M., Varsavsky, I., Cacucci, F., and Wills, T.J. (2019). Coordinated emergence of hippocampal replay and theta sequences during post-natal development. *Curr. Biol.* 29, 834.e4, 840.e4.
- Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* 7, 1951–1968.
- O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34, 171–175.
- Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L.F., Dalgleish, H., Carandini, M., and Harris, K.D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at bioRxiv. <https://doi.org/10.1101/061507>.
- Pachitariu, M., Stringer, C., and Harris, K.D. (2018). Robustness of spike deconvolution for neuronal calcium imaging. *J. Neurosci.* 38, 7976–7985.
- Palacios-Filardo, J., and Mellor, J.R. (2019). Neuromodulation of hippocampal long-term synaptic plasticity. *Curr. Opin. Neurobiol.* 54, 37–43.
- Robinson, N.T., Descamps, L.A., Russell, L.E., Buchholz, M.O., Bicknell, B.A., Antonov, G.K., Lau, J.Y., Nutbrown, R., Schmidt-Hieber, C., and Häusser, M. (2020). Targeted activation of hippocampal place cells drives memory-guided spatial behavior. *Cell* 183, 1586–1599.
- Rolls, E.T. (2007). An attractor network in the hippocampus: theory and neurophysiology. *Learn. Mem.* 14, 714–731.
- Rolotti, S.V., Ahmed, M.S., Szoboszlai, M., Geiller, T., Negrean, A., Blockus, H., Gonzalez, K.C., Sparks, F.T., Solis Canales, A.S., Tuttman, A.L., et al. (2022). Local feedback inhibition tightly controls rapid formation of hippocampal place fields. *Neuron* 110, 783.e6, 794.e6.
- Roxin, A., and Fusi, S. (2013). Efficient partitioning of memory systems and its importance for memory consolidation. *PLoS Comput. Biol.* 9, e1003146.
- Royer, S., Zemelman, B.V., Losonczy, A., Kim, J., Chance, F., Magee, J.C., and Buzsáki, G. (2012). Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat. Neurosci.* 15, 769–775.
- Schapiro, A.C., Turk-Browne, N.B., Botvinick, M.M., and Norman, K.A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 372, 20160049.
- Schoenfeld, G., Carta, S., Rupprecht, P., Ayaz, A., and Helmchen, F. (2021). In vivo calcium imaging of ca3 pyramidal neuron populations in adult mouse hippocampus. *eNeuro* 8, ENEURO.0023-21.2021.
- Sheffield, M.E., and Dombeck, D.A. (2019). Dendritic mechanisms of hippocampal place field formation. *Curr. Opin. Neurobiol.* 54, 1–11.
- Sheffield, M.E.J., Adoff, M.D., and Dombeck, D.A. (2017). Increased prevalence of calcium transients across the dendritic arbor during place field formation. *Neuron* 96, 490.e5, 504.e5.
- Silva, D., Feng, T., and Foster, D.J. (2015). Trajectory events across hippocampal place cells require previous experience. *Nat. Neurosci.* 18, 1772–1779.
- Souza, B.C., Pavão, R., Belchior, H., and Tort, A.B.L. (2018). On information metrics for spatial coding. *Neuroscience* 375, 62–73.
- Takahashi, H., and Magee, J.C. (2009). Pathway interactions and synaptic plasticity in the dendritic tuft regions of ca1 pyramidal neurons. *Neuron* 62, 102–111.
- Takeuchi, T., Duzkiewicz, A.J., Sonneborn, A., Spooner, P.A., Yamasaki, M., Watanabe, M., Smith, C.C., Fernández, G., Deisseroth, K., Greene, R.W., and Morris, R.G.M. (2016). Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature* 537, 357–362.
- Teles-Grilo Ruivo, L.M., and Mellor, J.R. (2013). Cholinergic modulation of hippocampal network function. *Front. Synaptic Neurosci.* 5, 2.
- Van Strien, N.M., Cappaert, N.L., and Witter, M.P. (2009). The anatomy of memory: an interactive overview of the parahippocampal-hippocampal network. *Nat. Rev. Neurosci.* 10, 272–282.
- Villette, V., Chavarha, M., Dimov, I.K., Bradley, J., Pradhan, L., Mathieu, B., Evans, S.W., Chamberland, S., Shi, D., Yang, R., et al. (2019). Ultrafast two-photon imaging of a high-gain voltage indicator in awake behaving mice. *Cell* 179, 1590.e23, 1608.e23.
- Warren, R.A., Zhang, Q., Hoffman, J.R., Li, E.Y., Hong, Y.K., Bruno, R.M., and Sawtell, N.B. (2021). A rapid whisker-based decision underlying skilled locomotion in mice. *eLife* 10, e63596.
- Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* 261, 1055–1058.
- Zaremba, J.D., Diamantopoulou, A., Danielson, N.B., Grosmark, A.D., Kaifosh, P.W., Bowler, J.C., Liao, Z., Sparks, F.T., Gogos, J.A., and Losonczy, A. (2017). Impaired hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion. *Nat. Neurosci.* 20, 1612–1623.
- Zhao, X., Hsu, C.-L., and Spruston, N. (2022). Rapid synaptic plasticity contributes to a learned conjunctive code of position and choice-related information in the hippocampus. *Neuron* 110, 96.e4, 108.e4.
- Zhao, X., Wang, Y., Spruston, N., and Magee, J.C. (2020). Membrane potential dynamics underlying context-dependent sensory responses in the hippocampus. *Nat. Neurosci.* 23, 881–891.
- Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gamal, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* 16, 264–266.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
rAAV1.Syn.GCaMP6f.WPRE.SV40	This paper	Addgene 100837-AAV1
Experimental models: Organisms/strains		
Mouse: C57BL/6J	This paper	Jackson Laboratory Strain 000664
Software and algorithms		
Suite2p	Pachitariu et al., 2017	https://github.com/cortex-lab/Suite2P
OASIS	Friedrich et al., 2017	https://github.com/j-friedrich/OASIS
Custom python code	This paper	Zenodo: https://doi.org/10.5281/zenodo.6348097
Python 3.7	https://www.python.org	N/A
SIMA	Kaifosh et al., 2014	https://github.com/losonczylab/sima

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to the lead contact Attila Losonczy (al2856@columbia.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Data generated in this study are available from the [lead contact](#) upon reasonable request.
- All original code has been deposited at Zenodo and is publicly available as of the date of publication. DOIs are listed in the [key resources table](#).
- Any additional information is available from the [lead contact](#) upon reasonable request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All experiments were conducted in accordance with the NIH guidelines and with the approval of the Columbia University Institutional Animal Care and Use Committee. Experiments were performed with adult (8-16 weeks) male C57Bl/6 mice (Jackson Laboratory).

METHOD DETAILS

Behavior and imaging

Viruses

Pyramidal cell imaging experiments were performed by injecting a recombinant adeno-associated virus (rAAV) encoding *GCaMP6f* (rAAV1-Syn-GCaMP6f-WPRE-SV40, Addgene/Penn Vector Core) into male wild-type mice.

Surgical procedure

Methods for viral delivery and surgical implant of imaging window and headposts were largely identical to previous work ([Lovett-Baron et al., 2014](#)). Briefly, mice were anesthetized under isoflurane and the virus was injected in dorsal CA1 (-2 mm AP; -1.5 ML, -1.2 DV relative to bregma; 500 nL) using a Nanoject syringe. Mice recovered in their home cage for 3 days following viral infusions. We then aspirated the cortex overlying the left dorsal hippocampus and implanted a 3 mm glass-bottomed stainless steel cannula for imaging access, and cemented a titanium headpost to the skull for head-fixation. For all surgeries, monitoring and analgesia (buprenorphine or meloxicam as needed) was continued for 3 days postoperatively.

Behavioral apparatus and virtual reality system

Mice were head-fixed above a low-friction, lightweight running wheel ([Warren et al., 2021](#)). The axle of the running wheel was coupled to a rotary encoder, which connected to a circuit that decoded and buffered the quadrature data from the encoder and transmitted these position updates to an Intel NUC i5 mini-PC that was used to control the behavior system. As described previously ([Kaufman](#)

et al., 2020), the experiments were managed through custom software on the PC that would send and receive instructions from a custom GPIO circuit on the behavior apparatus, which managed the water delivery system, lick port sensor and synchronization with the imaging system, as well as from the position tracking circuit and from the virtual reality system. All system elements were connected using high-speed Ethernet and communicated via UDP message passing.

The running wheel was surrounded by 5 LCD computer monitors (Acer SB230 23" IPS screens) arranged in a half-octagon, covering approximately 220° of the visual field of the mouse. Each monitor was connected to an individual ODROID-C2 single board computer running the Android operating system (version 6.0.1) which rendered a fraction of the VR scene on each display. Each ODROID received continuous instructions over Ethernet from the behavior control computer to update the VR environment as the position of the animal advanced. Virtual reality scenes were designed using the Unity game engine.

Behavior training

Starting 7 days after implant surgery, mice were habituated to handling and head fixation as previously described (Lovett-Barron et al., 2014). After two days of acclimation and free running on the wheel, we began exposing the animal to the 3 m virtual environment. The environment used during this training period would later become the "Familiar" context for the main experiment. When the mice reached the end of the virtual track, the screens were momentarily blanked for a 2 sec inter-trial interval, after which they were instantly teleported back to the beginning of the track.

At this point, mice were water deprived to 85-90% of their starting weight. Over a period of 1-2 weeks, we trained mice to run forward through the virtual environment and lick for small volume sucrose solution rewards (5% sucrose, $\sim 4\mu\text{m}$ per reward). Rewards were initially dispersed randomly throughout the environment to encourage running, and we slowly reduced the reward count to two fixed reward zones (located at ~ 1.2 and 2.7 m on the track) over the training period. We used two reward zones as this tended to lead to faster training and more reliable running behavior during the recordings (i.e. animals would rarely pause/slow down outside of the reward zones). Mice reached training criteria when they could consistently run > 130 trials using the two fixed reward zones in under an hour. Mice were given additional water as needed daily to maintain weights.

Familiar-novel context switching paradigm

For context-switch experiments, mice ran through alternating blocks of trials in the Familiar (training) context and a Novel context that was previously unseen prior to the start of the imaging experiments. Each recording session was organized into 4 blocks: 40 trials in Familiar, 30 trials in Novel, 30 trials in Familiar, and 30 trials in Novel. We additionally repeated this experiment with the same contexts across two additional days, and so a complete experiment sequence represented 3 days of recording. In all trials and contexts, the distance to the sucrose rewards remained fixed at 1.2 and 2.7 m. For each mouse, we also repeated the entire experiment sequence with a second Novel context (the familiar context remained the same for both sequences).

2-photon microscopy

Mice were habituated to the imaging apparatus (e.g. microscope/objective, laser, sounds of resonant scanner and shutters) during the training period. All imaging was conducted using a 2-photon 8 kHz resonant scanner (Bruker) and 16x NIR water immersion objective (Nikon, 0.8 NA, 3 mm working distance). For excitation, we used a 920 nm laser (50-100 mW at objective back aperture, Coherent). Green (GCaMP6f) fluorescence was collected through an emission cube filter set (HQ525/70 m-2p) to a GaAsP photomultiplier tube detector (Hamamatsu, 7422P-40). A custom dual stage preamp (1.4×10^5 dB, Bruker) was used to amplify signals prior to digitization. All experiments were performed at 1.2-2x digital zoom, acquired as 512×512 pixels images at 10 Hz.

QUANTIFICATION AND STATISTICAL ANALYSIS

Image preprocessing

Imaging data was organized using the SIMA software package (Kaifosh et al., 2014). Data was motion corrected in Suite2p (Pachitariu et al., 2017) using the non-rigid registration mode. ROIs were also detected using Suite2p (using "sparse mode"), followed by Suite2p's standard fluorescence extraction and neuropil correction. Identified ROIs were curated post-hoc using Suite2p's graphical interface to exclude non-somatic components.

Neural data analysis

Event detection

All fluorescence traces were deconvolved to detect putative spike events, using the OASIS implementation of the fast non-negative deconvolution algorithm (Friedrich et al., 2017). As in Ahmed et al. (2020), we discarded any events whose amplitude was below 4 median absolute deviations of the raw trace, and binarized the resulting signal for all subsequent analysis to indicate whether the neuron was active in a given frame. We note that we do not claim to uncover true underlying spike times, but rather use deconvolution as a tool for denoising and reducing the impact of the calcium autocorrelation on our analysis.

Calculating spatial tuning curves

Spatial tuning curves were calculating for each neuron on each trial. The virtual track was discretized into 100 evenly spaced bins (3 cm), which were used to compute a histogram of neural events. After normalizing for animal occupancy, the histogram was convolved with a Gaussian kernel ($\sigma = 9$ cm) to obtain a smooth activity rate estimate.

Spatial decoding

In order to assess position coding and context discrimination, we trained a decoder to predict the spatial location of the animal from neuronal population activity, using an ensemble of support vector machines (SVM, “one-vs-one” multi-class decoding [Ahmed et al., 2020]). For each day in Figure 1E, we trained the decoder on each context block separately and then tested the decoder on all remaining context blocks to generate the error matrices. Performance was quantified as the absolute error in predicted position. In Figure 1G, we tested decoding accuracy trial-by-trial for the first 10 trials in each context block, training the decoder on the remaining trials in the block. For each context switch on each day, we compared trial-by-trial performance between the Familiar and Novel context (Figures 1H and S1D).

Population cross-correlation analysis

We used a cross-correlation approach to identify spatial drift in neural tuning across trials at the level of the neural population. For each pair of trials a and b in a context block, we can compute their population vector correlation by concatenating the spatial tuning curves of all neurons in the population. Similarly, we obtained the spatial cross-correlation between the two trials across a range of spatial lags by first shifting each neuron’s spatial tuning in trial b and then recomputing the population vector correlation with a . Here a peak in the spatial cross-correlation at a positive lag indicates that the activity of a is ahead of b in space (since b must be shifted forward to maximize the correlation), and conversely, a peak at a negative lag indicates that activity in a is generally at locations behind b . We summarized the directionality and distance of the shift between trial pairs by the center-of-mass (COM) of the spatial cross-correlation curve.

The pattern of spatial drift across trials can be visualized by plotting the spatial shifts for all pairs of trials as a matrix. The matrix is necessarily antisymmetric (since the cross-correlation between a and b is the mirror image of that between b and a). Different neural dynamics predict qualitatively different shift matrices: a population of continually drifting place cells would exhibit a diagonally-banded shift matrix (since shift is then a monotonic function of the time between trials), while a more transient population shift (for example, during a period with many place fields forming via BTSP) would appear as vertical/horizontal bands at the affected trials (see Figure S3 for simulations).

The shifting tendency of each individual trial relative to the whole trial block can be determined by averaging the rows (or columns) of the symmetrized shift matrix (i.e. by multiplying the matrix lower triangle by -1). Symmetrizing is necessary so that all shift comparisons for a given trial consistently indicate the direction of movement. For example, if the neural population is consistently shifting backward, trials before trial n will be ahead and trials after will be behind. Averaging the n th row of the symmetrized matrix is equivalent to averaging all trial comparisons in the matrix upper triangle that include trial n . By our convention, a positive average shift for trial n indicates it is part of a generally backward trend, while a negative average shift indicates a forward trend, relative to the surrounding trials in the block. Occasionally activity on some trials correlated poorly with the rest of the block (possibly due to lapses in attention or behavior), and so we excluded any trial pairs where the peak of the spatial cross-correlation was less than 0.1. This method is used to calculate the average shifts in Figures 2C and S2, where we additionally omitted any trials in an experiment where $> 1/3$ rd of trial pair comparisons were missing due to this exclusion criteria.

Simulations of place field accumulation

We validated the population cross-correlation on simulated datasets of accumulating place fields, where a sub-fraction of fields exhibited BTSP-like characteristics (first-trial shift and gain), or that exhibited linear drifting over trials. Each simulated population comprised 300 neurons, where each neuron had a 0.2 probability of acquiring a place field. Fields accumulated over 30 trials according to a geometric process with mean 6, which is approximately the observed average first trial for place fields during the first exposure to the novel context during Day 1 of the experiment. Place field centers were sampled uniformly over the environment, and activity was modeled on each trial as a Gaussian bump ($\sigma = 5$ bins) at the sampled location.

For simulations of BTSP, we fixed a probability of BTSP for place fields in each simulation. If a field was drawn as BTSP, we shifted its first trial activity by a random distance sampled uniformly between 0 and 0.25 of the environment length and scaled its first trial activity by a random gain sampled uniformly between 1 and 3. We then considered the analysis results for different BTSP probabilities.

For simulation of continuously drifting place fields, we fixed a probability of drifting for place fields in each simulation. If a field was drawn as drifting, we sampled a slope (distance-by-trial) for the drift uniformly from 0 to 3 cm/trial, and shifted each trial’s activity according to that drift function. We then considered the analysis results for different drifting probabilities.

For all neurons, we additionally added out-of-field noise to individual trials with a probability of 0.2. This was simulated as additional Gaussian bumps added to that trial’s activity, centered at a random location in the environment. We additionally applied random scaling noise to each trial, with each trial’s gain drawn $\sim \mathcal{N}(1, 0.05)$.

Place field detection

We detected place fields separately in each context trial block (4 per session). Our detection scheme works by finding locations in the virtual environment where a neuron was more active than expected by chance, given its average firing rate and the animal’s spatial occupancy. We constructed a null distribution of spatial tuning curves for every neuron, by circularly shifting each trial’s activity independently by a random distance, and recomputing the smoothed, trial-averaged tuning curve as described previously. This procedure was repeated 1000 times, and we determined the 95th percentile of null tuning values at every spatial bin (i.e. the threshold for $p < 0.05$ spatial tuning). Segments of space where the true spatial tuning curve exceeded this null threshold were marked as candidate place fields, and the place field width was calculated as the distance between where the true tuning curve first exceeded and then

again fell below the threshold curve. To restrict our analysis to neurons with unambiguous firing fields, we additionally required that place fields have a width < 1 m (i.e. $(1/3)$ of the track), and that the neuron is active within the bounds of the place field on at least 15 trials.

Ensemble overlap

For each day, we considered a binary vector for each context block with elements $1 \dots N$ where N is the number of neurons recorded that day. The i th element was equal to 1 if that neuron had a significant place field at any location during that context block, and 0 otherwise. We calculated the ensemble overlap between each pair of context blocks on each day as the Jaccard similarity between these vectors. Since the raw overlap can be influenced by differences in the number of place cells identified in each block, we z-scored the overlap according to a shuffle distribution whose samples were drawn by randomly permuting each context block's place cell vector independently. This procedure was repeated 10000 times to produce the distributions in [Figures S1E](#) and [S1F](#).

Spatial tuning score

We computed a pseudo-probability mass function $p(x_i)$ for each spatial tuning curve, where x_1, x_2, \dots, x_n are the discrete position bins, by normalizing the tuning curve so that it sums to 1 over positions. We then compute the un-normalized tuning score \hat{s} as the KL-divergence between p and the uniform distribution u :

$$\hat{s} = D_{KL}(p||u) = \log_2 n + \sum_{i=1}^n p(x_i) \log_2 p(x_i)$$

Intuitively, uniform activity over space will give the minimum score $s = 0$, while having all activity concentrated in a single position bin will yield the maximum score of $s = \log_2(n)$. Inhomogenous spatial tuning can also arise simply from very sparse, noisy activity, a regime in which using raw information-theoretic metrics can give misleading results ([Souza et al., 2018](#)). To control for the effects of firing rate and sampling, we used a normalized tuning score s by standardizing \hat{s} relative to a null distribution, formed by calculating the tuning score on all null tuning curves (generated as described in the place field detection procedure):

$$s = \frac{\hat{s} - \mu(\hat{S}_{\text{null}})}{\sigma(\hat{S}_{\text{null}})}$$

Detecting the trial of place field formation

For our analysis of place field characteristics, it was necessary to determine the first trial that a neuron fired within its place field(s) in a given trial block. Following [Sheffield et al. \(2017\)](#), we identified for each place field the first window of 5 trials where the neuron was active within the place field boundaries on at least 3 of those trials, and called the first active trial within that window the place field formation trial. A trial was considered active if there was activity within the place field boundary on that trial with amplitude of at least 5% of the neuron's peak activity level across all trials. Our results however did not strongly depend on the particular choice of thresholds.

Formation trial activity displacement

We estimated the spatial shift between activity on the trial of place field formation and activity on the remaining trials within the place field. This was computed as the difference between the location of peak firing on the first trial, and the location of peak firing in the average tuning of all subsequent trials (as in the population shift analysis, a positive shift indicates the formation trial's activity was at positions ahead of the later trial's activity).

Formation trial gain

Evidence for elevated activity or burst-firing during the formation trial was assessed by computing the formation trial gain, defined as the peak activity rate on the formation trial within the place field divided by the peak activity rate in the averaging tuning of all subsequent trials that contained activity within the place field (i.e. excluding any silent trials, so that the gain does not simply reflect unreliable place fields). A gain > 1 indicates elevated firing during the place field formation trial, relative to later traversals through the place field.

Perifield formation velocity

For each place field, we computed the average velocity of the animal within the place field boundaries during the place field formation trial.

Velocity-field width correlation

For each trial block on each day, we fit a linear model to predict the width of place fields from the perifield formation velocity. We compared the slope of the linear fits between Familiar and Novel contexts by calculating a Δ slope (Novel – Familiar) for each context switch. We calculated the probability that this Δ would be observed under random permutations of the context labels for place fields, recomputing Δ slope for 10000 shuffles in each condition. A p-val was obtained from the cumulative density of the Gaussian fit to the resulting null distribution. We reported the significance relative to the null distribution as the negative \log_{10} (p-val). In [Figures 3F](#) and [3G](#), we pooled place fields across all experiments from a given day and trial block.

Network model of BTSP-mediated place field formation

Our characterizations of BTSP at the level of the population and single neurons are largely premised on the shift in activity from the first trial on which a place field appears relative to the remaining trials, which occurs due to the temporal asymmetry of the BTSP plasticity rule ([Bittner et al., 2017](#); [Zhao et al., 2020, 2022](#); [Rolotti et al., 2022](#)). We used a simple model of BTSP-mediated place field

formation to explore how detectable this shift could be under different recording conditions, since our recordings are relatively slow (10 Hz) and indirect (somatic calcium vs spiking).

We simulated a CA1 neuron as a linear-nonlinear-Poisson cascade model (LNP) (Gerstner et al., 2014). The input to the neuron was the firing rate vector $\mathbf{r}(x)_{\text{CA3}}$ of 300 CA3 place cells, with elements:

$$r_i(x(t))_{\text{CA3}} = A * \exp \left[-\frac{1}{2} \left(\frac{x(t) - c_i}{\sigma} \right)^2 \right]$$

where the maximum firing rate $A = 30$ Hz, Gaussian place field width $\sigma = 10$ cm, $x(t) \in [0, 3]$ m is the position of the animal at time t , and c_i is the place field center for the i th input neuron, which were distributed evenly across the environment. The total input current to the CA1 neuron at time t was given by the scalar product:

$$h(t) = \mathbf{r}(x(t))_{\text{CA3}} \cdot \mathbf{w}$$

where \mathbf{w} is the vector of synaptic weights. This was used to calculate the instantaneous firing rate of the CA1 neuron:

$$r(t)_{\text{CA1}} = \exp[k * h(t)]$$

where k is a scale factor that enforced a maximum instantaneous firing rate of 30 Hz, computed post hoc from the maximum current throughout the simulation. For simplicity, we assumed a constant velocity within each simulation. For each velocity we considered, we simulated the instantaneous CA1 firing rate across 16 trials (1 preplasticity, 15 post-plasticity) in time steps of 1 msec (incorporating a 2 sec inter-trial interval between trials as in the experiment, during which the CA3 inputs were silent). We then used this firing rate trace to simulate 200 unique ground-truth spike trains assuming Poisson spiking.

Spike trains were converted to fluorescence by convolution with an idealized calcium kernel ($\tau_D = 0.7$ sec, $\tau_R = 0.07$ sec) (Pachitariu et al., 2018). Traces were then scaled to unit variance, sub-sampled to the same sampling rate as our experimental data (10 Hz), and corrupted with different levels of additive Gaussian noise $\mathcal{N}(0, \sigma^2)$. We considered the signal-to-noise ratio as $\text{SNR} = \sigma^{-1}$, and calculated noisy fluorescence traces at $\text{SNR} \in \{0.5, 1, 2, 4, 8, 16\}$ for each spike train. Finally, we passed these traces through the spike deconvolution workflow used in the main paper. We analyzed and compared the simulation results on both the noisy calcium and resulting inferred spikes.

For modeling BTSP, we fixed a plateau onset location at $\hat{x} = 200$ cm for all simulations. For all post-plasticity trials, we set the synaptic weights from CA3 place cells according to an asymmetric, double-exponential plasticity kernel as parameterized in Bittner et al. (2017). For inputs active at time t :

$$w(\hat{t} - t) \propto \Theta(\hat{t} - t) * \exp[(t - \hat{t}) * 0.69] + \Theta(t - \hat{t}) * \exp[(\hat{t} - t) * 1.31]$$

where Θ is the Heaviside step function and \hat{t} is the plateau onset time. Since velocity is constant, we approximate the post-plasticity weight vector by evaluating this equation at each CA3 place field center c_i as $w(\hat{x} - c_i / v)$ where v is the simulation velocity (Figure S4A). The weight vector was z-scored to equalize the weight distribution across different simulation parameters. These updated weights were then used to drive CA1 output on all post-plasticity trials (1-15). On trial 0 prior to learning, CA3 weights were homogeneously depressed. Spiking activity on trial 0 was instead driven by a 0.5 sec step function when the animal reached the plateau onset location, to simulate plateau-mediated burst spiking.

In Figure S4, we surveyed the measured shift in activity from the plateau trial to the average of subsequent trials, calculated as the center of mass of the neural tuning curve on each trial. This was measured for the noisy fluorescence trace and inferred spike for each choice of SNR, for each of $n = 200$ unique spike trains, for each choice of velocity $v \in \{10, 15, 20 \dots 60\}$. We then compared these shifts to the shift measured in the ground-truth spike train.

Non-negative matrix factorization

We aimed to identify subsets of place fields in the dataset that exhibited different kinds of trial-to-trial dynamics (e.g. BTSP-like shift in the formation trial activity, drifting responses, or amplitude modulation over trials). To this end, we first aligned all place field tuning profiles by extracting a 150 cm (50 bins) window of activity from each trial around the place field center, and concatenating these windows for the first 10 trials from the formation trial of the place field (e.g. Figure 4A). This yielded a 500 element vector \mathbf{x} for every place field, which we stacked into a matrix \mathbf{X} so that each row was the spatiotemporal profile (trial \times position) of a single place field.

Our goal then was to cluster the rows of \mathbf{X} in order to identify groups of place fields with similar trial-to-trial dynamics, but clustering samples directly in high-dimensional spaces is generally a poor strategy. Instead, we first reduced the dimensionality of \mathbf{X} using non-negative matrix factorization:

$$\mathbf{W}, \mathbf{H} = \underset{\mathbf{W}, \mathbf{H}}{\operatorname{argmin}} \|\mathbf{X} - \mathbf{W}\mathbf{H}\|_{\text{Fro}}^2$$

Where \mathbf{W} and \mathbf{H} are rank n matrices, for some $n \ll 500$. \mathbf{H} is an $n \times 500$ matrix, where each of the n rows is a spatiotemporal pattern. NMF models each row (place field) in \mathbf{X} as a weighted sum of these n patterns, given by \mathbf{W} . As with any dimensionality reduction method, since n is much less than the number of place fields in the dataset, the model is forced to identify spatiotemporal patterns in \mathbf{H} that are shared between many place fields. The strict non-negativity of \mathbf{H} and \mathbf{W} in NMF however can often give particularly

interpretable decompositions, due to its parts-based reconstruction of \mathbf{X} (i.e., since all elements in the weighted sum for each place field are non-negative, the different patterns in \mathbf{H} cannot “cancel out”). Each row of \mathbf{X} was mean-normalized prior to decomposition, so the optimization was not dominated by inhomogeneous activity scales across neurons or experiments.

We inspected NMF decompositions of the data for a range of n , and found that $n = 11$ corresponded to a prominent “elbow” between two linear regimes in the loss function (i.e. the rate of improvement slows when adding any additional components > 11 , [Figure 4C](#)). Additionally, $n = 11$ consistently gave the most interpretable spatiotemporal patterns in \mathbf{H} ([Figures 4D and S5](#)). In particular, this model produced components that separately represented the place field on each of the 10 trials, plus an additional forward-shifted component on the first trial, reminiscent of plateau-driven place field formation. The weights of the two first-trial components among place fields were strongly anti-correlated ([Figure S5](#)).

Clustering place fields in NMF space

We used K-means clustering with $K = 2$ to partition place fields into “BTSP-like” and “other” groups, using the 11 NMF patterns as the feature space for clustering. We found that 2 clusters were sufficient to reliably segregate place fields with BTSP-like characteristics. In particular, we found that increasing the number of clusters had negligible effects on the BTSP-like cluster; additional clusters formed largely through additional subdivisions of the “other” group ([Figure S5](#)).

BTSP fraction

We computed the fraction of place fields in each experimental condition (day, trial block) that were assigned to the BTSP-like group. In [Figure 5E](#), we computed this per experimental session. In [Figures 5F and 5G](#), this was computed by mouse (pooling each mouse’s data from the two repetitions of the experiment).

Decoding

We validated the NMF-clustering analysis by attempting to decode the BTSP/Other labels of place fields based on secondary characteristics of their activity profiles. We computed 4 features for every place field in the dataset: the shift in its tuning from the first trial to remaining trials, the width of the place field, the stability of its place field (correlation between even and odd trials), and its first trial activity gain. We then trained a support vector machine with a linear kernel to classify place fields as BTSP or Other in this 4-dimensional feature space ([Figure S5](#)). We reported an averaged cross-validated decoding accuracy for each mouse, by randomly partitioning the data into 10 50/50% training/test splits, stratified by BTSP label. Samples were weighted during training to be inversely proportional to label frequency to account for the greater number of Other cells. We additionally compared the cross-validated results to a null distribution constructed by rerunning the cross-validated decoding analysis on copies of the dataset where the BTSP classifications of place fields were randomly permuted. This procedure was repeated 1000 times, and significance thresholds were computed from a 95% interval on the resulting distribution of null accuracies.

Modeling the effects of velocity on place field density

We fit a linear model to predict place field density at each position as a function of the animal’s spatial velocity profile. The predictors in the model were exponentially filtered versions of the spatial velocity profile ($\tau \in \{0.01, 0.025, 0.05, 0.1, 0.25\}$ spatial bins), to account for any short-term history effects, e.g. those induced by the calcium autocorrelation. The velocity filters plotted in [Figure 6C](#) are obtained by a weighted sum of the exponential kernels, with the weights given by the regression coefficients.

Next-trial field distance

To test for history dependence in the locations of accumulated place fields, we computed for every place field n its average distance to place fields that formed on trial $t + 1$, where t is the formation trial for field n . These distances were computed only between fields of the same BTSP classification. We then standardized this distance according to a null distribution, obtained by randomly permuting the trial of place field formation between all fields of a given BTSP classification within each trial block. In this way, we disrupt the correlations between adjacent trials while maintaining the marginal distributions of place fields over positions and trials. This procedure was repeated 200 times for each place field. The resulting σ measures the next-trial field distance relative to the distance expected simply from random sampling according to the marginal distributions.