


Crucial role for CA2 inputs in the sequential organization of CA1 time cells supporting memory

Christopher J. MacDonald^{a,b,1} and Susumu Tonegawa^{a,b,c,1} 

^aRIKEN–MIT Laboratory for Neural Circuit Genetics, The Picower Institute for Learning and Memory, Department of Biology, Massachusetts Institute of Technology, Cambridge, MA 02139; ^bDepartment of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139; and ^cHHMI, Massachusetts Institute of Technology, Cambridge, MA 02139

Contributed by Susumu Tonegawa, December 7, 2020 (sent for review October 2, 2020; reviewed by Thomas J. McHugh and Richard G. M. Morris)

There is considerable evidence for hippocampal time cells that briefly activate in succession to represent the temporal structure of memories. Previous studies have shown that time cells can be disrupted while leaving place cells intact, indicating that spatial and temporal information can be coded in parallel. However, the circuits in which spatial and temporal information are coded have not been clearly identified. Here we investigated temporal and spatial coding by dorsal hippocampal CA1 (dCA1) neurons in mice trained on a classic spatial working-memory task. On each trial, the mice approached the same choice point on a maze but were trained to alternate between traversing one of two distinct spatial routes (spatial coding phase). In between trials, there was a 10-s mnemonic delay during which the mouse continuously ran in a fixed location (temporal coding phase). Using cell-type-specific optogenetic methods, we found that inhibiting dorsal CA2 (dCA2) inputs into dCA1 degraded time cell coding during the mnemonic delay and impaired the mouse's subsequent memory-guided choice. Conversely, inhibiting dCA2 inputs during the spatial coding phase had a negligible effect on place cell activity in dCA1 and no effect on behavior. Collectively, our work demonstrates that spatial and temporal coding in dCA1 is largely segregated with respect to the dCA2–dCA1 circuit and suggests that CA2 plays a critical role in representing the flow of time in memory within the hippocampal network.

hippocampus | time cells | place cells | memory | CA2

A large body of work has shown that the hippocampus (HPC) is crucial for remembering events in the context of where and when they occurred (1, 2). Since the discovery of place cells in the HPC (3, 4), there is compelling evidence supporting their role in coding cognitive maps of space that emerge and stabilize while animals explore new environments (reviewed in refs. 5 and 6). Individual place cells exhibit spatial coding; they selectively fire when an animal is in a particular location and sets of place cells can collectively tile environments to represent a distinct spatial context (7, 8). More recently, a striking example of temporal coding in the HPC has been identified in the form of time cells. When animals repeatedly experience distinct event sequences that have a consistent temporal structure, many cells in the HPC reliably fire for brief periods at specific times during the sequence independently of space (9–11). In this way, sets of such time cells activate in succession and collectively tile intervals to represent a distinct temporal context for a specific experience. Both place and time cells share several commonalities for how repeated experiences are coded during spatially or temporally structured memory tasks that require an intact HPC for successful performance (10, 12, 13). Given the prevalence of spatial and temporal coding in these tasks, place and time cells are often studied as an underlying mechanism for hippocampal-dependent memories.

A critical open issue is to what degree temporal information is segregated from spatial information within the HPC. There is little known about the specific hippocampal subcircuits that support temporal coding and hippocampal-dependent memory, as well as their relationship to the circuits that generate spatial coding. Specific to this issue, some studies have shown that time

cells in the dorsal hippocampal subregion CA1 (dCA1) can be disrupted under conditions where dCA1 place cells remain intact (14, 15), while another study reported the opposite effect (16). Though these studies differed in several aspects, all of them employed methods to disrupt brain activity that leave open the question of whether specific inputs into dCA1 may preferentially carry temporal information. In order to investigate this issue, we focused our attention on the dCA2 subregion as a source of temporal information into the dCA1. Previous physiological work has identified a prominent time-varying signal in the dCA2 network (17–21) and additional studies have demonstrated that dCA2 inputs are the most potent excitatory drivers of dCA1 pyramidal cells in comparison to alternative excitatory inputs into dCA1 (22, 23). Moreover, temporal processing deficits have been reported in mice with a targeted disruption of the vasopressin 1b receptor (Avpr1b) gene (24), which is predominantly expressed in the CA2 subregion (25).

We adapted a classic hippocampal-dependent spatial working-memory task (delayed spatial alternation or DSA task) for mice that required them to alternate between taking a distinct left and right spatial route (spatial coding phase) on successive trials (26, 27). In between trials, mice were required to continuously run on a treadmill for an extended temporal delay (9, 28). Therefore, to make a correct subsequent choice, the mice had to maintain trial-specific information across a mnemonic delay while their spatial location was fixed (temporal coding phase). Previous work has shown that the inclusion of a mnemonic delay that

Significance

The hippocampus (HPC) is critical for remembering “where” and “when” past events occurred, as though our memories are bound to a neural representation defined by space and time. In support of this idea, the HPC contains place and time cells whose firing patterns exhibit spatial and temporal coding, respectively. There is evidence that space and time are coded in parallel, but the underlying hippocampal subcircuits have not been determined. We combined cell-type-specific inhibition with *in vivo* extracellular recordings to demonstrate that spatial and temporal coding is segregated in the final hippocampal output (region CA1) with respect to a major input (region CA2). These data suggest a selective role for CA2 temporally organizing CA1 neural activity for a hippocampal-dependent memory.

Author contributions: C.J.M. and S.T. designed research; C.J.M. performed research; S.T. contributed new reagents/analytic tools; C.J.M. and S.T. analyzed data; and C.J.M. and S.T. wrote the paper.

Reviewers: T.J.M., RIKEN Center for Brain Science; and R.G.M.M., University of Edinburgh.

The authors declare no competing interest.

Published under the PNAS license.

¹To whom correspondence may be addressed. Email: tonegawa@mit.edu or cjm@mit.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2020698118/-DCSupplemental>.

Published January 11, 2021.

segments the trials renders the task hippocampal dependent (9, 29). The explicit separation of trials into a spatial and temporal coding phase enabled us to use cell-type-specific optogenetic methods to selectively inhibit dCA2 inputs into the dCA1 during these time-restricted trial periods. Using this approach, we found that inhibition of dCA2 inputs into dCA1 during the temporal coding phase impaired the mouse's subsequent choice and degraded the precision and stability of dCA1 time cells. In contrast, we found that dCA2 inhibition while the mouse traversed a spatial route prior to the mnemonic delay had a comparatively smaller effect on spatial coding, nor did it affect the mouse's behavioral choice. Taken together, our results indicate that dCA2 is a crucial source of temporal information into dCA1 and supports the sequential organization of time cells while a memory is maintained across a mnemonic delay.

Results

Inhibition of dCA2 Inputs into dCA1 during the Temporal Coding Phase Impairs the Mouse's Subsequent Memory-Guided Choice. For one group of CA2-specific Cre knockin mice [Mapk315-Cre, hereafter referred to as CA2-Cre mice (23)], we used a Cre-dependent virus (30) to express the outward proton pump enhanced ArchT3.0 (hereafter ArchT) in Cre⁺ cells within the hippocampal dCA2 subregion (Fig. 1A). These mice (CA2-ArchT; $n = 6$) were food-restricted and trained on a modified version of a DSA task using a rectangular figure-eight maze (Fig. 1B) (9, 28). On each trial, the mouse progressed toward a choice point at the end of the maze's central alleyway. There the mouse committed to traversing the left or right return arm that brought it back to the beginning of the central alleyway. The mice were only rewarded if on each trial they switched between taking left and right spatial routes upon reaching the choice point. Before being allowed to access the choice point, the mice were required to run on a treadmill in the middle of the central alleyway for 10 s during which their spatial location was largely fixed. Therefore, the task requires the mouse to process spatial and reward information while traversing one of the return arms (spatial coding phase) and maintain trial-specific information across a mnemonic delay while their spatial location was fixed (temporal coding phase) to correctly make a choice.

We first tested whether dCA2 inputs into dCA1 were crucial for performance by inhibiting dCA2 during the temporal coding phase on a pseudorandom subset of trials in two testing sessions (*Materials and Methods*). On trials in which dCA2 input to dCA1 was inhibited (light-on trials), CA2-ArchT mice subsequently made a higher percentage of choice errors compared to trials in which the light was not turned on (light-off trials) (Fig. 1C, *Left*), as confirmed by a two-way repeated-measures ANOVA (rmANOVA) that indicated a main effect of the factor Light [$F(1, 20) = 21.67, P = 0.002$]. There was no effect of the factor Session [$F(1, 20) = 0.44, P = 0.52$] and no Light \times Session interaction [$F(1, 20) = 2.51, P = 0.13$]. A separate control cohort of CA2-Cre mice (CA2-EYFP; $n = 4$) that bilaterally expressed EYFP alone in Cre⁺ dCA2 cells (*Materials and Methods*) showed no difference in the percentage of choice errors between light-off and light-on trials (Fig. 1C, *Right*), as supported by the rmANOVA. There was no main effect of the factor Light [$F(1, 12) = 0.27, P = 0.61$], Session [$F(1, 12) = 0.0001, P = 0.99$], and no interaction [$F(1, 12) = 1.77, P = 0.22$]. We also considered the possibility that the impairment was secondary to an acute anxiety-like effect caused by inhibition of dCA2. After testing on the DSA task was complete, the CA2-ArchT mice were put through an open-field test as a coarse assessment of this possibility (31). The mice were placed in the arena for a total duration of 6 min and dCA2 inputs into dCA1 were inhibited during the final 3 min. There was no difference in the total distance traveled during the light-off and light-on periods (paired t test; $t_5 = 0.32, P = 0.76$) (Fig. 1D, *Left*) or time spent

within the center of the arena (paired t test; $t_5 = 0.68, P = 0.53$) (Fig. 1D, *Right*). Therefore, light-mediated inhibition of dCA2 into dCA1 does not alter locomotor activity or a willingness to explore a new arena.

Does inhibiting dCA2 inputs into dCA1 increase the percentage of choice errors if the extended mnemonic delay is removed? In order to answer this question, a new cohort of CA2-ArchT mice ($n = 5$) was tested on a nondelayed alternation (NDA) task (26). In this version, mice freely moved through the central alleyway and were not required to run on the treadmill to access the choice point. In this task, light-on trials consisted of turning on the green light when the mouse entered the central alleyway and then turning the light off when it reached the end of the alleyway. In this protocol, behavior appeared stereotyped because the mice quickly moved down the central alleyway. The average duration of inhibition in these experiments was 5.06 ± 0.85 s (mean \pm SD). Under these conditions, there was no difference between the percentage of error trials on light-off and light-on trials (Fig. 1E, *Left*); there was no main effect of the factors Light [$F(1, 16) = 0.26, P = 0.62$] or Session [$F(1, 16) = 0.10, P = 0.74$] and no Light \times Session interaction [$F(1, 16) = 0.79, P = 0.39$]. Finally, we wondered if inhibition of dCA2 inputs into dCA1 during the spatial coding phase would impair the mouse's choice after the mnemonic delay. In a new cohort of CA2-ArchT mice ($n = 5$) trained on the DSA task, inhibition took place on a pseudorandom subset of trials while they traversed one side of the maze. The average duration of inhibition during these trials was 25.47 ± 11.78 s (mean \pm SD). We found no evidence for a choice impairment (Fig. 1E, *Right*) because the rmANOVA indicated no effect of Light [$F(1, 16) = 0.49, P = 0.49$] or Session [$F(1, 16) = 0.097, P = 0.76$], and no interaction between the two factors [$F(1, 16) = 0.66, P = 0.43$].

The results from these behavioral experiments suggest that, in the context of this task, the dCA2 projection into dCA1 is not required for processing spatial information per se. This claim is supported by finding that inhibition of dCA2 inputs into dCA1 during the spatial coding phase had no effect on the mouse's subsequent choice. Therefore, spatial information can be processed by alternative circuits during this period to support normal performance. Additional support for the claim is provided by the experiment in which the treadmill running period was removed. In the NDA task, inhibition of dCA2 inputs into dCA1 while mice traversed the central alleyway had no effect on their subsequent choice. This result suggests that dCA2 projections into dCA1 are dispensable for the decision-making process that underlies response selection. Taken together, these data suggest that the dCA2→dCA1 circuit is required when there is an extended temporal interval in between trials and contributes to maintaining trial-specific information across the mnemonic delay (see also ref. 32).

Inhibition of dCA2 Decreases the Stability and Precision of Temporal Coding in dCA1. How does the delay-dependent behavioral impairment map onto dCA1 activity while the animal is performing the task? In order to address this question, we tested how inhibition of dCA2 inputs into dCA1 during the temporal coding phase alters dCA1 firing patterns. We bilaterally expressed ArchT in the dCA2 of CA2-Cre mice ($n = 6$) and then chronically implanted a multitetrode recording drive targeting dCA1 in one hemisphere in addition to optic fibers bilaterally placed above dCA1 (*Materials and Methods*). This strategy enabled for us to bilaterally inhibit dCA2 inputs while monitoring extracellular activity from large ensembles of dCA1 cells in mice performing the DSA task. As before, experiments consisted of light-off trials in addition to light-on trials that were interspersed throughout the recording session. A total of 629 putative pyramidal cells were isolated across 15 recording sessions using waveform and spiking characteristics (33) (*Materials and Methods*

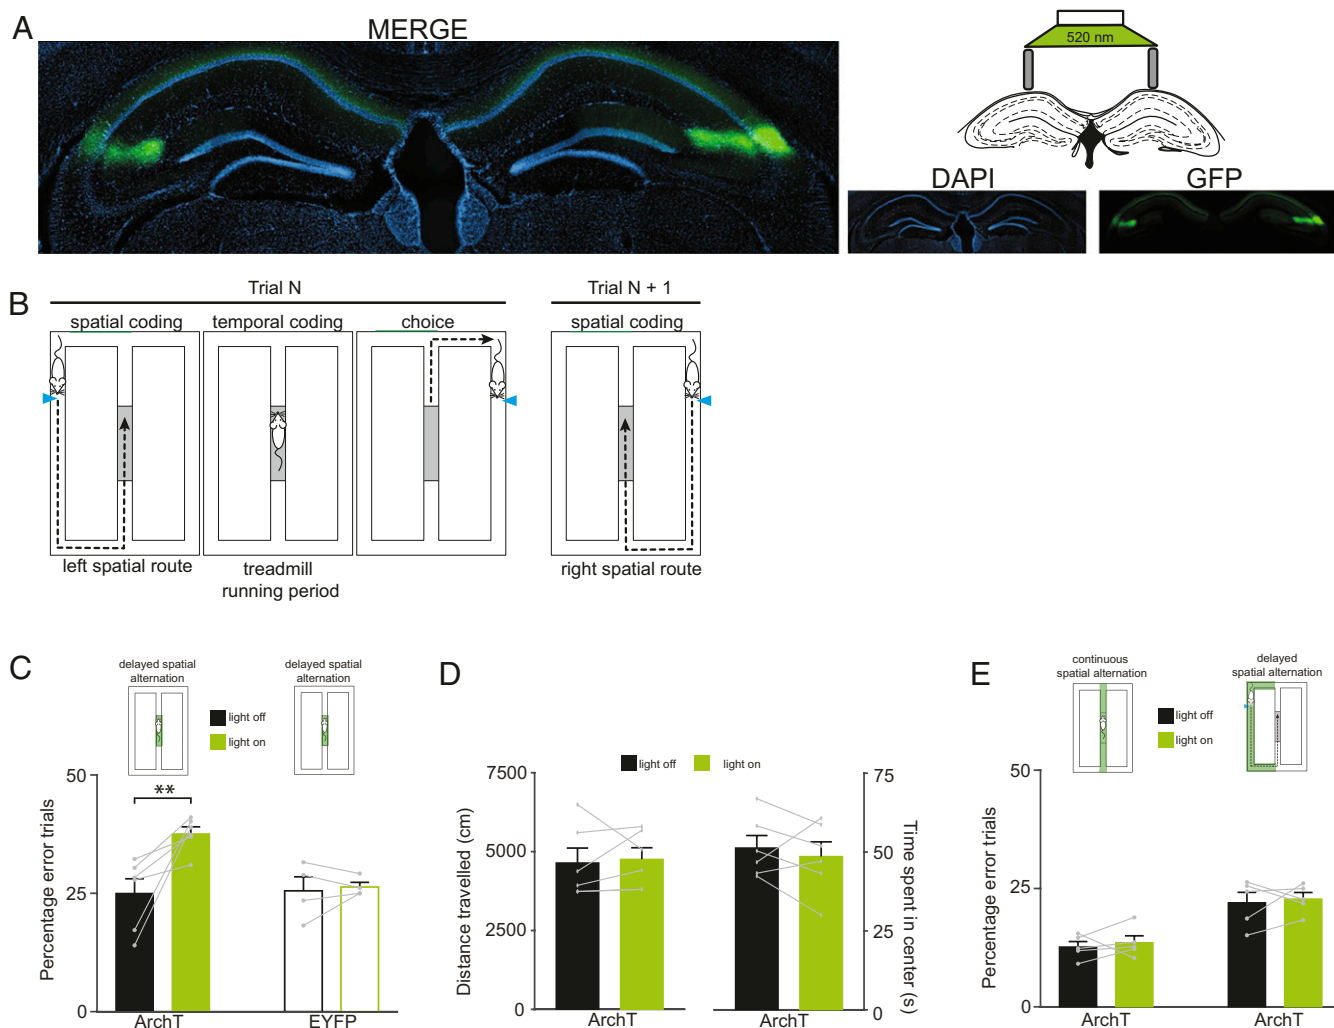


Fig. 1. dCA2 inputs into dCA1 are involved with memory maintenance across an extended temporal delay. (A) CA2-specific (CA2-Cre) Cre knockin mice were injected with a Cre-dependent virus to express the ArchT-EYFP in Cre⁺ cells comprising dCA2. (Magnification: 10 \times .) (B) The DSA task: Mice are trained to alternate between making left- and right-side choices on successive trials. A single trial begins with a spatial coding phase during which the mouse traverses one side of the maze and progresses up the central alleyway to approach a treadmill that is built into the maze. The temporal coding phase begins with the treadmill turning on and the mouse running for 10 s, during which the mouse's position is largely fixed. After the temporal coding phase, the mouse leaves the treadmill area, approaches the choice point, and commits to traversing the left or right side of the maze. Once the mouse commits to choosing a side, the spatial coding phase for the subsequent trial begins. (C) Bilateral inhibition of dCA2 inputs into dCA1 during the temporal coding phase (light on, green) increases the percentage of error trials compared to trials in which dCA2 inputs were not inhibited (light off, black). CA2-ArchT group corresponds to CA2-Cre mice expressing ArchT in Cre⁺ cells in dCA2. The EYFP group corresponds to CA2-Cre mice expressing EYFP alone in Cre⁺ dCA2 cells. Two-way rmANOVA with factors Session and Light. ** $P < 0.01$, main effect or factor Light. (D) Bilateral inhibition of dCA2 inputs into dCA1 in an open arena does not change the amount of locomotion (Left) or the time spent in the center of the arena (Right). (E, Left) Bilateral inhibition of dCA2 inputs into dCA1 does not increase the percentage of error trials on the NDA task with no treadmill running period. Here the dCA2 inhibition period occurred while the mouse traversed the central stem. (Right) Bilateral inhibition of dCA2 inputs into dCA1 during the spatial coding phase does not increase the percentage of error trials. All bar plots represent the mean \pm SEM across 2 testing days and the individual data points for this average from light-off and light-on trials are shown for each mouse. The gray connecting lines are included to visualize the paired comparisons.

and Table 1, "Treadmill ArchT" condition). Trial-averaged firing rates were estimated during the temporal coding phase for each cell in the light-off and light-on conditions. For each cell, we used permutation testing to assess if the estimated firing rate in each time bin was significantly greater than the average firing rate across the entire temporal coding phase (see *Materials and Methods* for details). Using this approach, in many cells we identified distinct firing fields, which were periods during which firing rates were significantly elevated above the cell's overall activity and flanked on at least one side by periods with low activity levels. A cell was considered to exhibit temporal coding if at least one firing field was identified with this permutation test and the peak firing rate was greater than 0.75 Hz. We found that

nearly half of the neural population ($n = \sim 50\%$; 313 of 629) showed such temporal coding. As these cells reliably activated at the same specific moment during the temporal coding phase of each trial, we refer to them as time cells. As a population, time cells activated in succession to collectively span the entire mnemonic delay (Fig. 2A).

Next, we compared time cell activity from light-off trials to activity from light-on trials. Using an independent sample permutation test (*Materials and Methods*), for each time cell we tested if there was a difference in its firing rate between light-off and light-on trials in at least one time bin during the temporal coding phase (*Materials and Methods*). We found that almost half of the time cell population (48%; 149 of 313) exhibited different

Table 1. Summary of the number of recording sessions and total cells recorded for each mouse used in the three extracellular recording experiments

Condition	Mouse	No. of recording sessions	Total no. of cells recorded
Treadmill ArchT	CA2-ArchT-01	3	123
	CA2-ArchT-02	3	112
	CA2-ArchT-03	2	113
	CA2-ArchT-04	2	68
	CA2-ArchT-05	2	84
	CA2-ArchT-06	3	129
	Total cells		629
Treadmill EYFP	CA2-EYFP-01	1	51
	CA2-EYFP-02	2	54
	CA2-EYFP-03	2	79
	Total cells		184
Maze ArchT	CA2-ArchT-07	1	30
	CA2-ArchT-08	1	30
	CA2-ArchT-09	3	53
	CA2-ArchT-10	1	22
	CA2-ArchT-11	1	32
	Total cells		167

firing patterns between these trial types (Fig. 2*B*). The average firing rate of the time cell population was largely the same across the temporal coding phase during both light-off and light-on trials (Fig. 2*C*). However, at the level of individual cells we observed a broad range of differing activity patterns in response to dCA2 inhibition. We observed a robust decrease in activity during light-on trials in many time cells (18%; $n = 57$), as would

be expected by removing a major source of excitatory drive (Fig. 2*D, Upper*). In many other cases (30%; $n = 97$ overall), time cells showed increased levels of activity during light-on trials compared to light-off trials (Fig. 2*D, Lower*). However, firing rates were often elevated because the time cell's tuning curve broadened, resulting in less precise temporal coding. In fact, the proportion of time cells showing elevated firing rates during

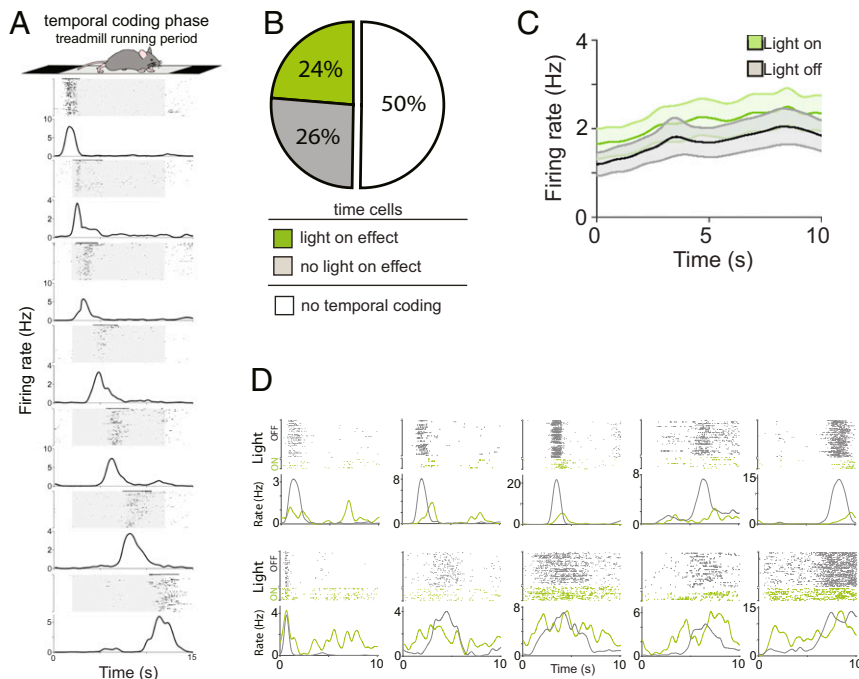


Fig. 2. Inhibition of dCA2 altered the firing rates of a large proportion of dCA1 cells that exhibited temporal coding. (A) Peristimulus time histograms (PSTH) and raster plots depicting spiking activity from seven time cells that were recorded at the same time. The shaded area shown on the raster plot corresponds to the temporal coding phase during which the mouse was running on the treadmill for 10 s. Each cell selectively fires at a particular time during the temporal coding phase but differs according to the time coded such that they collectively activate in sequence. The black bars running along the time axis above the raster plots demarcate the cell's firing field. (B) Proportion of total recorded CA1 cells that were classified as time cells and showed altered temporal firing patterns (green: light-on effect) or no change (gray: no light-on effect) during trials in which dCA2 was inhibited. The remaining cells showed no evidence of temporal coding. (C) Average firing rate across all time cells during light-off and light-on trials. Shaded areas correspond to 99% confidence intervals. (D) PSTH and raster plots depicting activity from 10 cells during the temporal coding phase and subdivided into light-off (gray) and light-on (green) trials. The top row depicts cases in which CA2 inhibition decreased firing rates while the bottom row depicts cases in which CA2 inhibition impaired the precision of temporal coding through the elevation of firing rates outside the temporal firing field.

light-on trials was greater than the proportion of cells showing a decrease ($\chi^2 = 13.83$, $P = 0.0002$).

We identified firing fields during light-off trials to assess the effect of inhibition on in-field and out-of-field firing rates across the time cell population (*Materials and Methods*). In addition, we quantified each cell's information content and sparsity, both of which reflect the precision of a cell's temporal coding, and firing stability, which reflects the consistency of a cell's response from trial-to-trial in relation to its trial-averaged firing pattern. While dCA2 inhibition did not change in-field firing rates in a coherent direction across the population ($P = 0.57$; Wilcoxon signed-rank test) (Fig. 3A), there was a global increase in out-of-field firing rates detected ($P = 1.37 \times 10^{-31}$; Wilcoxon signed-rank test) (Fig. 3B). Inhibition also decreased information content ($P = 0.007$; Wilcoxon signed-rank test) (Fig. 3C) and increased sparsity in the time cell population ($P = 0.003$; Wilcoxon signed-rank test) (Fig. 3D), which indicates more diffuse firing in time (i.e., less temporal selectivity). Finally, dCA2 inhibition decreased firing stability across the cell population ($P = 0.001$; Wilcoxon signed-rank test) (Fig. 3E).

In a separate cohort of CA2-Cre mice ($n = 3$), we expressed EYFP in Cre⁺ CA2 cells, and implanted a recording drive along with optic fibers above the left and right dCA1. These CA2-EYFP mice were trained on the DSA task, then tested in five recording sessions during which we intermingled light-on trials, as in the CA2-ArchT group. We isolated 184 putative pyramidal cells (Table 1, "Treadmill-EYFP" condition), of which 78 exhibited temporal coding and were classified as time cells. The proportion of time cells in CA2-EYFP mice (43%) was not different from the proportion found in CA2-ArchT mice (50%; $\chi^2 = 2.39$, $P = 0.12$). A comparison of light-off and light-on activity confirmed that there was a minimal effect of light on time cell-firing patterns; the proportion of time cells affected in CA2-EYFP animals (6%; 6 of 78) was significantly less than the proportion in CA2-ArchT animals (48%; $\chi^2 = 39.92$, $P = 2.65 \times 10^{-10}$). We also quantified and compared in- and out-of-field firing rates, information, sparsity, and stability measures between

light-off and light-on trials in these mice (Fig. 3A–E) and found no differences (all $P > 0.23$; Wilcoxon signed-rank test).

Altogether, the results from these analyses reveal that dCA1 time cell firing patterns become less temporally precise and less reliable while dCA2 inputs are inhibited. Interestingly, during light-on trials we also observed drastic changes in the firing rates of some cells that were not classified as putative pyramidal cells but whose spiking characteristics were consistent with interneurons (33) (*Materials and Methods*). The examples shown in Fig. 3F fired at high, largely fixed rates during the temporal coding phase and showed sustained decreases in activity during dCA2 inhibition. Despite these observations, we found no reliable difference when we compared θ rhythm power during light-off and light-on trials across all tetrodes ($P = 0.115$; Wilcoxon signed-rank test).

Loss of dCA2 Input during the Temporal Coding Phase Decorrelates the Neural Ensemble Representation of Time in dCA1.

The individual cell correlates of dCA2 inhibition manifested at the population level as a progressive degradation of the time cell sequence that was normally observed during light-off trials (Fig. 4A). Fig. 4B illustrates the periods during which dCA2 inhibition significantly decreased (Fig. 4B, Left) or increased activity (Fig. 4B, Right) across the time cell population. These plots reinforce the fact that in cells for which dCA2 inhibition decreased firing rates, the effect was generally confined to inside its firing field. In contrast, when dCA2 inhibition increased cell firing rates, the effect largely occurred outside the firing field and often for extended periods. To further characterize the effect of dCA2 inhibition at the neural ensemble level, for each recording session we constructed a population vector (PV) indexing each neuron's trial averaged firing rate in successive 250-ms time bins from light-off and light-on conditions (Fig. 4C). In each corresponding time bin, we correlated the PVs from light-off trials with the PVs constructed from the same set of cells during the light-on trials to obtain a set of PV correlations (*Materials and Methods*). Next, we took the median PV correlation from the first half (0 to 5 s; time

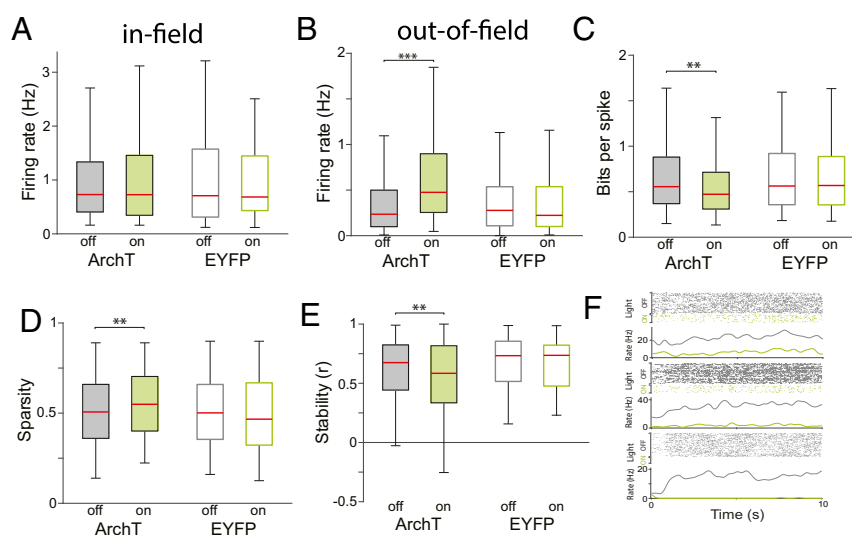


Fig. 3. Inhibition of dCA2 decreased the precision and stability of temporal coding in dCA1. (A) Average in-field firing rate (Hz) for time cells during light-off ("off") and light-on ("on") trials in CA2-ArchT (Left) and CA2-EYFP (Right) mice. (B) Average out-of-field firing rate (Hz) for time cells during light-off ("off") and light-on ("on") trials for CA2-ArchT (Left) and CA2-EYFP (Right) mice. (C) Average information (bits per spike) for time cells during light-off ("off") and light-on ("on") trials in CA2-ArchT (Left) and CA2-EYFP (Right) mice. (D) Average sparsity for time cells during light-off ("off") and light-on ("on") trials in CA2-ArchT (Left) and CA2-EYFP (Right) mice. (E) Average stability (measured as a correlation value, r) for time cells during light-off ("off") and light-on ("on") trials in CA2-ArchT (Left) and CA2-EYFP (Right) mice. (F) PSTH and raster plots depicting activity from three putative interneurons during the temporal coding phase of light-off and light-on trials. For A, the box-plots indicate the median (red line) as well as the 25th and 75th percentiles. The whiskers extend to 1.5 \times interquartile range. All comparisons were Wilcoxon signed-rank tests where $**P < 0.01$ and $***P < 0.001$.

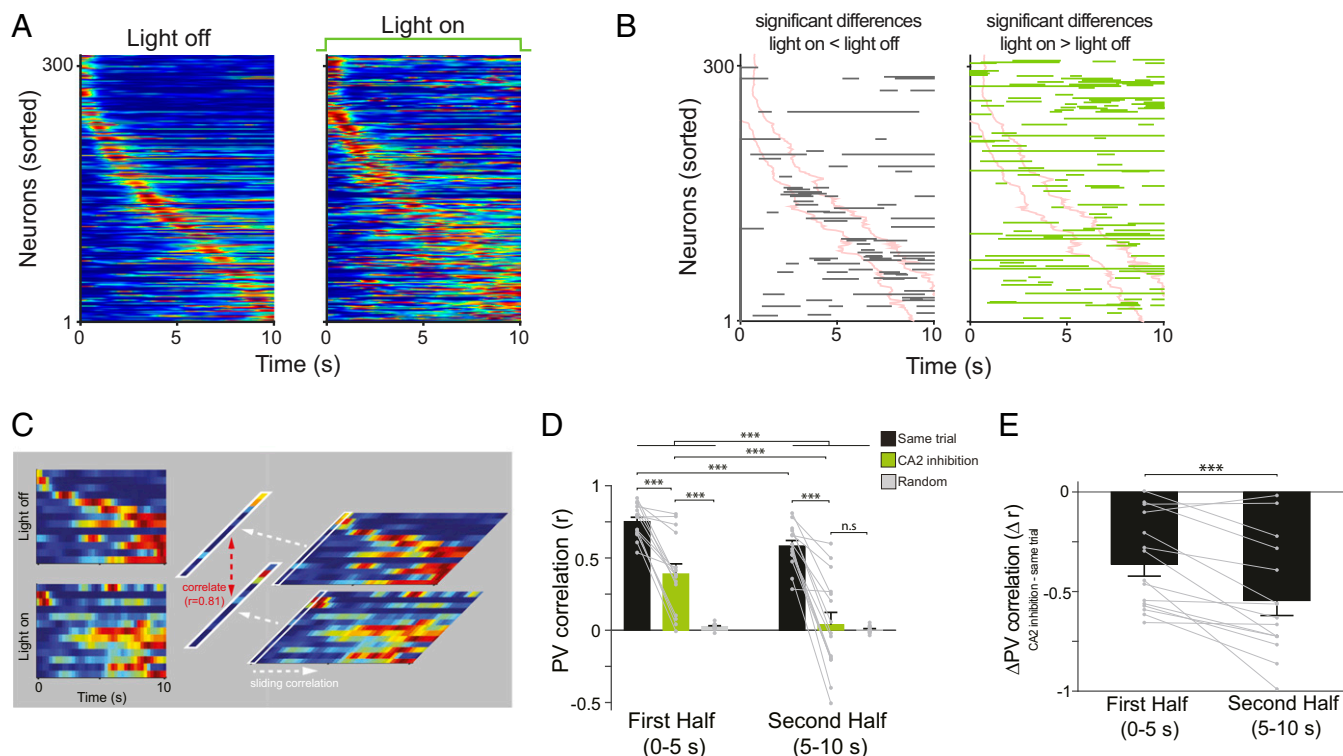


Fig. 4. Loss of dCA2 input during the temporal coding phase decorrelates the neural ensemble representation of time in dCA1. (A) Each panel shows normalized (range: 0 to 1) trial-averaged firing rates of time cells during the temporal coding phase of light-off (Left) and light-on (Right) trials. Cells are sorted by increasing latency to its maximum firing rate in light-off trials. (B) Each panel shows the periods during the temporal coding phase in which a cell's firing rate was significantly decreased in light-on trials (Left, gray segments) or significantly increased in light-on trials (Right, green segments). Cells are sorted in these panels as in A. The red bands running down the diagonal of each panel traces out the half-maximum width from the cell's peak firing rate in light-off trials (compare with Left in A). (C) Depiction of how the PV correlation is computed. The panels on the left-hand side show trial averaged activity from an ensemble of 13 time cells recorded during a single experiment and sorted as in A. The x axis corresponds to time and is segmented into 250-ms bins. Each row on the top panel shows one cell's average activity during light-off trials while each row on the bottom panel shows the same corresponding cell's average activity during light-on trials. In this example, the PV from the first time bin in the light-off condition is correlated with the first time bin in the light-on condition to yield a PV correlation of 0.81. This is done for each time bin to obtain a set of 40 PV correlation values. The median PV correlation is taken from the first half (time bins 1 to 20) and second half (time bins 21 to 40) to obtain the first and second PV correlation for this recording session. This process is carried out for each recording session. (D) Depiction of the average PV correlation (mean \pm SEM) during the first and second half of the temporal coding phase for the CA2 inhibition, same trial, and random comparison (see main text). The PV correlation values computed for each recording session are overlaid on the bar plots. To assist visualization, we used gray lines to connect data points from the same recording session with respect to same trial and CA2 inhibition PV correlations. (E) Illustration of the difference between the PV correlation for CA2 inhibition and the same trial comparison in the first and second half (mean \pm SEM). The difference is more negative in the second half indicating that CA2 inhibition significantly decorrelated firing patterns more in the second half compared to the first half. As before, gray lines connect data points obtained from the same recording session. Post hoc comparisons are shown and Bonferroni-corrected where $***P < 0.001$. See main text for details.

bins 1 to 20) and second half (5 to 10 s; time bins 21 to 40) of this set. Thus, for each experiment we obtained two PV correlation values that measured the similarity between light-off and light-on time cell firing patterns with respect to the first or second half of the temporal coding phase (hereafter referred to as the "CA2 inhibition" comparison).

In order to properly interpret this measure, we evaluated PV correlations for two additional control comparisons. For one control comparison, we correlated PVs between even and odd trials from the temporal coding phase of light-off trials. This correlation provides an upper bound for the PV correlation because it measures the self-consistency of the ensemble pattern across trials of the same trial-type ("same trial" comparison). For the second control comparison, we took a resampling approach to estimate a PV correlation between time cell ensembles that randomly fired during the temporal coding phase of each trial (*Materials and Methods*). For each recording session, we randomly split all trials from the light-off condition into two trial sets and randomly shifted (circularly) each trial's temporal firing pattern from each cell before computing a trial average. Then we

obtained a PV correlation from the first and second half of the temporal coding phase, just as we did for the CA2 inhibition comparison. This process was repeated 999 more times and the mean PV correlation of this resampled distribution for the first and second half was taken as the PV correlation for the "random" comparison in the recording session. This comparison yields a lower bound for the PV correlation, as it measures the similarity between PVs that were randomly rearranged in time.

The results from these analyses are shown in Fig. 4D. A two-way rmANOVA comparing PV correlations across Halves ("first half" and "second half") and Comparison (same trial, CA2 inhibition, and random) revealed a main effect of Halves [$F(1, 14) = 96.52$, $P = 1.16 \times 10^{-7}$], Comparison [$F(1, 14) = 62.79$, $P = 1.15 \times 10^{-7}$], and an interaction between these two factors [$F(2, 28) = 34.36$, $P = 9.79 \times 10^{-7}$]. The significant interaction term indicated that the difference between the PV correlation in each half depended on the Comparison and prompted us to make several specific statistical comparisons using different contrasts in the rmANOVA model. First, the PV correlation for CA2 inhibition was less than the same trial correlation in the first

($P = 9.49 \times 10^{-5}$; Bonferroni-corrected) and second half ($P = 1.22 \times 10^{-5}$; Bonferroni-corrected) of the temporal coding phase. Furthermore, the PV correlation decreased from the first to second half of the temporal coding phase in the same trial ($P = 9.76 \times 10^{-7}$; Bonferroni-corrected) and CA2 inhibition comparison ($P = 3.49 \times 10^{-5}$; Bonferroni-corrected). However, the degree in which the PV correlation decreased across Halves depended on the CA2 inhibition or same trial comparison [two-way rmANOVA that excludes the random comparison: Halves \times Comparison interaction: $F(1, 14) = 15.37$, $P = 0.002$]. Specifically, the PV correlation showed a larger decrease across Halves for CA2 inhibition compared to the same trial comparison (Fig. 4E) ($P = 0.0009$; Bonferroni-corrected). The PV correlation for the same trial comparison was greater than the random comparison in each half (both $P < 1.16 \times 10^{-9}$; Bonferroni-corrected). However, while the PV correlation for the CA2 inhibition and “error” comparison was different in the first half of the temporal coding phase ($P = 0.004$; Bonferroni-corrected), they were not different in the second half ($P = 1$; Bonferroni-corrected). Finally, we conducted this analysis on the CA2-EYFP group and confirmed a decrease in the PV correlation between Halves [$F(1, 4) = 15.34$, $P = 0.017$] but there was no difference between Comparisons [$F(1, 4) = 0.75$, $P = 0.435$] and no interaction between these factors [$F(1, 4) = 1.58$, $P = 0.28$].

These analyses reveal that time cell firing patterns are more variable in the second half of the temporal coding phase under normal conditions (light-off trials), which is consistent with a decrease in the precision of temporal coding with elapsed time (10, 34). Despite this greater variability, temporal coding is preserved in the second half and considerably greater than what one would expect if cells fired randomly over time. Conversely, dCA2 inhibition significantly decorrelated dCA1 temporal firing patterns from the patterns observed under normal conditions. In addition, this effect was more pronounced in the second half and during this period the PV correlation was comparable to the value one would obtain using temporally random firing patterns. Combined with our single-cell analyses, these data suggest that the loss of dCA2 input degrades temporal coding and destabilizes firing patterns in dCA1.

Inhibition of dCA2 Marginally Degrades Spatial Coding in dCA1.

Given that dCA2 inhibition during the spatial coding phase does not impair the mouse’s behavior, we tested if dCA1 place cell activity remained intact during inhibition. To test this, we recorded cells from dCA1 in five mice (seven recording sessions) while inhibiting dCA2 inputs during the spatial coding phase on a pseudorandom subset of trials in which the mouse traversed the right side of the maze. Firing rates were estimated as a function of the mouse’s linearized position during the spatial coding phase of light-off and light-on trials. We identified cells exhibiting spatial coding by using the same analysis that was used to confirm temporal coding (*Materials and Methods* and Table 1, “Maze ArchT” condition). Using this approach, we found over 33% of the cells from which we recorded (33%; $n = 55$ of 167) showed robust spatial coding during right-side trials. The spatial firing patterns in 15% ($n = 8$ of 55) of these place cells significantly differed between light-off and light-on trials (Fig. 5A), but this proportion was much lower than the proportion of time cells that was affected by dCA2 inhibition (48%; $\chi^2 = 19.89$, $P = 8.19 \times 10^{-6}$). There were no differences between light-off and light-on trials regarding population averaged measures of in- and out-of-field firing rates, information, sparsity, and firing stability (all $P > 0.12$; Wilcoxon signed-rank test).

Next, we quantified the effect of dCA2 inhibition on spatial coding using the PV correlation analysis described earlier. In this case, we separated the spatial coding phase into a first and second half with respect to the mouse’s linearized position on the

right-side of the maze (Fig. 5B). A rmANOVA comparing PV correlations across Halves and Comparison (Fig. 5C) identified a main effect of Comparison [$F(2, 12) = 364.97$, $P = 2.62 \times 10^{-7}$] but there was no effect of Halves [$F(1, 14) = 0.87$; $P = 0.39$] and no interaction between Comparison and Halves [$F(1, 14) = 0.35$; $P = 0.68$]. Expectedly, the PV correlation for the CA2 inhibition comparison was different from the “error” comparison in the first ($P = 6.84 \times 10^{-5}$; Bonferroni-corrected) and second ($P = 3.59 \times 10^{-7}$; Bonferroni-corrected) half, and this was also the case for the same trial Comparison (first half: $P = 0.002$; Bonferroni-corrected and second half: $P = 2.91 \times 10^{-7}$; Bonferroni-corrected). To confirm if the main effect of Comparison was largely driven by the inclusion of the “error” comparison in our model, we directly compared the PV correlations from the same trial and CA2 inhibition comparison. This analysis confirmed a modest but reliable difference between the same trial and CA2 inhibition PV correlations [main effect of Comparison: $F(1, 6) = 19.29$, $P = 0.005$], though there was no evidence for an effect of Halves [$F(1, 6) = 0.60$, $P = 0.48$] or a Halves \times Comparison interaction [$F(1, 6) = 0.71$, $P = 0.43$].

Finally, we directly compared the PV correlations for the CA2 inhibition comparison from our time cell and place cell experiments (Fig. 5D). We used a mixed ANOVA to compare the PV correlation for the CA2 inhibition between Condition (between-subject factor: “place cell” or “time cell”) and Halves (within-subject factor: “first half” and “second half”), and found a main effect of Condition [$F(1, 20) = 15.91$, $P = 0.0007$], Halves [$F(1, 20) = 17.02$, $p = 0.0005$], and a Halves \times Condition interaction [$F(1, 20) = 24.02$, $p = 8.63 \times 10^{-5}$]. Altogether, these data indicate that dCA2 inhibition had a substantially larger effect on time cells compared to place cells and place cell disruption did not worsen while dCA2 was inhibited, unlike in time cells.

Discussion

There is a great deal of interest in defining a functional role of the hippocampal CA2-CA1 circuit with regard to learning and memory (35, 36) and a growing literature indicates a major contribution of CA2 to processing social memories (37, 38). However, this function depends on information flow from dCA2 to vCA1 rather than dCA2’s interactions with dCA1 (38, 39). Our present work combined cell-type-specific optogenetics with large-scale extracellular recordings of hippocampal dCA1 neurons as mice underwent the hippocampal-dependent DSA task. The study resulted in a discovery of a functional role of dCA2 projections into dCA1 for the organization of time in a hippocampal-dependent memory.

The dCA2→dCA1 Projection Is Crucial for DSA Memory. We found that selective and reversible inhibition of dCA2 inputs into dCA1 during the treadmill running period of the DSA task caused mice to make more errors in the subsequent choice of two alternative outer arms within the trial (Fig. 1C). When the delay period was reduced or the treadmill removed (NDA task), the inhibition of dCA2 inputs during the remaining delay period or the central track running period was ineffective in impairing the animal’s performance (Fig. 1E, Left). These results are consistent with prior work demonstrating a requirement for the HPC in DSA tasks but not in NDA tasks (29). But we substantially advanced our knowledge on the requirement of the HPC for the DSA task by identifying not only the specific dCA2→dCA1 circuit, but also the specific phase of the DSA task during which it is crucial. The latter finding was even more specified by demonstrating that inhibition of the dCA2→dCA1 circuit activity is largely dispensable during the pre-delay period of the outer arm running within the trial (Fig. 1E, Right). Taking these data together, we find that a crucial factor that determines a requirement for the dCA2→dCA1 circuit to support normal performance in the DSA

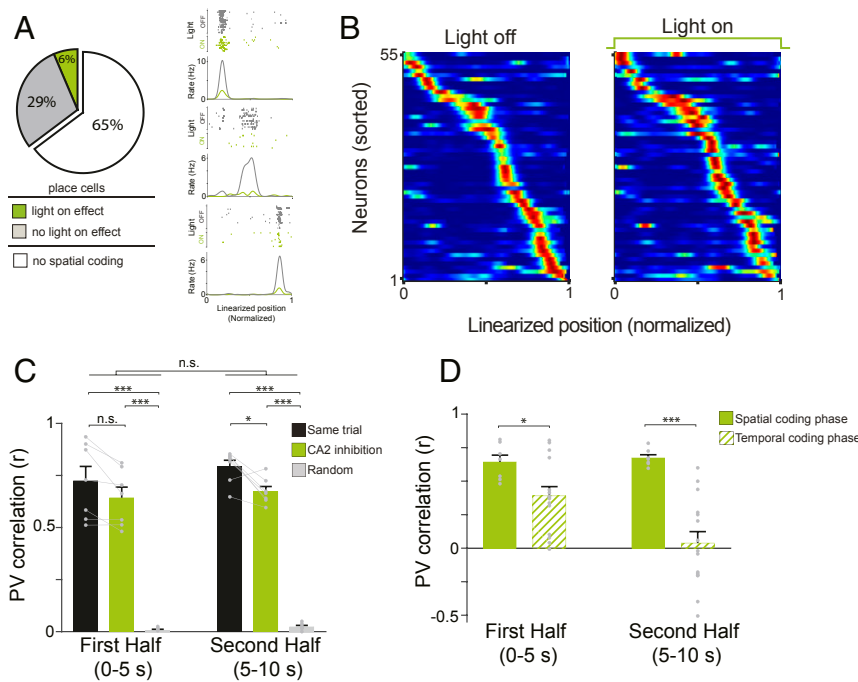


Fig. 5. Inhibition of dCA2 has a marginal effect on spatial coding compared to temporal coding in dCA1. (A, Left) The proportion of cells exhibiting spatial coding on the right-hand side of the maze as well as those whose firing patterns were significantly affected by inhibition of dCA2. (Right) Depiction of the raster plot and trial averaged firing rate of three place cells with respect to the mouse's linearized position on the right-hand side of the maze during light-off (black) and light-on (green) trials. These cells exhibited robust spatial coding and were inhibited during light-on trials. (B) Each panel shows normalized trial averaged firing rates of the cells identified as having spatial coding in A during light-off (Left) or light-on (Right) trials. The cells are sorted with respect to the linearized position at which the cell fired maximally during light-off trials. (C) The average PV correlation (mean \pm SEM) is shown for the cell population identified in A in the first and second half of the spatial coding phase for the same trial, CA2 inhibition, and random comparison (see main text). (D) The PV correlation (mean \pm SEM) for the "CA2 inhibition" comparison in the spatial coding phase (C) is directly compared to the values obtained from the temporal coding phase illustrated in Fig. 4D. All comparisons were Bonferroni corrected such that * = $P < 0.05$, *** = $P < 0.001$, and n.s. = $P > 0.05$. As in Fig. 4, gray lines connect data points obtained from the same recording session. See main text for details.

task is whether the mouse experiences a distinct mnemonic delay before being allowed to make a memory-guided choice. This finding suggests that the behavioral impairment is related to the mouse's failure to maintain trial-specific information across an extended time interval. A role of CA2 in temporal processing was previously studied employing constitutive *Avpr1b* knockout mice (24) in a battery of tasks to test for social and cognitive deficits. However, the lack of a high level of spatial and temporal restriction of the knockout prevented a precise interpretation of the observed behavioral impairments as the present study.

The Loss of dCA2 Input Disrupts Temporal Coding in dCA1 during the Mnemonic Delay. After clarifying the nature of the behavioral impairment, we determined how dCA2 inhibition altered temporal coding in the dCA1 network during the mnemonic delay. In our task, temporal coding was isolated from spatial coding during this period because the mouse's spatial location was fixed. Consistent with previous work, we identified a robust form of temporal coding in the dCA1, as shown by time cells that reliably activated for brief periods in succession to span the entire mnemonic delay (Fig. 24). Firing patterns were also more variable for cells that coded for times in the last half of this trial phase (Fig. 4D), which agrees with several previous time cell studies that report decreasing accuracy in firing times as elapsed time grows (10, 11, 28, 34, 40). Importantly, we found that inhibition of dCA2 inputs during this trial phase drastically altered time cell firing patterns in dCA1 (Figs. 2 B and D and 4A). The effect of inhibition was to reduce firing rates within the main firing field of some cells. However, in many other cells, activity was elevated outside of its main firing field (Fig. 4B). One

potential explanation for elevated firing rates is the presence of feed-forward inhibitory transmission between dCA2 and interneurons within the dCA1 (41). We observed several putative dCA1 interneurons to decrease their firing rates in response to optogenetic dCA2 inhibition (Fig. 3F), so it is possible that the increase in a time cell's out-of-field firing rate reflects a disinhibitory effect.

Overall, the reduction of dCA2 inputs decreased the precision and stability of dCA1 time cell firing (Fig. 3 A–E), and degraded the fidelity of the dCA1 time cell sequence that unfolded while animals were behaving normally (Fig. 4 A, D, and E). That these changes are associated with a behavioral deficit is analogous to previous studies that have identified associations between the precision and stability of dCA1 place fields and other measures of spatial memory (42–45). Our results provide strong evidence for a link between the temporal organization of neural activity during a mnemonic delay and the accuracy of the mouse's eventual memory-guided choice (11, 14, 15, 46, 47). They are also consistent with some aspects of hippocampal network models of spatial alternation behavior (48, 49) in which separate inputs for spatial and temporal information are combined within CA1, and the prefrontal cortex (PFC) interacts with hippocampal output to guide behavioral response selection (i.e., left or right choice) (see also refs. 50 and 51). Although we do not tease apart these CA1-PFC circuits in the present study, based on such previous work we also assume that there are at least two types of memory supporting performance on this DSA task: One that is a hippocampal-dependent memory and the other a PFC-dependent "rule memory" that is acquired gradually with repeated experience and guides response selection. In this respect, when output from the HPC is unstable and temporally imprecise, as is

the case when dCA2 inputs are inhibited, the PFC cannot properly guide response selection.

Spatial Coding in dCA1 Is less Reliant on dCA2 Input and Can Be Supported by Other Circuits. We also optogenetically inhibited dCA2 inputs and monitored spatial coding in dCA1. We found a comparatively smaller effect of dCA2 inhibition on the sequential activity of place cells while mice moved through the maze (Fig. 5) (see also ref. 52). Therefore, spatial coding is largely preserved during dCA2→dCA1 inhibition in contrast to a much stronger effect on temporal coding, and this finding mirrors our behavioral results. A similar dissociation between place and time cells in dCA1 was reported in one study (14) that found a disruption of dCA1 time cells, but not place cells, during muscimol-mediated inhibition of the medial septum, which is thought to affect cells throughout the HPC and surrounding cortices. In another study (15), optogenetic inhibition of the medial entorhinal cortex (MEC) had minimal effects on dCA1 place cell activity but had a lasting disruption on time cell firing patterns that persisted into light-off periods. In this case, the viral construct used to express the inhibitory opsin was not selective for excitatory or inhibitory neurons in the MEC, nor for cells in a particular layer of the MEC. Such a strategy may have led to an enduring, general disruption of cellular activity in the HPC and surrounding cortices that nevertheless had a selective impact on temporal coding. Conversely, our use of the CA2-Cre mouse line enabled for ArchT expression in dCA2 pyramidal cells as defined by CA2-specific markers (23). While our manipulation of dCA2 altered the firing stability and dynamics of dCA1 time cells, these changes were confined to the temporal coding phase of light-on trials and did not carry over into light-off trials (compare light-off and light-on trials from neurons depicted in Fig. 2D). Another relevant and recent study of DSA behavior in rats found a relatively low percentage of time cells during the mnemonic delay, and that excitotoxic lesions of the MEC had an effect on spatial coding instead of temporal coding despite a delay-dependent behavioral impairment (16). There was no treadmill involved with these experiments so the authors considered that the differences may have been because the rats were not constrained to continuously run in place during the temporal coding phase. However, temporal coding in dCA1 has been reported in rats without such requirements in another variant of a DSA task (53). Altogether, these studies emphasize the importance of future work taking into account cell-type specificity, chronic and acute circuit manipulations, as well the behavioral task used to assess memory.

In addition to identifying dissociations between place and time cells in functionally compromised circuits, there is some additional evidence from animal and human work in support of spatial and temporal information being coded in parallel within the HPC (17; reviewed in ref. 13). One recent study that provides an interesting complement to our own reported that spatial coding in dCA1 depends on direct input from the dorsal hippocampal CA3 subregion (dCA3) (54). On the other hand, an earlier study reported intact dCA1 place cells when intra-hippocampal inputs were permanently severed that spared spatial memory acquisition and impaired recall (55). Time cells have also been reported in the dCA3 (40) and some computational models CA1 time cells assume an important role for CA3 (14, 56). Therefore, it will be important to test if the HPC's intrinsic circuitry per se is important for temporal coding in dCA1, and if CA2 and CA3 make unique contributions to temporal and spatial coding to support learning and memory.

Relationship to Idiopathic Cue Processing and Path Integration. Unlike time cells, place cells that signaled locations in the second half of the spatial coding phase were not less accurate than those

signaling locations in the first half (compare “same comparison” in Figs. 4D and 5C) (40). Moreover, the marginal effect of optogenetic dCA2 inhibition on spatial coding was not stronger in the second half of this trial phase despite inhibition lasting well over twice as long as the temporal coding phase (Fig. 5C and D). This difference between time and place cells may reflect a sensitivity of time cells to accumulation of error. Place cells are thought to contribute to maintaining a representation of the organism's current location by integrating distance and direction traveled over time (i.e., path integration) (57–60). Time cell activity can be influenced by distance traveled (28, 61) and may reflect a temporal integration process that tracks the continuous change of elapsing time within an episode (10, 62–64). For integration of any variable over time to be useful to an organism, there must also be mechanisms in place to counter the rapid accumulation of error. There is evidence for such mechanisms operating in the spatial domain, such as when an animal encounters visual boundaries, sharp turns, and other landmarks (65–67). In the temporal domain, resets at the field potential level have been reported in a variety of nonspatial memory tasks too (68–70). Viewed from this perspective, the dCA2 may be particularly crucial for the temporal organization of cellular activity in dCA1 when surrounding cues are largely static and there are less salient environmental cues to provide error correction, as is the case during the temporal coding phase (see also ref. 14).

dCA2→dCA1 May Play a Crucial Role for the Temporal Aspect of Episodic Memory. The HPC is critical for the encoding and recalling of episodic memory in humans (1, 2, 71) and episodic-like memory in animals, which takes place in spatial and temporal contexts (72–74). Therefore, understanding how hippocampal circuits process space and time in service of memory is essential (13, 75). Indeed, the existence of spatial and temporal coding in various forms throughout the HPC and surrounding cortices (76, 77) indicates that space and time are prominent continuous dimensions of our experience. One emerging view is that spatial and temporal coding organizes memory by mapping our experience to a representation of context that is defined by both space and time (13). Our present study employed an HPC-dependent spatial working-memory task (DSA task) in which spatial and temporal coding are separated into distinct phases. This permitted us to identify a hippocampal subcircuit, dCA2→dCA1, that plays a crucial role in the temporal organization of cellular activity representing an HPC-dependent memory. We speculate that the dCA2→dCA1 circuit broadly supports the temporal organization of experience including episodic-like and even episodic memory in which spatial and temporal coding are often less separated. Indeed, the HPC and associated cortices may be comprised of multiple microcircuits that collectively function to organize our memories in relation to any dimension that frames our ongoing experience (75, 78–80).

Materials and Methods

Animals. All procedures relating to mouse care and treatment conformed to the MIT institutional and NIH guidelines. Our laboratory previously generated a CA2-specific Cre knockin mouse (*Map3k15-Cre*) (23) that enables for the selective expression of opsins in Cre+ CA2 cells by means of Cre-dependent viruses. A total of 34 male Cre+ CA2-Cre mice were used in these studies. All mice were housed individually after the viral injection surgery.

Treadmill-Based DSA Task. We adapted a previous task (28) for use with mice. A rectangular figure-eight maze (Fig. 1C) was constructed from black fiber glass. The maze was 100-cm long and 50-cm wide and bisected width-wise by a 100-cm-long central stem. A 35-cm segment in the middle of the central stem was replaced with a commercially available motorized treadmill (Columbus Instruments) that could be turned on and off at desired speeds through custom software programmed in Matlab. The width of the track used throughout the entire maze was 5 cm and 4-cm walls extended

upwards from each side of the track. Additional details about the training protocol and testing are found in *SI Appendix, Materials and Methods*.

Open-Field Test. As described previously (31), an automated video-tracking system (Ethovision by Noldus) was used to track the amount of time spent in the center of an open metal chamber (Accuscan system) compared to the edges, as well as the total distance traveled across a session (31).

Surgery, Recording and Tracking Procedures, Spike Sorting, and Optogenetics. The details about these procedures are in *SI Appendix, Materials and Methods*.

Histological Procedures and Immunohistochemistry. For the in vivo physiology experiments, after the completion of experiments the mice were deeply anesthetized, and small lesions were made near the tips of each tetrode by passing current (30 mA for 10 s) for subsequent electrode track position confirmation. All mice were deeply anesthetized and transcardially perfused with cold phosphate-buffered saline (PBS) followed by 4% paraformaldehyde in cold PBS. Their brains were extracted for histology using standard procedures. Additional details are in *SI Appendix, Materials and Methods*.

Statistical Analyses of Spiking Activity and Local Field Potential. All analyses were performed using Matlab. For each neuron, each spike was assigned a time in relation to the onset of the spatial or temporal coding phase in each trial as well as a linearized spatial coordinate in relation to mouse's position during left- or right-side trials. θ -Power on each tetrode during the temporal coding phase was estimated using multitaper spectral methods available as part of the Chronux toolbox (81) (chronux.org). The function `mtspectrumc.m` was used to compute the trial-averaged spectrum and θ -power was taken as the average value between 6 and 12 Hz. The time-bandwidth product was set to 3 and the number of tapers used was five. This computation was done separately for light-off and light-on trials and compared across all tetrodes.

To identify time cells, for each neuron we separated light-off from light-on trials, smoothed the trial averaged firing rates with a Gaussian kernel (SD = 200 ms), and these rates were down-sampled to 1,000 bins. For light-off and light-on trials, we used a one-sample paired permutation method to test if the trial-averaged firing rate in each time bin was different from the overall firing rate across the entire temporal coding phase. We accounted for multiple comparisons by adjusting the P values with respect to a family-wise error rate of 0.05 (82). The details of this analysis and Matlab code are found in Groppe et al. (83). For each cell, if at least one time bin significantly exceeded the overall firing rate in the entire temporal coding phase using $P < 0.05$ corrected for multiple comparisons, then it was considered a candidate "time cell." Next, we identified all time bins for which the firing rate exceeded the overall firing rate using a $P < 0.05$ that was not corrected for multiple comparisons, and bins that passed this criterion and less than 250 ms apart from one another were merged to obtain firing fields. Candidate time cells that had a maximum trial averaged firing rate > 0.75 Hz for light-off or light-on trials were considered time cells. The same approach was used to identify place cells during the spatial coding phase except that firing rates were occupancy-normalized with respect to the mouse's linearized position and only data during which the mouse was moving at a speed greater than 4 cm/s were used. Spatial bins that were less than 5 cm apart were merged.

A two independent-sample permutation method was used to test if a neuron's firing rate in light-off and light-on trials was different from zero. Here again, P values were adjusted according to a family-wise error rate of 0.05. The details of this analysis and Matlab code can be found in Groppe et al. (83). To visualize the periods during which firing rates between conditions significantly differed (Fig. 4B), we took the same approach used to identify firing fields described above.

In-Field and Out-Of-Field Firing Rates. Firing fields during the temporal or spatial coding phase were identified during light-off trials as described above. The firing rates inside and outside this firing field were computed separately for each cell.

Firing Stability. In order to assess firing stability for each neuron in a given condition, we 1) computed the trial average firing rate of first five trials and then 2) computed the trial average of all remaining trials in the condition. Finally, 3) these firing rate vectors were correlated using Matlab's `corr` function. The five-trial window was advanced by one trial (i.e., trials 2 to 6) and steps 1 to 3 were repeated. This process was iterated until the sliding five-trial block reached the last trial to create a list of correlation values. The firing stability for the cell was taken as the median value within the list.

Information. Information (84) was computed for time or space as:

$$\Sigma \left[\hat{\lambda}_x \log_2 \left(\frac{\hat{\lambda}_x}{\lambda} \right) * p_x \right], \text{ where } \hat{\lambda}_x \text{ is the estimated firing rate for the neuron at}$$

spatial or temporal bin x , λ is the mean firing rate of the neuron's tuning curve across all n bins of x , and p_x is the probability of the animal being in bin x .

PV Correlations to Measure Time Cell Ensemble Pattern Similarity between Conditions.

We adapted an analysis that is often used in studies of place cell populations (85). Regarding Fig. 4, for each recording session trials were split into light-off and light-on conditions and the trial-averaged firing rate for each time cell was computed using consecutive, nonoverlapping 250-ms bins. The cell population firing pattern in each condition was organized as a two-dimensional matrix with each row indexing a neuron and each column indexing a firing rate in a 250-ms bin during the temporal coding phase (Fig. 4B). Therefore, a PV was defined for each 250-ms time bin in a condition. Using Matlab's `corr` function, we correlated the PVs from the same time bin in each condition to obtain an r value. As there were 40 time bins in the temporal coding phase, we obtained a set of 40 r values. Finally, we took the median r value from the first (time bin 1 to 20) and second half (time bin 21 to 40) to yield a PV correlation for the CA2 inhibition condition in a recording session. To compute a same-trial PV correlation, we correlated the PVs representing the ensemble's trial averaged firing rates from even and odd trials from the light-off condition. To compute a random PV correlation, 1) the light-off trials were randomly divided into two sets of trials, 2) the spike train from each neuron in each trial was circularly shifted by a value that was selected randomly for each neuron and trial, then 3) the trial average was computed for each neuron using these two shifted trial sets to construct two PVs for each time bin in the temporal coding phase. Each PV for a given time bin was correlated and the median value taken from each half was taken as the r value. This process was repeated 999 more times and the mean PV correlation from this resampled distribution was used to measure similarity in the random condition. The PV correlations were compared using an rmANOVA and post hoc comparisons were Bonferroni-corrected. All P values were Greenhouse-Geisser-corrected. The same procedure was used to measure and compare PVs from the spatial coding phase (Fig. 5) by subdividing the linearized position into 40 spatial bins. All trial-averaged firing rates were occupancy-normalized and only data during which the mouse was moving > 4 cm/s was used.

Data Availability. All study data are included in the article and supporting information.

ACKNOWLEDGMENTS. We thank Mike Ragion, Anthony Moffa, Carl Twiss, and Jayson Derwin for technical assistance; Quentin Ferry, Shruti Muralidhar, and Chen Sun for comments about the manuscript before submission; and the rest of the S.T. laboratory for their support. This work was supported by the RIKEN Center for Brain Science, the Howard Hughes Medical Institute, and the JPB Foundation (S.T.).

1. W. B. Scoville, B. Milner, Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* **20**, 11–21 (1957).
2. F. Vargha-Khadem et al., Differential effects of early hippocampal pathology on episodic and semantic memory. *Science* **277**, 376–380 (1997).
3. J. O'Keefe, J. Dostrovsky, The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175 (1971).
4. J. O'Keefe, L. Nadel, *The Hippocampus as a Cognitive Map* (Oxford University Press, 1978).
5. T. Hartley, C. Lever, N. Burgess, J. O'Keefe, Space in the brain: How the hippocampal formation supports spatial cognition. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **369**, 20120510 (2013).

6. G. Buzsáki, E. I. Moser, Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat. Neurosci.* **16**, 130–138 (2013).
7. T. J. Willits, C. Lever, F. Cacucci, N. Burgess, J. O'Keefe, Attractor dynamics in the hippocampal representation of the local environment. *Science* **308**, 873–876 (2005).
8. S. Leutgeb et al., Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619–623 (2005).
9. E. Pastalkova, V. Itskov, A. Amarasingham, G. Buzsáki, Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).

10. C. J. MacDonald, K. Q. Lepage, U. T. Eden, H. Eichenbaum, Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron* **71**, 737–749 (2011).
11. C. J. MacDonald, S. Carrow, R. Place, H. Eichenbaum, Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *J. Neurosci.* **33**, 14607–14616 (2013).
12. H. Eichenbaum, Time cells in the hippocampus: A new dimension for mapping memories. *Nat. Rev. Neurosci.* **15**, 732–744 (2014).
13. H. Eichenbaum, On the integration of space, time, and memory. *Neuron* **95**, 1007–1018 (2017).
14. Y. Wang, S. Romani, B. Lustig, A. Leonardo, E. Pastalkova, Theta sequences are essential for internally generated hippocampal firing fields. *Nat. Neurosci.* **18**, 282–288 (2015).
15. N. T. M. Robinson *et al.*, Medial entorhinal cortex selectively supports temporal coding by hippocampal neurons. *Neuron* **94**, 677–688.e6 (2017).
16. M. Sabariego *et al.*, Time cells in the hippocampus are neither dependent on medial entorhinal cortex inputs nor necessary for spatial working memory. *Neuron* **102**, 1235–1248.e5 (2019).
17. E. A. Mankin, G. W. Diehl, F. T. Sparks, S. Leutgeb, J. K. Leutgeb, Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts. *Neuron* **85**, 190–201 (2015).
18. L. Lu, K. M. Igarashi, M. P. Witter, E. I. Moser, M.-B. Moser, Topography of place maps along the CA3-to-CA2 axis of the hippocampus. *Neuron* **87**, 1078–1092 (2015).
19. H. Lee, C. Wang, S. S. Deshmukh, J. J. Knierim, Neural population evidence of functional heterogeneity along the CA3 transverse axis: Pattern completion versus pattern separation. *Neuron* **87**, 1093–1105 (2015).
20. G. M. Alexander *et al.*, Social and novel contexts modify hippocampal CA2 representations of space. *Nat. Commun.* **7**, 10300 (2016).
21. M. E. Wintzer, R. Boehringer, D. Polygalov, T. J. McHugh, The hippocampal CA2 ensemble is sensitive to contextual change. *J. Neurosci.* **34**, 3056–3066 (2014).
22. V. Chevaleyre, S. A. Siegelbaum, Strong CA2 pyramidal neuron synapses define a powerful disinaptic cortico-hippocampal loop. *Neuron* **66**, 560–572 (2010).
23. K. Kohara *et al.*, Cell type-specific genetic and optogenetic tools reveal hippocampal CA2 circuits. *Nat. Neurosci.* **17**, 269–279 (2014).
24. L. M. DeVito *et al.*, Vasopressin 1b receptor knock-out impairs memory for temporal order. *J. Neurosci.* **29**, 2676–2683 (2009).
25. W. S. Young, J. Li, S. R. Wersinger, M. Palkovits, The vasopressin 1b receptor is prominent in the hippocampal area CA2 where it is unaffected by restraint stress or adrenalectomy. *Neuroscience* **143**, 1031–1039 (2006).
26. E. R. Wood, P. A. Dudchenko, R. J. Robitsek, H. Eichenbaum, Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–633 (2000).
27. L. M. Frank, E. N. Brown, M. Wilson, Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* **27**, 169–178 (2000).
28. B. J. Kraus, R. J. Robinson 2nd, J. A. White, H. Eichenbaum, M. E. Hasselmo, Hippocampal “time cells”: Time versus path integration. *Neuron* **78**, 1090–1101 (2013).
29. J. A. Ainge, M. A. A. van der Meer, R. F. Langston, E. R. Wood, Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. *Hippocampus* **17**, 988–1002 (2007).
30. J. Yamamoto, J. Suh, D. Takeuchi, S. Tonegawa, Successful execution of working memory linked to synchronized high-frequency gamma oscillations. *Cell* **157**, 845–857 (2014).
31. S. Ramirez *et al.*, Activating positive memory engrams suppresses depression-like behaviour. *Nature* **522**, 335–339 (2015).
32. H. Eichenbaum, The role of the hippocampus in navigation is memory. *J. Neurophysiol.* **117**, 1785–1796 (2017).
33. J. Csicsvari, H. Hirase, A. Czurkó, A. Mamiya, G. Buzsáki, Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. *J. Neurosci.* **19**, 274–287 (1999).
34. M. W. Howard *et al.*, A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *J. Neurosci.* **34**, 4692–4707 (2014).
35. S. M. Dudek, G. M. Alexander, S. Farris, Rediscovering area CA2: Unique properties and functions. *Nat. Rev. Neurosci.* **17**, 89–102 (2016).
36. S. J. Middleton, T. J. McHugh, CA2: A highly connected intrahippocampal relay. *Annu. Rev. Neurosci.* **43**, 55–72 (2020).
37. F. L. Hitti, S. A. Siegelbaum, The hippocampal CA2 region is essential for social memory. *Nature* **508**, 88–92 (2014).
38. T. Okuyama, T. Kitamura, D. S. Roy, S. Itoharu, S. Tonegawa, Ventral CA1 neurons store social memory. *Science* **353**, 1536–1541 (2016).
39. T. Meira *et al.*, A hippocampal circuit linking dorsal CA2 to ventral CA1 critical for social memory dynamics. *Nat. Commun.* **9**, 4163 (2018).
40. D. M. Salz *et al.*, Time cells in hippocampal area CA3. *J. Neurosci.* **36**, 7476–7484 (2016).
41. K. Nasrallah *et al.*, Routing hippocampal information flow through parvalbumin interneuron plasticity in area CA2. *Cell Rep.* **27**, 86–98.e3 (2019).
42. C. A. Barnes, B. L. McNaughton, J. O’Keefe, Loss of place specificity in hippocampal complex spike cells of senescent rat. *Neurobiol. Aging* **4**, 113–119 (1983).
43. T. J. McHugh, K. I. Blum, J. Z. Tsien, S. Tonegawa, M. A. Wilson, Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell* **87**, 1339–1349 (1996).
44. C. G. Kentros, N. T. Agnihotri, S. Streater, R. D. Hawkins, E. R. Kandel, Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* **42**, 283–295 (2004).
45. A. S. Bahar, P. R. Shirvalkar, M. L. Shapiro, Memory-guided learning: CA1 and CA3 neuronal ensembles differentially encode the commonalities and differences between situations. *J. Neurosci.* **31**, 12270–12281 (2011).
46. C. D. Harvey, P. Coen, D. W. Tank, Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
47. K. Rajan, C. D. Harvey, D. W. Tank, Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
48. M. E. Hasselmo, H. Eichenbaum, Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Netw.* **18**, 1172–1190 (2005).
49. Y. Katz, W. L. Kath, N. Spruston, M. E. Hasselmo, Coincidence detection of place and temporal context in a network model of spiking hippocampal neurons. *PLOS Comput. Biol.* **3**, e234 (2007).
50. M. W. Jones, M. A. Wilson, Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol.* **3**, e402 (2005).
51. M. W. Jones, M. A. Wilson, Phase precession of medial prefrontal cortical activity relative to the hippocampal theta rhythm. *Hippocampus* **15**, 867–873 (2005).
52. R. Boehringer *et al.*, Chronic loss of CA2 transmission leads to hippocampal hyperexcitability. *Neuron* **94**, 642–655.e9 (2017).
53. P. R. Gill, S. J. Y. Mizumori, D. M. Smith, Hippocampal episode fields develop with learning. *Hippocampus* **21**, 1240–1249 (2011).
54. H. Davoudi, D. J. Foster, Acute silencing of hippocampal CA3 reveals a dominant role in place field responses. *Nat. Neurosci.* **22**, 337–342 (2019).
55. V. H. Brun *et al.*, Impaired spatial representation in CA1 after lesion of direct input from entorhinal cortex. *Neuron* **57**, 290–302 (2008).
56. E. T. Rolls, P. Mills, The generation of time in the hippocampal memory system. *Cell Rep.* **28**, 1649–1658.e6 (2019).
57. B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser, M.-B. Moser, Path integration and the neural basis of the ‘cognitive map’. *Nat. Rev. Neurosci.* **7**, 663–678 (2006).
58. F. Savelli, J. J. Knierim, Origin and role of path integration in the cognitive representations of the hippocampus: Computational insights into open questions. *J. Exp. Biol.* **222**(suppl. 1), jeb188912 (2019).
59. E. I. Moser, M. B. Moser, B. L. McNaughton, Spatial representation in the hippocampal formation: A history. *Nat. Neurosci.* **20**, 1448–1464 (2017).
60. A. Samsonovich, B. L. McNaughton, Path integration and cognitive mapping in a continuous attractor neural network model. *J. Neurosci.* **17**, 5900–5920 (1997).
61. C. Haimerl *et al.*, Internal representation of hippocampal neuronal population spans a time-distance continuum. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 7477–7482 (2019).
62. M. E. Hasselmo, Arc length coding by interference of theta frequency oscillations may underlie context-dependent hippocampal unit data and episodic memory function. *Learn. Mem.* **14**, 782–794 (2007).
63. M. E. Hasselmo, *How We Remember: Brain Mechanisms of Episodic Memory* (The MIT Press, 2012).
64. V. Itskov, C. Curto, E. Pastalkova, G. Buzsáki, Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus. *J. Neurosci.* **31**, 2828–2834 (2011).
65. K. Hardcastle, S. Ganguli, L. M. Giocomo, Environmental boundaries as an error correction mechanism for grid cells. *Neuron* **86**, 827–839 (2015).
66. R. P. Jayakumar *et al.*, Recalibration of path integration in hippocampal place cells. *Nature* **566**, 533–537 (2019).
67. D. Derdikman *et al.*, Fragmentation of grid cell maps in a multicompartment environment. *Nat. Neurosci.* **12**, 1325–1332 (2009).
68. D. S. Rizzuto *et al.*, Reset of human neocortical oscillations during a working memory task. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 7931–7936 (2003).
69. M. J. Jutras, P. Fries, E. A. Buffalo, Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 13144–13149 (2013).
70. M. E. Hasselmo, C. E. Stern, Theta rhythm and the encoding and retrieval of space and time. *Neuroimage* **85**, 656–666 (2014).
71. R. G. M. Morris *et al.*, Elements of a neurobiological theory of the hippocampus: The role of activity-dependent synaptic plasticity in memory. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **358**, 773–786 (2003).
72. N. S. Clayton, A. Dickinson, Episodic-like memory during cache recovery by scrub jays. *Nature* **395**, 272–274 (1998).
73. N. J. Fortin, S. P. Wright, H. Eichenbaum, Recollection-like memory retrieval in rats is dependent on the hippocampus. *Nature* **431**, 188–191 (2004).
74. R. G. M. Morris, Episodic-like memory in animals: Psychological criteria, neural mechanisms and the value of episodic-like tasks to investigate animal models of neurodegenerative disease. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **356**, 1453–1465 (2001).
75. H. Eichenbaum, N. J. Cohen, Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron* **83**, 764–770 (2014).
76. A. Tsao *et al.*, Integrating time from experience in the lateral entorhinal cortex. *Nature* **561**, 57–62 (2018).

77. B. J. Kraus *et al.*, During running in place, grid cells integrate elapsed time and distance run. *Neuron* **88**, 578–589 (2015).
78. H. Eichenbaum, N. J. Cohen, *From Conditioning to Conscious Recollection* (Oxford University Press, 2004).
79. A. D. Ekstrom, C. Ranganath, Space, time and episodic memory: The Hippocampus is all over the cognitive map. *Hippocampus* **28**, 680–687 (2017).
80. H. Eichenbaum, P. Dudchenko, E. Wood, M. Shapiro, H. Tanila, The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron* **23**, 209–226 (1999).
81. H. Bokil, P. Andrews, J. E. Kulkarni, S. Mehta, P. P. Mitra, Chronux: A platform for analyzing neural signals. *J. Neurosci. Methods* **192**, 146–151 (2010).
82. R. C. Blair, J. J. Higgins, W. Karniski, J. D. Kromrey, A study of multivariate permutation tests which may replace Hotelling's T2 test in prescribed circumstances. *Multivariate Behav. Res.* **29**, 141–163 (1994).
83. D. M. Groppe, T. P. Urbach, M. Kutas, Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review. *Psychophysiology* **48**, 1711–1725 (2011).
84. E. J. Markus, C. A. Barnes, B. L. McNaughton, V. L. Gladden, W. E. Skaggs, Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus* **4**, 410–421 (1994).
85. S. Leutgeb, J. K. Leutgeb, A. Treves, M.-B. Moser, E. I. Moser, Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* **305**, 1295–1298 (2004).