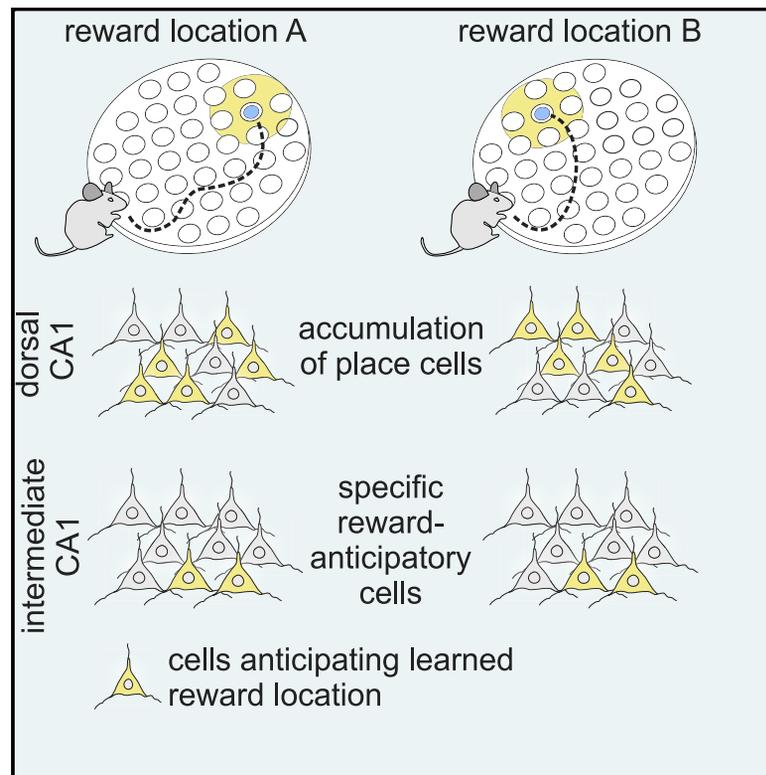


Different encoding of reward location in dorsal and intermediate hippocampus

Graphical abstract



Authors

Przemyslaw Jarzebowski,
Y. Audrey Hay, Benjamin F. Grewe,
Ole Paulsen

Correspondence

op210@cam.ac.uk

In brief

Jarzebowski et al. show that the memory of reward location is differently encoded along the hippocampus. In dorsal CA1, population activity increases in anticipation of reward, but the pool of active cells changes with the reward location. In intermediate CA1, the same cells anticipate multiple reward locations.

Highlights

- Ca^{2+} imaging from mice revealed spatial coding in dorsal and intermediate CA1
- Dorsal and intermediate CA1 encoded the memory of reward locations differently
- Dorsal, but not intermediate, CA1 encoded the memory by increasing population activity
- Intermediate, but not dorsal, CA1 encoded the memory in specific reward location cells



Report

Different encoding of reward location in dorsal and intermediate hippocampus

Przemyslaw Jarzebowski,¹ Y. Audrey Hay,¹ Benjamin F. Grewe,² and Ole Paulsen^{1,3,*}¹Department of Physiology, Development and Neuroscience, Physiological Laboratory, Cambridge CB2 3EG, UK²Institute of Neuroinformatics, University and ETH Zürich, Winterthurerstrasse 190, Zürich, Switzerland³Lead contact*Correspondence: op210@cam.ac.uk<https://doi.org/10.1016/j.cub.2021.12.024>

SUMMARY

Hippocampal place cells fire at specific locations in the environment. They form a cognitive map that encodes spatial relations in the environment, including reward locations.¹ As part of this encoding, dorsal CA1 (dCA1) place cells accumulate at reward.^{2–5} The encoding of learned reward location could vary between the dorsal and intermediate hippocampus, which differ in gene expression and cortical and subcortical connectivity.⁶ While the dorsal hippocampus is critical for spatial navigation, the involvement of intermediate CA1 (iCA1) in spatial navigation might depend on task complexity⁷ and learning phase.^{8–10} The intermediate-to-ventral hippocampus regulates reward-seeking,^{11–15} but little is known about the involvement in reward-directed navigation. Here, we compared the encoding of learned reward locations in dCA1 and iCA1 during spatial navigation. We used calcium imaging with a head-mounted microscope to track the activity of CA1 cells over multiple days during which mice learned different reward locations. In dCA1, the fraction of active place cells increased in anticipation of reward, but the pool of active cells changed with the reward location. In iCA1, the same cells anticipated multiple reward locations. Our results support a model in which the dCA1 cognitive map incorporates a changing population of cells that encodes reward proximity through increased population activity, while iCA1 provides a reward-predictive code through a dedicated subpopulation. Both of these location-invariant codes persisted over time, and together they provide a dual hippocampal reward location code, assisting goal-directed navigation.^{16,17}

RESULTS

To track the activity of the same dCA1 and iCA1 cells when mice learned different reward locations, we imaged calcium fluorescence of excitatory cells in Thy1-GCaMP6f mice¹⁸ with a head-mounted, miniature microscope^{19,20} over multiple days (Figures 1A–1C). We imaged daily 170 ± 19 dCA1 cells from seven animals and 70 ± 11 iCA1 cells from six animals and matched the identity of active cells between days (Figures S1A–S1D).

iCA1 and dCA1 encoded comparable spatial information despite iCA1 place fields being larger

First, to confirm that calcium signals yield similar place cell properties as reported using tetrode recordings,^{21–24} we characterized place-specific neuronal responses during foraging on the same maze as that subsequently used for reward location learning (Figure 1D; Video S1). Mice foraged for liquid rewards baited in randomly selected wells of the 120-cm-diameter maze. Because hippocampal lesions can increase mobility,²⁵ we confirmed that the dCA1 and iCA1 implanted mice ran in a similar fraction of the trials and with similar speeds (Figures S1E–S1F).

We identified $46\% \pm 3\%$ of dCA1 and $43\% \pm 4\%$ of iCA1 cells as place cells (inconclusive evidence for difference: $BF_{10} = 0.45$, $CI = [-15\%, 11\%]$; Figure 1E). In agreement

with tetrode recordings,^{21–23} the place fields were larger in iCA1 (strong evidence: $BF_{10} = 12$, $CI = [24\%, 204\%]$; Figure 1F). The mean count of place fields per place cell was similar in dCA1 and iCA1 (inconclusive evidence: $BF_{10} = 0.80$, $CI = [-16\%, 4\%]$; Figure 1G). Studies using tetrode recordings reported that spatial information decreases from dorsal to ventral CA1.^{21,23,24} Here, the spatial information relative to the cell's spatial information expected by chance did not differ between dCA1 and iCA1 place cells (moderate evidence for no difference: $BF_{10} = 0.25$, $CI = [-31\%, 65\%]$; Figure 1H). However, calcium imaging underestimates spatial information and could have failed to capture the difference.²⁶ Within-day stability did not differ between dCA1 and iCA1 place cells (moderate evidence for no difference: $BF_{10} = 0.22$, $CI = [-0.12, 0.19]$; Figure 1I), and both cell populations encoded the animal's position with similar accuracy (inconclusive evidence for difference in median decoding error: $BF_{10} = 0.8$, $CI = [-14 \text{ cm}, 3 \text{ cm}]$; Figure 1J).

dCA1, but not iCA1, place cells accumulated at reward locations

Next, we compared how the activity of dCA1 and iCA1 place cells changed with learning.^{2,3,5,27–29} Mice learned sets of two fixed reward locations in daily sessions of 8 trials. The learning period spanned 5 days for the first set of locations and 2 days



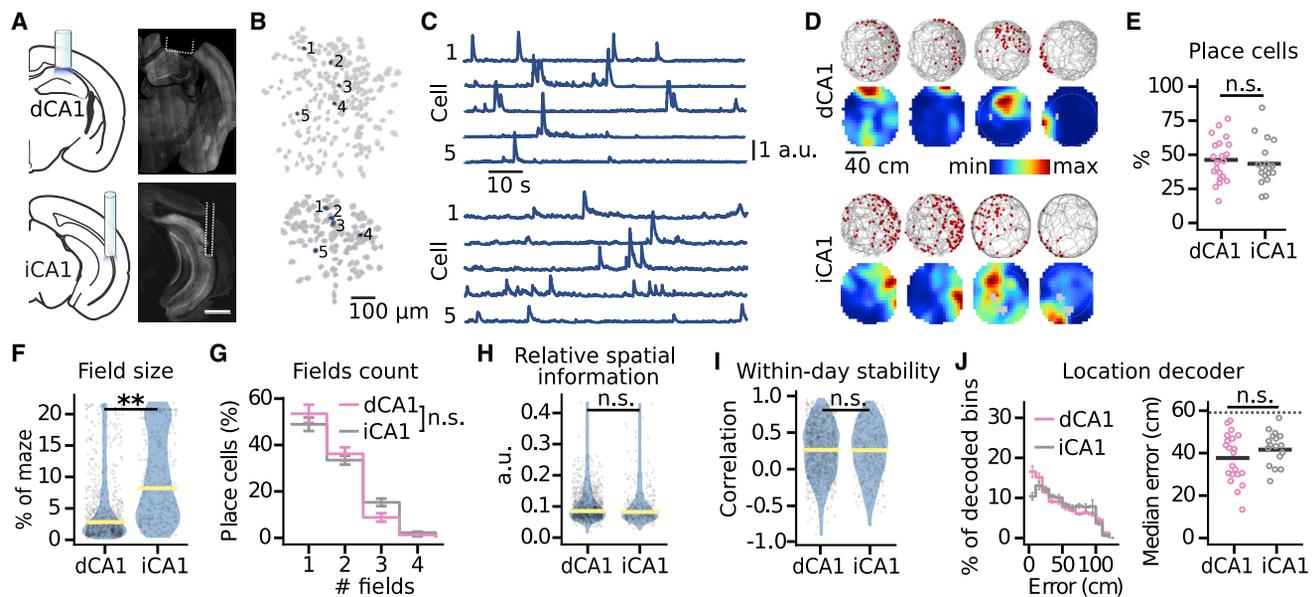


Figure 1. Similar spatial information encoded in dCA1 and iCA1 during foraging

(A) Location of the GRIN lens implanted above dCA1 (top) and iCA1 (bottom) pyramidal cells expressing GCaMP6f. Scale bar, 1 mm.
 (B) ROIs with putative neurons detected in a single day from a dCA1- (top) and iCA1-implanted (bottom) mouse.
 (C) Background-subtracted fluorescence traces from the blue cells in (B).
 (D) Examples of dCA1 (top) and iCA1 place cells (bottom) recorded during single-day foraging. Locations of calcium events marked with a red dot are overlaid over mouse paths (top); place maps are shown below. Gray pixels represent unsampled locations.
 (E) Percentage of the dCA1 and iCA1 cells identified as place cells during foraging.
 (F) Field sizes of the dCA1 and iCA1 place cells.
 (G) Field counts per place cell.
 (H) Spatial information of place cells relative to the cell's spatial information expected by chance.
 (I) Within-day stability of place cells measured as a correlation between place maps from early and late trials on the same day.
 (J) Accuracy of decoding location from neural activity in the dCA1 or iCA1. Distribution of decoding errors (left) and the median error (right) calculated using cross-validation on single-day activity. Decoders were trained and evaluated on 30 sampled cells to match the cell counts in dCA1 and iCA1; the sampling was repeated fifty times. Linear mixed-effects model: $F_{(1, 11)} = 2.4$, $p = 0.15$; $BF_{10} = 0.8$, $CI = [-14 \text{ cm}, 3 \text{ cm}]$; $n = 39$ sessions.
 Distribution of the values shown on violin plots has the width proportional to density; horizontal bars mark the means; individual data points are overlaid on top of the violin plots. Error bars mark \pm SEM. To avoid double-counting cells sampled on different days, (F), (H), and (I) show data for place cells recorded in the last-day foraging. The statistical tests in (E), (G), and (J) used linear and in (F) and (H) used log-linear mixed-effects models. For (E)–(I) statistics, see Table S1.
 See also Figure S1 and Video S1.

for subsequent sets, each with one reward translocated to a pseudo-random location, for a total of three or four sets (Figure 2A). Mice took progressively shorter paths to the rewards (Figure 2B; Video S2). Their memory was tested in unbaited test trials the day after learning each set. Mice crossed the reward zones (20-cm-radius disks centered on the learned reward locations) $64\% \pm 7\%$ more times in the first 120 s of the test trials compared with the same zones during foraging (Figures 2C and 2D).

To investigate how the reward location memory and goal-directed behavior affect spatial coding, we compared place fields during test trials with pre-learning foraging trials (Figures 2E and 2F). The fraction of place cells with fields inside the reward zones increased in dCA1 (by $65\% \pm 11\%$) but not in iCA1 (Figures 2G and 2H). We verified that increased sampling of the reward locations did not account for the accumulation of dCA1 place cells (Figure S2) and that the iCA1 place fields did not enlarge toward the rewards (log-linear mixed-effects model; $F_{(1, 973)} = 35$, $p = 10^{-8}$; $BF_{10} = 10^6$, $CI = [-31\%, -17\%]$). The different effects on dCA1 and iCA1 place fields were not

explained by speed differences between the two mouse groups (Figure S2C).

dCA1 place cells increased their population activity in anticipation of reward

To understand how reward location memory affects the population activity, we analyzed dCA1 and iCA1 activity as mice approached the reward. In late learning (the last day of learning a set of reward locations), the mean dCA1 activity increased by 0.09 ± 0.01 SD when mice approached the reward ($BF_{10} = 10^{11}$, $CI = [0.08, 0.13]$ SD; Figures 3A and 3B, left; Figures S3A and S3C). The effect was absent on the first learning day ($BF_{10} = 0.18$, $CI = [-0.03, 0.05]$ SD) and when the mice stopped at non-rewarded locations ($BF_{10} = 0.16$, $CI = [-0.03, 0.01]$ SD). The fraction of active cells (cells whose activity exceeded a Z-score of 0.5) increased $33\% \pm 4\%$ among place cells ($BF_{10} = 10^{10}$, $CI = [21\%, 46\%]$) and $15\% \pm 3\%$ among non-place cells ($BF_{10} = 168$, $CI = [6\%, 22\%]$; Figure 3B, middle; Figure S3E). The cells active at a particular reward consisted of a repeatedly activated cell population ($7\% \pm 1\%$ of place cells and $4\% \pm 1\%$

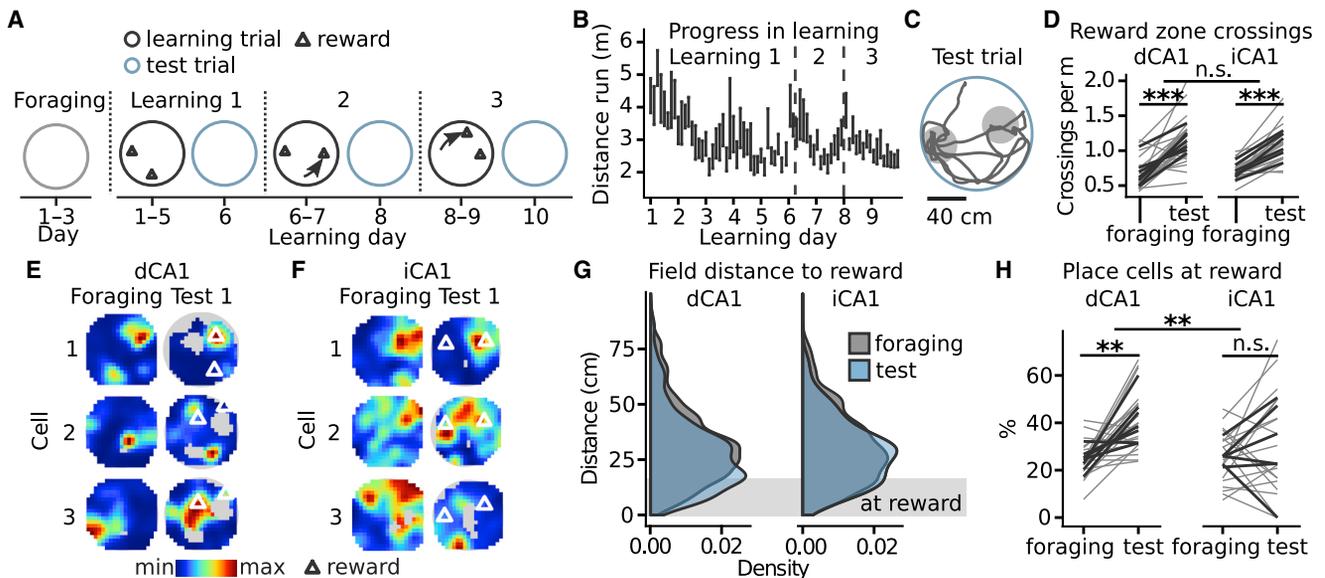


Figure 2. dCA1, but not iCA1, place cells accumulated at learned reward locations

(A) Timeline for the learning and test sessions showing when and how the reward locations (triangles) changed.
 (B) Progress in learning measured by distance run per trial. Vertical bars mark the mean distance \pm SEM, dashed lines mark the time of reward translocations.
 (C) Example running path of a mouse during the test trial. Gray discs show the extent of the reward zone used for analysis.
 (D) Number of reward zone crossings during the first 120 s of the test trials compared with the crossings of the zones centered on the same locations during foraging. Thin lines connect values for a single set of reward zones, thick lines connect animal averages. Effect of learning: $F_{(1, 61)} = 105$, $p = 10^{-14}$; $BF_{10} = 6 \times 10^{10}$, $CI = [45\%, 85\%]$; moderate evidence for lack of learning \times implant location interaction: $F_{(1, 16)} = 0.03$, $p = 0.87$; $BF_{10} = 0.18$, $CI = [-13\%, 11\%]$; $n = 44$ trials.
 (E) Examples of dCA1 place fields for the same cell during foraging and the test after learning 1. Triangles mark reward locations.
 (F) As in (E) but for iCA1 cells.
 (G) Distribution of distances from test trial place fields to the closer of the two reward locations compared with distances to the same locations from foraging place fields.
 (H) Proportion of place cells with a place field inside a reward zone. Thin lines connect values for a single set of reward locations, thick lines connect animal averages. Significant trial-type \times implant location interaction: $F_{(1, 13)} = 16.2$, $p = 0.001$; increase in dCA1: $t_{(12,1)} = 4.8$, $p = 0.002$, $BF_{10} = 42850$, $CI = [34\%, 96\%]$; inconclusive evidence for change in iCA1: $t_{(13,2)} = 1.27$, $p = 0.65$, $BF_{10} = 0.37$, $CI = [-18\%, 65\%]$; $n = 44$ trials.
 Linear mixed-effects model tested for the effects of learning, implant location, and their interaction in (D) and (H). Post-hoc test on least-square means of the model tested the activity change in dCA1 and iCA1 in (H).
 See also [Figure S2](#) and [Video S2](#).

of non-place cells were active in more than half of the approaches) and a changing, broader cell population: $50\% \pm 2\%$ of place cells were active at reward at least once, meaning that a large fraction of these cells fired outside of their place field ([Figure S3F](#)). The increased activity of place cells was visible in plots of mean activity against distance to reward ([Figure S3G](#)) and correlated with day-mean performance ([Figure 3C](#)).

The fraction of active place cells, but not non-place cells, increased at the learned reward locations also in unbaited test trials (place cells: $BF_{10} = 7.4$, $CI = [6\%, 33\%]$; non-place cells: $BF_{10} = 0.09$, $CI = [-9\%, 5\%]$; [Figure 3B](#), right; [Figure S3H](#)). Therefore, the increased number of active place cells was not exclusively caused by reward-associated olfactory cues, which might have contributed to the higher baseline activity in baited trials.

iCA1 non-place cells decreased their activity in anticipation of reward

Changes in iCA1 population activity contrasted with those in dCA1. In late learning, the mean iCA1 activity decreased by 0.09 ± 0.01 SD when mice approached the reward ($BF_{10} = 39$,

$CI = [-0.15, -0.4]$ SD; [Figures 3D](#) and [3E](#), left; [Figures S3B](#) and [S3C](#)). This effect was absent on the first learning day ($BF_{10} = 0.21$, $CI = [-0.11, 0.5]$ SD) and when the mice stopped at non-rewarded locations ($BF_{10} = 0.11$, $CI = [-0.02, 0.05]$). The opposite direction of change in dCA1 and iCA1 activity was not explained by mice approaching the reward with different speeds ([Figure S3D](#)). The fraction of active iCA1 cells did not change among place cells ($BF_{10} = 0.08$, $CI = [-15\%, 23\%]$) but decreased by $32\% \pm 3\%$ among non-place cells ($BF_{10} = 10^6$, $CI = [-44\%, -21\%]$; [Figure 3E](#), middle; [Figure S3E](#)). Large fractions of place cells and non-place cells were inactive 0–1 s before all approach to a particular reward ($46\% \pm 3\%$ of the place cells, $50\% \pm 4\%$ of the non-place cells; [Figure S3F](#)). The activity decrease was visible in plots of mean activity against distance to reward ([Figure S3G](#)), and it correlated with day-mean performance ([Figure 3F](#)).

The fraction of active iCA1 cells did not change at the learned reward locations in unbaited test trials (place cells: $BF_{10} = 0.08$, $CI = [-18\%, 12\%]$; non-place cells: $BF_{10} = 0.24$, $CI = [-26\%, 3\%]$; [Figure 3E](#), right; [Figure S3H](#)). This suggests that the decrease in iCA1 activity during learning was related to the

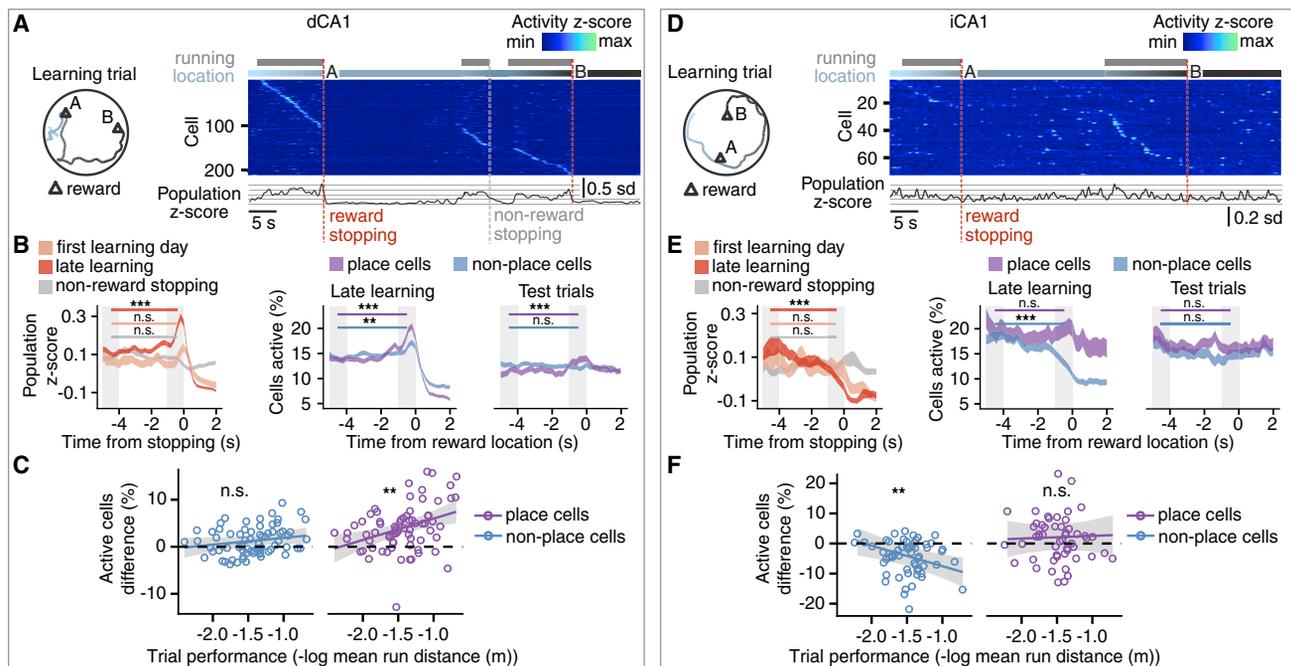


Figure 3. dCA1 population activity ramping-up and iCA1 activity ramping-down as mice approach the reward

(A) Example single-trial path of a mouse (left) together with dCA1 activity (right). Each row of the raster shows the Z-scored activity of a single cell. The cells are sorted by the time of their maximum activity. Periods when the mouse ran are marked above the raster with gray. Blue-colored gradient above the raster indicates the color-matched spatial location on the left; the trace further on shows population activity Z-score. Dashed vertical lines show the time when the mouse stopped at rewards or a non-rewarded location.

(B) dCA1 population activity as mice approached the reward. After learning, the activity increased before the reward (left), and the percentage of active place cells and non-place cells increased during approach to the reward location (middle). The percentage of active place cells increased during the approach in test trials (right). The traces have a width of \pm SEM; gray rectangles mark 1-s-long periods used for the statistical comparison. Data compared with post-hoc tests on least-square means of linear mixed-effects models for the effects of learning stage, reward proximity, and their interaction. Statistics presented in Table S2.

(C) Change in the number of active cells from the 4–5 s before reward approach to 0–1 s shown as a function of day-mean learning performance. The black line shows the slope of modeled regression together with its credibility interval in gray. Linear mixed-effects model, correlation of performance with place cells: $F_{(1, 66)} = 10$, $p = 0.002$, $BF_{10} = 7.2$, slope CI = [0.8, 7.2]; correlation with non-place cells: $F_{(1, 36)} = 2.4$, $p = 0.13$; $BF_{10} = 0.63$, slope CI = [−0.8, 3.1]; $n = 68$ day sessions.

(D) As in (A) but for iCA1.

(E) As in (B) but for iCA1. Statistics presented in Table S2.

(F) As in (C) but for iCA1. Linear mixed-effects model, correlation with non-place cells: $F_{(1, 55)} = 8$, $p = 0.005$, $BF_{10} = 7.0$, slope CI = [−10.5, −1.2]; correlation with place cells: $F_{(1, 49)} = 0.9$, $p = 0.77$; $BF_{10} = 0.31$, slope CI = [−5.4, 6.6]; $n = 58$ day sessions.

See also Figure S3.

reward-associated olfactory cues or speed profile when mice approached reward location (Figures S3D and S3I).

The same iCA1, but different dCA1 cells, tracked changing reward location

To investigate whether some cells signaled location-independent anticipation of reward,^{4,30} we compared their activity between test trials after the mice had learned different reward locations. Of the cells active on the first test trial, $60\% \pm 6\%$ were active again on the second and $50\% \pm 5\%$ on the third test trial (Figure S4A). We followed the remapping of place cells present in two subsequent test trials. Of the 89 dCA1 place cells with a reward field (place field within 20 cm of reward location) at the previous reward location, 25% retained their place field and 31% remapped to either of the current reward locations. However, their place fields were not closer to the current reward locations than those of cells previously without a reward field (moderate evidence for no difference: $BF_{10} = 0.16$, CI = [−6, 2] cm; Figures 4A and S4B).

In contrast, iCA1 cells with a reward field at the previous reward location had place fields closer to the current reward locations than the cells previously without a reward field (strong evidence: $BF_{10} = 122$, CI = [6, 19] cm; Figures 4A and S4B). The effect was not due to different place field sizes in the two groups of iCA1 place cells (moderate evidence for no difference in size: $BF_{10} = 0.32$; Figure S4C).

The subpopulations of iCA1 place cells with zero or multiple reward fields were larger than expected by chance—one remapped avoiding and the other remapped tracking reward locations. To quantify this, we looked at cells present in at least two test trials and classified them as place cells in more than half of them. Of 106 iCA1 place cells, 47.1% had zero reward fields, and 4.7% had reward fields at more than half of reward locations, significantly more than the respective 40.0% and 2.9% expected by chance (Figures 4B and S4C). In comparison, of 423 dCA1 place cells, 33.8% had zero reward fields, and 2.1% had reward fields at more than half of reward locations—fractions similar to the respective 34.1% and 2.5% expected by chance (Figures 4B and S4C).

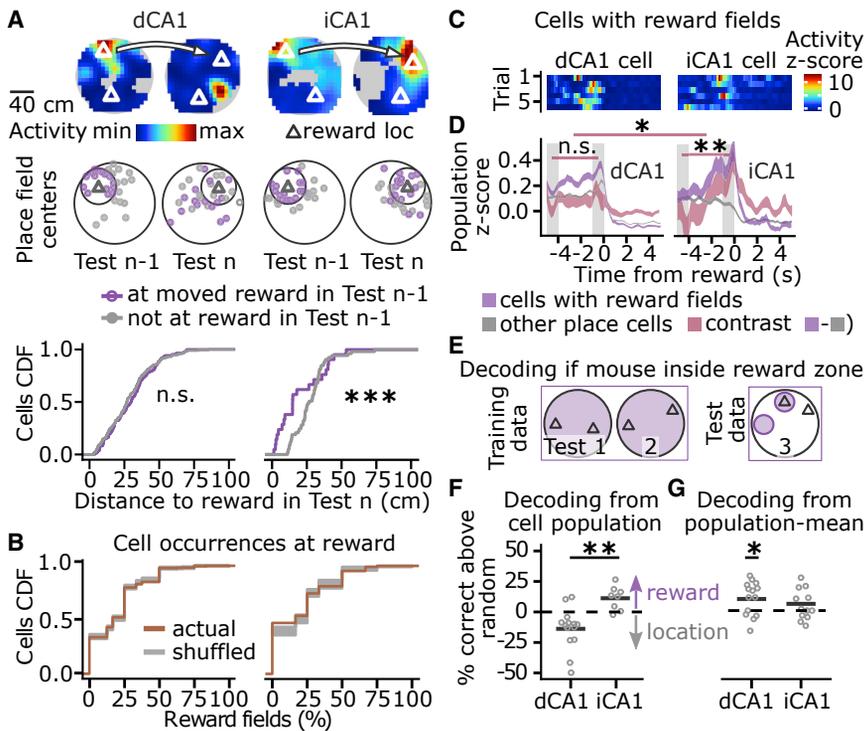


Figure 4. The same iCA1, but different dCA1, place cells tracked changing reward location and were active in anticipation of the reward

(A) Remapping of place cells after learning a changed reward location. iCA1 place cells with a field at the previous reward location had their fields subsequently closer to the current reward locations than the cells previously without a reward field. Top: example place maps of the same cells from the previous (Test n-1) and current test trial (Test n). Middle: place field centers of randomly sampled cells. Visualization preserves distances to the reward location, but the exact locations differ. Bottom: cumulative distribution function (CDF) of distances from the place field centers to the closer of the reward locations. Log-linear mixed-effects model, dCA1: $F_{(1, 378)} = 0.94$, $p = 0.33$, $BF_{10} = 0.16$, $CI = [-6, 2]$ cm; $n = 279$ cells; iCA1: $F_{(1, 90)} = 17$, $p = 0.0001$; $BF_{10} = 122$, $CI = [6, 19]$ cm; $n = 74$ cells.

(B) CDF for the frequency with which cells had a reward field during the test trials. For example, a cell with fields at two reward locations during the three trials had a frequency of 33%.

(C) Activity in single-day learning trials of example dCA1 and iCA1 cells that had a reward field in the next-day test trial. Each row shows activity in a single trial centered on the time the mouse arrived at the reward.

(D) Population activity in two cell groups: cells with reward fields in the next-day test trial and the cells without a reward field, as well as their mean within-trial difference (contrast). Activity shown around the time of mice arriving at reward. Linear mixed-effects model comparing the contrast at 4–5 s and 0–1 s, recording site \times reward proximity interaction: $F_{(1, 456)} = 5.9$, $p = 0.016$; post-hoc test showed no change in dCA1: $t_{(455)} = -0.2$, $p = 0.90$, $BF_{10} = 0.10$, $CI = [-0.1, 0.1]$ SD; and a change in iCA1: $t_{(456)} = -3.3$, $p = 0.006$, $BF_{10} = 21.2$, $CI = [0.1, 0.4]$ SD; $n = 278$ learning trials. The trace has a width of \pm SEM. Gray rectangles mark 1-s-long periods used for the statistical comparison.

(E) Training and test data used for binary decoders predicting whether the mouse was running inside a reward zone. The decoders were trained on the activity from test trials on two different days. They were tested on activity from another day when the decoder had to flip its prediction for the two tested locations: the previously rewarded location was unrewarded and vice versa.

(F) Accuracy of decoding from the activity of all cells whether the mouse was running inside a reward zone, shown as the difference from random predictions. Decoders were evaluated on data from test trials with different reward locations than in training. For the same location, the decoder had to give the opposite answer to the training data. Accuracy below the random level means the decoder predicted location rather than reward zone. Linear mixed-effects model: $F_{(1, 20)} = 12$, $p = 0.002$, $BF_{10} = 7.1$, $CI = [3\%, 33\%]$; $n = 20$ trials.

(G) Accuracy of decoding from population-mean activity, shown as the difference from random prediction. $t_{(13)} = 2.7$, $p = 0.02$, $BF_{10} = 3$, $CI = [1\%, 16\%]$; $n = 14$ dCA1 trials.

Individual points in (F) and (G) show decoding accuracy per test trial. Horizontal bars mark the means. See also Figure S4.

We investigated whether cells with reward fields during test trials were active in anticipation of reward during the previous-day learning trials. dCA1 and iCA1 cells with reward fields in test trials had higher activity than other place cells as the mice approached either of the two rewards (Figures 4C, 4D, and S4E). The difference between the two populations increased as mice approached the reward in iCA1, but not in dCA1, suggesting that iCA1 cells form distinct subpopulations (no change in dCA1: $BF_{10} = 0.10$, $CI = [-0.1, 0.1]$ SD; increase in iCA1: $BF_{10} = 21.2$, $CI = [0.1, 0.4]$ SD; Figures 4D and S4F).

The iCA1 cell activity and dCA1 population-mean activity predicted reward location with stable and location-independent code

Lastly, we assessed the similarity of the hippocampal encoding of memory for different reward locations. We created a binary decoder, which predicted from the instantaneous activity of

the dCA1 or iCA1 cell population during test trials whether the mouse was running inside a reward zone or not (Figure 4E). When tested on the same dataset as used for training, both decoders performed well-above random predictions (Figure S4G). To investigate whether the activity generalized to changed reward locations, we evaluated the decoders on data from a different test trial at one new and one of the previous reward locations (Figure 4E). Decoding from iCA1 had an accuracy of $10\% \pm 3\%$ above and from dCA1 of $11\% \pm 4\%$ below random predictions ($BF_{10} = 7.1$, $CI = [3\%, 33\%]$; Figure 4F). Thus, while the iCA1 decoder predicted the reward location, the dCA1 decoder predicted the same location as in the training dataset although the reward was moved, which means it decoded the mouse location rather than the reward location.

Because the number of active dCA1 cells ramped up when mice approached the learned reward location (Figures 3A–3C), we tested another decoder based on the population-mean

activity. The decoder used two inputs: the fraction of active place cells and the fraction of active non-place cells. Decoding performed with an accuracy of $10\% \pm 4\%$ above chance ($BF_{10} = 3$, $CI = [1\%, 16\%]$; Figure 4G). Thus, the ramping-up population activity of dCA1 place cells encoded reward location independently of their spatial coding.

DISCUSSION

We found that both dCA1 and iCA1 activity predict reward location; however, they do so using different codes: (1) The dCA1 anticipated the reward with increased population activity whose strength correlated with learning performance. The activity engaged place cells that changed with the reward locations, allowing independent reward and spatial coding. (2) In iCA1 the same cells were active in anticipation of the reward. Their population provided a code for learned reward location that persisted across different reward locations and time.

Comparison of reward location coding in dCA1 and iCA1

Consistent with previous reports, dCA1 place cells accumulated at the learned reward locations^{2,3,5,28,29} (Figure 2F) and, as mice approached the reward, the number of active place cells ramped up (Figure 3B). Ramping dCA1 activity was previously reported during reward anticipation in immobile animals.^{31,32} Therefore, it can predict the reward location independently of the spatial representations during movement.

Our findings suggest that dCA1 place cells active at reward locations are part of a flexible spatial rather than a dedicated reward-coding population. We found evidence against the hypothesis that dCA1 cells with place fields close to the reward remapped to track the translocated reward better than other place cells (Figure 4D). A random subset of place cells accumulating at reward locations accounted for the total number of reward fields per dCA1 place cell (Figure 4E), apparently at odds with the conclusion of Gauthier and Tank.⁴ There are several differences between the two studies: (1) we compared the cell activity across multiple days and reward locations and in the absence of a reward; (2) we used a two-dimensional maze, and the mice were freely moving; and (3) the null hypothesis we tested takes into account the accumulation of place cells at reward locations. The evidence we present suggests that dCA1 place cells are attracted to the reward stochastically, although with probabilities that might differ between them.^{33,34}

In contrast, the density of iCA1 place fields was unaffected by the memory of reward location. In a study where mice alternated between two marked reward locations, the iCA1 place cells accumulated at the reward locations and were sensitive to reward value.²⁴ Possibly, the iCA1 place cells accumulate at the reward during stereotypical running or in some form of value association. Heterogeneity among the intermediate-to-ventral CA1 cells^{30,35} could also contribute to the difference. In our study, although non-place cells decreased their activity as the mice approached the learned reward location, a subpopulation of place cells increased their activity and remapped to track the changing reward locations (Figures 3E, 4D, and 4E), similar to the goal-encoding cells suggested to exist in dCA1.⁴

Function of reward-predictive encoding

Both dCA1 and iCA1 predicted learned reward location using time- and location-invariant codes (Figures 4G–4I). Their signal might direct the animal during navigation by increasing their activity in the proximity of a goal¹⁶ or by signaling reward expectation.¹⁷

The different encoding of reward anticipation in dCA1 and iCA1 affects how the signal can be relayed downstream. The reward-anticipatory subpopulation in iCA1 could include nucleus accumbens-projecting neurons controlling appetitive memory^{11–13,15,36} and exclude those controlling aversion or fear.^{36–38} In dCA1, the ramping-up population activity resembles that seen in the dopaminergic system.³⁹ Such signal could indiscriminately excite downstream targets of dCA1 or recruit specific projection neurons, including the nucleus accumbens-projecting neurons that support spatial appetitive memory.⁴⁰

Further studies are required to determine how the reward-anticipatory signals in dCA1 and iCA1 affect activity downstream of the hippocampus. Hippocampal reward-predictive signals could be important for learning and choosing appropriate actions during reward-guided navigation as they are in reinforcement learning models.^{17,41}

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
 - Mice
- METHOD DETAILS
 - Surgery
 - Histological processing
 - Cheeseboard maze task
 - Calcium imaging
 - Calcium signal processing
 - Place cell detection and analysis
 - Calculation of place fields at reward location
 - Bayesian decoders
 - Downsampled data comparison
 - Population activity on reward approach
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.cub.2021.12.024>.

ACKNOWLEDGMENTS

We thank Dr Ben Phillips for experimental advice and Dr. Julija Krupic and Mr. Roman Boehringer for their feedback on early versions of the manuscript. We gratefully acknowledge the Cambridge Advanced Imaging Centre for their support and assistance in this work. We also thank Mr. Max Schwiening and Dr. Christof Schwiening for 3D printing parts of the experimental equipment. This work was supported by the Biotechnology and Biological Sciences

Research Council (O.P., BB/P019560/1), the Swiss National Science Foundation (B.F.G., CRSII5-173721 and 315230_189251), and ETH project funding (B.F.G., ETH-20 19-01). P.J. was supported by a Biotechnology and Biological Sciences Research Council Doctoral Training Programme studentship.

AUTHOR CONTRIBUTIONS

All authors designed research; P.J. performed research and analyzed data; P.J. wrote the paper with support from the other authors; and all authors approved the final version of the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: September 7, 2021

Revised: November 5, 2021

Accepted: December 8, 2021

Published: January 10, 2022

REFERENCES

- O'Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Oxford University Press).
- 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.
- Boccarda, C.N., Nardin, M., Stella, F., O'Neill, J., and Csicsvari, J. (2019). The entorhinal cognitive map is attracted to goals. *Science* *363*, 1443–1447.
- Gauthier, J.L., and Tank, D.W. (2018). A dedicated population for reward coding in the hippocampus. *Neuron* *99*, 179–193, e7.
- 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.
- Fanselow, M.S., and Dong, H.-W. (2010). Are the dorsal and ventral hippocampus functionally distinct structures? *Neuron* *65*, 7–19.
- Contreras, M., Pelc, T., Llofriu, M., Weitzenfeld, A., and Fellous, J.M. (2018). The ventral hippocampus is involved in multi-goal obstacle-rich spatial navigation. *Hippocampus* *28*, 853–866.
- de Hoz, L., Knox, J., and Morris, R.G.M. (2003). Longitudinal axis of the hippocampus: both septal and temporal poles of the hippocampus support water maze spatial learning depending on the training protocol. *Hippocampus* *13*, 587–603.
- Bast, T., Wilson, I.A., Witter, M.P., and Morris, R.G.M. (2009). From rapid place learning to behavioral performance: a key role for the intermediate hippocampus. *PLOS Biol* *7*, e1000089.
- Ruediger, S., Spirig, D., Donato, F., and Caroni, P. (2012). Goal-oriented searching mediated by ventral hippocampus early in trial-and-error learning. *Nat. Neurosci.* *15*, 1563–1571.
- Britt, J.P., Benaliouad, F., Mcdevitt, R.A., Stuber, G.D., Wise, R.A., and Bonci, A. (2012). Synaptic and behavioral profile of multiple glutamatergic inputs to the nucleus accumbens. *Neuron* *76*, 790–803.
- LeGates, T.A., Kvarta, M.D., Tooley, J.R., Francis, T.C., Lobo, M.K., Creed, M.C., and Thompson, S.M. (2018). Reward behaviour is regulated by the strength of hippocampus–nucleus accumbens synapses. *Nature* *564*, 258–262.
- Zhou, Y., Zhu, H., Liu, Z., Chen, X., Su, X., Ma, C., Tian, Z., Huang, B., Yan, E., Liu, X., and Ma, L. (2019). A ventral CA1 to nucleus accumbens core engram circuit mediates conditioned place preference for cocaine. *Nat. Neurosci.* *22*, 1986–1999.
- Riaz, S., Schumacher, A., Sivagurunathan, S., Van Der Meer, M., and Ito, R. (2017). Ventral, but not dorsal, hippocampus inactivation impairs reward memory expression and retrieval in contexts defined by proximal cues. *Hippocampus* *27*, 822–836.
- Ito, R., Robbins, T.W., Pennartz, C.M., and Everitt, B.J. (2008). Functional interaction between the hippocampus and nucleus accumbens shell is necessary for the acquisition of appetitive spatial context conditioning. *J. Neurosci.* *28*, 6950–6959.
- Burgess, N., and O'Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus* *6*, 749–762.
- Foster, D.J., Morris, R.G.M., and Dayan, P. (2000). A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus* *10*, 1–16.
- Dana, H., Chen, T.-W., Hu, A., Shields, B.C., Guo, C., Looger, L.L., Kim, D.S., and Svoboda, K. (2014). Thy1-GCaMP6 transgenic mice for neuronal population imaging in vivo. *PLoS One* *9*, e108697.
- Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., Gamal, A. El, and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* *16*, 264–266.
- Ghosh, K.K., Burns, L.D., Cocker, E.D., Nimmerjahn, A., Ziv, Y., Gamal, A.E., and Schnitzer, M.J. (2011). Miniaturized integration of a fluorescence microscope. *Nat. Methods* *8*, 871–878.
- Jung, M.W., Wiener, S.I., and McNaughton, B.L. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *J. Neurosci.* *14*, 7347–7356.
- Maurer, A.P., VanRhoads, S.R., Sutherland, G.R., Lipa, P., and McNaughton, B.L. (2005). Self-motion and the origin of differential spatial scaling along the septo-temporal axis of the hippocampus. *Hippocampus* *15*, 841–852.
- Keinath, A.T., Wang, M.E., Wann, E.G., Yuan, R.K., Dudman, J.T., and Muzzio, I.A. (2014). Precise spatial coding is preserved along the longitudinal hippocampal axis. *Hippocampus* *24*, 1533–1548.
- Jin, S.-W., and Lee, I. (2021). Differential encoding of place value between the dorsal and intermediate hippocampus. *Curr. Biol.* *31*, 3053–3072, e5.
- Bannerman, D.M., Yee, B.K., Good, M.A., Heupel, M.J., Iversen, S.D., and Rawlins, J.N.P. (1999). Double dissociation of function within the hippocampus: a comparison of dorsal, ventral, and complete hippocampal cytotoxic lesions. *Behav. Neurosci.* *113*, 1170–1188.
- Climer, J.R., and Dombeck, D.A. (2021). Information theoretic approaches to deciphering the neural code with functional fluorescence imaging. *eNeuro* *8*, ENEURO.0266-21.2021.
- 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.
- 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*, 1018–1026, e4.
- Sato, M., Mizuta, K., Islam, T., Kawano, M., Sekine, Y., Takekawa, T., et al. (2020). Distinct mechanisms of over-representation of landmarks and rewards in the hippocampus. *Cell Rep.* *32*, 107864.
- Ciocchi, S., Passecker, J., Malagon-Vina, H., Mikus, N., and Klausberger, T. (2015). Selective information routing by ventral hippocampal CA1 projection neurons. *Science* *348*, 560–563.
- Hok, V., Lenck-Santini, P.-P., Roux, S., Save, E., Muller, R.U., and Poucet, B. (2007). Goal-related activity in hippocampal place cells. *J. Neurosci.* *27*, 472–482.
- Duvelle, É., Grieves, R.M., Hok, V., Poucet, B., Arleo, A., Jeffery, K.J., and Save, E. (2019). Insensitivity of place cells to the value of spatial goals in a two-choice flexible navigation task. *J. Neurosci.* *39*, 2522–2541.
- Danielson, N.B., Zaremba, J.D., Kaifosh, P., Bowler, J., Ladow, M., and Losonczy, A. (2016). Sublayer-specific coding dynamics during spatial navigation and learning in hippocampal area CA1. *Neuron* *91*, 652–665.
- Lee, J.S., Briguglio, J.J., Cohen, J.D., Romani, S., and Lee, A.K. (2020). The statistical structure of the hippocampal code for space as a function of time, context, and value. *Cell* *183*, 620–635, e22.

35. Gergues, M.M., Han, K.J., Choi, H.S., Brown, B., Clausing, K.J., Turner, V.S., Vainchtein, I.D., Molofsky, A.V., and Kheirbek, M.A. (2020). Circuit and molecular architecture of a ventral hippocampal network. *Nat. Neurosci.* *23*, 1444–1452.
36. Shpokayte, M., McKissick, O., Yuan, B., Rahsepar, B., Fernandez, F.R., Ruesch, E.A., Grella, S.L., White, J.A., Liu, X.S., and Ramirez, S. (2020). Hippocampal cells multiplex positive and negative engrams. *bioRxiv*. <https://doi.org/10.1101/2020.12.11.419887>.
37. Xu, C., Krabbe, S., Gründemann, J., Botta, P., Fadok, J.P., Osakada, F., Saur, D., Grewe, B.F., Schnitzer, M.J., Callaway, E.M., and Lüthi, A. (2016). Distinct hippocampal pathways mediate dissociable roles of context in memory retrieval. *Cell* *167*, 961–972, e16.
38. Jimenez, J.C., Su, K., Goldberg, A.R., Luna, V.M., Biane, J.S., Ordek, G., Zhou, P., Ong, S.K., Wright, M.A., Zweifel, L., et al. (2018). Anxiety cells in a hippocampal-hypothalamic circuit. *Neuron* *97*, 670–683, e6.
39. Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* *275*, 1593–1599.
40. Trouche, S., Koren, V., Doig, N.M., Ellender, T.J., El-Gaby, M., Lopes-dos-Santos, V., Reeve, H.M., Perestenko, P.V., Garas, F.N., Magill, P.J., et al. (2019). A hippocampus-accumbens tripartite neuronal motif guides appetitive memory in space. *Cell* *176*, 1393–1406, e16.
41. Sutton, R.S., and Barto, A.G. (2018). *Reinforcement Learning: An Introduction*, Second Edition (MIT Press).
42. Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B.L., Koay, S.A., Taxis, J., Najafi, F., Gauthier, J.L., Zhou, P., et al. (2019). CalmAn an open source tool for scalable calcium imaging data analysis. *eLife* *8*, e38173.
43. Resendez, S.L., Jennings, J.H., Ung, R.L., Nambodiri, V.M.K., Zhou, Z.C., Otis, J.M., Nomura, H., McHenry, J.A., Kosyk, O., and Stuber, G.D. (2016). Visualization of cortical, subcortical and deep brain neural circuit dynamics during naturalistic mammalian behavior with head-mounted microscopes and chronically implanted lenses. *Nat. Protoc.* *11*, 566–597.
44. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* *21*, 1281–1289.
45. Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* *89*, 285–299.
46. Harland, B., Contreras, M., Souder, M., and Fellous, J.-M. (2021). Dorsal CA1 hippocampal place cells form a multi-scale representation of mega-space. *Curr. Biol.* *31*, 2178–2190, e6.
47. Markus, E.J., Barnes, C.A., McNaughton, B.L., Gladden, V.L., and Skaggs, W.E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. *Hippocampus* *4*, 410–421.
48. Treves, A., and Panzeri, S. (1995). The upward bias in measures of information derived from limited data samples. *Neural Comput.* *7*, 399–407.
49. Panzeri, S., Senatore, R., Montemurro, M.A., and Petersen, R.S. (2007). Correcting for the sampling bias problem in spike train information measures. *J. Neurophysiol.* *98*, 1064–1072.
50. Akrami, A., Kopec, C.D., Diamond, M.E., and Brody, C.D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* *554*, 368–372.
51. Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychon. Bull. Rev.* *14*, 779–804.
52. Keysers, C., Gazzola, V., and Wagenmakers, E.-J. (2020). Using Bayes factor hypothesis testing in neuroscience to establish evidence of absence. *Nat. Neurosci.* *23*, 788–799.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Processed calcium imaging and behavioral data	This paper	Mendeley Data: https://doi.org/10.17632/km4cdcvyfs.1
Experimental models: Organisms/strains		
C57BL/6J-Tg(Thy1-GCaMP6f)GP5.5Dkim/J	The Jackson Laboratory	Cat. #: 024276; RRID: IMSR_JAX:024276
Software and algorithms		
R 3.6.3	R Project for Statistical Computing	RRID: SCR_001905
MATLAB 2019b	MathWorks	RRID: SCR_001622
CalmAn	⁴²	RRID: SCR_021533
Original code for calcium signal processing	This paper	Zenodo: https://doi.org/10.5281/zenodo.5716279
Original code for calculating spatial metrics and implementing a Bayesian decoder	This paper	Zenodo: https://doi.org/10.5281/zenodo.5716283
Original code for data analysis	This paper	Zenodo: https://doi.org/10.5281/zenodo.5716271

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Ole Paulsen (op210@cam.ac.uk).

Materials availability

This study did not generate new unique reagents.

Data and code availability

All original code used for the analysis is publicly available: (1) for processing calcium signal using CalmAn from Zenodo: <https://doi.org/10.5281/zenodo.5716279>; for calculating spatial metrics and implementing a Bayesian decoder from Zenodo: <https://doi.org/10.5281/zenodo.5716283> for generating the figures from Zenodo: <https://doi.org/10.5281/zenodo.5716271>. Processed calcium imaging and behavioral data are publicly available from Mendeley Data: <https://doi.org/10.17632/km4cdcvyfs.1>. Due to their large size, raw data files will be shared by the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Mice

Thirteen adult male Thy1-GCaMP6f transgenic mice were used in this study (Jax: 024276). Mice were housed with 2-4 cage-mates in cages with running wheels in a 12:12 h reverse light cycle. All animal experiments were performed under the Animals (Scientific Procedures) Act 1986 Amendment Regulations 2012 following ethical review by the University of Cambridge Animal Welfare and Ethical Review Body (AWERB) under personal and project licenses held by the authors.

METHOD DETAILS

Surgery

Mice underwent two surgeries: the first one to implant a GRIN lens directly above the cells of interest, and the other to fix an aluminum baseplate above the GRIN lens for later attachment of the miniature microscope. The procedures followed the protocol as described in Resendez et al.⁴³

Surgeries were carried out following minimal standard for aseptic surgery. Meloxicam (2 mg kg⁻¹ intraperitoneal) was administered as analgesic 30 min prior to surgery initiation. Mice were anesthetized with isoflurane (5% induction, 1-2% maintenance, Abbott Ltd, Maidenhead, UK) mixed with oxygen as carrier gas (flow rate 1.0-2.0 L min⁻¹) and placed in a stereotaxic frame (David Kopf Instruments, Tujunga, CA, USA). The skull was exposed after skin incision and Bregma and Lambda were aligned horizontally. A

craniotomy was drilled above the implantation site. For the dCA1, the craniotomy was 1.5–2 mm in diameter. The cortical tissue and 2 layers of corpus callosum fibers above the hippocampal implantation site were aspirated. Saline was applied throughout the aspiration to prevent desiccation of the tissue. A GRIN lens (1 mm diameter, 4.3 mm length, 0.4 pitch, 0.50 numerical aperture, Grintech) was stereotaxically lowered at coordinates -1.75 AP, 1.75 ML, 1.35–1.40 DV (in mm from Bregma) and fixed to the skull surface with ultraviolet-light curable glue (Loctite 4305) and further fixed with dental adhesive (Metabond, Prestige Dental) and dental acrylic cement (Simplex Rapid, Kemdent). A metal head bar was attached to the cranium using dental acrylic cement for head-fixing the animal during the microscope mounting. For the iCA1 implanted mice, a 0.9 mm diameter hole was drilled, and no tissue was aspirated. The GRIN lens (0.6 mm diameter, 4.95 mm length, 1.0 pitch, 0.5 numerical aperture, Grintech) was lowered inside a 21 gage needle using a custom-made stereotaxic guide that allowed a precise placement of the lens. The lens was placed at coordinates -3.16 AP, 3.6–3.8 ML, 3.40–3.70 DV and the needle guide was retracted allowing for fixation of the lens to the skull surface. After the surgery, the mice were monitored daily for 5 days and given oral Meloxicam as analgesic.

If the GCaMP6f expression was visible in the implanted mice, 4 weeks later the animals were anesthetized for the purpose of attaching a baseplate for the microscope above the top of the GRIN lens. The baseplate was cemented into place and the miniscope was unlocked and detached from the baseplate.

Histological processing

Following the behavioral experiments, animals were terminally anesthetized by intraperitoneal injection of pentobarbital (533 mg kg⁻¹) and then transcardially perfused with phosphate-buffered saline (PBS) followed by 4% paraformaldehyde (PFA). Brains were removed and post-fixed for 24–48 hours, then rinsed and subsequently cryoprotected overnight in 30% (w/v) sucrose dissolved in phosphate-buffered saline (PBS). Coronal sections of the hippocampus were cut using a microtome (Leica) with 80–100 μm thickness.

After rinsing in PBS, the sections were mounted in Fluoroshield with DAPI (Sigma). Sections were examined with a Leica Microsystems SP8 confocal microscope using the 10× and 20× magnification objectives.

Cheeseboard maze task

The mice performed a rewarded spatial navigation task on a 120 cm diameter cheeseboard maze² with 177 evenly spaced wells. The rewarded wells were baited with ~100 μL of condensed milk mixed 1:1 with water.

For the first three days, the mice foraged for rewards baited in randomly selected wells. The mice explored the cheeseboard in three or four trials for a total of 30 minutes per day. A different, random set of wells was baited in each trial.

Next, we performed a spatial learning task. The mice had to learn two locations with baited wells. The baited wells had fixed locations that were at least 40 cm apart chosen pseudo-randomly for each mouse. Mice started the trial in one of three locations on the maze: south, east or west. The maze was rotated and wiped with a disinfectant (Dettol) between each trial to discourage the use of intra-maze cues. Landmarks of black and white cues were installed on the walls surrounding the maze. The trials were terminated once the mice had consumed both rewards, or after 300 s, whichever was sooner. Each learning day consisted of 8 trials with 2–4-minute-long breaks between the trials.

After the first 5-day-long learning period, memory retention was tested on the next day in a 4-to-5-minute-long unbaited test trial. The trial was started from a previously unused starting position (north). The performance was measured by the number of reward zone crossings counted when the mouse crossed a circular zone within 20 cm from either of the reward locations. The number of crossings was normalized by the total distance the animal had moved.

Following the learning sessions and memory retention test for the first set of locations, one of the two reward locations was translocated. The new location was pseudo-randomly chosen to be at least 40 cm away from the current and previous reward locations. The learning of the new sets of locations was performed over two days and tested in an unbaited test trial the following day as described above.

The trials were recorded with an overhead webcam video camera at 24 Hz frame rate. The mouse body location was tracked with DeepLabCut software,⁴⁴ and custom-written software was written to map the mouse coordinates to the relative location on the maze. The extracted tracks were smoothed by applying locally weighted scatterplot smoothing (LOWESS) which used a moving average of coordinates in 15 video frames. Periods of running were identified when the running speed smoothed with a moving average 0.5 s window exceeded 4 cm/s.

Calcium imaging

Calcium imaging was acquired using Miniscope – a head-mounted microscope²⁰ (v3 and v4 Miniscope). A blue LED was used for excitation (~470 nm spectral peak) with power adjusted to approximately match the mean brightness of the image across animals. Fluorescence was passed through an emission filter (bandpass filter, 525/50 nm) and collected by a CMOS imaging sensor. Before the start of the recording, the mouse was head-fixed on a running wheel to attach the microscope and adjust its focal plane so it matched the field of view from the previous recordings. Afterwards, the mouse with the Miniscope attached was placed in a start box for 3–5 minutes before the recording session started. The calcium imaging was acquired at 20 Hz and started synchronously with WebCam camera recording.

Calcium signal processing

CalmAn software was used to motion-correct any movements between the calcium imaging frames, identify the cells and extract their fluorescence signal from the video recordings.⁴² The method for cell and signal detection was based on constrained non-negative matrix factorization.⁴⁵ CalmAn extracted background-subtracted calcium fluorescence values and deconvolved the signal. The deconvolved signal can be interpreted as a scaled probability of a neuron being active. The calcium imaging videos recorded in the same-day trials were motion corrected to a common template frame and were concatenated. Signal extraction and further processing was performed on the resulting long video, allowing the detection of cells and signals present across the trials. To improve the computational performance, the videos were cropped to a rectangle containing the imaged cells and the video width and height was downsampled by a factor of 2.

The identified putative cells were automatically filtered using CalmAn. The results were visually inspected and the filtering parameters adjusted to exclude non-cell like shapes and traces from the filtered components. The criteria used for the filtering included a threshold for signal to noise ratio of the trace, the minimum and maximum size of the component's region of interest (ROI), threshold for consistency of the ROI at different times of the component's activation, and a threshold for the component's resemblance to a neuronal soma as evaluated by a convolutional neural network provided with CalmAn software.

The identity of cells between the recordings on different days was matched using a registration algorithm implemented in CalmAn. The algorithm aligned the image with ROI of cells from all days to the image from the reference day and matched the cells when their centers of mass were closer than 10 μm .

The deconvolved traces were smoothed in time with a Gaussian kernel ($\sigma = 75$ ms). The trace was time binned by averaging the values in 200 ms bins.

For the comparison between dCA1 and iCA1 activity, calcium event rates are reported. A calcium event was detected whenever the cell's deconvolved signal crossed 20% of its maximum value.

Place cell detection and analysis

To assess how spatial locations modulated the activity of a cell, we considered periods of running as described above and calculated place maps — mean neural activity per spatial bin. The total activity inside 6 x 6 cm bins was summed from the smoothed deconvolved signal. The mean neural activity in the spatial bin was then calculated as the ratio of the total activity to the total occupancy in the bin after both maps were smoothed across the space using a 2D Gaussian kernel with $\sigma = 12$ cm. The place map was filtered to include spatial bins with total occupancy that exceeded 1 s (5 time bins, thresholded on unsmoothed total occupancy). Because the size of place fields scales with the environment,⁴⁶ to facilitate comparison with other studies, we report field size as percentage of the maze area.

Spatial information of a cell's activity was calculated using the place map values. Spatial information was defined as:⁴⁷

$$\text{Spatial information} = \sum_{i=1}^N p_i \frac{\lambda_i}{\bar{\lambda}} \log_2 \left(\frac{\lambda_i}{\bar{\lambda}} \right) \quad (\text{Equation 1})$$

where $\bar{\lambda}$ represents the mean value of the neural signal, p_i represents probability of the occupancy of the i -th bin, and λ_i represents the bin's mean neural activity. Dividing by $\bar{\lambda}$ ensures the metric is independent of the cell's average activity. The units of spatial information calculated on calcium fluorescence can be reported as bits per action potential.²⁶ However, because the actual action potentials were not measured, we report them as arbitrary units.

Spatial information was compared with the value expected by chance. The chance level was calculated by circularly shifting the activity with regards to the actual location. For each cell, the activity was circularly shifted within the trial by a time offset chosen randomly (minimum offset 10 s for baited and 20 s for unbaited trials). If the cell's spatial information exceeded 95% values calculated on 1000 random shifts of its activity, it was defined as a place cell.

A limited number of neuronal responses sampled per spatial bin can lead to an upward bias in the estimated spatial information.⁴⁸ To correct for this bias, we report spatial information values relative to the mean spatial information from the time-shifting procedure used for place cell detection. This calculation does not require binning of the neuronal responses from the calcium imaging as required by analytical estimation,⁴⁹ and has been used previously to estimate mutual information bias.⁵⁰

We defined the field size as the fraction of a place map with values exceeding half the maximum value. Centers of place fields were identified in the place map by finding local maxima exceeding half the global maximum. The local maxima were restricted to be at least 25 cm apart and have at least one adjacent spatial bin exceeding half the global maximum.

Calculation of place fields at reward location

The center of mass for the field was calculated and used to report the field's distance from the reward locations. For place cells with multiple place fields, the distance to the closest reward was used. Fields ≤ 20 cm from the reward location were referred to as reward fields. For the count of reward locations where a place cell had a reward field, only cells that were classified as a place cell in at least half of the test trials were considered. The distribution of the expected counts was generated by a process that shuffled cell identities assigned to test trial place maps. The count of reward fields was summed for each resulting cell. A significantly higher fraction of cells with zero or many fields at reward location means their count exceeded that in 95% of the shuffles.

Bayesian decoders

Two Bayesian decoders were constructed from the neural activity: the first decoded spatial location of the running mouse, the second decoded whether the mouse was running inside a reward zone.

The decoders used binarized background-subtracted calcium fluorescence values F . The binarized trace had value 1 when the fluorescence exceeded the 90th percentile of the cell's values for that day (active cell); otherwise, the binarized trace had value 0 (inactive cell).

The Bayesian decoder assumed that the activity of the cells was independent given the output, and it chose the output to maximize posterior probability given the neural data:

$$\hat{s} = \operatorname{argmax}_s P(s) \prod_{i=1}^{n_{\text{cells}}} P(r_i | s) \quad (\text{Equation 2})$$

For decoding the mouse location during running, s represents the spatial bin, $P(s)$ represents the prior occupation probability in the spatial bin s , and $P(r_i | s)$ represents the probability of the i th cell being active in the spatial bin s .

For decoding whether the mouse was running in the proximity of learned reward location, s represents if the mouse was inside the reward zone (within 20 cm from the reward), $P(s)$ represents prior occupation probability inside or outside of the reward zone, and $P(r_i | s)$ represents the probability of the i -th cell being active inside or outside of the reward zone.

The decoders were trained and evaluated on two non-overlapping datasets:

- (1) The decoder for the spatial location was trained, and evaluated using a cross-validation method as follows: The day's session was split into five equal parts. A single part was reserved for evaluation and the others for training the decoder. The decoder was trained and evaluated, and the process was repeated five times, each time with a different part of the data reserved for evaluation. The decoder was compared with a baseline random decoder which predicted spatial location based on prior occupancy probabilities. The decoder errors were reported as a distance between the actual and the predicted spatial bin. Because fewer cells were recorded in the iCA1 than in the dCA1, we compared decoders trained on equally sized populations by randomly sampling 30 cells from each recorded session (72% of the iCA1 recordings had more than 30 cells). The spatial decoding from equally sized neuronal populations was repeated 50 times with different cell samples.
- (2) The decoder predicting whether the mouse was inside a reward zone was trained on data from two unbaited test trials, which were performed on different days and shared a single learned reward location. The training dataset was filtered to times when the mouse was in proximity of the learned reward location (distance ≤ 15 cm), or the mouse was well-away from the reward location (distance ≥ 40 cm). The decoder was evaluated on data from another unbaited test trial. In this trial, one of the learned reward locations was different from those in the training dataset, and one of the learned reward locations was missing from the current ones. Only data from the proximity of either of these two locations was used for evaluation (the reward zone vs the previous reward zone). The evaluation was restricted to trials that shared at least 10 cells with the training trials. The decoder assumed equal prior $P(s)$ of the zones. The resulting decoder's performance was compared with a baseline random decoder.

Downsampled data comparison

To verify that differences in maze occupancy between foraging and test trials were not the reason for the observed accumulation of place cells at goal location, we randomly downsampled the data. For each spatial bin in the two sessions, an equally sized subset of timestamps was selected to match the lower of the two occupancies. The selected timestamps were used to construct place maps and to identify place cells. The random downsampling procedure was repeated 100 times, and the statistics about the place field locations and their distance to reward locations aggregated.

Population activity on reward approach

To analyze the population activity during an approach to reward locations, periods of running that exceeded a minimum duration of 3 s were used. In the baited test trials, the running bouts were aligned to the time when the tracked mouse body stopped within 7 cm from the reward. For the bouts stopping at non-rewarded locations, only the stops at distance >24 cm from the reward were included. In the unbaited trials, the running bouts were included if they crossed a location <18 cm from a learned reward location and covered a distance >12 cm. The deconvolved z-scored activity was aligned to the timestamp when the mouse was the closest to the learned reward location. The mean population z-scored activity was calculated for 1 s-long bins and the activity at 4–5 s before the bout finished was compared with the activity at 0–1 s.

QUANTIFICATION AND STATISTICAL ANALYSIS

Results are reported using two statistical methods. First, we estimated p-values using null hypothesis significance testing. The p-values are low for small effects assessed on large sample sizes; they depend on unseen data, and on the plan for how many animals to test experimentally.⁵¹ Therefore, we also report Bayes Factors⁵² — a measure of relative evidence for two competing hypotheses. It is calculated as a ratio of posterior probabilities: the probability of the alternative hypothesis given the observed data over the probability of the null hypothesis given the observed data. We assumed equal prior probabilities of the alternative and null hypothesis. In addition to providing further statistical support to significant p-values, Bayes Factor analysis gives evidence for the absence of

differences where the effects are non-significant.⁵² The magnitude of evidence was graded as inconclusive, moderate or strong following Jeffreys' thresholds.⁵²

Mixed-effects models were used for the statistical analysis to allow for unbalanced sampling and correlated samples. Both apply to these data, for example due to correlations between the samples of cell activity recorded at the same timestamp, or recordings from the same mouse on different trials. The effects were assessed with linear and log-linear mixed-effects models. The fixed effects were the statistically tested effects such as implant location (dCA1 vs iCA1) or cell type (place cell vs non-place cell); the random effects were modeled as mouse-specific and session-specific random variables. The random effects were also included in the estimation of the linear regression intercept.

For the frequentist approach, the model coefficients were estimated using the restricted maximum-likelihood method. The residual errors were checked for linear model assumptions: zero mean, no correlation with the predicted values and homoscedasticity. To satisfy these assumptions, some models used a log-linear transformation of the response variable. The significant effects and their interactions were reported and the post-hoc tests were performed on differences in least-square means of the paired groups. The tests used Satterthwaite estimation of degrees of freedom and adjusted p-values using Holm-Bonferroni correction.

For the Bayes Factor analyses, the mixed-effects models mirrored the frequentist models and had the same fixed and random effects. The priors were specified as Cauchy distribution with $\sqrt{2}/2$ scale for fixed effects and 0.5 scale for random effects. These priors follow the expectation that the differences between mice are smaller than the effects of interest. Bayes Factor for the effect of interest was calculated as probability of the full model over the probability of the model excluding the tested effect.

The effect sizes were reported with 95% credibility intervals (CI; equal-tailed interval). The interval can be interpreted as a range within which the effect falls with 95% probability given the evidence from the observed data. Credibility intervals were estimated from the samples of the model's posterior distribution.

Statistical analysis was performed in R version 3.6.3. The linear mixed-effects models were built in R with package 'lme4' and p-values for the fixed effects were obtained using Satterthwaite estimation of degrees of freedom implemented in the 'lmerTest' R package. Least-square means were calculated and tested with 'lsmeansLT' function from the same package. Bayesian linear mixed-effects models were created using 'BayesFactor' R package and 'lmBF' function.

Data are reported as mean \pm SEM unless otherwise stated. p-values < 0.05 were considered significant. Statistical details can be found in figure legends or supplemental tables.