

Object-centered population coding in CA1 of the hippocampus

Highlights

- Up to 620 CA1 cells were recorded simultaneously during foraging in rats
- The majority of cells changed their firing pattern when an object was introduced
- Population activity reorganized around the object in a distance-dependent manner
- The response to the object was distributed and did not depend on specific cells

Authors

Anne Nagelhus,
Sebastian O. Andersson,
Soledad Gonzalo Cogno,
Edvard I. Moser, May-Britt Moser

Correspondence

soledad.g.cogno@ntnu.no (S.G.C.),
edvard.moser@ntnu.no (E.I.M.),
may-britt.moser@ntnu.no (M.-B.M.)

In brief

Objects are crucial for guiding navigation. Nagelhus, Andersson, et al. show that CA1 cells change their spatial firing pattern in response to salient objects in the environment. At the population level, these changes are widely distributed and smoothly organized according to the animal's distance from the object.



Article

Object-centered population coding in CA1 of the hippocampus

Anne Nagelhus,^{1,3} Sebastian O. Andersson,^{1,2,3} Soledad Gonzalo Cogno,^{1,4,5,*} Edvard I. Moser,^{1,4,5,*} and May-Britt Moser^{1,4,5,6,*}

¹Kavli Institute for Systems Neuroscience and Centre for Algorithms in the Cortex, Norwegian University of Science and Technology (NTNU), Trondheim, Norway

²Present address: Max Planck Institute for Brain Research, Frankfurt, Germany

³These authors contributed equally

⁴These authors contributed equally

⁵Senior author

⁶Lead contact

*Correspondence: soledad.g.cogno@ntnu.no (S.G.C.), edvard.moser@ntnu.no (E.I.M.), may-britt.moser@ntnu.no (M.-B.M.)
<https://doi.org/10.1016/j.neuron.2023.04.008>

SUMMARY

Objects and landmarks are crucial for guiding navigation and must be integrated into the cognitive map of space. Studies of object coding in the hippocampus have primarily focused on activity of single cells. Here, we record simultaneously from large numbers of hippocampal CA1 neurons to determine how the presence of a salient object in the environment alters single-neuron and neural-population activity of the area. The majority of the cells showed some change in their spatial firing patterns when the object was introduced. At the neural-population level, these changes were systematically organized according to the animal's distance from the object. This organization was widely distributed across the cell sample, suggesting that some features of cognitive maps—including object representation—are best understood as emergent properties of neural populations.

INTRODUCTION

Objects and landmarks are crucial for guiding navigation and must be part of the brain's cognitive map of space.¹ A long history of work has shown that objects influence spatial firing in the hippocampus in a variety of ways. In place cells,^{2,3} the number, size, position, and directionality of firing fields may be changed when a salient object is introduced in the environment.^{4–9} The hippocampus has also been reported to have cells with specific firing at the location of the object.^{5,10} These “object cells” consistently move their firing fields along with the object when the object is moved.

One synapse upstream of the hippocampus, in the medial entorhinal cortex (MEC), specialized cells fire whenever the animal is at a fixed distance from the object in a given direction (i.e., the cell's firing location is given by a vector between the animal and the object).^{11,12} These “object-vector cells” fire independently of the object's shape, size, and identity. A similar set of cells, “landmark-vector cells,” has been reported in the hippocampus^{13,14} but these cells only began to fire at fixed distances and directions from objects after the animal gained experience with the object.¹³ However, limited amounts of data made it difficult to characterize the nature of these cells in a systematic manner.

Despite numerous studies on object coding in the hippocampus, the exact nature of object representation remains unclear. Early studies used limited quantitative criteria to identify object

cells and landmark-vector cells. Moreover, cells in the hippocampus often show complex responses that lack a simple relationship with the object location.^{7,15} In the present study, we considered the possibility that information about the presence and location of the object is expressed in the collective activity of the hippocampal population in ways that cannot be extracted from the firing of individual neurons.^{16–18} Thus, with access to spike activity from almost 1,200 unique principal neurons and more than 600 neurons recorded simultaneously, we set out to determine how the presence of objects influences population activity in the CA1 area of the hippocampus in freely moving rats. We report that population activity shows a strikingly smooth organization according to the animal's distance from the object.

RESULTS

Object cells and object-vector cells in the hippocampus

To investigate single-cell and population coding of objects in the hippocampus, we implanted tetrodes or Neuropixels probes in the CA1 of rats (7 and 2 rats, respectively; [Figure S1](#)) and recorded activity daily for four trials as the rats foraged in a 150 cm × 150 cm open field arena ([Figure 1A](#)). Each session began with an empty box trial (no object [NO]). Then, on the second trial (object [O]), a multi-colored Duplo tower 52 cm × 10 cm × 10 cm (height × width × length) was placed in the arena. On the third trial (object moved [OM]), the object was



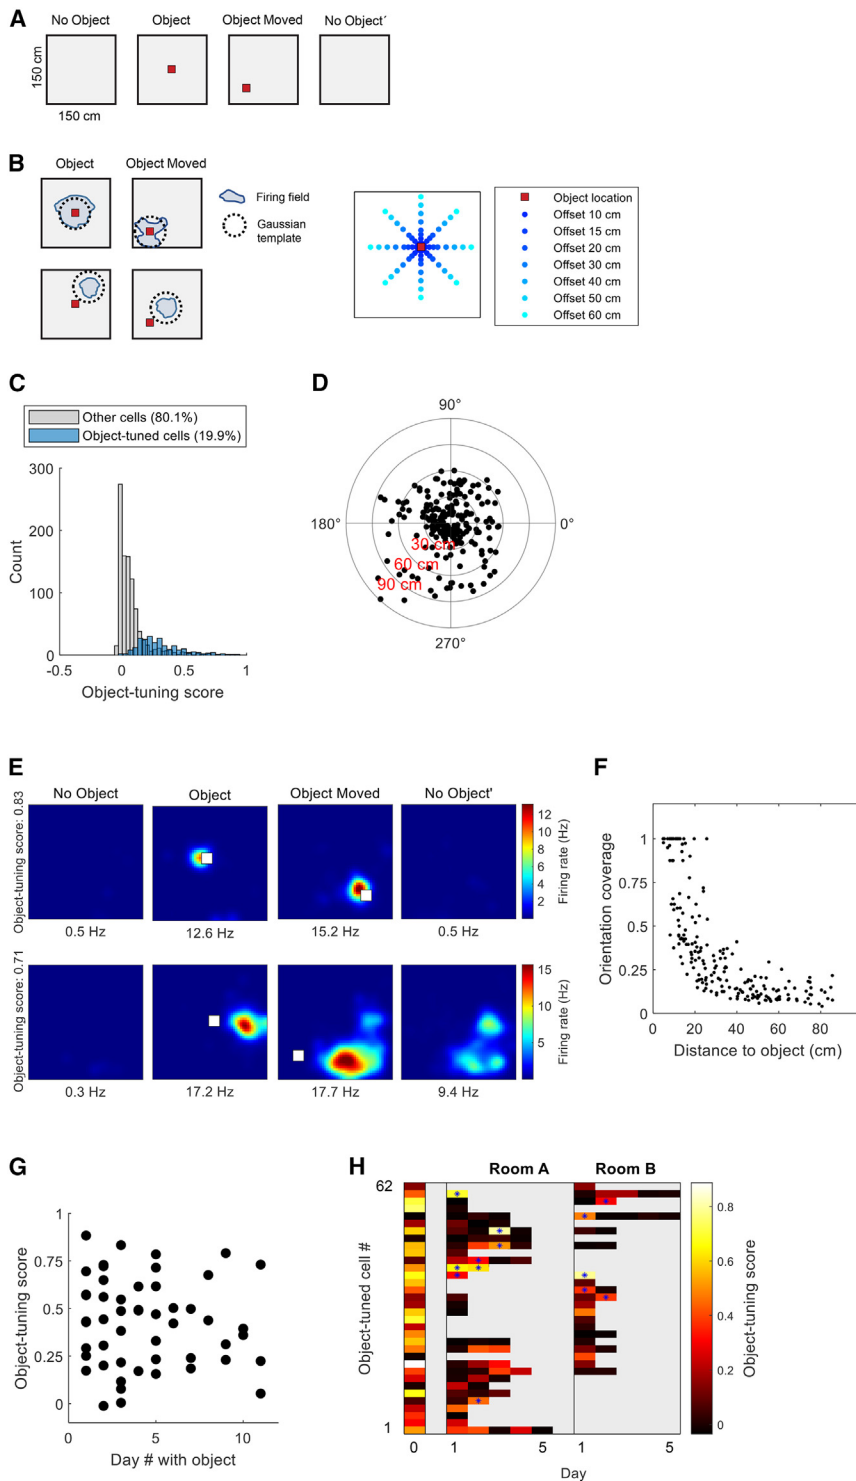


Figure 1. Object-tuned cells in CA1

(A) Schematic of experimental protocol. Each recording session consisted of four consecutive trials where rats foraged freely in the same 150 cm × 150 cm open field arena. In the two middle trials (object [O] and object moved [OM]), a Duplo tower was placed into the environment (red squares).

(B) Schematic of iterative template matching procedure that we used to identify cells that fired at the object (object cells; top row) or at fixed distances and directions from it (object-vector cells; bottom row). (Left) In each iteration, we placed a Gaussian template at a specific location in the O trial, moved it to the same location relative to the object in the OM trial, and then calculated the Pearson correlation between rate map and template in the two trials separately. (Right) We repeated this procedure with the template offset from the object location in different directions (8 offsets from 10 to 60 cm and 8 cardinal directions) and for 8 different variances of the Gaussian firing field (not shown). For each repetition, the minimum of the two Pearson correlations was taken as an “object-tuning score”.

(C) Distribution of object-tuning scores ($n = 1,189$ unique CA1 cells). Blue bars, cells classified as object-tuned cells; gray bars, cells not classified as object-tuned.

(D) Polar scatter plot showing orientation (polar axis, in degrees) and distance (radial axis, in cm, red) of object-tuned cells (one dot per cell).

(E) Color-coded firing rate maps from example object cell (top) and example object-vector cell (bottom). Object locations are indicated by the white square. Object-tuning score is indicated at the left and peak firing rate at the bottom of each rate map.

(F) Orientation coverage of firing fields of object-tuned cells as a function of distance to object. The nearer the cell’s firing field is to the object, the larger is the orientation coverage.

(G) Object-tuning score as a function of the number of days the animal had been exposed to the object (day #) for all cells recorded across different days and that were identified as object-tuned at least in 1 day ($n = 62$). Each dot indicates a cell.

(H) Color-coded matrix showing object-tuning scores as a function of recording day in two different rooms for 35 object-tuned cells that were recorded on 2 days or more. Left column, object-tuning scores the 1st day the cells were recorded (“day 0”). Middle and right column, object-tuning scores for the same cells on additional days in the same room (room A) or in a different room (room B). Only 11 of the 35 cells (31.4%) passed the criteria for object-tuned cells on more than 1 day (cells indicated by blue asterisks). Gray regions indicate that the cell was not recorded on the specific day. Bar, range of object-tuning scores (lighter colors, high score; darker colors, low score).

moved to a new, semi-random location, before it was taken out from the box for the last trial (no object’ [NO’]).

We first set out to identify individual cells with activity related to the presence of the object. We developed an iterative template

matching procedure (Figure 1B) that identified cells with object-related activity (“object-tuned cells”) regardless of whether they fired at the location of the object (object cells)^{5,10} or at fixed distances and directions from the object (object-vector or

landmark-vector cells).^{11–14} In this procedure, a Gaussian template was placed successively at 65 locations and at 8 different widths in each location in the O and OM trials (Figure 1B; STAR Methods). For each iteration, the correlation between rate map and template was determined for each trial. The smallest of the two correlations was then taken as the iteration's object-tuning score. We classified a cell as object-tuned cell if at least one of the 65×8 scores was larger than the 99th percentile of a shuffled distribution created by randomly choosing the template location in both trials.

Using the template matching procedure, 19.9% of CA1 cells were classified as object-tuned (237/1,189 cells) (Figures 1C, S2A, and S2B). This fraction was significantly higher than a chance level determined by running the same template matching procedure on control trials without objects (11/219 cells, 5.03%; 19.9% compared to chance level of 5.03%, $p < 5.73 \times 10^{-141}$, binomial test). The firing fields of the object-tuned cells covered a wide range of distances from the object (range: 4.8–80.7 cm; mean \pm SD: 37.7 ± 16.0 cm) and the entire azimuthal range (Figure 1D). We did not observe a correlation between the proximodistal recording location and the percentage of cells classified as object-tuned (Figures S1 and S2C). Object-related activity was present also in CA3^{13,19} (Figures S1 and S3).

To enable a rough comparison between the present dataset and previous literature, we estimated that 4.1% of all recorded cells were object cells (49/1,189 cells, object-tuned cells with firing fields <15 cm from the object, Figure 1E top) and 13.3% were object-vector cells (158/1,189 cells, object-tuned cells with firing fields ≥ 15 cm distance from the object, Figure 1E bottom) (note that in 30 object-tuned cells, firing fields could not be detected because the object-related activity was not confined to contiguous regions—these 30 cells were therefore not classified as object cells or object-vector cells; see STAR Methods for details). While object cells have fields covering all orientations around the object (Figure 1F),^{5,10} the data suggest a continuum between object cells and object-vector cells (Figures 1E and 1F). We observed no significant effect of day of exposure to the object on the object-tuning score (Figure 1G) ($r = -0.080$, $p = 0.57369$, Pearson correlation).

We next asked if the object responses were stable across time and environments. A total of 35 object-tuned cells were recorded on more than 1 day (in room A and room B) (Figure 1H). Only 13 of these cells (31.4%) had significant object-tuning scores on more than 1 day (Figure 1H). However, object-tuned cells were more stable than expected by chance, with higher probability of passing the classification on later days compared to other cells (13/35 object-tuned cells vs. 36/370 other cells, $p = 3.35 \times 10^{-4}$, binomial test). In addition, during the course of a single trial, cells with larger object-tuning scores tended to be more stable than cells with lower scores (Figure S2E) (Pearson correlation between within-trial stability and object-tuning score, $r = 0.45$, $p = 3.65 \times 10^{-13}$), pointing to the existence of a discrete cell class with transient object-tuning.

The majority of hippocampal cells change spatial firing patterns in the presence of objects

Many CA1 cells that did not pass criteria for object-tuned cells still responded in some way or other to the object (Figure 2A). In

the cell sample as a whole, introducing the object led to larger changes in spatial firing patterns (i.e., smaller Pearson correlations of rate maps) than in pairs of control trials with no object (Figures 2B, 2C, and 2D, top left; NO-O vs. CON1-CON2: $z = 12.17$, $p = 4.75 \times 10^{-34}$, Wilcoxon rank-sum test). While firing locations changed, the overall rate change, estimated as the mean firing rate in the less active trial divided by mean firing rate in the more active trial (rate overlap), was not different between trial pairs with objects (NO-O) and control trials (CON1-CON2) (Figure 2D, top right; $z = 0.57$, $p = 0.57$, Wilcoxon rank-sum test). Using a shuffling procedure to determine chance levels (see STAR Methods), we estimated that 67.5% of the cells ($n = 803/1,189$ cells) expressed either spatial change, rate change, or both between NO and O. Considerable change was also observed between the O and OM trials (Figure 2D, second row; O-OM vs. CON2-CON3; spatial correlation: $z = 13.23$, $p = 5.87 \times 10^{-40}$; rate overlap: $z = 3.43$, $p = 0.001$, Wilcoxon rank-sum tests). In these trials, 69.7% of the cells changed their spatial firing pattern, firing rate, or both when the object was moved ($n = 829/1,189$ cells). Major changes were evoked also when we removed the object (Figure 2D, third row; OM-NO' vs. CON3-CON4; spatial correlation: $z = 11.85$, $p = 2.11 \times 10^{-32}$; rate overlap: $z = 1.79$, $p = 0.074$, Wilcoxon rank-sum tests). Upon removal of the object, many cells retained firing fields from the preceding trial, consistent with previous observations of hysteresis in CA1 place cells^{13,20,21} (Figures 2A and S4). In line with this observation, major changes in spatial firing were also observed between NO and NO' (Figure 2D, fourth row; NO-NO' vs. CON1-CON4; spatial correlation: $z = 11.85$, $p = 2.11 \times 10^{-32}$, Wilcoxon rank-sum test).

We observed extensive changes in firing locations and rates also in CA3 (Figures S3D–S3F).

Decoding the presence of an object from neural-population activity

To determine the extent to which information about object location is present in the collective firing of hippocampal neurons, we first reasoned that if hippocampal population activity contains object-related information, we would be able to use population activity to decode whether an object is present in the environment. For all decoding analyses, we used recordings from 15 tetrode sessions with 25–39 cells in each recording and 2 NeuroPixels sessions with 620 and 135 cells, respectively, unless otherwise indicated (total of 4 rats). Building upon an existing Bayesian decoding algorithm,²² we developed an approach where we fed data from the NO trial and the O trial into a decoder, which classified testing data into either of these two types of trial (Figure 3A). The decoding accuracy was above chance (50%) in all experimental sessions (Figure 3B), showing that we can infer the presence of the object using population activity. The decoder's performance increased approximately logarithmically with the bin size of the testing data vector (Figure 3C) and the success of decoding increased with the number of cells included (Figure 3D). We chose 1 s bins for the main analyses because performance saturated around this level (Figure 3C). As expected from the trace analysis, we were also able to decode the NO trial from the NO' trial using population activity (Figure S5A).

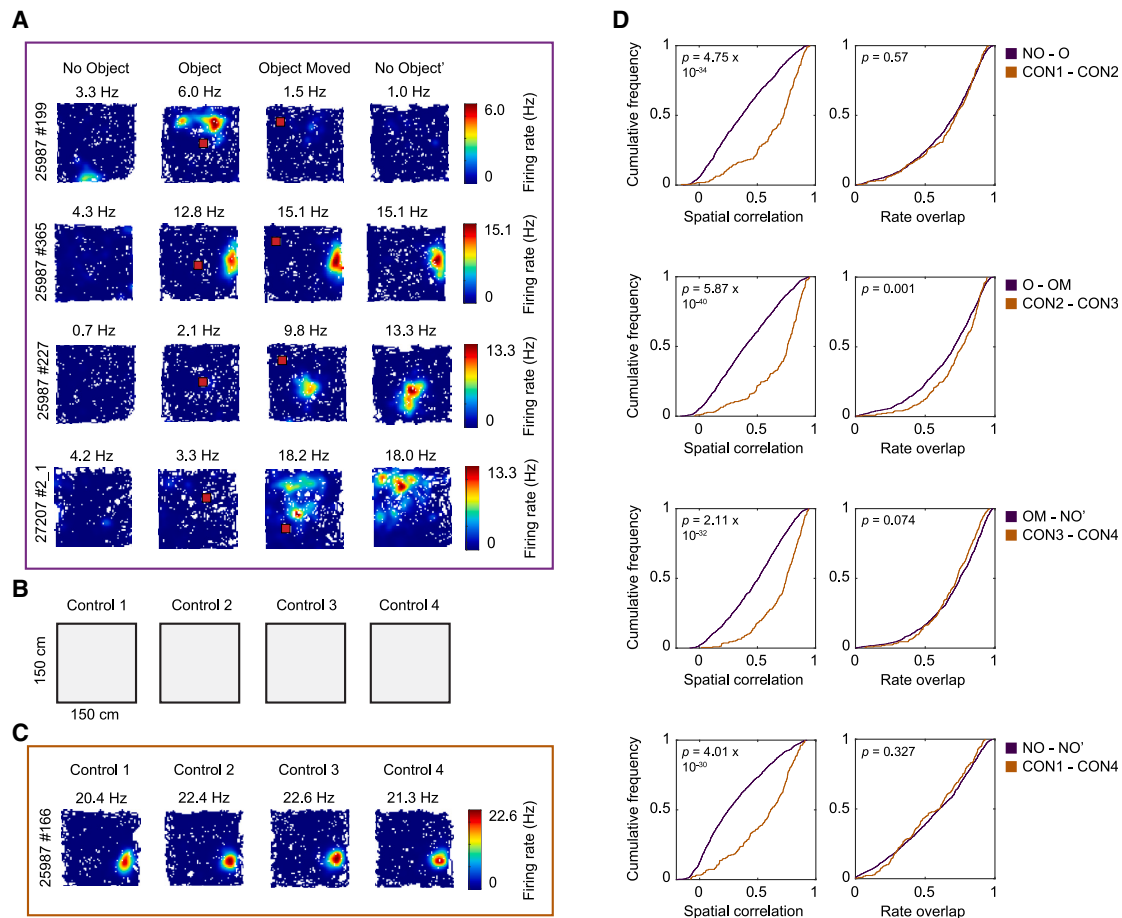


Figure 2. CA1 cells change spatial firing patterns in response to objects

(A) Color-coded firing rate maps of 4 example cells recorded in CA1 (experimental protocol as in Figure 1A). The red square indicates the object. The first cell changes its spatial firing pattern only upon introduction of the object (top row). The second cell changes its spatial firing pattern when the object is introduced but it preserves the firing pattern during the next trials, leaving a trace of the original firing field (second row). The third cell changes its firing pattern when the object is moved and maintains a trace on the last trial (third row). The fourth cell changes its firing pattern when the object is moved, but the spatial firing pattern is more diffuse than in the other cells (last row). Rat ID and cell ID are indicated to the left of the firing rate maps, and peak firing rate in each trial is indicated at the top of each rate map. Bars, color-coded firing rate.

(B) Control experiments used to compare object-induced changes to changes that occur due to passage of time. Four consecutive trials were run with rats foraging freely in the same 150 cm × 150 cm empty open field arena.

(C) Color-coded firing rate map from an example CA1 cell in the control experiment.

(D) Cumulative normalized frequency distributions showing spatial correlation (left column) and rate overlap (right column) between pairs of trials. Each row shows a given pair of trials in sessions with objects (as in Figure 1A; dark purple lines; $n = 1,189$ principal cells recorded from 9 rats) and the corresponding pair of control trials (as in Figure 2B; brown lines; $n = 219$ principal cells recorded from 3 rats). NO, no object; O, object; OM, object moved; NO', no object'; CON 1–4, control 1–4.

Having established as a proof-of-concept that we can use population activity to decode the presence of the object, we next asked if decoding accuracy differed as a function of the animal's spatial relationship with the object. We divided the environment into 5 distance bins, consisting of concentric rings radiating out from the object (Figure 4A). By feeding data from each distance bin into the decoder, we found that decoding performance decreased gradually with increasing distance from the object (Figure 4B) ($H = 44.559$, $p = 4.91 \times 10^{-9}$, Kruskal-Wallis test, $n = 16$ experimental sessions; significantly different decoding accuracies for distance bins 1 vs. 2; 2 vs. 3; and 4 vs. 5, $W \geq 111$, $p \leq 0.0049$, Wilcoxon signed-rank tests). Such a decrease

was also seen in the individual experiments (Figure 4C). In the main analyses, we optimized distance bins for each experiment so that we had an equal number of samples in each bin. We found the same decrease in decoding accuracy using fixed distance bins rather than optimized bins (Figure S5B) (fixed bins of 0–20, 20–40, 40–60, 60–80, and 80–100 cm).

To test if the reorganization was equally large in all directions relative to the object, we divided the environment into 5 angular bins, with each bin covering a different set of allocentric angles relative to the object (Figure 4D). The decoding accuracy was the same in different angular bins (Figures 4E and 4F) ($H = 3.75$, $p = 0.4416$, Kruskal-Wallis test, $n = 16$ experimental

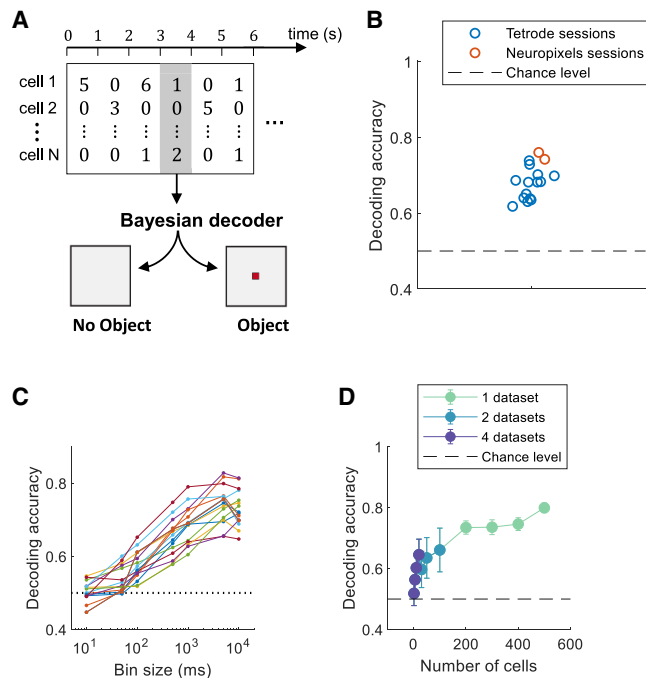


Figure 3. Decoding the presence of an object in the animal's environment

(A) Schematic of the conceptual framework. A Bayesian decoder infers whether a 1 s sample of population data (highlighted in gray) occurred in the absence or presence of an object (bottom left and right, respectively). (B) Decoding accuracy in all sessions ($n = 16$). The decoding accuracy is given by the fraction of samples that the decoder classified correctly. (C) Decoding accuracy as a function of bin size (ms), expressed on a logarithmic scale. The bin size was 10^1 , 5×10^1 , 10^2 , 5×10^2 , 10^3 , 5×10^3 , and 10^4 ms. Each color shows a different experimental session. (D) Decoding accuracy as a function of number of cells used to perform the decoding. Cells were randomly subsampled 100 times for each specific n and dataset. Data points show means and error bars show standard deviations (calculated on decoding accuracies pooled from all datasets). Stippled lines indicate the chance level (0.5). Color scale shows how many datasets (i.e., rats) contribute to each datapoint. 2 datasets are from tetrode rats (between 25 and 39 cells) and 2 datasets are from Neuropixels rats (135 and 620 cells, respectively).

sessions). This shows that the distance-dependent reorganization of population activity is general and not specific for certain compass directions.

To ensure that the result reflects a genuine phenomenon in the population activity we performed several controls. First, we observed no decrease in decoding accuracy across the same distance bins when attempting to decode the first half of the NO trial from the second half of the NO trial (Figure S5C) or when performing decoding in NO control trials (Figures 2B and S5D). Second, we randomly shuffled the trial labels of the spike count vectors (labels being “NO” and “O”), destroying the association between spike count vectors and trials. On the shuffled data, the decoder performed with a mean accuracy of 50% in all distance bins (Figure S5E). Third, we partitioned the data into training and testing in a random manner (in the standard decoder the first half of the session was used for training and the second half for testing and then vice versa). The results

from this decoder showed the same decline in decoding accuracy as a function of distance from the object (Figure S5F). Finally, across all cells, firing fields were widely distributed across the environment and not concentrated at the object location (Figures S5G and S5H), whereas firing rates were similar between trials, being largest at middle distances from the object (peaking in distance bin 3; Figure S5I). These observations indicate that firing field locations and firing rates cannot trivially explain our findings.

To further visualize the hippocampal object representation, we took the entire stack of rate maps from the NO trial and O trial and computed the pairwise Pearson correlation between the two stacks. This gives a correlation value (from -1 to 1) for each spatial bin (size: $3 \text{ cm} \times 3 \text{ cm}$). The correlation maps had low correlations in the central regions near the object, but high correlations in peripheral regions further away from the object (Figure 4G). Next, we binned the environment into 5 distance bins and collected the correlation values in each bin (number of correlation values for each bin: 489 ± 27 ; mean \pm SEM; $n = 16$ experimental sessions). This showed a stepwise linear increase in correlation values as a function of distance from object, with saturation at larger distances (Figure 4H) ($H = 30.1$, $p = 4.67 \times 10^{-6}$, Kruskal-Wallis test; significantly different correlation values for distance bins 1 vs. 2; 2 vs. 3; and 3 vs. 4, $W \leq 19$, $p \leq 0.0113$, Wilcoxon signed-rank tests). Altogether, these findings suggest that the object acts as a pivot around which the hippocampal population activity reorganizes in a distance-dependent manner. The internal walls of the recording environment did not play a similar role as the free-standing object, since position decoding accuracy and population vector correlations were comparable near and far away from the walls (Figures S5J and S5K).

Behavioral analyses

We next reasoned that if the animals' behavior is different between the NO and O trials, and increasingly different near the object, this could also produce decoding accuracy that falls with object distance (Figure 4B). To ask if such a behavioral confound was present in the previous analyses, we developed decoders that attempted to infer the trial (NO vs. O) based on the animal's behavior (see STAR Methods). Decoders for speed, head direction, and occupancy performed only slightly above 50% correct, with accuracies independent of object distance (Figures S6A, S6C, and S6E). We verified that all behavioral decoders worked: when simulating speed, head direction, or occupancy that was clearly different between the NO and O trials in all distance bins, the decoders performed with virtually 100% accuracy (Figures S6B, S6D, and S6F, left). Similarly, when simulating speed, head direction, or occupancy that was increasingly different near the object, the decoders had distance-dependent accuracies (Figures S6B, S6D, and S6F, right). That the behavioral decoders show flat accuracies in the experimental data (Figures S6A, S6C, and S6E) is therefore a strong indication that variations in behavior are not sufficient to explain our findings (Figures 4B and 4C).

We further reasoned that we could test for any role of behavior by artificially removing any behavioral differences between the trials prior to running the decoder on the neural data (example

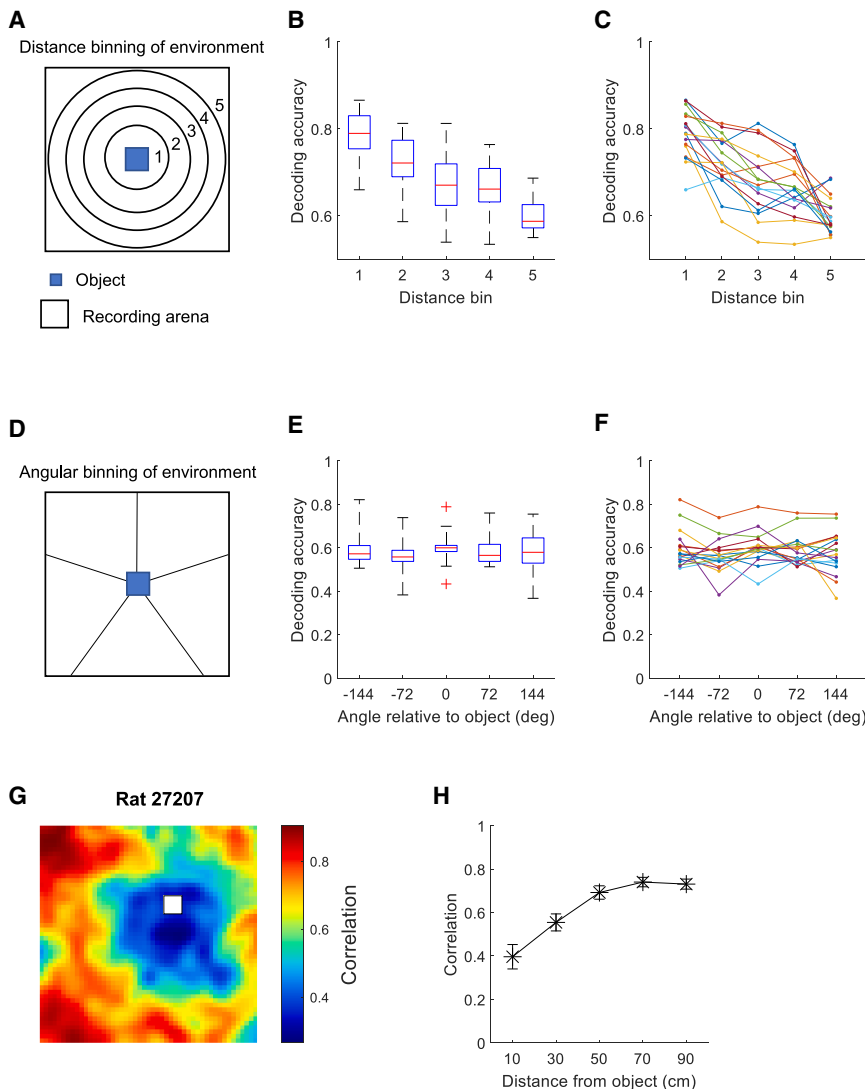


Figure 4. Graded reorganization of hippocampal population activity

(A) Schematic of partitioning of the environment for the main analysis. We divided the recording arena (150 cm × 150 cm) into 5 distance bins, forming circles around the object. We then applied the decoder to data from each distance bin to assess how decoding accuracy scales with distance from object.

(B) Boxplot of decoding accuracy as a function of distance from object. The central dashed line indicates the median, and the top and bottom lines indicate upper and lower and quartiles, respectively. The length of the whiskers indicates 1.5 times the interquartile range. Distance bin 1 is nearest the object; distance bin 5 furthest away. Each distance bin had an equal number of samples.

(C) Decoding accuracy as a function of distance from object for each individual experiment (n = 16). Data are the same as in the previous panel, but each color shows a different experimental session. Note that all experimental sessions have decreasing curves.

(D) Schematic of the partitioning of the environment for angular binning. We divided the arena into 5 angular bins, each bin spanning 72° and applied a decoder to each bin.

(E) Boxplot of decoding accuracy as a function of angle relative to object. Outliers exceeding 1.5 times the interquartile range are shown as crosses.

(F) Similar to (C) but for angle relative to object. All experimental sessions have flat curves.

(G) Correlation map showing similarity of time-averaged population firing on trials with and without the object (rat 27207). The correlations are closer to zero (near 0.3) in the vicinity of the object (white square) and more positive (beyond 0.8) in the periphery.

(H) Correlation between stacks of rate maps (NO trial vs. O trial) as a function of distance from the object. For each experimental session, we calculated a correlation map as in the previous panel. Then, we binned all values in the correlation map into 5 distance bins (based on distance to object) and calculated the mean correlation inside each distance bin. Plot shows mean ± SEM for all experimental sessions (n = 16) from all animals.

for speed distributions before and after removing behavioral differences in Figure S6G). We still observed a graded decrease in decoding accuracy for the neural activity decoder after matching distributions for speed (Figure S6H), head direction (Figure S6I) and occupancy (Figure S6J), confirming that behavioral differences between the trials were not necessary to obtain our results. The standard decoder used a prior that included information about the animal’s occupancy (see STAR Methods for details). As expected from the behavioral analyses, removal of this prior had no effect on our results (Figure S6K).

Finally, in the presence of the object there will be some regions that the animal cannot explore, preventing certain neural activity patterns from occurring in the O trial, which might lead to high decoding success near the object. To test that possibility, we included only spatial bins in which the animal had spent a mini-

mum amount of time (either 100, 500, 1,000, or 2,000 ms) in both the NO and O trials (i.e., excluding all spatial bins that the animal had not occupied for that duration in both trials). Regardless of the threshold value, the result was the same as before, with accuracy declining along with object distance (Figure S6L; note that to perform this analysis, the rate map bin size had to be increased from 3 to 10 cm²).

Objects are encoded by distinct patterns of joint activity in CA1 cells

We next asked if specific features of the population activity become different when the object is introduced. To parametrize the population activity, we introduced the concepts of “words” and “counts” as used previously to study population coding.^{23–25} We defined words as unique patterns of spiking and

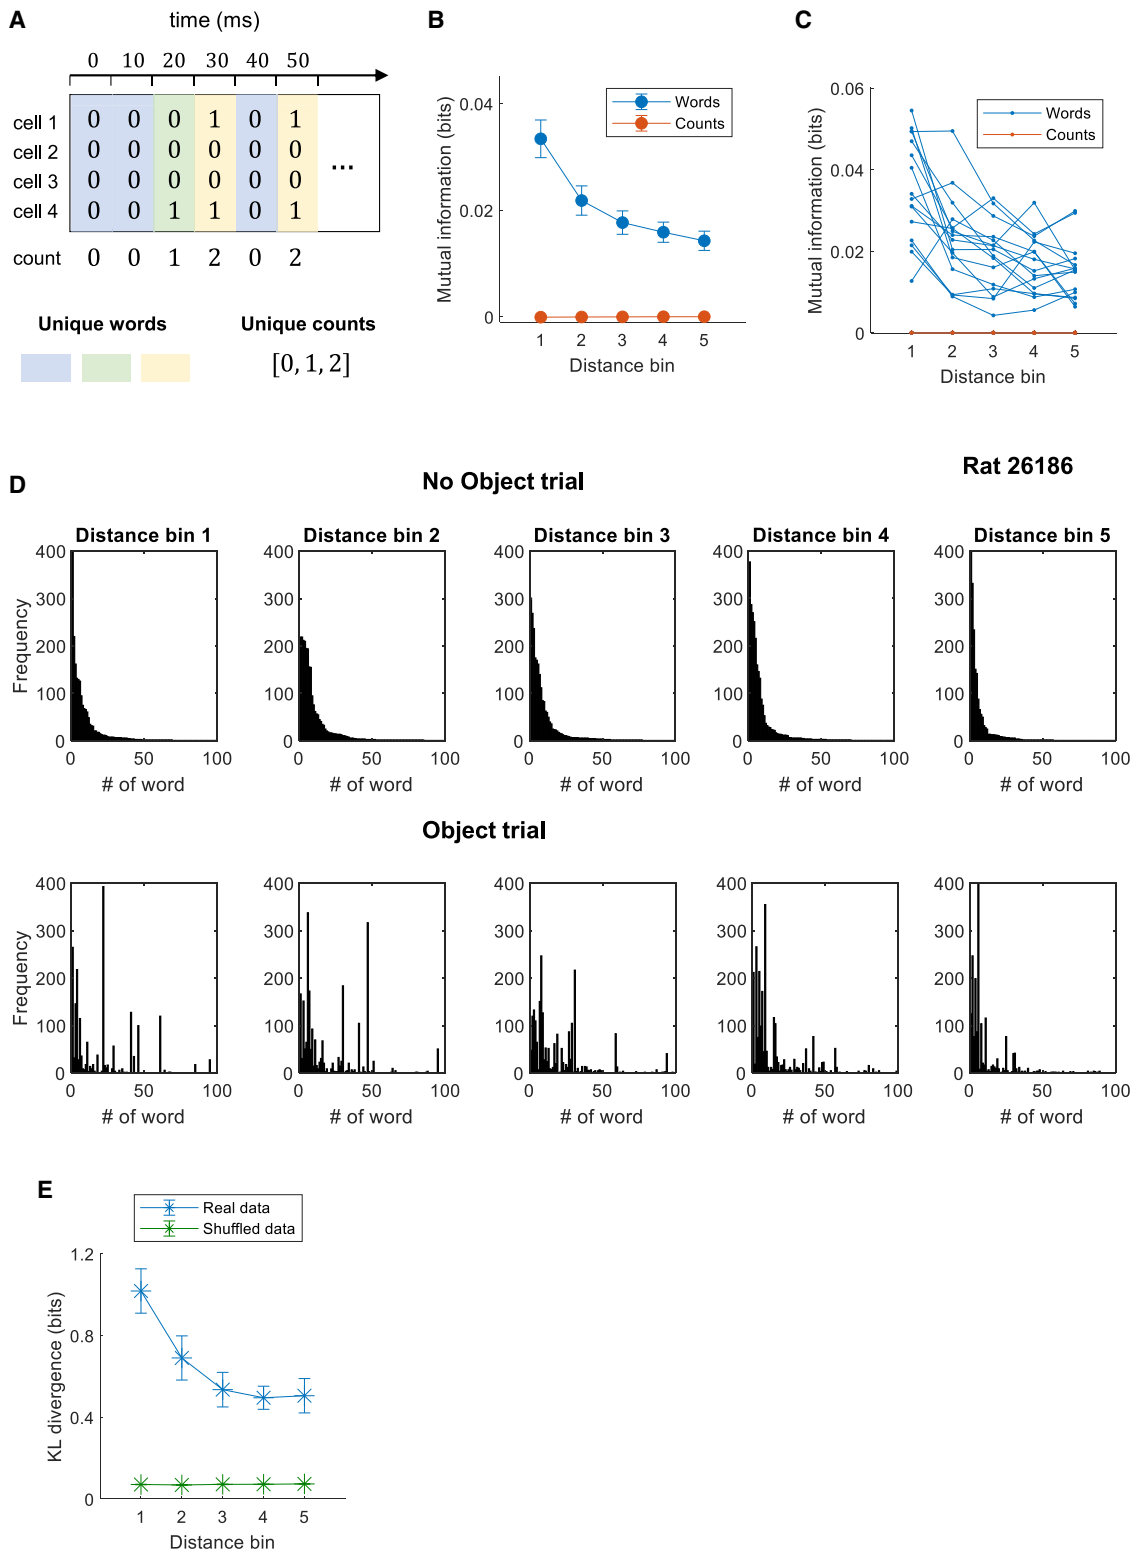


Figure 5. Objects are represented by distinct patterns of joint activity in the CA1 cell population

(A) Framework for analyzing population activity. We represent neural data in an $N \times T$ spike matrix, where N is the number of recorded cells and T is the number of time points in the trial sampled at 10 ms. We define a specific pattern of spiking across the neural population as a word. In the schematic, each unique word is

(legend continued on next page)

silence across neurons in the population in a single time bin (Figure 5A, time bins of 10 ms) and counts as the total sum of spikes across the population in a single time bin (Figure 5A, bottom row). We then calculated the mutual information (MI) between words and trial, and between counts and trial (see STAR Methods). Words were much more informative than counts (Figure 5B) ($U \geq 439$, $p \leq 1.19 \times 10^{-6}$, Wilcoxon rank-sum tests). Importantly, the amount of information gained from words decreased as a function of distance from the object (Figures 5B and 5C) ($H = 24.29$, $p = 6.98 \times 10^{-5}$, Kruskal-Wallis test), reproducing the pattern of the decoder. In contrast, the amount of information from counts did not depend on object distance ($H = 4.08$, $p = 0.3947$, Kruskal-Wallis test) and was not significantly above zero ($t \leq 1.46$, $p \geq 0.1636$ for all comparisons, t tests).

To visualize the word representation, we calculated how many times each word occurred in each trial and sorted the words by frequency in the NO trial (Figure 5D, top). By applying the same sorting in the O trial (Figure 5D, bottom) we could see which words were up- and down-regulated with objects present. Many low-frequency words in the NO trial were high-frequency words in the O trial, standing out as spikes (Figure 5D, bottom). No such spikes were present in shuffled versions of data (Figure S7A), suggesting that the NO and O trials recruit different patterns of neural activity.

Word frequency distributions appeared different near the object (Figure 5D, distance bin 1) but more similar further away from the object (Figure 5D, distance bins 2–5; note that the same pattern was obtained when we sorted both histograms according to word frequencies in the O trial rather than the NO trial; Figure S7B). Consistent with this, the Kullback-Leibler (KL) divergence between the NO and O distributions became smaller with increasing object distance (Figure 5E, blue curve) ($H = 16.22$, $p = 0.0027$, Kruskal-Wallis test; KL divergences in 1st distance bin vs. 5th distance bin: $W = 102$, $p = 6.10 \times 10^{-4}$, Wilcoxon signed-rank test). No decrease in the KL divergence was present in shuffled data (Figure 5E, green curve). In contrast to the identity of words, the number of words was independent of trial type and distance from the object (Figures S7C and S7D).

Because the number of observed states is much larger for words than counts, the entropy was higher for words than counts

(entropy of words: $1.65 \text{ bits} \pm 0.06$; entropy of counts: $0.84 \text{ bits} \pm 0.02$; mean \pm SEM). However, the difference in entropy could not trivially explain why words were more informative than counts, because after downsampling the number of states to be equal for words and counts, words were still much more informative (Figure S7E).

Distance-specific reorganization of the hippocampal map is widely distributed across the CA1 population

We next wondered to what extent the distance-based reorganization of the hippocampal map is distributed across the cell population. The reorganization may be driven by a dedicated subpopulation of cells with “strong tuning” to distance from the object, or it could be distributed across most of the cells, as an emergent property of the network, independently of the responses at the single-cell level.

To understand how many cells are necessary to see the reorganization, we first asked if it is present already at the level of single cells. To estimate the distance modulation of individual cells, we first ran the decoder (Figure 3A) with one cell at a time ($n = 1,189$ cells and runs of the decoder). Individual cells showed no obvious decrease in decoding accuracy as a function of distance, with medians just slightly above 0.5 (Figure 6A, compare to Figure 4B). We then calculated

$$\text{distance tuning} = \text{decoding accuracy bin \#1} - \text{decoding accuracy bin \#5} \quad (\text{Equation 1})$$

as a proxy for distance modulation of individual cells, where a more positive value implies stronger distance modulation. CA1 cells had a slight positive bias in distance tuning (mean: 0.02; for reference, the same mean calculated from populations is 0.19; Figure 4C) and nearly all cells were contained between the bounds -0.2 and 0.2 (98.7% of cells). By shuffling data (circularly shifting spike trains) and re-calculating the distance tuning, we observed that some cells had significant positive distance tuning (12.2%, $n = 145/1,189$ cells) (Figure 6B, red bars) and some cells had significant negative distance tuning (3.8%, $n = 45/1,189$ cells) (Figure 6B, yellow bars). However, while these cells were individually informative, they were dispensable: after removing cells with significant positive distance tuning (12.2%), we still observed a gradual drop in decoding accuracy

labeled by a different color. We define counts as the total number of spikes in the neural population in any given time bin. The total count is illustrated in the bottom row.

(B) Mutual information between trial type (NO or O) and population response (either words, in blue, or counts, in orange) as a function of distance from object. Distance bins were constructed to contain the same number of samples, as in Figures 4B and 4C. Data points show mean \pm SEM ($n = 16$ experimental sessions, 14 from tetrodes data, 2 from Neuropixels data where the number of cells was subsampled to 30, see STAR Methods for further details). Distance bin 1 is nearest the object; distance bin 5 furthest away. Words but not counts are informative about trial type, and they are more informative if activity is sampled near the object.

(C) Mutual information from words (blue) or counts (orange) for each individual experiment. Data and analysis are as in the previous panel. Each curve shows a different experimental session. Notice that all orange curves overlap.

(D) Histograms of word frequencies from the NO trial (top) and the O trial (bottom) of an example session with 26 simultaneously recorded cells. Each histogram shows how many times each unique word appeared during each trial, in each distance bin. The words are indexed from 1 to 100 (only the 100 most frequent words in the NO trial are plotted). All distributions are sorted according to how frequently each word occurred in the NO trial. For this reason, distributions from the NO trial are decreasing (top) while distributions from the O trial have more arbitrary shapes (bottom). This also means that indices of words match between top and bottom (word #1 at the top is the same as word #1 at the bottom).

(E) Kullback-Leibler (KL) divergence between word distributions in the O (D, bottom) and NO (D, top) trial as a function of distance to object. Real data use the true trial labels of the words (blue curve); in shuffled data, trial labels were randomly assigned to each word (green curve). Word distributions were sorted by word frequencies in the NO trial (as described in the previous, D).

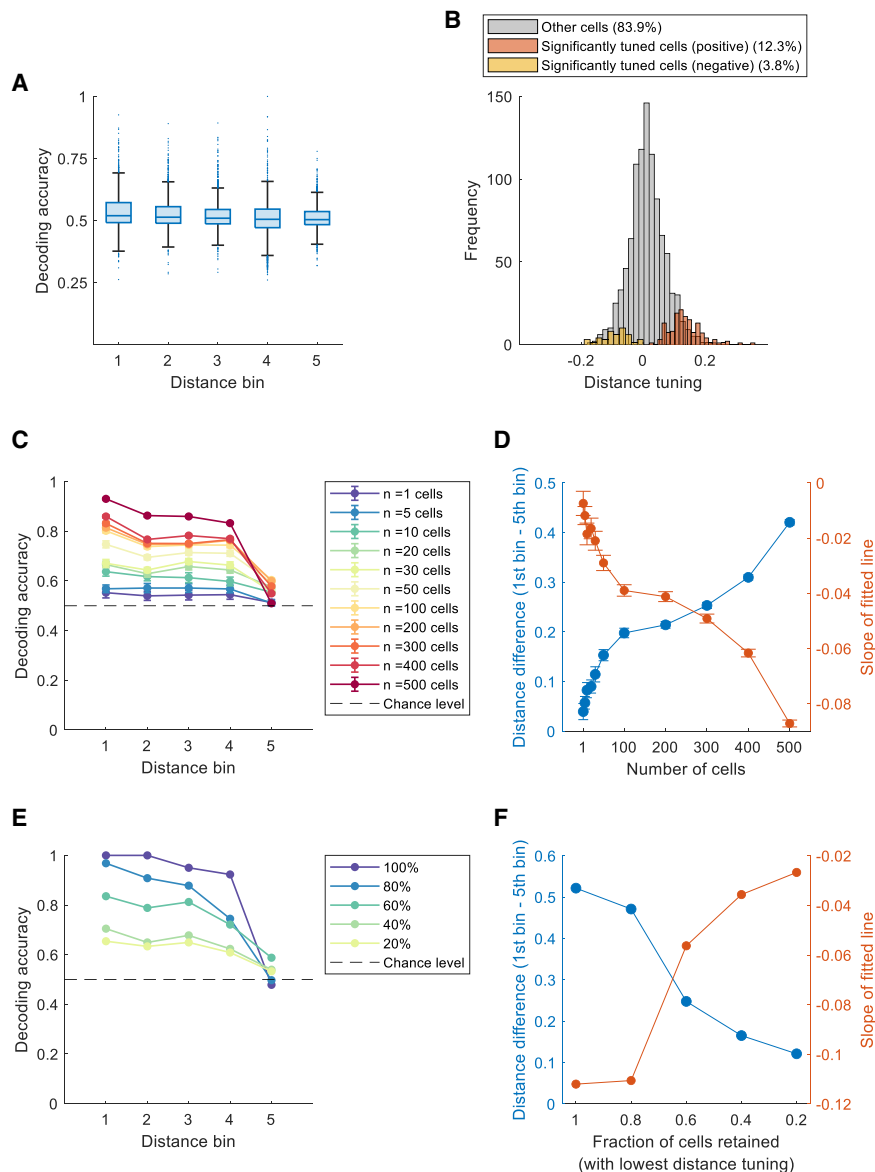


Figure 6. The distance-specific reorganization is widely distributed

(A) Decoding accuracy as a function of distance from object for individual cells (15 tetrode sessions and 2 Neuropixels sessions, $n = 1,189$ cells). Individual cells have mostly flat decoding accuracy. Boxplot conventions are as in Figures 4B. (B) Distance tuning of individual cells ($n = 1,189$ cells). Some cells had significant positive distance tuning (red bars) and other cells had significant negative distance tuning (yellow bars), compared to shuffled data. Note the slight positive bias of distance tuning values in the entire population. (C) Decoding accuracy as a function of distance from object, based on different subsets of cells of size n ($n = 1, 5, 20, 30, 50, 100, 200, 300, 400,$ and 500). Cells were randomly selected from a total population of $N = 620$ cells simultaneously recorded. Data points show means over 20 random realizations; error bars show SEM. When error bars cannot be seen it is because they are too small. Stippled line indicates chance level. (D) Distance reorganization as a function of number of cells used in the decoding. The left y axis (blue curve) shows the distance difference (decoding accuracy in bin #1 – decoding accuracy in bin #5) and the right y axis (orange curve) shows the slope of the least-square line (fitted to decoding accuracy vs. distance curves in the previous panel). These two measures are proxies for the amount of distance reorganization. Data points show means; error bars show SEM. Note the increasing distance difference, as well as the decreasing slope, as a function of the number of cells. (E) Decoding accuracy as a function of distance from object, depending on the fraction of cells with lowest distance tuning retained. We ran the decoder using the $k\%$ of cells with lowest tuning, with k being 100, 80, 60, 40, and 20. Keeping 20% of cells with the least tuning is sufficient to see reorganization of neural-population activity according to distance from the object. Data are from the same Neuropixels rat as in (C) and (D). (F) Distance reorganization as a function of the percentage of cells with the lowest distance tuning retained. Conventions are as in (D).

obtained with population data as a function of object distance (Figure S5L). This is consistent with the idea that many cells contribute to the smooth change in population activity according to object distance (Figures 4B and 4C).

We next asked how the reorganization scales with the number of cells used in the decoder. For this, we first used Neuropixels data from a rat implanted with two probes in the left hippocampus, yielding 620 CA1 cells with a wide variety of properties, including place cells, object cells, and object-vector cells. To simulate different population sizes, we picked random subsets of cells (subsets of size $n = 1, 5, 10, 30, 50, 100, 200, 300, 400,$ and 500) from the 620 cells. For each n we ran the decoder on 20 different randomly chosen subsets of cells. All curves showed reduced decoding accuracy as a function of distance, except the curve constructed using only a single cell (Figure 6C) ($n = 1$ curve:

$W = 2.48, p = 0.6489$; all other curves: $W \geq 11.89, p \leq 0.0182$, Wilcoxon signed-rank tests). The difference between bins #1 and #5 grew systematically with the number of cells (Figure 6D; left y axis, blue) ($H = 174.57, p = 3.14 \times 10^{-32}$, Kruskal-Wallis test) and reached its maximum at $n = 500$ cells. In parallel, the slope of the least-square line (fitted to the curves in Figure 6C) became increasingly negative (Figure 6D, right y axis, orange) ($H = 157.6, p = 1.011 \times 10^{-28}$, Kruskal-Wallis test) and reached its minimum at $n = 500$ cells, creating the appearance of two curves that were mirror-images of one another. First, this implies that even a small number of cells picked randomly is sufficient to see the distance modulation, independently of the individual cell properties (Figure 6C, blue and green curves). Second, the more cells we include, the stronger the distance-modulated reorganization of the map is (Figure 6D). This suggests that the factors

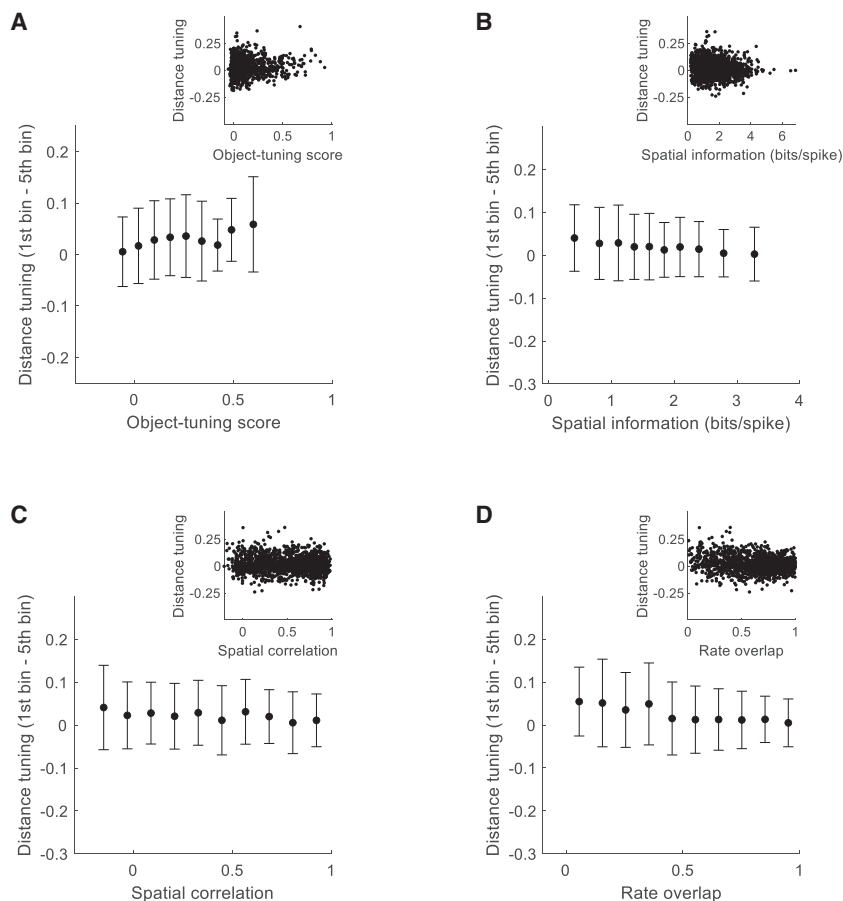


Figure 7. Single-cell properties are not predictive of distance tuning

(A) Relationship between each cell's object-tuning score and its distance tuning. Due to the large number of cells, we bin the data to aid visualization (10 bins, ranging from minimum object-tuning score to maximum object-tuning score). Data points show mean \pm SD for each bin. Inset shows scatterplot of original data before binning ($n = 1,189$ cells). Note lack of a strong correlation between object-tuning score and the distance tuning.

(B) Relationship between each cell's spatial information (bits/spike) (in the O trial) and its distance tuning ($n = 1,189$ cells).

(C) Relationship between each cell's spatial correlation (between rate maps from the NO and O trials) and its distance tuning ($n = 1,189$ cells).

(D) Relationship between each cell's rate overlap (between the NO and O trials) and its distance tuning ($n = 1,189$ cells). In all panels, note absence of relationships between single-cell properties and tuning to distance from the object.

underlying the reorganization are distributed across the entire population of hippocampal cells. The finding was upheld in data from animals with smaller numbers of cells (135 cells in a Neuropixels recording; 25–39 cells in tetrode recordings) (Figures S8A–S8D).

We next asked whether we can visualize the distance reorganization even when using cells that show the least “distance tuning.” After calculating the distance tuning of single cells (Figures 6A and 6B; Equation 1), we sorted the cells according to this variable, starting with the lowest value. We then subsampled the same Neuropixels data (Figures 6C and 6D), keeping either 100%, 80%, 60%, 40%, or 20% of cells with the lowest distance tuning. As expected, curves gradually became flatter when using smaller populations of cells (Figure 6E). However, even the 20% of cells with lowest tuning was sufficient to demonstrate the distance modulation, in the form of positive distance tuning as well as a negative slope (Figures 6E and 6F; right-most point on blue curve: 0.121; on orange curve: -0.0267). Thus even the “worst” cells (with lowest distance tuning) contribute to a distance-specific reorganization of the hippocampal map. As before, the finding also held in smaller-sized tetrode and Neuropixels data (Figures S8E–S8H).

Finally, we asked if any properties of single cells were predictive of distance tuning. We correlated the distance tuning of individual cells (Figure 6B; Equation 1) against their object-tuning

score, spatial information, spatial correlation, or rate overlap (Figures 7A–7D, insets). Distance tuning did not appear to depend strongly on any single-cell property. The object-tuning score had a significant positive correlation with distance tuning, but very weakly so ($r = 0.12$, $p = 1.95 \times 10^{-5}$). Spatial information had a significant negative correlation, but also this was very weak ($r = -0.13$, $p = 3.70 \times 10^{-7}$). Similarly, spatial correlation and rate overlap were also negatively but weakly correlated with distance tuning

(spatial correlation: $r = -0.08$, $p = 7.61 \times 10^{-4}$; rate overlap: $r = -0.18$, $p = 1.57 \times 10^{-9}$). We obtained similar results when, instead of distance tuning, we used the slope of the regression line fitted to decoding accuracy vs. distance as a measure of reorganization (object-tuning score: $r = -0.12$, $p = 3.05 \times 10^{-5}$; spatial information: $r = 0.10$, $p = 3.58 \times 10^{-5}$; spatial correlation: $r = -0.08$, $p = 7.98 \times 10^{-4}$; rate overlap: $r = 0.19$, $p = 1.13 \times 10^{-12}$). After removing all object-tuned cells (19.9%) from the data and re-running the decoder, we observed the same gradual fall in decoding accuracy as a function of distance (Figure S5M). Collectively the observations suggest that cells can contribute to the distance fall-off independently of whether they are object-tuned cells (Figure 7A), place cells (Figure 7B), or exhibit small or large changes in spatial selectivity (Figure 7C) or firing rate (Figure 7D). The dissociation of our findings from single-cell properties (Figures 7 and S5M) is supportive of the notion that distance-dependence is widespread throughout CA1 cells and manifests at the neural-population level (Figure 6).

DISCUSSION

By simultaneously recording up to 620 principal neurons in the hippocampal area CA1 of freely moving rats, we could evaluate single-cell and population-level perspectives on how local objects influence representations of space. Some cells fired at the

object (object cells), others at fixed distances and directions from it (object-vector cells, reminiscent of landmark-vector cells¹³). In addition, the majority of the CA1 cells changed their spatial firing patterns when the object was introduced, moved, or taken out from the environment. At the level of neural populations, the heterogeneous responses were coherently organized across cells so that the amount of reorganization of the hippocampal map depended on distance from the object in a graded manner. The reorganization of the hippocampal map in the presence of objects was widely distributed and did not require specific subsets of cells or the presence of object cells and object-vector cells. Although past studies have already raised the possibility that the hippocampal representation of objects may be non-factorized,¹⁵ we here demonstrate how by describing (1) a strikingly smooth change in population activity according to object distance, and (2) an emergent property of the code, where the distance-based representation becomes more evident as sample sizes are increased. Our findings hint that some aspects of hippocampal cognitive maps may not be understood through properties of dedicated functional cell types—such as object cells, object-vector cells, or place cells—but through emergent properties of cell populations whose activities appear random and disorganized when viewed individually.

We first set out to estimate how many object cells and object-vector cells are present in CA1 when animals explore two-dimensional spaces, in a dataset where the number of cells is an order of magnitude larger than in previous studies of such cells.^{5,10,13,14} We found both object cells and object-vector cells (4.1% and 13.3%, respectively, if using an arbitrary 15 cm distance cut-off between the two groups). Object-tuned cells had variable tuning across time and environments, unlike their MEC counterparts,¹¹ which are stable over weeks and across different tasks.^{11,26} A different type of “instability” across time was reported in hippocampal landmark-vector cells, which only began to fire at fixed distances and directions from objects once the rat had gained experience with the objects.¹³ In the present study we did not observe such experience-dependence. Other studies have reported stable object tuning at the level of individual cells,⁵ whereas we only find hints of stability during the course of a single trial (30 min). While this may reflect differences in training procedures (B. Rivard, personal communication), a broader lesson from our work may be that even if a stimulus (for example, an object and its location) is not represented via stable activities of individual cells, it may still be represented stably at the level of the network.

How exactly did objects influence the spatial firing patterns? Our findings are consistent with studies that report remarkable heterogeneity of object responses in the hippocampus.^{7,15} We observed appearance and disappearance of firing fields, movement of firing fields, object cells, object-vector cells, and firing fields that persisted after objects were removed. This heterogeneity points to a transformation of the object-centered spatial codes within the entorhinal-hippocampal circuit, where the representation of space is changed from a low-dimensional code in MEC with coherence between cells' responses across environments^{27–29} to a higher-dimensional code in the hippocampus that orthogonalizes distinct representations of space.^{20,30,31} The higher dimensionality of hippocampal responses is closely

linked to the phenomenon of global remapping, where distinct place-field codes are created for distinct environments.^{4,32,33}

In parallel with the large variability of responses at the level of single hippocampal cells, we observed a striking organization of the cells' collective activity. The object acted as a pivot for reorganizing hippocampal population activity: in the object's neighborhood, the reorganization was large, leading to high decoding accuracy, whereas further away from the object, the reorganization became increasingly smaller, leading to low decoding accuracy. Thus, the highly variable activity of the CA1 cells was brought together into a unified population code, where activity was largely determined by the state of a single variable: the animal's distance from the object. In that sense, the population response had a simple and predictable structure, in comparison with the disorganized activities of individual cells. The graded reorganization of the population response between NO and O trials could not be explained by differences in the animal's behavior: (1) we were unable to decode the trial type based on behavior and (2) downsampling data to match behavior between the trials had no effect on the main result.

Remapping restricted to subsets of place cells has been shown in many contexts.^{10,34–37} However, the reorganization of population activity described here does not correspond to forms of place-cell remapping described in the literature,^{4,32,33,35,38} one key difference being the smoothness of the change in the population representation. Remapping in place cells often appears in one of two forms: global remapping or rate remapping.³² During global remapping, place cells collectively reorganize their spatial firing patterns, providing an entirely new and orthogonal representation of space. This form of remapping is thought of as discrete,^{33,39–41} and during periods of ambiguity, there may even be flickering between distinct maps.^{40,42,43} Such discreteness is clearly different from the smooth reorganization seen in the present study. Rate remapping⁴⁴ can also be smooth during gradual changes of the environment.²¹ However, these smooth changes, expressed in the firing rates of individual cells, took place all over the environment. Those changes were clearly different from the smooth reorganization across space seen in the present study, where the amount of reorganization of the hippocampal map was a continuous function of the animal's distance from an object. The presence of this smooth reorganization, combined with the simultaneous presence of place cells and object-tuned cells, suggests that population-level representations co-exist with single-cell representations in CA1. Future work will have to determine if such co-existence arises *de novo* in the hippocampus or already at input stages such as the MEC. While responses of individual MEC cells may show less heterogeneity than their hippocampal counterparts (MEC cells fractionate into discrete functional cell types^{11,45–48}), we cannot exclude the possibility that similar population-level representations exist on top of the cells' individual firing profiles in MEC. This possibility would be consistent with topological analyses demonstrating population structure in grid cells not apparent in individual neurons.²⁷

A critical element of our findings is how widely distributed the object-centered population coding is across the recorded cells. The more cells we put into the decoder, the larger the improvement in decoding accuracy near the object (relative to far from

the object). In addition, we could still visualize some degree of distance-based reorganization when only the cells with the weakest distance tuning were retained in the analysis. While distance tuning is generally weak in individual cells, at the population level the responses transform into a smooth distance-dependent code. The transformation has parallels to classic work in the vertebrate retina, where weak pairwise correlations are transformed into strong correlations in population activity,²³ and supports recent work in the hippocampus showing that place cells and non-place cells encode information in a collective and distributed fashion.^{16,17} Taken together, the findings suggest that some neural representations may be best described as emergent properties of large numbers of cells.¹⁸

STAR★METHODS

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

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
 - Lead contact
 - Materials availability
 - Data and code availability
- **EXPERIMENTAL MODEL AND SUBJECT DETAILS**
- **METHOD DETAILS**
 - Surgery and electrode implantation
 - Behavioral procedures
 - Recording procedures
 - Histology and reconstruction of electrode placement
 - Spike sorting and cell-inclusion criteria
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Firing rate maps and identification of firing fields
 - Spatial information content
 - Shuffling of spike data
 - Template matching procedure to identify object-tuned cells
 - Object-vector score and object-vector cell criteria
 - Spatial correlation and rate overlap
 - Region-of-interest analysis for object traces
 - Bayesian decoding algorithm
 - Calculation of decoding accuracy as a function of distance and allocentric angle
 - Performance of decoder
 - Controls for results of decoder
 - Behavioral analyses I: Decoding
 - Behavioral analyses II: Subsampling
 - Subsampling for decoder
 - Measuring strength of distance-specific reorganization
 - Significance of distance tuning in individual cells
 - Population vector correlation
 - Definition of words and counts
 - Mutual information
 - Word distributions
 - Downsampling words
 - Kullback-Leibler divergence between word distributions
 - Statistical tests

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2023.04.008>.

ACKNOWLEDGMENTS

We thank Ø. Høydal and E. Ranheim Skytøen for discussions and R. Gardner for developing the Neuropixels data acquisition and analysis pipeline. We thank A.M. Amundsgård, S. Ball, K. Haugen, K. Jenssen, E. Kråkvik, I. Ulsaker-Janke, and H. Waade for technical assistance. The study was supported by funding from the Research Council of Norway (FRIPRO grant number 300394 to M.-B.M., Centre of Excellence grant number 223262 to M.-B.M. and E.I.M., and National Infrastructure grant number 295721 to E.I.M. and M.-B.M.); a grant from the K.G. Jebsen Foundation (grant number SKGJ-MED-022); a Synergy grant to E.I.M. from the European Research Council (“KILONEURONS,” grant agreement no. 951319); the Kavli Foundation (M.-B.M. and E.I.M.); and a direct contribution to M.-B.M. and E.I.M. from the Ministry of Education and Research of Norway.

AUTHOR CONTRIBUTIONS

A.N. and S.O.A. have equal right to list themselves first in bibliographic documents. A.N., E.I.M., and M.-B.M. designed and planned experiments; A.N. performed the experiments (surgeries, recordings, spike sorting); A.N. and S.O.A. performed single-cell analyses; S.O.A. performed neural-population analyses; S.O.A. and A.N. wrote the paper with inputs and edits from all authors; S.G.C., E.I.M., and M.-B.M. supervised the project; E.I.M. and M.-B.M. obtained funding.

DECLARATION OF INTERESTS

E.I.M. is on the Advisory Board of *Neuron*.

INCLUSION AND DIVERSITY

We support inclusive, diverse, and equitable conduct of research.

Received: July 7, 2022

Revised: December 22, 2022

Accepted: April 7, 2023

Published: May 5, 2023

REFERENCES

1. O’Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Clarendon Press; Oxford University Press).
2. O’Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175.
3. O’Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Exp. Neurol.* **51**, 78–109. [https://doi.org/10.1016/0014-4886\(76\)90055-8](https://doi.org/10.1016/0014-4886(76)90055-8).
4. Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* **7**, 1951–1968.
5. Rivard, B., Li, Y., Lenck-Santini, P.-P., Poucet, B., and Muller, R.U. (2004). Representation of objects in space by two classes of hippocampal pyramidal cells. *J. Gen. Physiol.* **124**, 9–25. <https://doi.org/10.1085/jgp.200409015>.
6. Battaglia, F.P., Sutherland, G.R., and McNaughton, B.L. (2004). Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. *J. Neurosci.* **24**, 4541–4550. <https://doi.org/10.1523/JNEUROSCI.4896-03.2004>.
7. Lenck-Santini, P.P., Rivard, B., Muller, R.U., and Poucet, B. (2005). Study of CA1 place cell activity and exploratory behavior following spatial and

- nonspatial changes in the environment. *Hippocampus* 15, 356–369. <https://doi.org/10.1002/hipo.20060>.
8. Burke, S.N., Maurer, A.P., Nematollahi, S., Uprety, A.R., Wallace, J.L., and Barnes, C.A. (2011). The influence of objects on place field expression and size in distal hippocampal CA1. *Hippocampus* 21, 783–801. <https://doi.org/10.1002/hipo.20929>.
 9. 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. <https://doi.org/10.1523/JNEUROSCI.21-05-01635.2001>.
 10. Gothard, K.M., Skaggs, W.E., Moore, K.M., and McNaughton, B.L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci.* 16, 823–835. <https://doi.org/10.1523/JNEUROSCI.16-02-00823.1996>.
 11. Høydal, Ø.A., Skytøen, E.R., Andersson, S.O., Moser, M.B., and Moser, E.I. (2019). Object-vector coding in the medial entorhinal cortex. *Nature* 568, 400–404. <https://doi.org/10.1038/s41586-019-1077-7>.
 12. Andersson, S.O., Moser, E.I., and Moser, M.-B. (2021). Visual stimulus features that elicit activity in object-vector cells. *Commun. Biol.* 4, 1219. <https://doi.org/10.1038/s42003-021-02727-5>.
 13. Deshmukh, S.S., and Knierim, J.J. (2013). Influence of local objects on hippocampal representations: landmark vectors and memory. *Hippocampus* 23, 253–267. <https://doi.org/10.1002/hipo.22101>.
 14. Geiller, T., Fattahi, M., Choi, J.S., and Royer, S. (2017). Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nat. Commun.* 8, 14531. <https://doi.org/10.1038/ncomms14531>.
 15. GoodSmith, D., Kim, S.H., Puliyadi, V., Ming, G.L., Song, H., Knierim, J.J., and Christian, K.M. (2022). Flexible encoding of objects and space in single cells of the dentate gyrus. *Curr. Biol.* 32, 1088–1101.e5. <https://doi.org/10.1016/j.cub.2022.01.023>.
 16. Meshulam, L., Gauthier, J.L., Brody, C.D., Tank, D.W., and Bialek, W. (2017). Collective behavior of place and non-place neurons in the hippocampal network. *Neuron* 96, 1178–1191.e4. <https://doi.org/10.1016/j.neuron.2017.10.027>.
 17. Stefanini, F., Kushnir, L., Jimenez, J.C., Jennings, J.H., Woods, N.I., Stuber, G.D., Kheirbek, M.A., Hen, R., and Fusi, S. (2020). A distributed neural code in the dentate gyrus and in CA1. *Neuron* 107, 703–716.e704. <https://doi.org/10.1016/j.neuron.2020.05.022>.
 18. Yuste, R. (2015). From the neuron doctrine to neural networks. *Nat. Rev. Neurosci.* 16, 487–497. <https://doi.org/10.1038/nrn3962>.
 19. Manns, J.R., and Eichenbaum, H. (2009). A cognitive map for object memory in the hippocampus. *Learn. Mem.* 16, 616–624. <https://doi.org/10.1101/lm.1484509>.
 20. Leutgeb, S., Leutgeb, J.K., Treves, A., Moser, M.B., and Moser, E.I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305, 1295–1298. <https://doi.org/10.1126/science.1100265>.
 21. Leutgeb, J.K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C.A., McNaughton, B.L., Moser, M.B., and Moser, E.I. (2005). Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron* 48, 345–358. <https://doi.org/10.1016/j.neuron.2005.09.007>.
 22. Posani, L.A.-O., Cocco, S., Ježek, K., and Monasson, R. (2017). Functional connectivity models for decoding of spatial representations from hippocampal CA1 recordings. *J. Comput. Neurosci.* 43, 17–33.
 23. Schneidman, E., Berry, M.J., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440, 1007–1012. <https://doi.org/10.1038/nature04701>.
 24. Osborne, L.C., Palmer, S.E., Lisberger, S.G., and Bialek, W. (2008). The neural basis for combinatorial coding in a cortical population response. *J. Neurosci.* 28, 13522–13531. <https://doi.org/10.1523/JNEUROSCI.4390-08.2008>.
 25. Schneidman, E., Puchalla, J.L., Segev, R., Harris, R.A., Bialek, W., and Berry, M.J. (2011). Synergy from silence in a combinatorial neural code. *J. Neurosci.* 31, 15732–15741. <https://doi.org/10.1523/JNEUROSCI.0301-09.2011>.
 26. Obenaus, H.A., Zong, W., Jacobsen, R.I., Rose, T., Donato, F., Chen, L., Cheng, H., Bonhoeffer, T., Moser, M.B., and Moser, E.I. (2022). Functional network topography of the medial entorhinal cortex. *Proc. Natl. Acad. Sci. USA* 119, e2121655119. <https://doi.org/10.1073/pnas.2121655119>.
 27. Gardner, R.J., Hermansen, E., Pachitariu, M., Burak, Y., Baas, N.A., Dunn, B.A., Moser, M.B., and Moser, E.I. (2022). Toroidal topology of population activity in grid cells. *Nature* 602, 123–128. <https://doi.org/10.1038/s41586-021-04268-7>.
 28. Fyhn, M., Hafting, T., Treves, A., Moser, M.B., and Moser, E.I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446, 190–194. <https://doi.org/10.1038/nature05601>.
 29. Yoon, K., Buice, M.A., Barry, C., Hayman, R., Burgess, N., and Fiete, I.R. (2013). Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat. Neurosci.* 16, 1077–1084. <https://doi.org/10.1038/nn.3450>.
 30. Alme, C.B., Miao, C., Jezek, K., Treves, A., Moser, E.I., and Moser, M.B. (2014). Place cells in the hippocampus: eleven maps for eleven rooms. *Proc. Natl. Acad. Sci. USA* 111, 18428–18435. <https://doi.org/10.1073/pnas.1421056111>.
 31. McNaughton, B.L., and Morris, R.G.M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.* 10, 408–415. [https://doi.org/10.1016/0166-2236\(87\)90011-7](https://doi.org/10.1016/0166-2236(87)90011-7).
 32. Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309, 619–623. <https://doi.org/10.1126/science.1114037>.
 33. Colgin, L.L., Moser, E.I., and Moser, M.-B. (2008). Understanding memory through hippocampal remapping. *Trends Neurosci.* 31, 469–477. <https://doi.org/10.1016/j.tins.2008.06.008>.
 34. Tanila, H., Shapiro, M.L., and Eichenbaum, H. (1997). Discordance of spatial representation in ensembles of hippocampal place cells. *Hippocampus* 7, 613–623.
 35. Anderson, M.I., and Jeffery, K.J. (2003). Heterogeneous modulation of place cell firing by changes in context. *J. Neurosci.* 23, 8827–8835. <https://doi.org/10.1523/JNEUROSCI.23-26-08827.2003>.
 36. Knierim, J.J. (2002). Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. *J. Neurosci.* 22, 6254–6264. <https://doi.org/10.1523/JNEUROSCI.22-14-06254.2002>.
 37. Chen, G., King, J.A., Burgess, N., and O'Keefe, J. (2013). How vision and movement combine in the hippocampal place code. *Proc. Natl. Acad. Sci. USA* 110, 378–383. <https://doi.org/10.1073/pnas.1215834110>.
 38. Latuske, P., Kornienko, O., Kohler, L., and Allen, K. (2017). Hippocampal remapping and its entorhinal origin. *Front. Behav. Neurosci.* 11, 253. <https://doi.org/10.3389/fnbeh.2017.00253>.
 39. Wills, T.J., Lever, C., Cacucci, F., Burgess, N., and O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science* 308, 873–876. <https://doi.org/10.1126/science.1108905>.
 40. Jezek, K., Henriksen, E.J., Treves, A., Moser, E.I., and Moser, M.-B. (2011). Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478, 246–249. <https://doi.org/10.1038/nature10439>.
 41. Kubie, J.L., Levy, E.R.J., and Fenton, A.A. (2020). Is hippocampal remapping the physiological basis for context? *Hippocampus* 30, 851–864. <https://doi.org/10.1002/hipo.23160>.
 42. Kelemen, E., and Fenton, A.A. (2010). Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. *PLoS Biol.* 8, e1000403. <https://doi.org/10.1371/journal.pbio.1000403>.
 43. Renaudineau, S., Poucet, B., and Save, E. (2007). Flexible use of proximal objects and distal cues by hippocampal place cells. *Hippocampus* 17, 381–395. <https://doi.org/10.1002/hipo.20277>.

44. Markus, E.J., Qin, Y.L., Leonard, B., Skaggs, W.E., McNaughton, B.L., and Barnes, C.A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.* *15*, 7079–7094. <https://doi.org/10.1523/JNEUROSCI.15-11-07079.1995>.
45. Hafting, T., Fyhn, M., Molden, S., Moser, M.B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature* *436*, 801–806. <https://doi.org/10.1038/nature03721>.
46. Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.B., and Moser, E.I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* *312*, 758–762. <https://doi.org/10.1126/science.1125572>.
47. Solstad, T., Boccara, C.N., Kropff, E., Moser, M.B., and Moser, E.I. (2008). Representation of geometric borders in the entorhinal cortex. *Science* *322*, 1865–1868. <https://doi.org/10.1126/science.1166466>.
48. Kropff, E., Carmichael, J.E., Moser, M.B., and Moser, E.I. (2015). Speed cells in the medial entorhinal cortex. *Nature* *523*, 419–424. <https://doi.org/10.1038/nature14622>.
49. Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature* *551*, 232–236. <https://doi.org/10.1038/nature24636>.
50. Steinmetz, N.A., Aydın, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M., et al. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* *372*, eabf4588. <https://doi.org/10.1126/science.abf4588>.
51. Mankin, E.A., Sparks, F.T., Slayyeh, B., Sutherland, R.J., Leutgeb, S., and Leutgeb, J.K. (2012). Neuronal code for extended time in the hippocampus. *Proc. Natl. Acad. Sci. USA* *109*, 19462–19467. <https://doi.org/10.1073/pnas.1214107109>.
52. Mankin, E.A., Diehl, G.W., Sparks, F.T., Leutgeb, S., and Leutgeb, J.K. (2015). Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts. *Neuron* *85*, 190–201. <https://doi.org/10.1016/j.neuron.2014.12.001>.
53. Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* *42*, 283–295. [https://doi.org/10.1016/s0896-6273\(04\)00192-8](https://doi.org/10.1016/s0896-6273(04)00192-8).
54. Skaggs, W.E., McNaughton, B.L., and Gothard, K.M. (1993). An information-theoretic approach to deciphering the hippocampal code. In *Advances in Neural Information Processing Systems 5*, [NIPS Conference] (Morgan Kaufmann Publishers Inc.), pp. 1030–1037.
55. Agresti, A. (2015). *Foundations of Linear and Generalized Linear Models* (John Wiley & Sons Inc).
56. Cover, T.M., and Thomas, J.A. (2006). *Elements of Information Theory* (John Wiley & Sons, Inc). <https://doi.org/10.1002/047174882X.ch2>.
57. Shannon, C.E. (1948). A mathematical theory of communication. *Bell Syst. Tech. J.* *27*, 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.
58. 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. <https://doi.org/10.1152/jn.00559.2007>.
59. Kullback, S., and Leibler, R.A. (1951). On information and sufficiency. *Ann. Math. Statist.* *22*, 79–86. <https://doi.org/10.1214/aoms/1177729694>.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
Rat: Long-Evans	Bred inhouse at KISN	Kavli Institute for Systems Neuroscience
Software and algorithms		
Matlab 2020a	MathWorks	https://mathworks.com/products/matlab.html RRID: SCR_001622
Cheetah 6.3.1	Neuralynx	https://neuralynx.com/
MClust	A.D. Redish	https://redishlab.umn.edu/mclust
SpikeGLX	Jun et al. ⁴⁹	https://billkarsh.github.io/SpikeGLX/
Kilosort	Steinmetz et al. ⁵⁰	https://github.com/MouseLand/Kilosort
Phy	Jun et al. ⁴⁹	https://github.com/cortex-lab/phy
OptiTrack Motive	OptiTrack	https://optitrack.com/software/motive/
Custom software	This paper	https://doi.org/10.5281/zenodo.7797252
Other		
Axona Microdrives	Axona	https://www.axona.com
Neuralynx Acquisition System	Neuralynx	https://neuralynx.com/
Neuropixels Probes	Neuropixels	https://www.neuropixels.org/
Neuropixels Control System	Neuropixels	https://www.neuropixels.org/control-system
Optitrack Flex 13 USB cameras	OptiTrack	https://optitrack.com/cameras/
Zeiss AxioImager	Zeiss	N/A
Deposited data		
Datasets	This paper	https://doi.org/10.25493/Q0FG-1X8

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, May-Britt Moser (may-britt.moser@ntnu.no).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Data are publicly available as of the date of publication (<https://doi.org/10.25493/Q0FG-1X8>).
- Code has been deposited on Zenodo (<https://doi.org/10.5281/zenodo.7797252>).
- All information is provided in the [key resources table](#). Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Data were obtained from 10 male Long-Evans rats aged 2-3.5 months (320-470 grams) at the time of implantation. The rats were housed with 2-4 littermates prior to surgery. After implantation, the rats were housed individually on a single floor in large, dual-floor metal cages (95 x 63 x 61 cm) with access to food and water *ad libitum* throughout the entire experiment. The cages contained nesting material and were enriched with cardboard and plastic houses and contained a heightened platform. The rats were kept in a temperature and humidity controlled environment and kept on a 12 hour light/12 hour dark schedule. All testing occurred during

the dark phase. All experiments were performed in accordance with the Norwegian Animal Welfare Act and the European Convention for the Protection of Vertebrate Animals used for Experimental and Other Scientific Purposes and were approved by the Norwegian Food Safety Authority (permit numbers 7163 and 18011).

METHOD DETAILS

Surgery and electrode implantation

Tetrode experiments

Seven rats were implanted with one or two "microdrives," each containing a single bundle of four or eight tetrodes. Five of these rats were implanted with a single eight-tetrode microdrive targeting the CA1 and CA3 areas of the hippocampus, whereas two rats were implanted with one four-tetrode microdrive targeting CA1/CA3 in one hemisphere and one eight-tetrode microdrive targeting the medial entorhinal cortex (MEC) of the opposite hemisphere. Few cells were recorded on the MEC tetrodes and the data from these tetrodes was not used. Choice of hemisphere varied across rats. Tetrodes were constructed from four twisted 17- μm polyimide-coated platinum-iridium (90–10%) wires (California Fine Wire). The electrode tips were plated with platinum to reduce electrode impedances to between 120 and 220 k Ω at 1 kHz.

Anaesthesia was induced by placing the rat in a closed glass induction chamber filled with 5% isoflurane (air flow: 0.4 l/min; oxygen flow: 0.4 l/min). The rat was given subcutaneous injections of buprenorphine (Temgesic, 0.03 mg/kg), meloxicam (Metacam, 1 mg/kg), and atropine (0.05 mg/kg). The local anaesthetic bupivacaine (Marcain, 1 mg/kg) was injected subcutaneously before the incision was made. After injections the rat was moved to a raised platform with a mask and a stereotactic frame and the head was secured in place with a set of ear bars. The rat's body was resting on a heating pad (38°C) to ensure that the body temperature was maintained throughout the surgery. Isoflurane levels were gradually reduced to 0.5–1.5% depending on the physiological condition of the rat, which was evaluated by reflex responses and breathing patterns.

After the incision had been made, the periosteum was removed and the craniotomies were drilled. Tetrodes targeting CA1/CA3 were implanted without an angle at 3.2 mm lateral to the midline and 3.8–4 mm posterior to bregma. Tetrodes targeting MEC were implanted at 4.5 mm lateral to the midline and 0.1 mm anterior to the transverse sinus edge, with an anterior angle of 25° relative to the bregma/lambdoid horizontal reference. The initial depth of the tetrode tips varied between 1500 and 2000 μm relative to the dura mater. The microdrives were secured to the skull using an adhesive (OptiBond, Kerr), dental cement (Meliodent), and 2–5 Jeweller screws (M1.4) that were placed in vertically drilled holes in the parietal bones. One or two screws were placed in the occipital bone over the cerebellum and were connected to the drive grounds.

After the surgery, the rat was left in a heated chamber (30°C) for 30–60 minutes for the immediate recovery phase, after which it was transferred back to the home cage. Additional doses of buprenorphine were administered 8–12 and 24 hours after the first injection. Meloxicam was administered once every 24 hours for as long as was assessed necessary (usually 3–4 days). The rat was left to recover for 3–5 days before the experiments began.

Neuropixels experiments

Three rats were implanted bilaterally with four-shank Neuropixels 2.0 silicone probes^{49,50} targeting the CA1/CA3 in one hemisphere, and the MEC in the opposite hemisphere. The MEC data were not used in the present study due to the small number of cells and experimental sessions. Two rats were implanted with one probe in each hemisphere, whereas one rat was implanted with two probes in the left hippocampus and one probe in the right MEC. In this case, the two hippocampal-targeting probes were glued together at the probe base prior to surgery and the total eight probe shanks were implanted in the same craniotomy. The hippocampal probe tips were lowered to 7000–7500 μm relative to the dura mater. The stereotactic coordinates and the surgical procedure were the same as described above for tetrode experiments and have been described in detail elsewhere.²⁷ The probes were secured to the skull using Optibond and Venus (Kulzer) and protected by fitting a modified falcon tube. After surgery and the initial recovery in the heating chamber, the rat was kept in the home cage for approximately 3 hours before experiments began. Post-operative analgesia was administered as described above.

Behavioral procedures

Before surgery, the rats were habituated to being handled by the experimenter and familiarised with the experimental task and environment during training sessions. Each rat underwent 7–15 training sessions, lasting at least 20 minutes each on separate days. During each training session, the rat foraged for randomly scattered corn puff crumbs in a 150 cm x 150 cm square arena with 50 cm tall black walls. The floor mat was either matt black or dark blue/green. Tetrode rats were not familiarised with the object prior to surgery, whereas Neuropixels rats had six or seven extra training sessions before surgery where the object was placed into the arena in order to avoid novelty effects of the object during recordings.

Tetrode sessions were performed in one of two different rooms. In one room, the arena was encircled by a set of lightproof blue curtains. A large white paper sheet was attached to the curtain on one side as a distal cue. The arena in the second room had curtains hanging on only one side of the arena and various distal cues were visible to the rat on the remaining sides. Neuropixels sessions were performed in a single room, where the arena was surrounded by a set of lightproof blue curtains on three sides. On the fourth side, the curtains were open towards a white wall on which a shelf containing the recording apparatus was placed.

During experimental testing, each recording session consisted of either three or four trials. In all trials, the rat was freely foraging for randomly scattered corn puff crumbs. Each trial typically lasted 20–30 minutes, depending on how long it took the rat to fully cover the extent of the arena. The first trial was always an empty arena (No Object). This was followed by a trial in which a colourful Duplo tower (52 cm x 10 cm x 10 cm) was placed at or close to the middle of the arena (Object). In the third trial, the Duplo tower was moved to another, semi-random location in the arena, usually towards one of the four corners (Object Moved; approximate range of distance from the corner: 35–55 cm). A fourth trial in which the object was removed was used to check for carry-over (hysteresis) of firing at the former object location (No Object'). In the few cases where the rat would stop exploring the arena, the fourth trial was omitted from the experiment (this happened to 8.5% of the cells). Between each trial, the rat was placed on a towel in a flowerpot on a pedestal next to the arena for 1–3 minutes and given a cookie with vanilla cream and access to a water bottle.

It has previously been shown that the activity of CA cells can change dynamically over time without changes in context^{51–53}. To control for this, a separate set of experiments was performed, where cells were recorded as the rat foraged freely in the same open field arena during four consecutive trials without an object (Control sessions; Figure 2B). All other parameters were the same as for sessions with the object.

In the tetrode experiments, recordings typically started 3–5 days after implantation. The tetrodes were turned down in steps of 50 μm per day. To determine if the tetrodes were in the CA1 pyramidal cell layer, cells were recorded as the rat explored the empty open field arena and the presence of theta modulation together with the presence of at least 5–10 place cells was used as criteria. When the criteria were met for the first time, the full experimental protocol was started. The tetrodes were subsequently moved down in steps of 50 μm per day at the most, and the experiment was repeated several times over the course of 1–2 months (up to 25 recording sessions per animal). In two of the rats, the tetrodes were subsequently lowered to the CA3 and experiments were repeated as for the CA1 recordings.

The Neuropixels data were obtained from a total of four recording sessions from the three rats. Two of the rats were only recorded once, whereas the third was recorded two times (first in CA1 and then in CA3) with a 3-day interval between the two recording sessions. All recording sessions were performed within 72 hours after recovery from surgery.

Recording procedures

Tetrode experiments

During electrophysiological recordings, the microdrives (with tetrodes) were connected to a Neuralynx data acquisition system (Neuralynx, Tucson, AZ; Neuralynx Digital Lynx SX) through a multichannel, impedance matching unity gain head stage. The output of the head stage was conducted via a lightweight multi wire tether cable and a slip-ring commutator. Unit activity was amplified by a factor of 3000–5000 and band-pass filtered from 600 to 6000 Hz. Spike waveforms above a threshold of 30 or 40 μV were time-stamped and digitized at 32 kHz for 1 ms. Local field potential (LFP) signals were recorded from each tetrode, amplified by a factor of 250–1000, low-pass filtered at 300–475 Hz and sampled at 1800–2500 Hz. An overhead camera recorded the position of two light-emitting diodes (LEDs) on the head stage in order to track the rat's movements at a sampling rate of 25 Hz. The diodes were separated by 4 cm.

Neuropixels experiments

Signals were recorded using a Neuropixels acquisition system as described previously.^{27,49,50} Briefly, the spike band signal was amplified at a gain of 80, low-pass-filtered at 0.5 Hz, high-pass-filtered at 10 kHz, and digitalised at 30 kHz on the probe circuit board. The signal was then multiplexed and transmitted to a Neuropixels PXIe acquisition module via a 5 m tether cable made using twisted pair wiring. A 3D motion capture system consisting of six OptiTrack Flex 13 cameras and Motive software was used to track the rat's movement at a sampling rate of \sim 120 Hz. Five reflective markers glued to a rigid body was attached to the rat's implant during recording. Randomised sequences of digital pulses were generated by an Arduino microcontroller. In order to synchronise the time-stamps of the two recording systems, the pulses were sent to the Neuropixels acquisition system as direct TTL input and to the OptiTrack system via infrared LEDs that were placed on the edge of the arena.

Histology and reconstruction of electrode placement

The rat was placed in a closed glass induction chamber filled with 5% isoflurane (air flow 1.2 l/min) to induce anaesthesia. It was then given an overdose of pentobarbital injected intraperitoneally and perfused intracardially with saline followed by 4% formaldehyde. The brain was removed and stored in formaldehyde. Frozen 30 μm sagittal sections were cut, mounted on glass, and stained with Cresyl violet (Nissl).

The final position of the tip of each tetrode or the placement of the probe shank was identified on photomicrographs obtained with an Axio Scan.Z1 microscope and ZEN software (Zen 2.6; Carl Zeiss). The tip of the tetrode was used to determine the anatomical location of the recorded cell. For the Neuropixels experiments, the depths of the probe shank tips were measured in ZEN software. The known position of active recording sites relative to the tip was then used to determine the location of the recorded cells. CA1 and CA3 recording locations were distinguished based on depth and activity profiles along the probe shanks. Recorded cells that were judged to be in the cortex overlying the CA1 or in the dentate gyrus were excluded from all analyses.

Spike sorting and cell-inclusion criteria

Spike sorting of tetrode data was performed offline using the graphical cluster-cutting software MClust (A.D. Redish, <https://redishlab.umn.edu/mclust>). Spikes were clustered manually in 2D projections of the multidimensional parameter space (consisting of waveform amplitudes and waveform energies), using autocorrelation and crosscorrelation functions as additional separation tools and separation criteria. To follow cells across recording sessions, clusters of successive sessions were compared and identified to be the same unit if the locations of the spike clusters in the multidimensional parameter space were stable. Spike sorting of Neuropixels data was performed using a version of KiloSort 2.5⁵⁰ with customisations as described previously.²⁷

For both types of experiments, all trials in a session were clustered together. Single units were discarded if more than 1% of their inter-spike intervals were less than 2 ms. Further, cells were excluded if they had a mean spike rate of less than 0.05 Hz or greater than 10 Hz across the full recording session.

QUANTIFICATION AND STATISTICAL ANALYSIS

Data analyses were performed with custom-written scripts in MATLAB (Version 2020a). The study did not involve any experimental subject groups; therefore, random allocation and experimenter blinding did not apply and were not performed.

Firing rate maps and identification of firing fields

2D rate maps that displayed the firing rate as a function of the position of the animal in the arena were calculated for each trial separately. First, position estimates were binned into a square grid of 2.5 cm x 2.5 cm bins. Next, the firing rate in each position bin was calculated as the number of spikes recorded in the bin, divided by the time the rat spent in the bin. The 2D rate map was smoothed with a Gaussian kernel with standard deviation of 2 bins in both the *x* and the *y* direction. Only time epochs in which the animal was moving at a speed above 2.5 cm/s were used for constructing the rate maps and for spatial analyses.

Firing fields in the rate map were detected by iteratively applying the Matlab "contour" function (Mathworks), starting from the cell's peak firing rate until it reached 3 × median absolute deviation (MAD) of the firing rates of all bins in the rate map. Firing fields were defined as contiguous areas within a contour with a size of at least 225 cm² and with a minimum firing rate of at least 1 Hz.

Spatial information content

To calculate the spatial information content in each cell's firing, the rate map was used. First, the spatial information rate⁵⁴ was computed as

$$\sum_{i=1}^N P_i \frac{\lambda_i}{\lambda} \log_2 \left(\frac{\lambda}{\lambda_i} \right),$$

where *N* is the total number of bins (3200 bins in total), λ_i is the firing rate in the *i*-th bin of the rate map, λ is the mean firing rate in the trial, and P_i is the probability of the rat being in the *i*-th bin, calculated as the occupancy in the *i*-th bin (time spent in the bin) divided by the total duration of the trial. Spatial information content (in bits per spike) was then obtained for each trial by dividing the spatial information rate by the mean firing rate of the cell in that trial.

Shuffling of spike data

To determine significance for spatial information content, a per-trial shuffling procedure was implemented for each cell separately. For each trial, the sequence of spike timestamps was circularly shifted (with the end of the trial wrapped to the beginning of the trial) by a displacement that took random values between 20 seconds and 20 seconds less than the duration of the trial. Time shifts varied randomly between trials and between cells. The shuffling procedure was repeated 1000 times per trial per cell. For each shuffle iteration, the rate maps were recalculated and the spatial information content was computed using the position of the rat, which remained unchanged, and a shuffled distribution of spatial information content values was built. A cell was judged to be spatially modulated in the trial if the spatial information content calculated on the experimental data was above the 99th percentile of the shuffled distribution.

Template matching procedure to identify object-tuned cells

We developed a template matching procedure to identify cells firing at the location of objects ("object cells") and cells firing in fixed distances and directions from objects ("object-vector cells"). Throughout the paper, we refer to such cells as "object-tuned cells."

First, we placed a Gaussian template at a specific location in the Object trial. The template was then transferred to the same location in the Object Moved trial, relative to the object. For example, the template's location (i.e. the mean of the Gaussian) could be 15 cm North from the object, or directly at the object (as in Figure 1B, left). The template location (relative to the object) was always the same in both trials.

Second, we calculated the Pearson correlation between the template and the cell's actual firing rate map in both the Object and Object Moved trials. This gave a measure of how strongly the cell's firing pattern resembled the template in each of the trials. Because an object-tuned cell should fire relative to the object consistently in both trials, we took the minimum of these two Pearson

correlations as an "object-tuning score." Taking the minimum ensures that we only classify cells as object-tuned cells if their firing patterns move along when the object moves.

$$\text{object tuning score} = \min(\text{corr}_{\text{object}}, \text{corr}_{\text{object moved}})$$

Parametric variations

Since cells can fire at a variety of distances and directions from the object, we repeated the calculation of object-tuning score with the template in different locations (9 offsets, including the template located at the object, and 8 directions, [Figure 1B](#), right). Similarly, we repeated the calculation with 8 different variances (σ^2) of the Gaussian template. Each combination of offsets, angles and variances gives one object-tuning score and constitutes a specific hypothesis about how the cell's firing could be related to the object.

The offsets between the object and the template's center were 0, 5, 10, 15, 20, 30, 40, 50 and 60 cm. The angles between the object and template varied between 0 and 2π , in steps of $\pi/4$. That is, the template could be shifted in 8 cardinal directions. The variances of the Gaussian template were 5, 10, 25, 75, 50, 100, 150 and 200 cm.

This meant that, for each cell, we calculated 520 object-tuning scores: 8 variances \times (8 angles \times 8 offsets + 1 offset at the object) = 520. In other words, we evaluated 520 different hypotheses of how the cell's firing could be related to the object.

To determine whether the obtained scores were statistically significant, we compared the actual object-tuning scores to scores calculated on shuffled data. We created shuffled data by choosing the template location randomly in both the Object and Object Moved trials, and then re-calculated the object-tuning score. We repeated this procedure 500 times. 8 shuffling distributions were created for each cell (1 for each variance of the Gaussian template), because we only compared actual scores to shuffled scores using the same variance of the Gaussian template. For example, if the template's variance was $\sigma^2 = 50$ when calculating the actual object-tuning score, the template's variance was also $\sigma^2 = 50$ for the shuffled scores. We therefore had a total of 8×1189 shuffling distributions, where 8 is the number of variances and 1189 is the number of cells. If the cell had one or more object-tuning scores larger than the 99th percentile of the shuffled distribution, we considered the cell object-tuned.

Most object-tuned cells (88.6%) had more than one significant object-tuning score, while only a small fraction (11.4%) had one significant score ([Figure S2B](#)). Since object-tuned cells typically fitted more than one template (with slight variations in template parameters), this means that an approximate fit between the template and the cell's actual firing pattern is sufficient for detection. The result also suggests that step sizes in our parameters were small enough.

To roughly estimate the number of object-tuned cells that were object cells and object-vector cells, we used the distance between the cell's firing field and the object. Cells with firing fields less than 15 cm from the object were considered object cells; cells with firing fields more than 15 cm from the object were considered object-vector cells. In a small number of object-tuned cells, the firing field detection algorithm did not identify any firing fields, either because the activity was too low or scattered to form firing fields (see [Firing rate maps and identification of firing fields for details about the algorithm](#)). These cells were therefore not separated into object cells or object-vector cells (30/1189 cells).

Chance level

To estimate the chance level in classification of object-tuned cells (5.03%), we repeated the same template matching procedure for cells recorded in control trials without objects ([Figure 2B](#)) ($n = 219$ CA1 cells). For this, we used the middle of the arena as the object location in the Object trial ($x = 75$ cm, $y = 75$ cm) and the upper-left corner of the arena as the object location in the Object Moved trial ($x = 35$ cm, $y = 120$ cm) (the most common object location in Object Moved trials).

Object-vector score and object-vector cell criteria

To confirm that the general template matching procedure (see above) identified object-vector cells, we applied the methods used in previous publications to identify object-vector cells.^{11,12} For this, we constructed "vector maps" that expressed each cell's firing as a function of the animal's distance and direction from the object. After constructing vector maps from the Object and Object Moved trials, we used the Pearson correlation between the two maps as an "object-vector score." To create shuffled data, we circularly shifted each cell's spike train 500 times (as described under [Shuffling of spike data](#)) and re-calculated the object-vector score. Object-vector scores from cells that were above the 99th percentile of the shuffled distribution were considered significant.

Before applying the criteria, we excluded cells with firing fields that would fall outside of the arena in the Object Moved trial (for example, when the object was moved to the south-east corner, when the cell fired at a south-east angle from the object) (28.8% of CA1 cells were excluded).

Cells were considered object-vector cells if the following criteria were met:

- 1) The spatial information content was above the 99th percentile of the shuffled distribution in both the Object and the Object Moved trials.
- 2) The object-vector score was above the 99th percentile of the shuffled distribution.
- 3) One or more firing fields were present in the Object and the Object Moved trials.
- 4) The firing field had to be at least 15 cm distance from the object (to select object-vector cells rather than object cells).

40 out of 1189 cells (3.4%) were classified by this method but not by the template matching procedure. Manual inspection showed that these cells were mostly false positives, confirming that the template matching procedure correctly classified object-vector cells.

Spatial correlation and rate overlap

In order to determine whether the firing activity of single cells changed or remained stable when the object was introduced, moved, or removed from the environment, two measures were considered:

- 1) Spatial correlation
- 2) Rate overlap

The spatial correlation between two trials was defined as the Pearson correlation between the overlapping parts of the two firing rate maps (excluding bins that contained no value in either rate map). The rate overlap between two trials was calculated by dividing the mean firing rate in the less active trial by the mean firing rate in the more active trial.

To classify a cell's response as changing or stable, a baseline stability of the cell's firing within a trial ("within-trial") was first assessed. The position data and the corresponding spike timestamps from the first trial (No Object) were binned into n sections of 1-second duration, where n is the full duration of the No Object trial in seconds. The bins were randomly assigned to two groups of size $n/2$. New rate maps for the two groups were then constructed, and the spatial correlation and the firing rate overlap between them was calculated. The procedure was repeated 500 times for each cell, with different bins randomly assigned to the two groups for each permutation.

For each between-trial comparison (e.g., No Object – Object), the same procedure as for the within-trial comparison was used. Position data and corresponding spiking data were randomly assigned to two groups of $n/2$ bins of 1 second duration for each of the two trials separately. Only one of the two groups (i.e., half of the data) from each trial was used and spatial correlation and rate overlap between the two trials were calculated as described above. The procedure was again repeated 500 times per between-trial comparison. A cell's response was classified as changing if the mean of the between-trial distribution fell below the 1st percentile of the within-trial distribution from the No Object trial, and if the mean ranks of the two distributions were judged to significantly different (assessed using a Wilcoxon signed-rank test).

Region-of-interest analysis for object traces

To investigate the presence of object-trace fields, we first pre-selected for cells that changed their spatial firing patterns or firing rates between No Object and No Object' (quantified with spatial correlation and rate overlap, using criteria defined in the previous section). For these pre-selected cells ($n = 919$), we identified firing fields in No Object'. Then, a circular ROI was defined in the same way as described in section [Region-of-interest analysis for object responses](#). The ROI was transferred to the same location in the other three trials (No Object, Object, and Object Moved). The firing rate overlap was calculated between the No Object' ROI and the other three ROIs. We used rate overlaps from all trials to test which trial combinations produced the largest rate overlap (and therefore the strongest object traces) ([Figure S4B](#)).

To determine the presence of object traces, we used the rate overlap with the Object Moved trial. For this, a shuffling procedure where the ROI identified in No Object' was placed at random locations in the Object Moved rate map ("Shuffled ROI") was applied. Bins of the Shuffled ROI that fell outside the boundaries of the rate map were wrapped around to the opposite site of the rate map. This was repeated 1000 times, and for each shuffle iteration, the rate overlap between the actual ROI in the No Object' trial and a Shuffled ROI in the Object Moved trial was calculated. All 1000 values of rate overlap were then used to build a shuffled distribution. If the rate overlap between the ROI in No Object' and Object Moved was greater than the 99th percentile of the shuffled distribution, the cell was classified as having an object-trace field ([Figures S4C and S4D](#)).

Bayesian decoding algorithm

For this and all other sections using population activity, we used $n = 15$ tetrode sessions with a minimum of 25 and a maximum of 39 cells recorded simultaneously (the threshold for including a tetrode session being that it had a minimum of 25 cells) and $n = 2$ Neuro-pixels sessions with 620 cells (from two probes) and 135 cells (from a single probe), respectively.

To assess whether population activity in the hippocampus contains information about objects in the environment, we used a decoding approach. Building upon an existing Bayesian decoder,²² we used rate maps to infer whether samples of neural activity occurred in the absence or presence of an object (No Object or Object trial). We divided the data into two groups, for training and testing (first half of data used for training and second half for testing; then vice versa in a separate iteration). While the training data was used to create rate maps, the testing data consisted of samples that the decoder classified as coming from the "No Object" or "Object" trial. Specifically, one sample of the testing data was a vector of spike counts from all hippocampal cells in a 1 second time window:

$$Data = \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_N \end{pmatrix}$$

where d_i is the number of spikes produced by each hippocampal cell in the 1 second time window and N is the total number of recorded cells in a session. Overall, our goal is to compare the likelihood that the “Data” are from the “No Object” trial against the likelihood that they are from the “Object” trial. This is given by the likelihood ratio:

$$\text{likelihood ratio} = \frac{P(\text{Data} | \text{No Object})}{P(\text{Data} | \text{Object})}$$

To calculate $P(\text{Data} | \text{No Object})$, because “Data” could have occurred when the animal was in any spatial bin, we started by introducing another term x on the right-hand side of the probability and asking how likely it is that the “Data” are from a specific spatial bin x (50×50 spatial bins, $x = 1, 2, \dots, 2500$) in the No Object trial: $P(\text{Data} | x, \text{No Object})$.

To calculate the likelihood $P(\text{Data} | x, \text{No Object})$, we assume that each cell fires according to a Poisson distribution with mean firing rate $r_i(x)$, where i is the cell index and x is the spatial location of the animal. Each $r_i(x)$ is obtained by going to the rate map for cell i and extracting the firing rate in spatial bin x . In that sense, rate maps constructed from training data are used as a look-up table. The likelihood was calculated by:

$$P(\text{Data} | x, \text{No Object}) = \prod_{i=1}^N \frac{r_i(x)^{d_i}}{d_i!} e^{-r_i(x)}$$

where d_i is the number of spikes produced by cell i and where the multiplication over Poisson distributions from N cells indicates the assumption of independence between firing of different cells. Note that the more consistent the set of spike counts $\{d_i\}$ are with the set of firing rates $\{r_i(x)\}$, the larger the value of the likelihood.

The likelihood was next multiplied by a prior,

$$P(x | \text{No Object})$$

given by the fraction of time the animal spent in spatial bin x during the No Object trial. This accounts for the intuition that the “Data” are more likely to come from a given spatial bin if the animal spent a lot of time in that bin. We then marginalised over all spatial bins x (50×50 bins),

$$P(\text{Data} | \text{No Object}) = \sum_x P(\text{Data} | x, \text{No Object}) P(x | \text{No Object})$$

to obtain the overall likelihood of the No Object trial. The same calculation is performed for the Object trial to obtain the overall likelihood of the Object trial, $P(\text{Data} | \text{Object})$. The ratio of these terms determines the decoder’s output. That is, if $\frac{P(\text{Data} | \text{No Object})}{P(\text{Data} | \text{Object})} > 1$, the decoder infers the “Data” as being from the No Object trial; otherwise, from the Object trial.

Since the length of an experimental trial was 1800 seconds, $1800/2 = 900$ seconds could be used as testing data. With a bin size of 1 second, this gave us 900 testing samples. Given the two iterations described above (using either the first or second half of the trial as testing data), we had $900 \times 2 = 1800$ testing samples. The decoding accuracy was the fraction of 1800 samples classified correctly. The chance level of the decoder was calculated through a shuffling procedure (see [Shuffling trial labels](#)).

Calculation of decoding accuracy as a function of distance and allocentric angle

To understand how decoding performance depended on the animal’s spatial relationship with the object, we asked how decoding accuracy changed as a function of the animal’s distance and allocentric angle relative to the object (Figures 4A–4C and 4D–4F, respectively). To examine the role of distance, we divided spike and position data into 5 distance bins, forming concentric rings around the object (Figure 4A). By creating training data and testing data for each specific distance bin and feeding these separately into the decoder, we could assess how decoding accuracy differed as a function of distance from object. For each experimental session, we optimized distance bins so that we had an equal number of testing samples within each bin. For example, if the animal spent little time near the object, distance bin 1 was made larger to obtain more samples. This helped to account for differences in the animal’s behavior in different experimental sessions. Specifically, the length of an experimental session was 1800 seconds, allowing $1800/2 = 900$ seconds to be used for testing data. This gave $900/5 = 180$ seconds of testing data per distance bin. Because the bin size was 1 second, we had 180 testing samples per distance bin. The decoding accuracy in each distance bin was the fraction of 180 testing samples classified correctly. To examine the role of allocentric angle, we used an analogous procedure. In this case, we divided spike and position data into 5 angular bins, forming sectors of a circle (Figure 4D), and created training and testing data separately for each angular bin.

Performance of decoder

Effect of bin size

To assess the performance of the decoder, we first asked how decoding performance depended on the bin size of testing data. For this analysis, we ran the decoder separately using bin sizes of 10, 50, 100, 500, 1000, 5000 and 10000 ms and calculated the decoding accuracy. Larger bin size corresponds to observing more spike counts from each hippocampal cell, which will lead to spike counts $\{d_i\}$ that are more similar to the mean firing rates $\{r_i(x)\}$. Consistent with this idea, performance increased with increasing

bin size. Specifically, performance increased rapidly up to 1000 ms, after which it saturated and gains of increasing bin size were lower (Figure 3C). Based on these results, we chose 1000 ms as bin size for all main analyses.

Effect of number of cells

To assess how decoding performance depended on the number of cells fed into the decoder (Figure 3E), we used a subsampling approach. We randomly selected a subset of k cells from a Neuropixels dataset with a total of $N = 620$ cells, iterating over $k = \{1, 5, 10, 20, 30, 50, 100, 200, 300, 400, 500\}$ cells. For each k , we ran the decoder 100 times, using different subsets of cells each time.

Controls for results of decoder

Fixed distance bins

To ensure that our findings did not depend on using distance bins with an equal number of samples (described in Calculation of decoding accuracy as a function of distance and allocentric angle), we repeated the main analysis with fixed distance bins. In this case, we divided spike and position data into distance bins with fixed values of 0-20 cm, 20-40 cm, 40-60 cm, 60-80 cm and 80-100 cm for all experimental sessions, ignoring differences in the animal's behavior.

Decoding 1st half vs. 2nd half of No Object trial

To ensure that our findings could not arise without the presence of an object, we attempted to classify data samples belonging to either the 1st and 2nd half of the No Object trial, as a negative control. For this, we divided the No Object trial at its midpoint. For example, a 1800 seconds long trial would be split at 900 seconds, with 450 seconds for training data and 450 seconds for testing data for each one of the halves. The two halves were then used as input to the decoder (rather than the No Object and Object trials).

Decoding in control experiments without objects

To further test whether our findings could arise without objects, we performed the decoding in control experiments without objects present (Figure 2B). For this, the trials Control 1 and Control 2 were used as input to the decoder (rather than the No Object and Object trials).

Shuffling trial labels

To ensure that our findings reflected a real relationship between population activity and trial type, we shuffled the trial labels of all testing data samples. That is, after shuffling, a spike count vector labelled "No Object" had 50% chance of being labelled "Object." To perform the shuffling, we first concatenated testing data from the "No Object" and "Object" trials into an $N \times T$ spike matrix, where N is the number of cells and T the total number of timepoints. To shuffle the trial labels, we first randomly permuted the set of integers $\{1, 2, 3, \dots, T\}$ using the MATLAB function `randperm`. For example, consider the simple case of $T = 10$. The original arrangement of the $N \times 1$ spike count vectors in the matrix would be $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$, where timepoints 1-5 would come from "No Object" and timepoints 6-10 would come from "Object." The permuted rearrangement might be $\{7, 1, 2, 6, 10, 4, 8, 3, 9, 5\}$. After permutation, the first $T/2$ spike count vectors were labelled "No Object" and the last $T/2$ spike count vectors were labelled "Object." This shuffling procedure destroys any real relationship between population activity and trial type.

Random partitioning into training and testing data

In the main analysis, we used the first half of the session as training data and the second half as testing data (and vice versa in a separate iteration). To confirm that our findings did not depend on such a partitioning, we used an alternative partitioning method where we randomly chose which set of timepoints would constitute training data. The random partitioning was analogous to the shuffling procedure above, but without concatenating spike matrices from the "No Object" and "Object" trials. For example, if data from "No Object" consisted of T timepoints, we generated a random permutation of the set of integers $\{1, 2, 3, \dots, T\}$ and used the first $T/2$ integers as timepoints for training data and the last $T/2$ integers for testing data (and vice versa in a separate iteration).

Removing object-tuned cells

To test whether object-tuned cells (object cells and object-vector cells) were necessary for the distance-dependent decrease in decoding accuracy (Figures 4B and 4C), we removed all object-tuned cells (237/1189 cells, 19.9%) (Figure 1C) and re-ran the decoder.

Removing effects of prior

The standard decoder used a prior that included information about the animal's occupancy (see above). To ensure that our findings (Figures 4B and 4C) did not depend on using this prior, we replaced it with a uniform probability distribution (a prior that assumes that the animal spent an equal amount of time in each bin) and re-ran the decoder.

Sampling analysis

In the Object trial, the object constitutes a physical barrier preventing the animal from exploring some places. This may prevent neural activity patterns associated with that place from occurring, possibly providing a trivial explanation for high decoding success near the object. In general, parts of the environment never or rarely explored in one trial (but explored in the other trial) could potentially lead to high decoding accuracy. To ensure that lack of sampling of some parts of the environment did not influence our results (Figures 4B and 4C), we identified spatial bins in which the animal had spent less time than a threshold value (100 ms, 500 ms, 1000 ms or 2000 ms) in either trial and excluded those spatial bins from both trials. That is, only spatial bins explored longer than the threshold value in both trials were included. After excluding spatial bins with low sampling, we re-ran the decoder. Note that for this analysis, we increased the rate map bin size from 3 cm^2 to 10 cm^2 , since otherwise almost no bins would be explored longer than the sampling threshold.

Behavioral analyses I: Decoding

To test whether behavioural confounds could explain the distance-dependent decrease in decoding accuracy (Figures 4B and 4C) (which could happen, for example, if the animal behaves differently between trials, and more different near the object), we developed decoders that tried to infer the trial of data samples based on the animal's behaviour. We focused on three key and independent behavioural features: speed, head direction and spatial occupancy. The behavioural decoders were based on the same framework as the standard decoder applied to neural data (see [Bayesian decoding algorithm](#)), used the same number of spatial bins (50×50 bins) and divided data into training and testing data the same way. However, there are three differences between the behavioural decoders and the standard decoder:

- The sample of testing data is now a scalar (speed and head direction) or 2-dimensional vector (occupancy), rather than a vector of spike counts.
- The likelihood no longer uses a Poisson distribution, but other distributions that are more appropriate models of each behaviour.
- The training data are constructed by calculating "behaviour maps" rather than firing rate maps.

Speed decoder

For the speed decoder, each sample of testing data was now a scalar representing the animal's instantaneous speed (cm/s) (sampled at 25 Hz). We estimated the animal's speed (in this section denoted by s) as the norm of the vector that has as components the speed in the x and y directions. These were calculated as the change in position of the animal in the x and y directions in consecutive time bins, divided by the size of the time bin and smoothed with a Lowess filter. All speed values were included, including zero speed (immobility). Probability distributions for speed typically had a large peak at low speed values, followed by rapid decay at larger speed values (see [Figure S6G](#) for examples). We therefore used a gamma distribution for the likelihood, which provides good fits for unimodal and exponentially decaying distributions (the exponential distribution being a special case of the gamma distribution). The gamma distribution was a reasonable model for speed: after having learnt model parameters in training data, we could accurately predict the animal's speed in testing data (statistics in [Table S1](#)). The model parameters ($a(x)$ and $b(x)$, see below) were estimated by maximum likelihood. The likelihoods of parameters had well-defined peaks and covered small intervals of the parameter space (size of interval covering 95% of the likelihood's area under the curve: for $a(x)$, 1.22 ± 0.85 ; for $b(x)$, 4.98 ± 1.07 ; median interval size \pm SD), compared to the range of inferred parameter values ($a(x)$: from 0.011 to 42.40; $b(x)$: from 2.69×10^{-4} to 36.38), suggesting that we were able to infer model parameters with reasonable precision. Using the gamma distribution, we calculated the likelihood as:

$$P(\text{Data} | x, \text{No Object}) = P(s | a(x), b(x), \text{No Object}) = \frac{s^{a(x)-1} e^{-\frac{s}{b(x)}}}{b^{a(x)} \Gamma(a(x))}$$

where s is the animal's instantaneous speed, x indexes the spatial bins (50×50 bins in total), $\Gamma(\cdot)$ is the gamma function, and $a(x)$ and $b(x)$ are shape and scale parameters that are functions of x because they describe the animal's speed within every spatial bin (analogous to how the mean parameter of the Poisson distribution described the cell's firing within every spatial bin).

To construct training data (using the animal's speed from half the trial), we divided speed observations in the No Object trial into spatial bins (based on the animal's location at the time of the speed measurement), so that each spatial bin contained a vector of speed values. For every spatial bin x we then fitted a gamma distribution to the observed speed values and estimated the parameters $a(x)$ and $b(x)$ via maximum likelihood. This way we obtained "maps" for $a(x)$ and $b(x)$ (analogous to firing rate maps). We repeated this procedure using data from the Object trial. Finally, the obtained parameters $a(x)$ and $b(x)$ could then be used to calculate the likelihood. As in the standard decoder, for every testing sample we then marginalise over spatial bins to obtain the overall likelihood of the No Object and Object trials, and finally classify the sample based on the ratio of likelihoods of the No Object and Object trials (see [Bayesian decoding algorithm](#) for details).

We verified that the speed decoder accurately identified speed differences between trials present in simulated data. For this, we either simulated (1) speed distributions that were different between the No Object and Object trials, or (2) speed distributions that became more and more different near the object. For (1), each speed measurement was drawn randomly from a gamma distribution with parameters $a = 2$, $b = 2$ for all x (No Object trial) or $a = 2$, $b = 10$ for all x (Object trial). For (2), speed measurements were drawn randomly from a gamma distribution with parameters $a = 2$, $b = 2$ (No Object trial, distance bins 1-5) or $a = 2$, $b = 10$ (Object trial, distance bin 1), $a = 2$, $b = 8$ (Object trial, distance bin 2), $a = 2$, $b = 6$ (Object trial, distance bin 3), $a = 2$, $b = 4$ (Object trial, distance bin 4) and $a = 2$, $b = 2$ (Object trial, distance bin 5). That is, in the Object trial, the b parameter approached the value of b in the No Object trial with further distance from the object, while the a parameter remained the same. Simulated speed values were sampled at 25 Hz (once per position sample). The animal's true position at the time of sampling determined the distance bin of the speed measurement (and therefore the parameters of the gamma distribution in simulation 2). Position data were taken from the 16 experimental sessions.

Head direction decoder

For the head direction (HD) decoder, each sample of testing data was a scalar representing the animal's instantaneous HD (radians) (sampled at 25 Hz). The animal's HD was determined from the relative positions of LEDs or reflective markers on the implant. Because HD distributions were bimodal or unimodal, we used a mixture of two Gaussians to calculate the likelihood. As before, after having learnt model parameters in training data (means, standard deviations and weights of the Gaussians, see below), the model could

accurately predict the animal's HD in testing data (statistics in Table S1). The model parameters (means, standard deviations and weights, see below) were estimated by maximum likelihood. The likelihood of model parameters had well-defined peaks and covered small intervals of the parameter space (size of interval covering 95% of the likelihood's area under the curve: for means of Gaussians, 0.20 ± 0.096 ; for standard deviations of Gaussians, 0.27 ± 0.23 ; for weights of Gaussians, 0.18 ± 0.037 ; median interval size \pm SD) compared to the range of inferred parameter values (means: from 0 to 2π ; standard deviations: from 6.06×10^7 to 6.18; weights: from 0.013 to 1), suggesting that we were able to infer model parameters with reasonable precision.

Using a mixture of two Gaussians, we calculated the likelihood as:

$$P(\text{Data} | x, \text{No Object}) = P(h | \mu_1(x), \sigma_1(x), w_1(x), \mu_2(x), \sigma_2(x), w_2(x), \text{No Object}) = w_1(x) \left(\frac{e^{-\frac{(h - \mu_1(x))^2}{2\sigma_1(x)^2}}}{\sigma_1(x)\sqrt{2\pi}} \right) + w_2(x) \left(\frac{e^{-\frac{(h - \mu_2(x))^2}{2\sigma_2(x)^2}}}{\sigma_2(x)\sqrt{2\pi}} \right)$$

where h is the animal's head direction, x indexes the spatial bin, $\mu_1(x)$ is the mean of the first Gaussian, $\sigma_1(x)$ is the standard deviation of the first Gaussian and $w_1(x)$ is the weight of the first Gaussian, and similarly for the second Gaussian. Again, the mean, standard deviation and weight of each Gaussian are functions of x , describing the animal's HD within every spatial bin. Note that the model is appropriate also when HD is unimodally distributed, because a mixture of two Gaussians includes one Gaussian as a special case (for example, the weight of one Gaussian can be set to zero, or the two Gaussians can be assigned the same mean and standard deviation). Finally, the model is also appropriate when the HD distribution is uniform, because the Gaussians can be assigned very large standard deviations.

To construct training data (using the animal's HD from half the trial), we divided HD observations in the No Object trial into spatial bins (based on the animal's location at the time of measurement), so that each spatial bin contained a vector of HD values. For every spatial bin x we then fitted a mixture of two Gaussians to the observed HD values and estimated the means, standard deviations and weights via maximum likelihood. This way we obtained "maps" for all 6 parameters (analogous to firing rate maps). We repeated this procedure using data from the Object trial. Finally, the obtained parameters could be used to calculate the likelihood. As in the standard decoder, we then marginalise over spatial bins to obtain the overall likelihood of the No Object and Object trials, and finally classify the trial based on the ratio of likelihoods of the No Object and Object trials (see Bayesian decoding algorithm for details).

We verified that the HD decoder accurately detected differences in HD between two trials in simulated data. For this verification, we simulated HD distributions from a single Gaussian since this is a special case of a mixture of two Gaussians. We either simulated (1) HD distributions that were different between the No Object and Object trials, or (2) HD distributions that became more and more different near the object. For (1), each HD sample was drawn randomly from a Gaussian distribution with parameters $\mu = 2$ and $\sigma = 0.5$ for all x (No Object trial) or $\mu = 5$ and $\sigma = 0.5$ for all x (Object trial). For (2), HD samples were drawn randomly from a Gaussian distribution with parameters $\mu = 2$ and $\sigma = 0.5$ (No Object trial, distance bin 1), $\mu = 2$ and $\sigma = 1$ (No Object trial, distance bin 2), $\mu = 2$ and $\sigma = 1.5$ (No Object trial, distance bin 3), $\mu = 2$ and $\sigma = 2$ (No Object trial, distance bin 4) and $\mu = 2$ and $\sigma = 10$ (No Object trial, distance bin 5). In the Object trial, the parameters were the same except that $\mu = 5$ for all distance bins. That is, the standard deviations of HD distributions from the No Object and Object trials increased further away from the object, so that they both approached uniform distributions (and could no longer be distinguished). Simulated HD values were sampled at 25 Hz (once per position sample). The animal's true position at the time of sampling determined the distance bin of the HD measurement (and therefore the standard deviation of the Gaussian in simulation 2). Position data were taken from the 16 experimental sessions.

Occupancy decoder

For the occupancy decoder, each sample of testing data represented the animal's instantaneous position in space (sampled at 25 Hz). Because the decoder itself calculates the likelihood of different spatial bins, and then marginalises over spatial bins (see Bayesian decoding algorithm), position will now appear on both sides of the probability expressions, with different roles. On the left-hand side is the animal's continuous location in space, as part of the testing data; on the right-hand side is the binned location in space, as part of the training data.

$$P(\text{Data} | \text{No Object}) = P(\text{pos} | \text{No Object}) = \sum_{ij} P(\text{pos} | x_i, y_j, \text{No Object}) P(x_i, y_j | \text{No Object})$$

where pos is the continuous location of the animal (a 2-dimensional vector of real numbers), as part of testing data; x_i is the i^{th} spatial bin for the x coordinate and y_j is the j^{th} spatial bin for the y coordinate (i and j are integers between 1 and 50), and both x_i and y_j are part of training data.

The occupancy decoder can now be simplified analytically, because each sample of testing data could only have occurred within a specific spatial bin (for example, the observation $\text{pos} = (149.2 \text{ cm}, 148.0 \text{ cm})$ could only have occurred when the animal was in the northeast corner of the environment (x_{50}, y_{50}) – that is, when the animal was in the "last" bin of an environment with 50 bins in the x coordinate and 50 bins in the y coordinate. In general, this means that:

$$P(\text{pos} | x_k, y_l, \text{No Object}) = 1$$

whenever k and l are the indices of the spatial bins that the observed position sample falls into. Also,

$$P(\text{pos} | x_i, y_j, \text{No Object}) = 0$$

whenever $i \neq k$ or $j \neq l$. Plugging these equations into the likelihood gives:

$$\begin{aligned} P(\text{Data} | \text{No Object}) &= \sum_{ij} P(\text{pos} | x_i, y_j, \text{No Object}) P(x_i, y_j | \text{No Object}) \\ &= P(\text{pos} | x_k, y_l, \text{No Object}) P(x_k, y_l | \text{No Object}) \\ &= (1) P(x_k, y_l | \text{No Object}) = P(x_k, y_l | \text{No Object}) \end{aligned}$$

That is, the probability that the sample of testing data comes from the No Object trial is proportional to the time that the animal spent in the relevant spatial bin. Therefore, the likelihood ratio becomes a ratio of occupancies in the relevant spatial bin:

$$\frac{P(\text{Data} | \text{No Object})}{P(\text{Data} | \text{Object})} = \frac{P(x_k, y_l | \text{No Object})}{P(x_k, y_l | \text{Object})}$$

so that we assign “No Object” if the animal spent more time in the spatial bin (x_k, y_l) in the No Object trial than the Object trial (and vice versa).

Given the simplification above, we constructed training data for the decoder by calculating occupancy maps (amount of time animal spent in every spatial bin, divided by total trial duration).

We verified that the occupancy decoder could accurately identify differences in occupancy between trials in simulated data. For this, we either simulated (1) occupancies that were different between the No Object and Object trials, or (2) occupancies that became more and more different near the object. We simulated occupancies by drawing random position samples either from a uniform distribution with edges from 0 to 150 cm (for both the x and y coordinates) or from a uniform distribution with edges that covered only half the environment (y edges from 0 to 75 cm in the No Object trial and y edges from 75 to 150 cm in the Object trial; x edges from 0 to 150 cm in both trials). That is, each position sample could either be drawn from the whole environment or from half the environment. To change how strongly occupancy differed between the trials, we varied the probability of drawing position samples from the whole or half environment. For (1), the probability of drawing from half the environment was 1. That is, all position samples were drawn from the lower (No Object trial) or upper (Object trial) half of the environment. For (2), in both trials, the probability of drawing from (half the environment, whole environment) was either (1, 0) (distance bin 1), (0.75, 0.25) (distance bin 2), (0.5, 0.5) (distance bin 3), (0.25, 0.75) (distance bin 4) or (0, 1) (distance bin 5). That is, in distance bin 5 position samples were drawn only from the whole environment (generating similar occupancy between the trials), while in distance bin 1 position samples were drawn only from half the environment (generating different occupancy between the trials). Also, for (2), edges of uniform distributions were adjusted so that (a) drawing from the whole environment was replaced by drawing from the whole distance bin; and (b) drawing from half the environment was replaced by drawing from half the distance bin. The number of position samples drawn was the same as the total number of position samples in real data.

Behavioral analyses II: Subsampling

To test whether the decoder applied to neural data (Figures 4B and 4C) made use of behavioural differences between the trials, we reran the decoder after matching distributions of behaviour. We specifically matched behaviour between trials in every distance bin to account for any distance-dependent behavioural effects (for example, if the animal ran had biased head direction near the object). To do this, we first calculated distributions of each behaviour (speed, head direction and occupancy) in every distance bin.

The distributions for speed had bins from 0 cm/s to 100 cm/s in steps of 5 cm/s. The distributions for head direction had bins from 0 to 2π radians in steps of $\pi/10$. The distributions for occupancy were two-dimensional and had both x and y bins from 0 cm to 150 cm in steps of 30 cm.

For each speed, HD or occupancy bin, we compared the number of observations from the No Object trial (n_{NO}) and the number of observations from the Object trial (n_O). If we found that:

$$n_{NO} > n_O$$

we threw away $n_{NO} - n_O$ random samples from the No Object trial, so that the number of samples matched between the trials. Conversely, if we found that,

$$n_O > n_{NO}$$

we threw away $n_O - n_{NO}$ random samples from the Object trial. This process was repeated for all bins, until distributions were completely matched (see example for speed in Figure S6G). After matching distributions within every distance bin, we ran the standard decoder (Bayesian decoding algorithm) on the subsampled neural data. The process was done separately for speed, head direction and occupancy.

Subsampling for decoder

To understand how widely distributed the object-centered population coding was across the recorded cells, we asked how many cells we needed to see the distance-based decrease in decoding accuracy (Figure 6). To answer this, we used a subsampling procedure before feeding data into the decoder. We used Neuropixels data from a rat implanted with two probes in left hippocampus ($N = 620$ cells). The subsampling was analogous to the shuffling described under Shuffling trial labels. We took the set of all cells $\{1, 2, 3, \dots, N\}$, where each integer corresponds to a different cell identity, and randomly permuted the set using the MATLAB function `randperm`. We then took the first k integers of the permuted set, corresponding to k cells, and used them to create the subset. The number of subsampled cells, k , took the values 1, 5, 10, 30, 100, 200, 300, 400 and 500. For each k , we ran the decoder 20 times, (20 times being sufficient to get negligible error bars), so that the results were independent of cell identities and cell properties. In this way, we could determine how the number of cells influenced the decoding results (Figure 5A) and show the distribution of distance coding across the entire hippocampal population.

Measuring strength of distance-specific reorganization

To estimate the strength of the distance-specific reorganization in populations of cells (and in individual neurons), we calculated two separate measures (Figures 6B, 6D, and 6F). Firstly, after having run the decoder and having obtained the curve showing decoding accuracy as a function of distance, we calculated

$$\text{distance tuning} = \text{decoding accuracy bin \#1} - \text{decoding accuracy bin \#5}$$

A strong distance-specific reorganization should be associated with large distance tuning. Secondly, we fitted a line to the same curve via least squares⁵⁵,

$$\hat{x} = (A^T A)^{-1} A^T b$$

where A was the design matrix:

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \\ 1 & 4 \\ 1 & 5 \end{pmatrix}$$

b was a vector containing the decoding accuracy of the cell in each distance bin i ,

$$b = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \end{pmatrix}$$

and \hat{x} was a 2-element vector containing the intercept (\hat{x}_1) and slope (\hat{x}_2) of the fitted line. A strong distance-specific organization should be associated with a more negative slope (\hat{x}_2). The results obtained using the two measures (distance tuning and slope) were consistent.

We also re-calculated these measures using distance bin #4 as the last distance bin (Table S1).

Significance of distance tuning in individual cells

To test whether distance tuning values in individual cells were significant (Figure 6B), we compared the actual values to shuffled data. To create shuffled data, we circularly shifted spike trains (as described under shuffling of spike data) and re-calculated distance tuning (100 repetitions per cell). Cells with actual values above the 99th percentile of the shuffled distribution were considered to have significant positive tuning (Figure 6B, red bars) while cells with actual values below the 1st percentile were considered to have significant negative tuning (Figure 6B, yellow bars).

Population vector correlation

To assess the similarity between populations of rate maps from the No Object and Object trials, we calculated the correlation between stacks of rate maps from each trial. For each experimental session, we stacked the rate maps from the No Object trial and Object trial separately, while preserving the cell ordering. This procedure generated two $N \times m \times n$ matrices, where N is the number of cells recorded in the session, m is the number of y-bins and n is the number of x-bins of the rate maps ($m = n = 50$ bins). We then calculated the Pearson correlation $r(x, y)$ between two N -dimensional vectors of firing rates, one from each trial, for every spatial bin. After computing the Pearson correlations, we obtained an $m \times n$ map showing Pearson correlations as a function of space (Figure 4G). After computing the correlation map, we allocated correlation values $r(x, y)$ to distance bins and calculated the mean correlation in every distance bin (5 bins in total, bin size = 20 cm). Repeating this analysis for all experimental sessions ($n = 16$ from both tetrode and Neuropixels data), we were able to determine how correlations between the stacks changed as a function of distance from object (Figure 4H).

Definition of words and counts

To understand which specific features of population activity explained the decrease in decoding accuracy as a function of distance from object, we considered two descriptions of population activity that we refer to as words and counts. These concepts have previously been used to study neural coding at the population level.^{23–25} First, we expressed population activity in an $N \times T$ matrix, where N is the number of cells and T the number of timepoints in one trial. Each entry of the matrix corresponds to the number of spikes generated by each cell in specific time bin. To obtain mostly ones and zeroes (in rare cases, 2 or 3) we used 10 ms bins. We defined a specific pattern of activity and silence across the population as a “word.” We had as many words as timepoints sampled at 10 ms. The length of each word was the number of cells in the population (tetraode data: 25–39 cells; Neuropixels data: subsampled to contain 30 cells). In contrast, we defined the total number of spikes generated by the population in a time bin as the “count.” Words are a detailed description of population activity that preserves the cell identity, while counts are a coarse description of population activity. To illustrate, in a population of 4 neurons, the word [1, 0, 0, 1] would mean that the 1st and 4th neuron fired one spike in that time bin, while the 2nd and 3rd neurons remained silent. The total count would be $2 = 1 + 1$. The word [0, 0, 0, 1] would mean that the 4th neuron fired one spike, while all others remained silent. The total count would be 1.

Mutual information

To quantify how much information words and counts provided about the presence of an object in the animal’s environment, we calculated the mutual information (MI)^{56,57} between the population response and the trial type (No Object and Object). The MI tells us how much information one variable (e.g., the words) provide about another variable (e.g., the trial type) independent of assuming any specific relationship between the two variables (e.g., the relationship could be completely non-linear). For words, the calculation was given by

$$MI_{\text{words}} = \sum_{\text{trial, word}} P(\text{trial, word}) \log_2 \left(\frac{P(\text{trial, word})}{P(\text{trial})P(\text{word})} \right)$$

where $P(\text{word})$ and $P(\text{trial})$ are the marginal probabilities of the words and trial type and $P(\text{trial, word})$ is the joint probability. Similarly, for counts the calculations was given by

$$MI_{\text{count}} = \sum_{\text{trial, count}} P(\text{trial, count}) \log_2 \left(\frac{P(\text{trial, count})}{P(\text{trial})P(\text{count})} \right)$$

where $P(\text{count})$ and $P(\text{trial})$ are the marginal probabilities of the counts and trial type and $P(\text{count, word})$ is the joint probability. All probabilities were estimated from the frequency distributions. When calculating the MI of words, the number of bins was the same as the number of words in both trials (No Object and Object). For counts, the number of bins was determined by the maximum counts observed in both trials. For example, if the maximum number of spikes in a single time bin was 14, bins would range from 1 to 14 (in steps of 1). Being highly frequent and uninformative, the all-zero vector was excluded from our analyses. For the same reason, the zero count was also excluded.

Bias-subtraction

Calculation of MI on datasets with small numbers of samples is well-known to lead to a bias where one overestimates the amount of information.⁵⁸ To ensure that our findings were not caused by sampling problems, we used a shuffling procedure to calculate the bias and subtract it from the mutual information. For this, we first shuffled trial labels as described for the decoder under Shuffling trial labels. After shuffling, a vector originally labelled “No Object” now had a 50% chance of being labelled “Object.” In other words, the shuffling destroys any relationship between the two variables (response and trial type) and shows the amount of “information” we should expect without such a relationship. The shuffling was performed 50 times, after which we calculated the mean MI across all 50 shuffles. This value (the “bias”) was subtracted from the MI calculated on real data.

Entropy

To estimate the entropy of words, we calculated

$$H(p) = - \sum_{i=1}^W p_i \log_2(p_i)$$

where H is the entropy function, $p = \{p_i\}$ is the entire probability distribution of words, p_i is the probability of word i occurring during the trial, and W is the total number of words. The calculation for the entropy of counts was done similarly, but with probabilities of words replaced by probabilities of counts.

Word distributions

To directly visualize the change of words between No Object and Object, we calculated frequency distributions of words. For this, we first created a spike matrix based on activity from both trials (as in Figure 5A, bin size = 10 ms, $n = 25$ –39 cells). We then calculated the occurrences of each word, plotting the result in a histogram (Figure 5D). To facilitate comparison between No Object and Object (for example, word #8 in No Object is the same as word #8 in Object), we applied the same sorting to the words in both trials. Specifically, we sorted words from No Object in order of descending frequency (Figure 5D, top) and applied the same sorting to the Object trial (Figure 5D, bottom).

To ensure that any differences in words between No Object and Object could not arise by chance, we randomly shuffled the trial labels (as described under “shuffling trial labels”) and calculated word distributions on the shuffled data (Figure S7A).

Downsampling words

Because the total number of words and counts differ by an order of magnitude (in a given trial, we usually observe a few hundred different words and approximately ten different counts), we wanted to assess how this difference influences the mutual information. For this, we downsampled the number of words to be equivalent to the number of counts (Figure S7E). In a given trial, we identified how many unique counts were present, c , and then kept the c words with the largest frequencies. The frequency distribution was then normalized, creating a downsampled probability distribution for words. Note that because we keep the most probable words, the total number of samples of words and counts are roughly the same even after downsampling (so that any difference in the number of samples should not have a large effect on the mutual information).

Kullback-Leibler divergence between word distributions

To quantify how strongly word distributions (as in Figure 5D) differed between the No Object and Object trials, we calculated the Kullback-Leibler (KL) divergence,^{56,59} which is given by,

$$D_{KL}(p||q) = \sum_{i=1}^W p_i \log_2 \left(\frac{p_i}{q_i} \right)$$

where p_i is the probability of word i occurring in the No Object trial, q_i is the probability of word i occurring in the Object trial, W is the total number of words from the No Object and Object trials, $p = \{p_i\}$ is the entire probability distribution of words from the No Object trial while $q = \{q_i\}$ is the entire probability distribution from the Object trial. Probabilities of the words were estimated in each trial as the normalized frequency distribution (number of times each word occurred/total number of timepoints). For visualization purposes, we plotted histograms of only the most 100 frequent words in each trial (Figure 5D). However, the KL divergence was calculated using all words.

Statistical tests

All statistical analyses were performed in MATLAB (version 2020a). P-values are indicated in figures and legends. Wilcoxon signed-rank tests were used for paired data and Wilcoxon rank sum tests for unpaired data. Kruskal-Wallis tests were used for comparing differences between multiple groups. Binomial tests were used to compare fractions of observed cells (for example, object cells) to chance levels. Correlations were determined using Pearson’s product-moment correlation coefficients (r). The significance level was set to $p = 0.05$. All statistical tests were two-sided.