Article

Cell Reports

Multiplexing of Information about Self and Others in Hippocampal Ensembles

Graphical Abstract

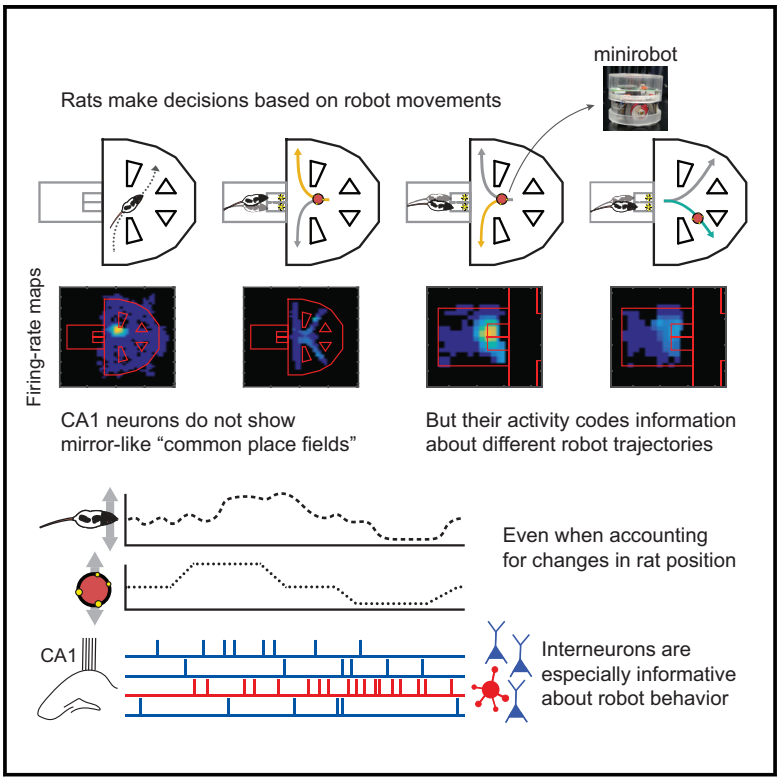

Authors

Jeroen J. Bos, Martin Vinck, Pietro Marchesi, ..., Jadin C. Jackson, Paul F.M.J. Verschure, Cyriel M.A. Pennartz

Correspondence

c.m.a.pennartz@uva.nl

In Brief

Bos et al. study hippocampal coding of an external agent's location using a minirobot and find no evidence of mirror-like neurons coding location. They discover that CA1 firing patterns (especially interneurons) carry information about robot behavior and highlight the importance of controlling for confounds due to changes in animal position.

Highlights

- Rats track the position of a robot to perform an observe-and-decide discrimination task
- CA1 neurons did not show any place fields common to both the rat and the robot
- Firing patterns of CA1 neurons were modulated by the behavior of the robot
- Fast-spiking interneurons were especially informative about robot movements

Bos et al., 2019, Cell Reports *29*, 3859–3871
December 17, 2019
https://doi.org/10.1016/j.celrep.2019.11.057

CellPress

Cell Reports
Article

Multiplexing of Information about Self and Others in Hippocampal Ensembles

Jeroen J. Bos,[1,2,6] Martin Vinck,[1,3,6] Pietro Marchesi,[1,2,6] Amos Keestra,[1] Laura A. van Mourik-Donga,[1,2] Jadin C. Jackson,[4] Paul F.M.J. Verschure,[5] and Cyriel M.A. Pennartz[1,2,7,*]
[1]Swammerdam Institute for Life Sciences, Center for Neuroscience, Faculty of Science, University of Amsterdam, Amsterdam, the Netherlands
[2]Research Priority Program Brain and Cognition, University of Amsterdam, Amsterdam, the Netherlands
[3]Ernst Strüngmann Institute for Neuroscience in Cooperation with Max Planck Society, Deutschordenstraße 46, 60528 Frankfurt, Germany
[4]Medtronic, 7000 Central Avenue NE, Minneapolis, MN 55432, USA
[5]Institute for Bioengineering of Catalonia (IBEC), Barcelona Institute of Science and Technology (BIST), Barcelona, Spain
[6]These authors contributed equally
[7]Lead Contact
*Correspondence: c.m.a.pennartz@uva.nl
https://doi.org/10.1016/j.celrep.2019.11.057

SUMMARY

In addition to coding a subject's location in space, the hippocampus has been suggested to code social information, including the spatial position of conspecifics. "Social place cells" have been reported for tasks in which an observer mimics the behavior of a demonstrator. We examine whether rat hippocampal neurons may encode the behavior of a minirobot, but without requiring the animal to mimic it. Rather than finding social place cells, we observe that robot behavioral patterns modulate place fields coding animal position. This modulation may be confounded by correlations between robot movement and changes in the animal's position. Although rat position indeed significantly predicts robot behavior, we find that hippocampal ensembles code additional information about robot movement patterns. Fast-spiking interneurons are particularly informative about robot position and global behavior. In conclusion, when the animal's own behavior is conditional on external agents, the hippocampus multiplexes information about self and others.

INTRODUCTION

The hippocampus plays a key role in the formation of episodic memory, defined by one's personal experiences set in space and time (Eichenbaum et al., 2012; O'Keefe and Conway, 1978; O'Keefe and Dostrovsky, 1971; Scoville and Milner, 1957; Tulving, 1983). Classically, hippocampal neurons are thought to code an organism's own position in space, as they exhibit an increase in firing rate when an animal visits a particular location in its environment (the "place field"; McNaughton et al., 2006; O'Keefe and Conway, 1978; O'Keefe and Dostrovsky, 1971; Scoville and Milner, 1957). However, previous studies suggest that the hippocampus may also play a role in social memory (Hitti and Siegelbaum, 2014; Kogan et al., 2000; Okuyama et al., 2016; Sliwa et al., 2016; Stevenson and Caldwell, 2014; Viskontas et al., 2009; von Heimendahl et al., 2012; but see Bannerman et al., 2001; Squires et al., 2006). For instance, recent studies in mice suggested important roles of CA2 and ventral CA1 neurons in social memory, more specifically, the ability to memorize the familiarity of a previously encountered conspecific (Hitti and Siegelbaum, 2014; Okuyama et al., 2016; Stevenson and Caldwell, 2014). Further evidence suggests that dorsal CA1 neurons may code information about the spatial position and movement of other agents (Danjo et al., 2018; Ho et al., 2008; Omer et al., 2018; Zynyuk et al., 2012).

Analogous to "mirror neurons," originally discovered in monkey ventral premotor cortex (area F5; di Pellegrino et al., 1992; Gazzola et al., 2007; Rizzolatti and Fabbri-Destro, 2008; Rizzolatti and Sinigaglia, 2016), Danjo et al. (2018) suggested that also CA1 neurons of rat hippocampus can display mirror-like firing. In a task in which a rat observed the trajectory of another rat (demonstrator) to decide whether it should subsequently run the same or an alternative trajectory, the authors found a subset of neurons showing an identical place field for the subject and demonstrator. Such a "common place field" can be considered a spatial equivalent of a "mirror neuron" in monkey F5. One difficulty in interpreting these results is that both the recorded and demonstrator rat could move simultaneously, making it difficult to disentangle whether neural responses should be ascribed to the demonstrator or the self, which may react to the other agent in various ways. In another recent study, Omer et al. (2018) also applied a spatial observation task, in which an observer bat was required to fly the same route to a target location as was demonstrated by another bat shortly beforehand. They also identified a subset of "social place cells" showing high similarity between place fields for the self and the other bat.

In the present study, we investigated the influence of another agent on hippocampal firing patterns, using not an "observe-and-mimic" paradigm as in Danjo et al. (2018) or Omer et al. (2018) but a design in which the subject first explores a maze in isolation and subsequently performs two "observe-and-decide" discrimination tasks based on the movements of a

remotely operated minirobot (an "e-puck;" Mondada et al., 2009). We used a minirobot instead of a conspecific because the robot's movements could be precisely controlled, thus being insensitive to the observer's behavior. Confined to an observation cage and thus separated from the robot, the rat was trained to choose between two directly apposed reward wells on the basis of the robot's movement pattern. This method allowed us, first, to study whether mirror-like "common place fields" are found in other multi-agent situations than in an observe-and-mimic paradigm.

Second, we assessed hippocampal firing sensitivity to dynamic variables related to the robot's behavioral trajectories on the maze. Specifically, we asked: (1) Are mirror-like "common place fields" generally found when another, behaviorally significant agent is present in the subject's vicinity, or are they rather exceptional? (2) Should robot-induced changes in CA1 firing rate be ascribed to the behavior of the external agent or to changes in the rat's own position? (3) What information do firing patterns of single cells and of neuronal populations convey about a movement pattern of the other agent when correcting for changes in rat position?

Our paradigm did not reveal "common place fields," and moreover, changes in rat position were significantly predictive of robot parameters, which indicates a serious confound that should be taken into account in future studies on social neural coding that may entail changes in the subject's behavior in response to another agent. However, when correcting for this confound, we found that CA1 firing patterns coded additional information about global robot behavior, robot position, and the animal's engagement in a task requiring robot tracking. Fast-spiking (FS) interneurons coded a remarkably large amount of information about the other agent.

RESULTS

Behavioral Task: Rat Observing Robot Movements

To study how an external agent affects CA1 firing patterns, we aimed to reduce potential confounds due to variability in the agent's behavior by using a minirobot ("e-puck;" Mondada et al., 2009) whose behavior was stereotyped and remotely controlled. Briefly, each recording session was divided into three epochs. During the first and last epochs (Rat-on-Maze 1 and 2; Figure 1A) the rat freely explored a maze consisting of several connected alleys. During the second epoch (Observation Period), the rat moved into a separate cage compartment and was exposed to the robot moving across the maze. During the Task Phase, the rat could obtain reward by tracking robot movements and making a nose poke into one of two closely apposed reward wells, where the rewarded side corresponded to the side of the (outbound) travel direction of the robot. The Task Phase presented two task types (Front Task [FT] and Mid Task [MT]) which differed in spatial movement patterns of the robot (in FT, moving toward the rat and to the side in the alley closest to the rat; in MT, moving away from the rat and to the side; see Figure 1B) but were identical in the behavioral response rules the rat had to follow to earn reward (i.e., the well delivering reward on a correct choice was always determined by the side of outbound robot movement). The Observation Period also contained a Free Roaming Control (FR-control) phase, during which the robot moved around on the maze without programmed consequences for the rat, and a No-Robot Control (NR-control) phase, during which the robot was removed from the maze while the rat remained in the observation cage. In both the Front Task and Mid Task, behavioral performance was above chance level (Figure S1).

Examining the Occurrence of "Common Place Fields" in Hippocampus

First, we examined whether a given hippocampal place cell, defined as coding a specific location visited by the animal, also codes the location of the robot when it enters the cell's place field while being observed by the rat in the cage (Figure 2A). Analysis of 608 CA1 neurons from three rats indicated that 68.8% of the cells (418 of 608) exhibited place fields as gauged from the rat's own spatial behavior during the Rat-on-Maze 1 and Rat-on-Maze 2 epochs. These classic place fields were compared with spatial maps of CA1 cell firing rate computed as a function of robot position. When considering data from the full Task Phase (FT and MT combined; see Figure 1A) we found 128 (21.1%) cells showing firing patterns with some selectivity for the robot's position on the maze ("firing fields;" see STAR Methods). Of these neurons, 82 cells (64.1%) also showed place fields during the Rat-on-Maze epochs. Of these 82 units, a subset of 43 cells (52.4%) showed peak firing rates across globally overlapping maze segments for self and other positions (Figure 2C).

Although this spatial overlap between self- and other-related firing patterns may suggest the occurrence of "common place cells," as previously reported by Danjo et al. (2018), the patterns we observed are different. First, robot firing fields encompassed a much larger fraction of space than found for the rat's own place fields during the Rat-on-Maze periods (Figure 2D; median and interquartile range for robot firing fields during Task Phase, 34.2% [18.4%–55.6%]; rat fields during Rat-on-Maze 1, 9.7% [5.4%–14.8%]; Rat-on-Maze 2, 8.7% [5.5%–15.8%]; throughout the text we report our findings as median and interquartile ranges unless stated otherwise; Friedman $\chi^2 = 67.5$, $p = 2.2 \times 10^{-15}$ with Tukey-Kramer's post hoc test: $p = 1.2 \times 10^{-9}$ and $p = 9.6 \times 10^{-10}$; Task Phase compared with Rat-on-Maze 1 and Rat-on-Maze 2, respectively).

Second, after removing non-specific firing fields (i.e., those covering >50% of visited space), only 8 of 43 units remained as potential "common place cells" (see Figure 2C for a neuron showing an atypical high overlap between rat and robot place fields [24.5%]). Overlap between rat and robot fields appeared to be coincidental, because the robot-based firing-rate maps showed a diffuse, spatially distributed pattern of activity (Figures 2C and 2E). A permutation test (see STAR Methods) confirmed that the overlap of 24.5% in the example in Figure 2C did not differ from chance, and the same held for all other neurons considered for the test. The maximum firing rates of classic place fields were higher than the equivalent fields for the robot (Figure 2E; Rat-on-Maze 1, 9.1 Hz [4.7–15.3 Hz]; Rat-on-Maze 2, 10.4 Hz [5.6–15.2 Hz]; for robot, Task Phase, 2.6 Hz [1.4–4.2 Hz]; Friedman $\chi^2 = 77.4$, $p = 1.6 \times 10^{-17}$ with Tukey-Kramer's post hoc test: $p = 1 \times 10^{-9}$ and $p = 9.6 \times 10^{-10}$ compared with Rat-on-Maze 1 and Rat-on-Maze 2, respectively). Because of these firing-rate differences, none of the neurons showed overlapping firing fields

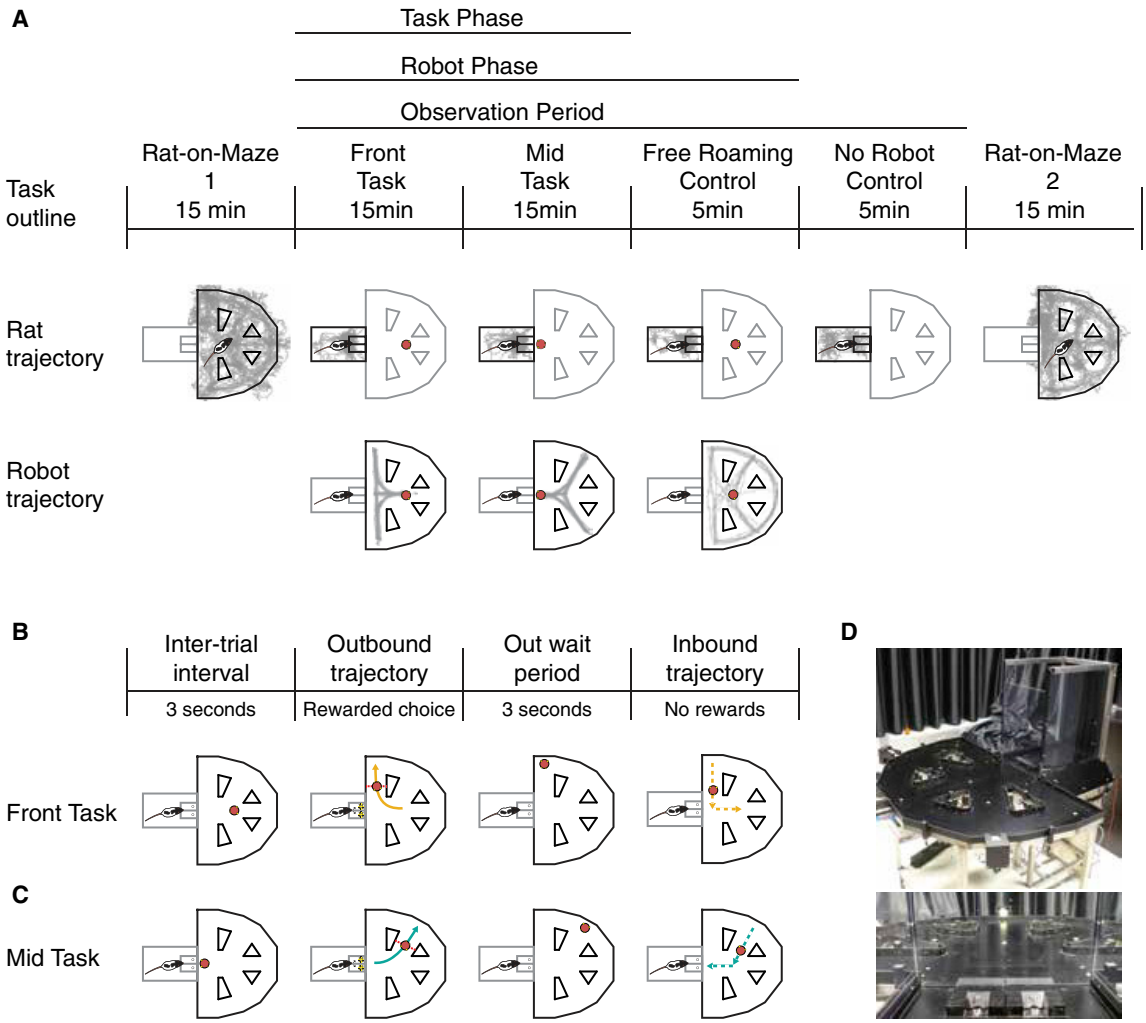

Figure 1. Behavioral Task

(A) For each behavioral phase, the trajectories of the rat and robot (if applicable) are shown in gray (example session). A recording session was divided into several epochs. During the first and last epochs, the rat walked on the maze collecting food pellets (Rat-on-Maze 1 and 2). During the Observation Period, the rat was positioned in a Plexiglas cage attached to the maze and performed two discrimination tasks based on movements of the robot (Front Task [FT] and Mid Task [MT]; see Figure S1 for task performance). The order of FT and MT was randomized across recording sessions. This Task Phase was followed by two control periods. During the Free Roaming control, the robot drove around on the maze, while during the No-Robot control, the robot was removed and the rat was exposed to an empty maze. Neither of these control phases resulted in reward delivery.

(B and C) Outline of a single trial in either the Front (B, ocher) or the Mid (C, cyan) task. After an intertrial interval, the robot moved into one of the side arms. Upon crossing an infrared light beam (dashed red lines), the lights above the reward wells turned on, and the rat made a choice by poking in one of the reward wells. At the end of the outbound trajectory, the robot was stationary for three seconds before moving back to the central position (inbound trajectory; dashed arrows).

(D) Photographs of the behavioral setup: overall configuration with maze and cage (top) and view of the cage with reward wells, as seen from the animal's potion in the cage facing the maze (bottom).

when applying the metric used by Danjo et al. (2018) (their Figure 3) for finding common place fields for self and other. Even when we applied less stringent firing-rate requirements, no mirror-like firing fields were found. In conclusion, the experimental design deployed here does not yield evidence for "common place fields" as described by Danjo et al. (2018).

Effects of Robot Presence on Hippocampal Coding

Despite this initial null result, we did observe that the firing rate of a subset of CA1 neurons varied in relation to the robot's wandering across the maze. To examine this in more detail, we first focused on the impact of the mere presence of the robot on hippocampal firing patterns by comparing the FR-control and NR-control stages. The corresponding firing-rate maps only included cage positions the rat visited in each of two contrasting conditions (see STAR Methods). Figures 3A and 3B show two example cells significantly discriminating between the presence and absence of the robot. Of 608 cells, 218 units (35.9%) showed place fields in at least one of the two conditions. Of these 218 units, 47 (21.6%) differentiated between the robot's presence

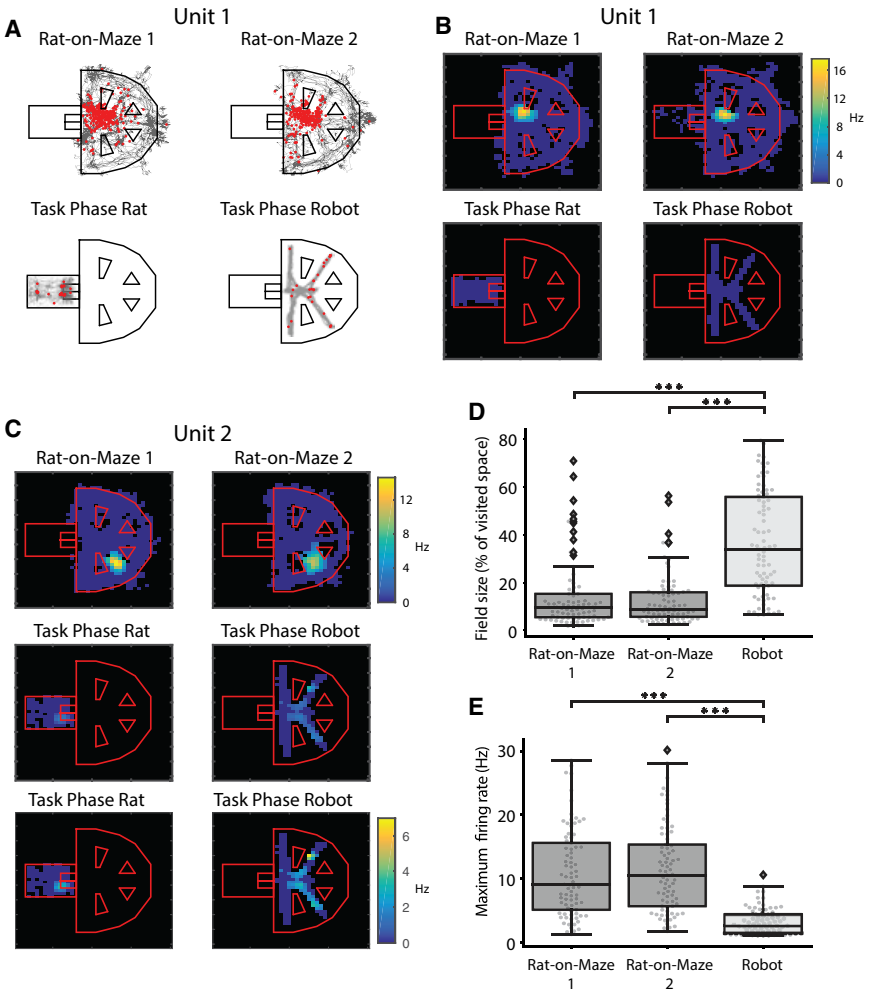

Figure 2. CA1 Firing Fields as a Function of Rat versus Robot Position

(A) Firing pattern of a CA1 unit (red: spikes) plotted as function of rat position for Rat-on-Maze 1 (top left), Rat-on-Maze 2 (top right), and Task Phase (bottom left; gray: rat positions). Bottom right: robot firing field was computed as a function of robot position by linking spike time stamps of the same cell to the corresponding positions of the robot.

(B) Same as (A), but now firing patterns are plotted as firing-rate maps (same color scale for all plots).

(C) CA1 unit that, relative to the total population of recorded CA1 cells, showed a high overlap (24.5%) between its classic place field (during Rat-on-Maze 1 and 2) and robot firing field during Task Phase. For Rat-on-Maze 1 and 2, "field" refers to a conventional place field. For "Task Phase Robot," the firing field was computed as a function of robot position. The two lowermost plots are calibrated to a color scale comprising the maximum firing rate for Task Phase Robot (lower right color scale), showing that there is no marked overlap with the classic place field even when rate remapping is considered.

(D) Boxplots of field size as a percentage of space visited for the 39 units that showed firing fields in three conditions (Rat-on-Maze 1, Rat-on-Maze 2, and Task Phase Robot). Throughout the figures, whiskers indicate the lowest and highest data points within 2 interquartile ranges from the low and high quartiles, respectively.

(E) Maximum firing rate reached within the cell's firing field was markedly lower when computed as a function of robot position relative to classic place fields.

***Significance at α = 0.001, determined using Friedman's test with Tukey-Kramer's as post hoc test.

and absence. This fraction was significantly above chance level (test of proportion, $Z = 6.1$; $p = 1.3 \times 10^{-9}$).

To assess if the presence of the robot could be decoded from CA1 population firing patterns, we trained a random forest decoder (RFD) to distinguish whether the robot was present on the maze (FR-control) or absent (NR-control). The RFD was trained on three types of input: neural data (in the form of binned spike counts), rat position, and the combination of neural data and rat position (see STAR Methods). Decoding on the basis of rat position was included as a control because the presence of the robot may cause the rat's position in the observation cage to change and thereby influence CA1 firing patterns. Although neural data distinguished robot presence versus absence above chance (0.5), the decoding accuracy when using rat position (or the combination of position and neural data) increased relative to neural data alone (Figure 3E; accuracy for neural data, 0.69 [0.63–0.74]; rat position, 0.87 [0.85–0.90]; both, 0.88 [0.86–0.89]; Wilcoxon's signed-rank test, Bonferroni-adjusted; $p = 3.7 \times 10^{-5}$, $p = 3.7 \times 10^{-5}$, and $p = 1.0$ for neural data versus rat position, neural data versus both, and rat position versus both, respectively; for control analyses using different decoders, see Figures S2A and S2B). These results underscore the importance of controlling for variable animal position as a confound in trying to identify influences of the position of another agent on CA1 firing patterns. During the NR-control, nothing eventful was occurring on the maze to capture the rat's attention, and thus the rat often entered a more quiescent state in this phase. In contrast, during FR-control, the robot's presence may constitute a significant motivational and attentional factor influencing the rat because it is globally associated with reward (even though no reward could be acutely obtained in this phase). We further verified whether the ability to distinguish behavioral epochs using the animal's position was due to minor repositioning around the reward spouts or to large locomotor movements in the cage. All decoding analyses were repeated selecting only time points at which the animal was positioned close to the reward wells; we found that the results were largely unaffected, indicating that large locomotor movements in the cage are not crucial to the decoding performance we observe (Figure S2C).

Effects of Engagement in a Task

Next, we asked whether hippocampal coding is sensitive to the subject's engagement in actively tracking the robot's

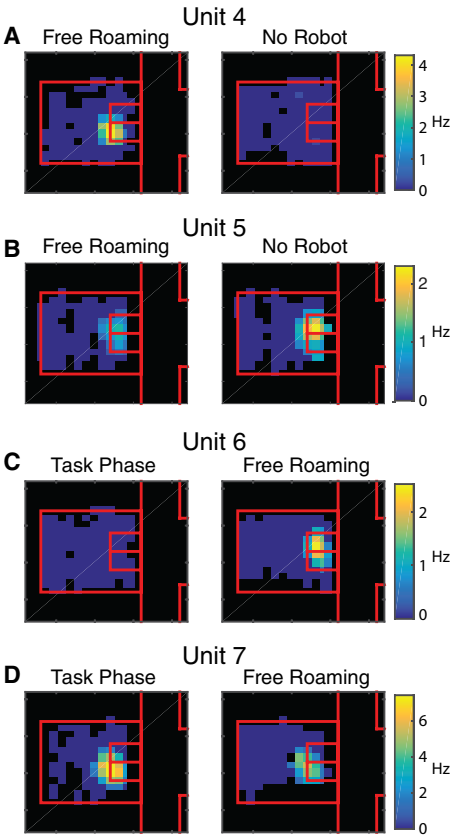

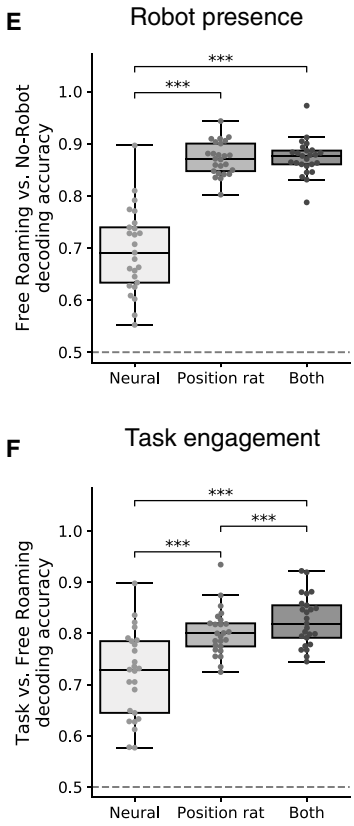

Figure 3. Sensitivity of CA1 Firing Patterns to Robot Presence and Task Engagement

(A) Firing-rate maps of a CA1 unit that discriminated between presence (Free Roaming) versus absence of the robot (No-Robot control), plotted as a function of rat position. Neither phase entailed task requirements. The plots are restricted to the cage compartment of the maze because the rat remained confined to this cage during both phases.

(B) Same as (A), for a CA1 cell showing a more gradual firing-rate modulation by robot presence.

(C and D) Example neurons discriminating between the presence and absence of a task requirement for the rat (Task Phase versus Free Roaming phase). In (C), a place field appears during the Free Roaming control, whereas in (D), firing rate is modulated by task engagement.

(E and F) Decoding robot presence (E) and task engagement (F) using a random forest decoder (RFD). For decoding using logistic regression and a feedforward neural network, see Figures S2A and S2B. The RFD was trained to classify time windows as belonging to one out of two phases of the recording session (for robot presence, Free Roaming versus No Robot; for task engagement, Task Phase versus Free Roaming), on the basis of spiking activity, rat position, and the combination of the two. In this way we determined whether population firing patterns contain information about robot presence and task engagement beyond the information that can be extracted from the movement of the animal. Dots indicate the decoding accuracy of individual recording sessions, averaged over five repetitions of 10-fold cross-validation.

Asterisks indicate significance at $^{*}\alpha = 0.05$, $^{**}\alpha = 0.01$, and $^{***}\alpha = 0.001$.

movements, thus contrasting the FR-control stage with the Task Phase. A significant proportion of units with place fields (31.9% [74 of 232 units]; $Z = 8.79$, $p = 0.000$, test of proportion) discriminated between task engagement and mere exposure to the robot, either via global remapping (Figure 3C) or rate remapping (Leutgeb et al., 2005; Figure 3D).

Population analysis (Figure 3F) showed that although neural data were sufficiently specific to distinguish between Task Phase and FR-control, decoding on the basis of rat position resulted in more accurate decoding of these two conditions. Combining neural data and rat position data modestly improved decoding (accuracy for neural data, 0.73 [0.64–0.78]; rat position, 0.80 [0.77–0.82]; both, 0.82 [0.79–0.85]; Wilcoxon's signed-rank test, Bonferroni-adjusted; $p = 4.2 \times 10^{-4}$, $p = 3.7 \times 10^{-5}$, and $p = 4.7 \times 10^{-4}$ for neural versus position, neural versus both, and position versus both, respectively; but see Figure S2B). Although it may seem surprising that neural data achieved accuracy of only 0.73 given the presence of 74 significantly discriminating units, these units were distributed across all recording sessions, and because decoding was performed separately for each session (and used all neurons), only a few of the discriminating units were available in each training of the decoder. It should be noted that task engagement but not FR-control was paired with reward consumption, significantly affecting the rat's motivational state.

Effects of Specific Robot Movement Patterns

We next focused on the contrast between the FT and MT stages, during which the animal was constantly engaged in discriminative behavior, applying the same set of task rules. Here, the difference between conditions was due to the dichotomy in spatial behavior of the robot. We found a significant fraction of cells (57 of 223 units with place fields [25.6%]; $Z = 7$, $p = 1.9 \times 10^{-12}$) which discriminated between the robot's Front Task versus Mid Task behavior (Figures 4A and 4B). Apparent firing-rate sensitivity to the robot's behavior may arise from accidental spike bursts, and thus we examined whether response patterns were reproducible across trials. To assess this, we analyzed firing rate as a function of linearized robot position and across trials (Figure 4C), and we observed consistent firing-rate differences between FT and MT. In the example cells (Figures 4A–4C), no marked differences between outbound versus inbound trajectories of the robot were found (but see Figure S3).

These findings were corroborated by population analysis using the RFD (Figure 4D). When classifying the two most behaviorally similar task periods (FT versus MT) of our paradigm, RFD performance based on rat position was better than for neural data. However, training the RFD with rat position and neural data combined significantly increased the accuracy above the level reached by position or neural data alone (accuracy for neural data, 0.66 [0.62–0.72]; rat position, 0.71 [0.69–0.74]; both, 0.75 [0.73–0.79]; Wilcoxon's signed-rank test, Bonferroni-adjusted;

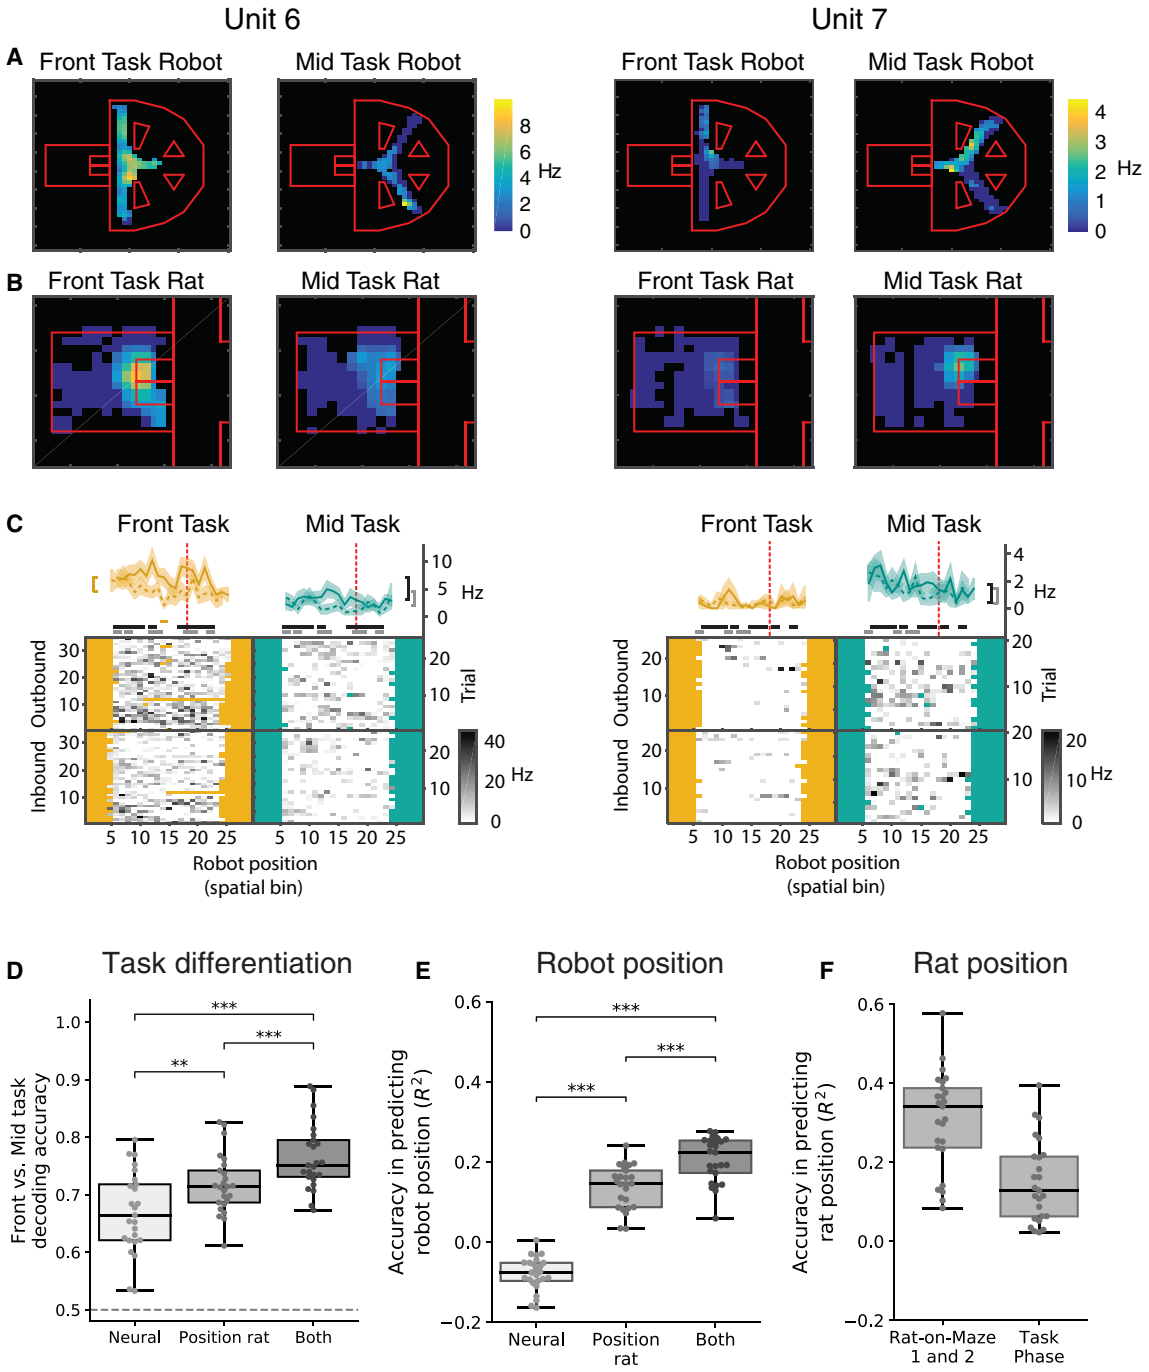

Figure 4. CA1 Neurons Distinguish between Different Patterns of Robot Movement

(A) Two units with firing fields as a function of robot position, showing discrimination between FT and MT. See Figure S3 for discrimination between inbound and outbound movements of the robot.

(B) Classic place field maps of same units, zoomed in on cage compartment. Place cell activity is rate-modulated as a function of task type.

(C) Histograms of binned firing rates as a function of linearized robot position. Same neurons as in (A) and (B) (FT and MT in ocher and teal, respectively). Red dashed line represents maze location that, when trespassed by the robot, marks future availability of reward if the rat makes a correct choice. Rasters: firing rate (gray) per spatial bin and per trial, ordered from bottom to top. Bins visited by the robot for less than 0.2 s (shown in color corresponding to task type) were excluded. Variation between trials is partially explained by rat moving in and out of the cell's place field. Top: firing-rate averages (solid, outbound; dashed, inbound trajectories) and SEs (shading). Horizontal lines below averages, bins significantly discriminating between two conditions (α = 0.05, false discovery rate [FDR] corrected). Ocher, significant outbound-inbound difference for Front Task (Wilcoxon's signed rank test). Black and gray, difference between FT and MT during outbound (black) or inbound trajectories (gray; Mann-Whitney U test). Vertical side bars, significant firing-rate difference across whole trajectories (α = 0.05).

(legend continued on next page)

$p = 9.6 \times 10^{-3}$, $p = 3.7 \times 10^{-5}$, and $p = 6.0 \times 10^{-5}$ for neural versus position, neural versus both, and position versus both, respectively). This suggests that the different robot trajectories in FT and MT induce modulations in both place field activity and rat position, which can be jointly exploited by the RFD to improve its prediction accuracy in classifying task type. In other words, CA1 firing patterns code information on specific robot behavior that cannot be extracted from rat position alone and cannot be attributed to incidental firing, as the improvement is consistent across sessions and decoding is performed with cross-validation.

We next tested to what extent robot position could be predicted from the neural data, rat position, or the combination of the two. Although neural data alone were generally unable to predict robot location (Figure 4E), the combination of neural data and rat position resulted in a modest but consistent improvement of the decoding of robot position relative to rat position alone. This indicates that robot location induces additional modulations in hippocampal coding, not accounted for by changes in rat position (goodness of fit R^2 for neural data = −0.08, range −0.10 to −0.05; rat position, 0.13 [0.09–0.18]; both, 0.21 [0.16–0.26]; Wilcoxon's signed-rank test, Bonferroni-adjusted; $p = 2.3 \times 10^{-9}$, $p = 2.3 \times 10^{-9}$, and $p = 2.7 \times 10^{-9}$ for neural versus position, neural versus both, and position versus both, respectively). By way of comparison, we verified that our data also allowed us to predict rat location, both on the maze and in the cage compartment (Figure 4F). A comparison of decoding quality showed that the position of the rat on the maze can be predicted with higher precision (explained variance $R^2 = 0.34$ [0.24–0.39]) than robot position ($R^2 = 0.21$ [0.16–0.26] using both position and neural data). Note that we used a measure of the explained variance in agent positioning, rather than the Euclidean error in centimeters, because this measure better captures the statistics of fine positional variability of the rat versus the robot (for decoding error expressed as Euclidean distance between true and predicted location, see Figure S4).

Mutual Information Analysis of Coding of Robot Movement Patterns

Next we examined whether CA1 information about robot behavior is coded by subsets of neurons or spread out across the entire population, focusing on FT and MT and correcting for animal position. First, we determined how well the firing of CA1 neurons distinguished between FT and MT by constructing receiver operating characteristic (ROC) curves (which measure the performance of a binary classifier as its classification threshold is varied) and their corresponding area under the curve (AUC) values at individual spatial bins, thus excluding the possibility that the difference in firing between FT and MT is explained solely by rat position (Figure 5A; see Figure S5 for parameter sensitivity of the ROC analysis). Of the 575 cells that satisfied the requirements to be included in the analysis, 312 (54.3%) were found to significantly discriminate between the two task types (Figure 5B).

We computed the mutual information (MI) between each neuron's firing activity and task type, and the conditional mutual information (cMI) between each neuron's firing activity and task type conditional on animal position (Figure 5C; see Figure S6 for parameter sensitivity of this analysis). The latter measure quantifies the average information contained in the firing rate about task type on top of the information that can be extracted from rat position. If the firing activity were only a (noisy) representation of the animal's position, then all the information that it conveys about task type would already be present in rat position, resulting in a cMI of zero. In contrast, a cMI significantly different from zero indicates that the firing rate contains additional information about task type.

Of 605 neurons with non-zero firing rates during the Task Phase, 76 neurons (12.6%; purple and green dots in Figure 5C) carried significant cMI about task type. AUC and cMI values similarly captured information about task type corrected for animal position (Figure 5D), and we found a strong correlation between the two measures, significantly stronger than the correlation between AUC and MI ($\rho = 0.64$ for AUC-cMI, $\rho = 0.50$ for AUC-MI, $p = 1.65 \times 10^{-6}$, Pearson and Filon's z test). Next, we followed up on the population coding of robot position by computing MI and cMI contained in each neuron's firing activity about robot position during the full Task Phase (Figure 5E, where cMI is again conditional on rat position). We found that 157 neurons (26%) carried significant information about robot position on top of the information contained in rat position.

Coding of Robot Behavior by Different Cell Types

To investigate whether different CA1 cell types convey distinct information about robot behavior, we identified 20 fast-spiking cells (FS; putative interneurons) on the basis of their spike waveform (Figures 6A and 6B; Barthó et al., 2004; Henze et al., 2000; Vinck et al., 2016) together with 474 broad-spiking (BS; putative excitatory) and 114 unclassified (UC) units. FS neurons coded more information about task type than BS and UC units, as determined both by AUC value and cMI (Figures 6D and 6E). We found that 95% of interneurons had significant AUC values, compared with 55.1% and 42.7% of BS and UC cells, respectively. Similarly, 55% of interneurons were found to carry significant cMI about task type, whereas this was the case for only 11.9% and 7.9% of BS and UC neurons, respectively (Figure 6G; see also Figure S6). The results were similar when considering information about robot position, where interneurons carried more information than BS and UC units and had a larger proportion of neurons with significant cMI (Figures 6F and 6H; see also Figure S6).

The marked difference between FS and BS neurons may be a result of the FS neurons' having higher overall firing rates, which can influence the information estimates. Thus, we repeated the

(D) Decoding of task type using RFD, following plotting conventions of Figure 3. For decoding using logistic regression and a feedforward neural network, see Figures S2A and S2B.

(E) Decoding of robot position averaged across FT and MT. R^2 represents the explained variance in robot position.

(F) Decoding of rat position using neural data during the Rat-on-Maze 1 and 2 periods and Task Phase. For decoding of rat and robot position expressed as the average Euclidean error, see Figure S4.

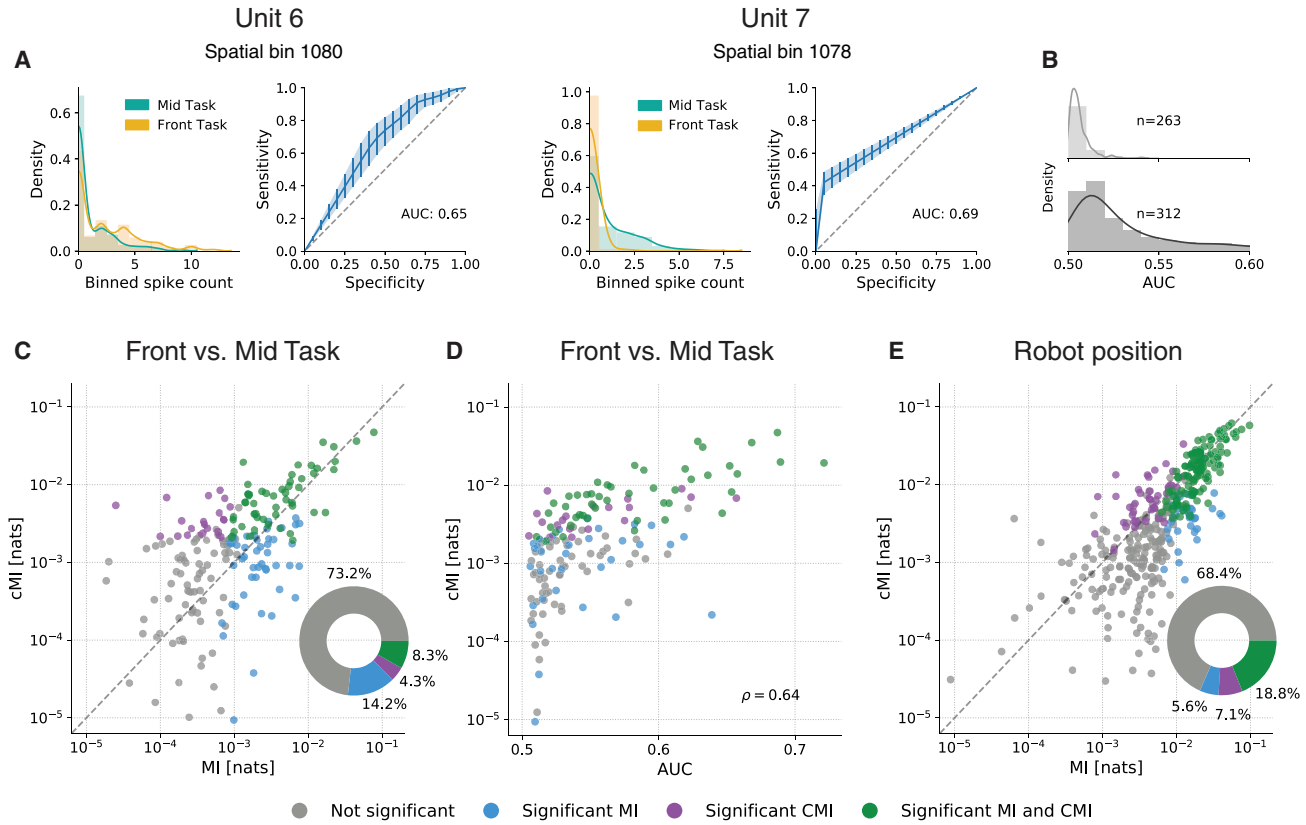

Figure 5. CA1 Neurons Distinguish between Task Types and Robot Position, Even When Correcting for Animal Position

(A) Distribution of binned spike counts for FT versus MT and ROC curves with 95% confidence intervals for the two example cells of Figure 4, computed for the individual spatial bin with maximum AUC for each unit. These cells differentiated between task types even when the animal was at the same location.

(B) Distribution of AUC values for all units with non-significant AUC (top) and significant AUC (bottom). For parameter sensitivity of the ROC analysis, see Figure S5.

(C) Comparison of mutual information (MI) and conditional mutual information (cMI) between neural activity and the binary variable indicating FT versus MT, for each unit (debiased values). Blue dots, units with significant MI (but non-significant cMI); purple dots, units with significant cMI (but non-significant MI); green dots, significant for both MI and cMI; gray dots, significant for neither MI nor cMI. The donut plot indicates the percentage of neurons that belong to each group (with respect to the total number of neurons considered for this analysis, N = 605). Debiased values can be slightly negative and thus not appear in this logarithmic plot. For the dependence of MI and cMI on the temporal binning, see Figure S6.

(D) Comparison of cMI about task phase and AUC values for neurons with significant AUC shows a strong correlation (Pearson's correlation, $\rho = 0.64$, $p = 4.7 \times 10^{-37}$), as both methods reflect information about task type corrected for rat position.

(E) The comparison of MI and cMI between neural activity and robot position identifies a population of neurons (purple and green) whose firing rate is correlated to robot position, even when correcting for rat position.

analysis after lowering the firing rate of FS cells to the average BS rate via random down-sampling (Figure S7): although cMI about robot position is no longer higher in interneurons relative to BS or UC cells, the information about task type, in terms of both AUC and cMI, remained higher for FS cells. This indicates that a bias due to higher firing rates does not fully account for the effect that we observe. Given the difference in information content between BS and FS neurons, we next asked whether the decoding of robot presence, task engagement, and task differentiation (Figures 3E, 3F, and 4D) was driven mostly by the latter group. We repeated the decoding analyses after excluding all putative interneurons and UC units (Figure S2D) and found that, although decoding accuracy was slightly lower (e.g., accuracy in decoding FT versus MT using neural data alone, 0.66 [0.62–0.72] with all cells versus 0.63 [0.60–0.71] for BS cells only, Wilcoxon's signed-rank test, $p = 2.5 \times 10^{-5}$), the results were largely unaffected, confirming that robot-induced modulations are distributed across the population, and are not a prerogative of interneurons.

DISCUSSION

Our main results can be summarized as follows. First, we showed that mirror-like common place fields are not found in rat hippocampus when the animal tracks another agent to make behavioral decisions but does not imitate this agent's behavior. Thus, previous reports (Danjo et al., 2018; Omer et al., 2018) on this phenomenon could not be generalized to the current behavioral paradigm. Second, we found instead that CA1 hippocampal neurons were influenced by the robot through remapping of place fields

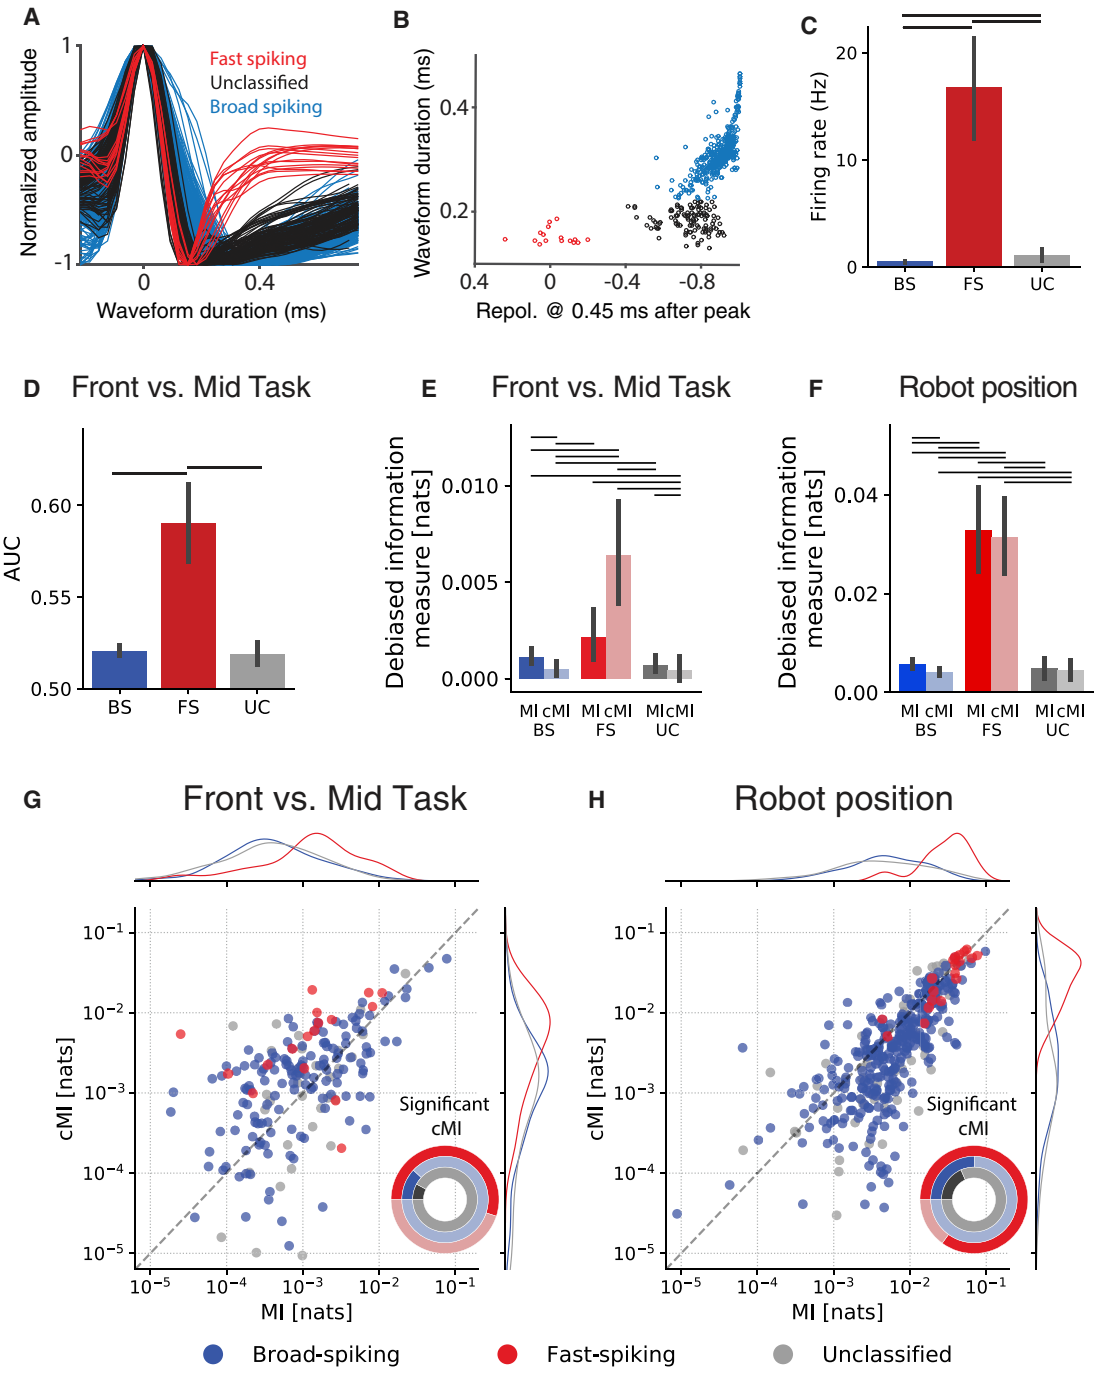

Figure 6. Coding of Task Phase and Robot Position by Different Classes of CA1 Neurons

(A) Normalized spike waveform amplitudes, grouped into fast-spiking (FS; red), unclassified (UC; black), and broad-spiking (BS; blue) neurons on the basis of two characteristics: repolarization level and waveform duration (see STAR Methods).

(B) Distribution of repolarization level against waveform duration.

(C) FS units showed higher firing rates than BS and UC neurons (Mann-Whitney U test, $p = 1.2 \times 10^{-11}$ and $p = 5.4 \times 10^{-10}$ compared with BS and UC respectively, $p = 8.2 \times 10^{-3}$ for BS versus UC). Here, as well as in (D)–(F), bars indicate 95% confidence interval (CI).

(D) Average AUC values show that FS units discriminate more strongly between FT and MT than BS and UC neurons (Mann-Whitney U test, $p = 5.5 \times 10^{-34}$ and $p = 3.0 \times 10^{-25}$ compared with BS and UC respectively, Bonferroni corrected). For parameter sensitivity of the ROC analysis, see Figure S5. For a procedure examining the effect of the higher basal firing rates of FS neurons, see Figure S7.

(legend continued on next page)

coding for the rat's own position. Both the robot's presence and the animal's engagement in the task were associated with consistent changes in CA1 place-cell firing, but information on these conditions could be deduced with greater accuracy from changes in rat position alone (Figure 3). For these two conditions, changes in rat behavior and CA1 firing may well cohere with changes in the animal's motivational and attentional state (Eichenbaum et al., 1987; Fenton et al., 2010; Hölscher et al., 2003; Lansink et al., 2012, 2016). Although manipulations of robot presence and task engagement are informative, the firing-rate modulations observed under these conditions should thus not be taken as firm evidence for an active role of hippocampal neurons in coding other-agent behavior.

In contrast, the identification of neuronal subsets discriminating FT and MT is particularly significant because these patterns were displayed while the animal operated under the same task rules and reward contingencies. Firing-rate modulations were not dependent on incidental bursting activity and were largely consistent across trials (Figure 4C). Although the rat's own position provided significant information on the robot's FT or MT behavior, the decoding of task type was significantly improved when neural and rat position were combined (Figure 4D). Using ROC analysis and information theoretic measures, we found that considerable fractions of cells conveyed significant information about task type and robot position, even after information already contained in rat position was removed (Figure 5).

Interestingly, FS neurons showed higher AUC and cMI values than other cell types, even when correcting for their high basal firing rates (Figure S7), suggesting a strong involvement of interneurons in socially induced firing-rate modulations (Figure 6). Nonetheless, it should be stressed that information about the robot was distributed across the wider population of cell types (Figure 6; Figure S2D). It remains unknown how social information may reach the hippocampus, specifically its FS cells (comprising a large majority of parvalbumin-positive [PV] interneurons; Pawelzik et al., 2002; Sik et al., 1995) in area CA1, but it is of note that chronic social isolation reduces the number and parvalbumin expression level of PV interneurons in the dorsal hippocampus, including area CA1 (Filipović et al., 2018; Ueno et al., 2017). Moreover, the functioning of hippocampal GABAergic cells is altered in disease conditions impairing social abilities, such as schizophrenia and major depression (Benes et al., 2008; Thompson Ray et al., 2011; cf. Piskorowski et al., 2016). The neuropeptide oxytocin, which has been implicated in regulating social discrimination and memory through its hippocampal targets (Raam et al., 2017), strongly depolarizes fast-spiking interneurons in both areas CA1 and CA2 (Owen et al., 2013; Tirko et al., 2018). Ventral hippocampal pyramidal cells and PV interneurons have been implicated specifically in retrieval of social memory (Deng et al., 2019; Okuyama et al., 2016). Whether the discriminatory firing behavior of FS cells in response to the other agent in our study may reflect social encoding, retrieval, or non-mnemonic signaling remains to be investigated in future studies.

A main point of discussion concerns the difference between the present results and those of Danjo et al. (2018), which share the use of rats as subjects. Danjo et al. (2018) found mirror-like place representations in 5% of their recorded neurons, whereas we found no convincing case. In contrast to our design, the observer and demonstrator rats in Danjo et al. (2018) shared the same task environment and ran the same maze trajectories. Possibly, "common place cells" emerge selectively when rats plan and execute locomotion actions in the same space as where the demonstrator is located. During the planning of a future trajectory, hippocampal assemblies may generate prospective sequences of spike trains (Johnson and Redish, 2007; Pfeiffer and Foster, 2013). A "common place field" could thereby be generated if the conventional place field of a given cell is also activated when the self observes and internally simulates a path that includes the place field under scrutiny. This potential confound deserves further testing in the paradigm used by Danjo et al. (2018), for instance by having the rat plan trajectories in the temporary absence of a demonstrator. Although our results show that an observer rat can actively track a demonstrator agent despite a physical barrier separating them, it remains possible that "common place fields" were not found in our study because a robot was used instead of a conspecific. Omer et al. (2018) tested for conspecific versus other-object effects in bats and found largely segregated subsets of CA1 neurons responding to these two classes. If such differential coding also applies to rats, one would not expect mirror-like coding to be completely absent in our dataset.

Furthermore, Omer et al. (2018) argued that even if an animal observing a demonstrator agent would generate prospective activity concerning its own position within planned future paths, this activity would still represent the spatial position of the demonstrator. Although we follow this argument, we note that this would equally apply to the coding of non-living objects or of spatial target locations to which the subject plans to travel, and thus it remains important to investigate how conspecifics are coded differently than other objects or target locations and if and how their social relevance specifically matters to neural coding. In conclusion, it remains to be examined whether mirror-like place cell phenomena should be reinterpreted as a consequence of cued path planning (while remapping of place fields reflects the more general effect of other-agent influences) or carry a specifically social signature.

Another potential confound concerning mirror-like hippocampal firing is that neural activity correlating to another agent is not

(E) Average MI (solid colors) and cMI (light colors) show that FS cells carry more information about task type than BS and UC neurons. We used Wilcoxon's signed rank test and the Mann-Whitney U test for comparisons within and across cell types, respectively (horizontal lines indicate significance at $\alpha = 0.01$ after Bonferroni correction). For the dependence of MI and cMI on the temporal binning, see Figure S6.
(F) Average MI (solid colors) and cMI (light colors) about robot position is also higher for FS neurons.
(G) Comparison of MI and cMI per cell type, with marginal probability densities. Donut plot, fraction of neurons carrying significant cMI about task type (solid color), irrespective of MI.
(H) Comparison of MI and cMI about robot position, showing that FS cells code information about robot position in the upper ranges of MI and cMI values.

representing a direct neural coding of the other's position but reflects an indirect effect due to the observer's behavioral reactions to the agent. Indeed, significant accuracy in decoding behaviorally contrasting conditions was achieved purely on the basis of changes in the observer's own positions (Figures 3 and 4). Even though the task rules were the same for FT versus MT performance, subtle changes in rat behavior correlated to task type (Figure 4D). These changes in self-behavior constitute a serious confound that has, to our knowledge, not or insufficiently been taken into account in previous studies. Nonetheless the neural data contributed significantly to FT-MT decoding and carried information about task type and robot position even when taking rat position into account (Figure 5). That an added value of neural data was not found when decoding robot presence versus absence (Figure 3E) may be explained by a greater variability of rat behavior during, and dissimilarity between, these two conditions. Although Danjo et al. (2018) did not examine this potential confound, Omer et al. (2018) did correct for large head movements of the observer bat.

These and other arguments are equally relevant when comparing less recent studies with the present results. In line with our findings, Mou and Ji (2016) did not report "common place fields" but also found no spatial tuning to a demonstrator. This difference may relate to the fact that the authors did not use a task to incentivize the observer animal to pay attention to the other agent, which may be an important factor for identifying spatial information pertaining to this agent as reported here. The lack of an incentivizing task in a study by von Heimendahl et al. (2012) might likewise be a factor explaining why they found only minor modulations of CA1 firing rate by conspecifics. Two other studies (Ho et al., 2008; Zynyuk et al., 2012) did reveal changes in CA1 place field activity by conspecifics or another moving object, but in both cases changes in hippocampal activity may have been due to changes in the subject's own positioning as it was allowed to be in proximity to the external agents.

The present findings shed new light on hippocampal coding of external agents in an animal's environment. Hippocampal mirror-like spatial coding may occur in particular social coding situations but is not generally found in situations in which another significant agent is active. We predict that apparent mirror-like firing will be restricted to observe-and-mimic situations because of path planning computations, rather than specifically reflecting social coding. The firing-rate modulations we observed rather resemble hippocampal remapping phenomena as found in other environmental manipulations (Lansink et al., 2012; Leutgeb et al., 2005; McNaughton et al., 2006) but are nonetheless informative on task engagement, task type, and the other agent's position (Figures 3, 4, and 5). How may such information be linked to processing in the mirror system? In monkeys, this system not only includes the ventral premotor cortex (area F5 in macaques) but also parietal motor areas PFG and AIP (Gazzola et al., 2007; Rizzolatti and Sinigaglia, 2016). The hippocampal system is indirectly connected to this system via, for instance, the superior temporal sulcus (STS) and inferior temporal lobe (Rizzolatti and Sinigaglia, 2016). When considering that animal behavior, in either a social or an interspecies context, often entails both self-motion and tracking of other agents, the question arises as to which neural mechanisms may underlie near simultaneous representations of self and others. The present findings suggest that primary coding of self-position by the CA1 network is multiplexed with rate-modulated coding of another agent's position. A question for future research is whether this combined coding generalizes to other brain areas implied in non-self-coding and may form the basis of joint distributed coding across premotor-parietal-temporal lobe networks.

STAR★METHODS

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

- KEY RESOURCES TABLE
- LEAD CONTACT AND MATERIALS AVAILABILITY
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
  - Behavioral setup and minirobot
  - Behavioral training and behavioral procedure during recordings
  - Surgical procedure and tetrode recordings
  - Data acquisition
- DATA ANALYSIS
  - Spike sorting and waveform quantification
  - Place fields and robot firing fields
  - Mirror-like place fields and overlap between rat place field and robot firing field
  - Cell sensitivity to contrasting behavioral conditions
  - Linearized peri-event space histograms
  - Population coding using Random Forest Decoders
  - Population coding using logistic regression and feed-forward neural networks
  - Receiver operating characteristic analysis
  - Conditional mutual information
- QUANTIFICATION AND STATISTICAL ANALYSIS
- DATA AND CODE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at https://doi.org/10.1016/j.celrep.2019.11.057.

ACKNOWLEDGMENTS

We would like to thank Romy Bakker, Barbara Samson, and Laury van Bedaf for their help during experiments and Santiago Brandi, Armin Duff, and the late Johan Lisman for helpful discussion. We acknowledge the software tools provided by Kenneth Harris (University College London; Klustakwik) and by A. David Redish (University of Minnesota, Minneapolis; MClust). We thank Ron Manuputy, Harry Beukers Jr., Gerrit Hardeman, Ed de Water, Daan Zwart, Cees van de Biggelaar, Eric Hennes, and Mattijs Bakker for providing the behavioral setup and hyperdrives. This work was supported by grants EU FP7 217148 ("Synthetic Forager" to P.F.M.J.V. and C.M.A.P.), EU FP7 ICT Grant 270108 ("Goal-Leaders"), and the European Union's Horizon 2020 Framework Programme for Research and Innovation under Specific Grant Agreement No. 785907 (Human Brain Project SGA2; both to C.M.A.P.).

AUTHOR CONTRIBUTIONS

J.J.B., M.V., P.F.M.J.V., and C.M.A.P. conceived and designed the experiment. J.J.B. and M.V. performed the tetrode recordings, with support from

L.A.v.M.-D. and J.C.J. J.J.B., M.V., P.M., A.K., and C.M.A.P. analyzed the data. C.M.A.P., J.J.B., and P.M. wrote the paper in collaboration with the other co-authors. C.M.A.P. supervised and coordinated the project.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: May 21, 2019
Revised: October 3, 2019
Accepted: November 13, 2019
Published: December 17, 2019

REFERENCES

Bannerman, D.M., Lemaire, M., Beggs, S., Rawlins, J.N., and Iversen, S.D. (2001). Cytotoxic lesions of the hippocampus increase social investigation but do not impair social-recognition memory. Exp. Brain Res. *138*, 100–109.

Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K.D., and Buzsáki, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. J. Neurophysiol. *92*, 600–608.

Benes, F.M., Lim, B., Matzilevich, D., Subburaju, S., and Walsh, J.P. (2008). Circuitry-based gene expression profiles in GABA cells of the trisynaptic pathway in schizophrenics versus bipolars. Proc. Natl. Acad. Sci. U S A *105*, 20935–20940.

Bishop, C.M. (2006). Pattern Recognition and Machine Learning (Springer).

Danjo, T., Toyoizumi, T., and Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. Science *359*, 213–218.

DeLong, E.R., DeLong, D.M., and Clarke-Pearson, D.L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. Biometrics *44*, 837–845.

Deng, X., Gu, L., Sui, N., Guo, J., and Liang, J. (2019). Parvalbumin interneuron in the ventral hippocampus functions as a discriminator in social memory. Proc. Natl. Acad. Sci. U S A *116*, 16583–16592.

di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., and Rizzolatti, G. (1992). Understanding motor events: a neurophysiological study. Exp. Brain Res. *91*, 176–180.

Eichenbaum, H., Kuperstein, M., Fagan, A., and Nagode, J. (1987). Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. J. Neurosci. *7*, 716–732.

Eichenbaum, H., Sauvage, M., Fortin, N., Komorowski, R., and Lipton, P. (2012). Towards a functional organization of episodic memory in the medial temporal lobe. Neurosci. Biobehav. Rev. *36*, 1597–1608.

Fenton, A.A., Lytton, W.W., Barry, J.M., Lenck-Santini, P.P., Zinyuk, L.E., Kubík, S., Bures, J., Poucet, B., Muller, R.U., and Olypher, A.V. (2010). Attention-like modulation of hippocampus place cell discharge. J. Neurosci. *30*, 4613–4625.

Filipović, D., Stanisavljević, A., Jasnić, N., Bernardi, R.E., Inta, D., Perić, I., and Gass, P. (2018). Chronic treatment with fluoxetine or clozapine of socially isolated rats prevents subsector-specific reduction of parvalbumin immunoreactive cells in the hippocampus. Neuroscience *371*, 384–394.

Gazzola, V., Rizzolatti, G., Wicker, B., and Keysers, C. (2007). The anthropomorphic brain: the mirror neuron system responds to human and robotic actions. Neuroimage *35*, 1674–1684.

Glaser, J.I., Chowdhury, R.H., Perich, M.G., Miller, L.E., and Kording, K.P. (2017). Machine learning for neural decoding. arXiv. https://arxiv.org/abs/1708.00909.

Glorot, X., Bordes, A., and Bengio, Y. (2011). Deep sparse rectifier neural networks. In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, PMLR, pp. 315–323.

Gray, C.M., Maldonado, P.E., Wilson, M., and McNaughton, B. (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. J. Neurosci. Methods *63*, 43–54.

Henze, D.A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K.D., and Buzsáki, G. (2000). Intracellular features predicted by extracellular recordings in the hippocampus in vivo. J. Neurophysiol. *84*, 390–400.

Hitti, F.L., and Siegelbaum, S.A. (2014). The hippocampal CA2 region is essential for social memory. Nature *508*, 88–92.

Ho, S.A., Hori, E., Kobayashi, T., Umeno, K., Tran, A.H., Ono, T., and Nishijo, H. (2008). Hippocampal place cell activity during chasing of a moving object associated with reward in rats. Neuroscience *157*, 254–270.

Hölscher, C., Jacob, W., and Mallot, H.A. (2003). Reward modulates neuronal activity in the hippocampus of the rat. Behav. Brain Res. *142*, 181–191.

Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. J. Neurosci. *27*, 12176–12189.

Kingma, D.P., and Ba, J. (2014). Adam: a method for stochastic optimization. arXiv. https://arxiv.org/abs/1412.6980.

Kogan, J.H., Frankland, P.W., and Silva, A.J. (2000). Long-term memory underlying hippocampus-dependent social recognition in mice. Hippocampus *10*, 47–56.

Kraskov, A., Stögbauer, H., and Grassberger, P. (2004). Estimating mutual information. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *69*, 066138.

Lansink, C.S., Bakker, M., Buster, W., Lankelma, J., van der Blom, R., Westdorp, R., Joosten, R.N., McNaughton, B.L., and Pennartz, C.M. (2007). A split microdrive for simultaneous multi-electrode recordings from two brain areas in awake small animals. J. Neurosci. Methods *162*, 129–138.

Lansink, C.S., Goltstein, P.M., Lankelma, J.V., and Pennartz, C.M. (2010). Fast-spiking interneurons of the rat ventral striatum: temporal coordination of activity with principal cells and responsiveness to reward. Eur. J. Neurosci. *32*, 494–508.

Lansink, C.S., Jackson, J.C., Lankelma, J.V., Ito, R., Robbins, T.W., Everitt, B.J., and Pennartz, C.M. (2012). Reward cues in space: commonalities and differences in neural coding by hippocampal and ventral striatal ensembles. J. Neurosci. *32*, 12444–12459.

Lansink, C.S., Meijer, G.T., Lankelma, J.V., Vinck, M.A., Jackson, J.C., and Pennartz, C.M. (2016). Reward expectancy strengthens CA1 theta and beta band synchronization and hippocampal-ventral striatal coupling. J. Neurosci. *36*, 10598–10610.

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.

Leutgeb, J.K., Leutgeb, S., Moser, M.-B., and Moser, E.I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science *315*, 961–966.

Lizier, J.T. (2014). JIDT: an information-theoretic toolkit for studying the dynamics of complex systems. Front. Robot. AI *1*, 11.

McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., and Moser, M.B. (2006). Path integration and the neural basis of the 'cognitive map'. Nat. Rev. Neurosci. *7*, 663–678.

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J.-C., Floreano, D., and Martinoli, A. (2009). The e-puck, a robot designed for education in engineering. In Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions, *1*, pp. 59–65.

Mou, X., and Ji, D. (2016). Social observation enhances cross-environment activation of hippocampal place cell patterns. eLife *5*, e18022.

O'Keefe, J., and Conway, D.H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. Exp. Brain Res. *31*, 573–590.

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.

Okuyama, T., Kitamura, T., Roy, D.S., Itohara, S., and Tonegawa, S. (2016). Ventral CA1 neurons store social memory. Science *353*, 1536–1541.

Omer, D.B., Maimon, S.R., Las, L., and Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. Science *359*, 218–224.

Owen, S.F., Tuncdemir, S.N., Bader, P.L., Tirko, N.N., Fishell, G., and Tsien, R.W. (2013). Oxytocin enhances hippocampal spike transmission by modulating fast-spiking interneurons. Nature *500*, 458–462.

Pawelzik, H., Hughes, D.I., and Thomson, A.M. (2002). Physiological and morphological diversity of immunocytochemically defined parvalbumin- and cholecystokinin-positive interneurones in CA1 of the adult rat hippocampus. J. Comp. Neurol. *443*, 346–367.

Paxinos, G., and Watson, C. (2006). The Rat Brain in Stereotaxic Coordinates (Elsevier).

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., and Dubourg, V. (2011). Scikit-learn: machine learning in Python. J. Mach. Learn. Res. *12*, 2825–2830.

Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. Nature *497*, 74–79.

Piskorowski, R.A., Nasrallah, K., Diamantopoulou, A., Mukai, J., Hassan, S.I., Siegelbaum, S.A., Gogos, J.A., and Chevaleyre, V. (2016). Age-dependent specific changes in area CA2 of the hippocampus and social memory deficit in a mouse model of the 22q11. 2 deletion syndrome. Neuron *89*, 163–176.

Raam, T., McAvoy, K.M., Besnard, A., Veenema, A.H., and Sahay, A. (2017). Hippocampal oxytocin receptors are necessary for discrimination of social stimuli. Nat. Commun. *8*, 2001.

Rizzolatti, G., and Fabbri-Destro, M. (2008). The mirror system and its role in social cognition. Curr. Opin. Neurobiol. *18*, 179–184.

Rizzolatti, G., and Sinigaglia, C. (2016). The mirror mechanism: a basic principle of brain function. Nat. Rev. Neurosci. *17*, 757–765.

Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., and Müller, M. (2011). pROC: an open-source package for R and S+ to analyze and compare ROC curves. BMC Bioinformatics *12*, 77.

Save, E., Nerad, L., and Poucet, B. (2000). Contribution of multiple sensory information to place field stability in hippocampal place cells. Hippocampus *10*, 64–76.

Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. J. Neurol. Neurosurg. Psychiatry *20*, 11–21.

Sik, A., Penttonen, M., Ylinen, A., and Buzsáki, G. (1995). Hippocampal CA1 interneurons: an in vivo intracellular labeling study. J. Neurosci. *15*, 6651–6665.

Sliwa, J., Planté, A., Duhamel, J.R., and Wirth, S. (2016). Independent neuronal representation of facial and vocal identity in the monkey hippocampus and inferotemporal cortex. Cereb. Cortex *26*, 950–966.

Sokal, R., and Rohlf, F. (1995). Biometry, 3rd (W. H. Freeman).

Squires, A.S., Peddle, R., Milway, S.J., and Harley, C.W. (2006). Cytotoxic lesions of the hippocampus do not impair social recognition memory in socially housed rats. Neurobiol. Learn. Mem. *85*, 95–101.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. *15*, 1929–1958.

Stevenson, E.L., and Caldwell, H.K. (2014). Lesions to the CA2 region of the hippocampus impair social memory in mice. Eur. J. Neurosci. *40*, 3294–3301.

Thompson Ray, M., Weickert, C.S., Wyatt, E., and Webster, M.J. (2011). Decreased BDNF, trkB-TK+ and GAD67 mRNA expression in the hippocampus of individuals with schizophrenia and mood disorders. J. Psychiatry Neurosci. *36*, 195–203.

Tirko, N.N., Eyring, K.W., Carcea, I., Mitre, M., Chao, M.V., Froemke, R.C., and Tsien, R.W. (2018). Oxytocin transforms firing mode of CA2 hippocampal neurons. Neuron *100*, 593–608.e3.

Tulving, E. (1983). Elements of Episodic Memory (Clarendon Press).

Ueno, H., Suemitsu, S., Murakami, S., Kitamura, N., Wani, K., Okamoto, M., Matsumoto, Y., and Ishihara, T. (2017). Region-specific impairments in parvalbumin interneurons in social isolation-reared mice. Neuroscience *359*, 196–208.

van Duuren, E., van der Plasse, G., Lankelma, J., Joosten, R.N., Feenstra, M.G., and Pennartz, C.M. (2009). Single-cell and population coding of expected reward probability in the orbitofrontal cortex of the rat. J. Neurosci. *29*, 8965–8976.

Vinck, M., Bos, J.J., Van Mourik-Donga, L.A., Oplaat, K.T., Klein, G.A., Jackson, J.C., Gentet, L.J., and Pennartz, C.M. (2016). Cell-type and state-dependent synchronization among rodent somatosensory, visual, perirhinal cortex, and hippocampus CA1. Front. Syst. Neurosci. *9*, 187.

Viskontas, I.V., Quiroga, R.Q., and Fried, I. (2009). Human medial temporal lobe neurons respond preferentially to personally relevant images. Proc. Natl. Acad. Sci. U S A *106*, 21329–21334.

von Heimendahl, M., Rao, R.P., and Brecht, M. (2012). Weak and nondiscriminative responses to conspecifics in the rat hippocampus. J. Neurosci. *32*, 2129–2141.

Zynyuk, L., Huxter, J., Muller, R.U., and Fox, S.E. (2012). The presence of a second rat has only subtle effects on the location-specific firing of hippocampal place cells. Hippocampus *22*, 1405–1416.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Male Lister Hooded rats, ordered at 7 weeks of age	Harlan	N/A
Software and Algorithms		
MATLAB Custom scripts	Authors	N/A
MATLAB Fieldtrip open source	www.fieldtriptoolbox.org	https://doi.org/10.1155/2011/156869
Tcl custom scripts	https://www.tcl.tk/	N/A
Python custom script	Authors	N/A
Java Information Dynamics Toolkit (JIDT)	https://github.com/jlizier/jidt	https://doi.org/10.3389/frobt.2014.00011
Other		
Custom experimental setup	Technology Center, University of Amsterdam	N/A
Custom 14 tetrode hyperdrives	Technology Center, University of Amsterdam	N/A
Nichrome tetrode wire 0.0005 inch diameter Stablohm 800 A (100-189)	California fine wire	CFW2013150
remotely controlled minirobot (e-puck1)	G-tec, Austria	N/A

LEAD CONTACT AND MATERIALS AVAILABILITY

This study did not generate new unique reagents. Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Cyriel Pennartz (c.m.a.pennartz@uva.nl).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Three male Lister Hooded rats (Harlan, Netherlands) were recorded in this experiment. All experiments were conducted according to the National Guidelines on Animal Experiments and with approval of the Animal Experimental Committee of the University of Amsterdam. During training the rats were housed with four in a cage. After implantation with a microdrive the animals were housed solitarily in high-walled, transparent cages (40 × 40 × 40 cm). Rats were maintained on a reversed 12:12 hour day/night cycle, lights on at 8 pm. To increase motivation the rats were maintained at 85% of their *ad libitum* weight, receiving 15-18 g food (standard lab chow) per day, delivered between 4 and 6 pm. Access to water was *ad libitum*.

METHOD DETAILS

Behavioral setup and minirobot

The animals were trained in a dimly lit environment on a custom-made, remotely controlled behavioral setup, consisting of a semi-circular maze compartment and a rectangular cage compartment (Figure 1). The maze (length: 62 cm; width: 103 cm) consisted of 5 radial arms connected at the ends. Each arm was equipped with infrared photobeam detectors to monitor the presence of either the rat or the robot. The cage (43 × 28 × 50 cm) overlooked the maze and was situated at the point where all arms converged. The cage and maze compartment were separated by a Plexiglas door, which could slide downward to allow the rat to move from one compartment to the other without interference by the experimenters. The door and the two side walls of the cage were made from transparent Plexiglas to allow the rat a full view on the maze compartment. The cage compartment included two reward wells (size 10 × 7 cm) right next to each other and placed along a left-right axis inside the cage. Reward (80-90 μL of a 15% sucrose solution in water) was delivered upon a correct nose poke. The reward wells were positioned such that the rat had visual access to the maze while sitting behind the reward wells and such that an implanted rat could make reward responses without hitting the door with its microdrive. The wells were equipped with LEDs and infrared photobeam detectors to monitor both nose pokes and licks. Photobeams were positioned in the middle of each individual arm of the maze to track the location of either the rat or the robot. During the Rat-on-Maze phases (Figure 1A), food pellets were distributed across the maze to stimulate exploration behavior. All behavioral events were time stamped by the recording software (tcl) and saved within a Neuralynx event file. For the Observation Period (Figure 1A), a

remotely controlled minirobot was used (e-puck; G-tec, Austria; height: 7 cm; cylindrical diameter: 7.7 cm; Mondada et al., 2009). Movements of the e-puck were remotely controlled by manual operation of keyboard presses, converted to e-puck commands via a custom-made MATLAB script. This e-puck was operated by a trained driver. The salience of the e-puck was enhanced by illuminating four LED lights on the minirobot's body. The step motor driving wheel movement produced an audible sound.

Behavioral training and behavioral procedure during recordings

During the first two weeks after arrival, rats were handled daily to get them accustomed to being handled by the experimenters. Following pre-training, rats were trained to perform the Front and Mid Task. During the Front Task, the robot moved from the middle position (inter-trial interval, ITI, 3 s) into either the left or right side arm closest to the cage, a movement that we label as ''outbound trajectory'' (Figure 1B). Upon reaching the point halfway of the side arms, the robot broke a photo-beam. One second after the beam was broken, the lights above the wells were turned on. With these lights on, the rats could acquire a reward by making a nose poke in the reward port (left or right) corresponding with the side of the maze the robot had traveled (left or right). For instance, upon a rightward movement of the robot (from the perspective of the rat facing the maze compartment), the rat obtained a reward by poking in the right reward well. Nose pokes into the other well were not rewarded. After either a correct or incorrect response the reward lights were turned off. Upon reaching the end of the side arm, the robot remained in place for 3 s before driving back to the middle of the maze. During this inbound trajectory, no rewards could be obtained (Figure 1B). The side chosen for the robot to travel to was pseudorandomly selected by the computer program (the number of trials consecutively to the same side was three at maximum). In addition to the Front Task, we also trained the animals on the Mid Task. This task was the same as the Front Task, except that the robot started right in front of the cage and moved into either the middle left or middle right arm (Figure 1C). In both versions of the task the rat always had to select the reward well corresponding to the (outbound) travel direction of the e-puck. Rats were trained daily on both the Front and Mid Task. The order of performing the two task versions was randomized across sessions. During the last 5 training sessions before implantation, rats were also accustomed to exploring the maze by allowing them 15 minutes of free exploration per day. To facilitate full sampling of the maze, rats were rewarded with food pellets (BioServ, dustless precision pellets 14 mg, Flemington, NJ) placed with random intervals at several locations throughout the maze.

During recording sessions, the behavioral procedure was as follows. A session commenced with a period of maze exploration (Rat-on-Maze 1; 15 min.), after which the rat entered the observation cage. In the subsequent Observation period, the rat engaged either in the Front or Mid Task (Figure 1A). The order of these task versions was pseudo-randomly determined. After performing one of the two tasks, the other version was performed (each task lasted 15 min.). During the Free Roaming control period (FR-control), the e-puck would be roaming across the maze, however without requiring the rat to perform any particular action (no rewards were applied; duration: 5 min.). This stage was followed by a period in which the e-puck was absent (No Robot Control, NR-control; no task and no rewards were applied; duration: 5 min.). The session was concluded with a second period of maze exploration (Rat-on-Maze 2; 15 min). Note that the main text presents some of the behavioral conditions in a different order than the behavioral sequence as applied during recording sessions.

Surgical procedure and tetrode recordings

Hippocampal area CA1 in the right hemisphere was targeted with a custom-made microdrive containing 14 individually movable tetrodes (Gray et al., 1995; coordinates: −4.1 mm posterior and +2.5 mm lateral relative to Bregma; Paxinos and Watson, 2006; tetrodes made from nichrome wires, California Fine Wire, lead diameter: 13 μm, gold-plated to 500-800 kΩ impedance at 1 kHz). We recorded neural activity with a 64-channel Digital Neuralynx Cheetah setup (Neuralynx, Bozeman MT). Twelve tetrodes were used for recordings and two served as reference electrodes. The design of the microdrive was similar to previously used versions (Lansink et al., 2007).

Prior to surgery rats received subcutaneous injections of Buprenorphin (Buprecare, 0.01 - 0.05 mg/kg), Meloxicam (Metacam, 2 mg/kg), and Baytril (5 mg/kg, Henry Schein Animal Health, Cuijk, the Netherlands). Rats were anesthetized using 3.0% (induction) and 1.0% - 3.0% (maintenance) isoflurane in pure oxygen and mounted in a stereotaxic frame. Body temperature was maintained between 35 and 36 ◦C using a heating pad. After the cranium was exposed, six holes were drilled to accommodate six surgical screws. The craniotomy, allowing implantation of the tetrode bundle, was approximately 2.3 mm in diameter. After removing the dura, the bundle was lowered onto the exposed cortex. The craniotomy was sealed with Kwik-Sil (World Precision Instruments, Friedburg, Germany) before the microdrive was fixed to six surgical screws in the skull (1.4 mm x 5 mm stainless steel, King Microscrews, Borculo, the Netherlands) using dental cement (Simplex Rapid, Dental Union, Nieuwegein, the Netherlands). A skull screw located on the caudal part of the parietal skull plate contralateral to the drive location served as ground. Next, the tetrodes were lowered 1 mm into the neocortex. Over the next seven days, rats were allowed to recover, with *ad libitum* food and water available. The recording tetrodes were gradually lowered to the CA1 pyramidal cell layer of the hippocampus over the course of 7-9 days after surgery. Electrode depths were continuously registered throughout the vertical descent. Positioning of the tetrodes was guided by keeping track of depths (estimated by the number of turns of the guide screws) and by online monitoring of LFPs and spike signals.

After the final recording session, current (12 μA for 10 s) was passed through one lead per tetrode to mark the end point of the tetrode track with a small lesion. At least 24 h after making the lesions, the animals were deeply anesthetized with Nembutal (sodium pentobarbital, 60 mg/ml, 1.0 mL i.p.; Ceva Sante Animale, Maassluis, the Netherlands) and transcardially perfused with a 0.9% NaCl solution, followed by a 4% paraformaldehyde solution (pH 7.4 phosphate-buffered). Following immersion post-fixation, coronal

sections of 40 μm were cut using a vibratome and stained with Cresyl Violet to reconstruct tetrode tracks and localize the endpoints. For all rats reported in this study, placement of tetrode endpoints in area CA1 was histologically verified.

Data acquisition

Signals recorded from tetrodes were passed through a unity-gain preamplifier headstage, a 64-channel commutator and band-pass filtered between 600-6000 Hz for spike recordings. One ms epochs of activity from all four leads were digitized at 32 kHz if a signal on any of the leads of a tetrode crossed a pre-set voltage threshold. Spike trains were sorted to isolate single units using a semi-automated clustering algorithm followed by manual refinement (KlustaKwik, K. Harris, and MClust 3.5, A. D. Redish). During recordings, rats were videotracked at 25 Hz, and an array of light-emitting diodes on the headstage allowed offline tracking of the rat's position and head direction.

DATA ANALYSIS

Spike sorting and waveform quantification

Automated and manual clustering of spikes was performed using the waveform peak amplitude, energy, and first derivative of the energy ("energyD1"). Clusters were accepted as single units when having no more than 0.1% of their inter-spike intervals shorter than 2 ms. During recordings, rats were videotracked at 25 Hz. An array of light-emitting diodes on the headstage allowed offline tracking of the rat's position and head direction.

To segment spike waveforms into groups corresponding to putative interneurons and excitatory cells, we quantified their waveform duration (peak to trough interval, in ms) and repolarization level (normalized amplitude at 0.45 ms after the spike peak, also known as decay of spike valley; Lansink et al., 2010). Waveforms were classified into three groups: fast-spiking cells (putative interneurons), broad-spiking cells (putative pyramidal neurons), and unclassified cells (Figures 6B and 6C; Vinck et al., 2016). The distribution of peak-to-trough durations was significantly bimodal (Hartigan's dip test, $p = 0.002$).

Place fields and robot firing fields

Place fields were classically defined as pertaining to subregions of space where the firing rate was selectively enhanced as a function of rat location. When studying firing rate as a function of robot position, we will refer to 'firing fields' as our results do not indicate discrete place fields for the robot. Classic place fields were determined using 3 × 3 cm spatial bins to segment the complete task space, including both the maze and observation cage. Only bins visited > 200 ms were included in the analysis. We next selected those bins having a non-zero firing rate and being surrounded by at least 3 other non-zero bins. In addition, the cluster of adjacent non-zero bins was required to comprise at least 9 bins (Save et al., 2000). Selected firing fields were smoothed with a 5-point Gaussian. Next, we computed the maximal firing rate obtained from this smoothed distribution. We defined the place field as the collection of bins falling within 80% of the maximal rate (Leutgeb et al., 2007), provided that this collection remained in excess of 9 bins and that the maximal firing rate was > 1 Hz. If coverage of a place field exceeded 80% of the total area visited, the neuron was considered to be a putative interneuron (see Figure 6 for additional criteria). Firing fields for the robot were determined using the same criteria used to determine place fields. To calculate robot firing fields, we used the neuronal data from the rat and the position data of the robot.

To examine the properties of rat place fields and robot firing fields, we contrasted the field size and maximum firing rate of each unit. The size of the field was expressed as the percentage of the total field (determined as above) activity divided by the total visited space. Differences between conditions, based on the same units, were tested using Friedman's test with Tukey-Kramer's post hoc test. For comparison across all units, a Kruskal-Wallis test was used ($\alpha = 0.05$).

Mirror-like place fields and overlap between rat place field and robot firing field

In order to examine whether hippocampal neurons coded place fields common to both the rat and robot ("mirror-like place fields") we determined place fields for each unit which was not deemed to be an interneuron, during three different periods: (1) place fields based on rat position during the first maze exploration period (Rat-on-Maze 1); (2) place fields during the second maze exploration period (Rat-on-Maze 2) and (3) firing fields based on robot position during the Task Phase (Figure 1A). The percentage of overlap between place field locations was determined by comparing Rat-on-Maze 1 and Rat-on-Maze 2. Overlap was calculated as the portion of bins visited in both conditions which are shared by the two firing fields, and cells were only included which displayed place fields in both conditions. To calculate the overlap between the combined Rat-on-Maze 1 and Rat-on-Maze 2 periods (yielding place fields for the rat) on the one hand, and the Task Phase (yielding robot firing fields) on the other hand, we concatenated the place field locations of Rat-on-Maze 1 and Rat-on-Maze 2 before determining the overlap with the firing fields in the Task Phase. Again, only cells showing place and firing fields in both conditions were included and overlap was restricted to areas visited in both conditions. Wilcoxon's signed rank test was used to compare the conditions ($\alpha = 0.05$).

To determine how much overlap between the robot firing field and the rat place field of a neuron is expected by chance, we generated a distribution of chance-level overlap by computing the overlap of the robot firing field with all recorded rat place fields. We then calculated a p value as the proportion of overlaps in the chance-level distribution which were larger than or equal to the observed overlap (i.e., the overlap between rat and robot fields of the same neuron). We repeated this computation for all neurons and

Bonferroni-adjusted the p values. For this analysis, we considered only cells which had a robot firing field according to our definition (see above), and which had a place field in at least one of the Rat-on-Maze periods. If a cell had a place field in both maze periods, we combined the place field occupancies before computing the overlap.

Cell sensitivity to contrasting behavioral conditions

When studying hippocampal sensitivity to two contrasting conditions (e.g., FR-control versus Task Phase), we focused on the subset of recorded CA1 cells showing place fields in one or both of these contrasting conditions. To determine whether a given CA1 neuron was sensitive to the contrast between two behavioral conditions, we performed a randomization test for each cell. For instance, to assess the statistical significance of Mid-Front Task differences for individual neurons, we randomly reallocated trials from both conditions to shuffled sets of trials and computed the difference in maximum firing rate within the place fields per permutation, thereby constructing a distribution of 1000 samples under the null hypothesis of having no difference between the two conditions. For those conditions lacking a trial structure, we subdivided every stage into segments of 30 s and randomly interchanged these between the two conditions, again constructing a permutation distribution. We then examined whether the ratio of firing rates reached in the two conditions exceeded the 2.5% or 97.5% percentile, indicating statistical significance at $p < 0.05$. We used a test of proportion to assess whether the fraction of significantly modulated units was different from chance ($\alpha = 0.05$; Sokal and Rohlf, 1995; van Duuren et al., 2009):

$$Zvalue = \frac{p_{cond} - \alpha}{\sqrt{\frac{p_{cond}(1-p_{cond})}{n}}}$$

Where p_{cond} is the proportion of significant units out of a population of n neurons.

Linearized peri-event space histograms

The consistency of CA1 neuronal responses to specific elements of robot behavior across trials was tested using peri-event space histograms plotting the binned firing of a single neuron as a function of linearized robot position (Figure 4C). Rasters were separately constructed for different spatial segments of the robot's trajectories, e.g., inbound versus outbound journeys in the Mid Task. The advantage of using these histograms is that they not only present spatial dependencies of average firing rate with respect to an event such as an outbound journey, but also spike density rasters across individual trials. Trajectories of the robot were linearized for each trial (using 3 cm bins) and firing rates were calculated for each bin. Firing rates of bins with an occupancy of less than 0.2 s (5 video frames) were excluded. This resulted in peri-event space histograms with robot position on the x axis instead of time. Firing rates are indicated using a gray tone code. Within these linearized peri-event space histograms, we compared bins from the Front versus Mid Task as well as outbound versus inbound trajectories. Comparisons between Front and Mid tasks were conducted using Mann–Whitney's U test with false discovery rate (FDR) corrections. Differences between outbound and inbound trajectories within a single task were tested using Wilcoxon's signed rank test. Again, FDR corrections were applied to control for multiple comparisons ($\alpha = 0.05$).

Population coding using Random Forest Decoders

Neurons were included in the analysis only if they fired at least 50 spikes during the segment of data considered (Glaser et al., 2017). Spikes were binned in 500 ms time bins, resulting in a neural activity feature vector $\boldsymbol{d}_i = (s_{i,0}, \ldots, s_{i,N})$ for every time bin i, where $s_{i,j}$ is the count of spikes of neuron j which fall in time bin i. To provide the random forest decoder (RFD) with information on neighboring time bins, as well as the current one, an expanded feature vector was generated by stacking neural activity vectors from subsequent time bins, as follows: $\boldsymbol{d}_i^E = (d_{i-k_b}, d_{i-k_b+1}, \ldots, d_i, d_{i+1}, \ldots d_{i+k_a})$, where k_a and k_b denote the number included preceding and subsequent bins respectively. In this way, each sample contains a short temporal trajectory of neural activity (Glaser et al., 2017). Rat position was linearly interpolated at the centers of the temporal bins used for spikes to obtain a feature vector of the stacked x and y position of the rat at time i, and every sample was expanded in the same way to include a short spatial trajectory of the animal. For the reported analysis, we used $k_a = k_b = 2$. Similar results were obtained when including only history $k_a = 2$, $k_b = 0$, or no expansion at all, $k_a = k_b = 0$ (data not shown).

The combination of neural data and position input is obtained by stacking the two feature vectors, $\boldsymbol{r}_i = (\boldsymbol{d}_i^E, \boldsymbol{p}_i^E)$. When training the RFD on either neural activity or rat position alone, we padded the feature vectors with zero in order to keep the dimensionality of the input constant. That is, we generated $\boldsymbol{r}_i^{neural} = (\boldsymbol{d}_i^E, 0)$, where 0 denotes a vector of zeros of the same size of $\boldsymbol{p}_i^E$, and $\boldsymbol{r}_i^{position} = (0, \boldsymbol{p}_i^E)$, where 0 denotes a vector of zeros of the same size of $\boldsymbol{d}_i^E$.

To discriminate between behavioral epochs we employed a random forest classifier with 200 trees. To ensure reproducibility, random forests were seeded. To predict robot position, we first interpolated the position tracking of the robot at the centers of the bins used to bin spikes, then used a random forest regressor to predict the x and y position of the robot at every time point. For both algorithms we relied on the implementation provided by Scikit-learn (Pedregosa et al., 2011).

Training and testing of the decoders was performed separately for each recording session, as each session yields a different population of neurons. When discriminating behavioral epochs, every dataset (corresponding to one recording session) was balanced with random undersampling (resampling the majority class without replacement to obtain as many samples as are present in the minority class), and classification accuracy (the fraction of correctly classified samples) was used as the performance metric.

When predicting robot and rat position, we quantified the goodness of fit as the fraction of explained variance (Glaser et al., 2017), given by

$$R^2 = 1 - \frac{\sum_i (\hat{y}_i - y_i)^2}{\sum_i (y_i - \bar{y})^2}$$

where $\hat{y}_i$ are the predictions, y_i is the actual robot/rat position, and $\bar{y}$ is the average robot/rat position. Note that in this formulation R^2 can be negative in case of overfitting (*cf.* Figure 4E). Expressing decoding performance in terms of variance explained (as opposed to a simple distance between true and predicted location, i.e., the Euclidean error) has the advantage of being more generalizable, as it accounts for the variance of the signal that is being predicted. This is particularly important since we want to compare the prediction errors of two signals which are quite different: while the movement of the robot is slow, regular, and restricted to a subregion of the whole arena (fixed trajectories during the Task Phase), the rat explorations are faster, unpredictable, and can cover the whole space of the arena, with the possibility that some locations are only visited sporadically or just once. Given that rat position covers a larger range of locations than robot position, a comparison of the Euclidean errors will likely be misleading. We computed a value of R^2 for the *x* and *y* components of robot position separately and reported their average. The reported decoding accuracies and R^2 were averaged over 5 repetitions of 10-fold cross-validations. In each fold, samples were assigned randomly to the training and test set. For a given session, decoding using different inputs (e.g., neural data and rat position) was performed on the same random folds, to prevent differences in the random partitioning of training and test data to cause differences in decoding performance. For the decoding of robot position, training and testing were carried out separately for Front and Mid Task, and the results averaged across the two epochs. For rat position, decoding was performed separately for Rat-on-Maze 1, Rat-on-Maze 2, Front Task, and Mid Task; we then averaged the two Rat-on-Maze periods and the two stages of the task.

To compare the accuracies (or R^2 values) of the three types of inputs (neural data, rat position, and the combination of the two) across sessions we used a Wilcoxon signed-rank test and report Bonferroni-adjusted p values (i.e., the p values are multiplied by the number of comparisons). For the boxplots data points which exceed the high and low quartiles by twice the interquartile range are graphically depicted as outliers, but all data points are always considered for the statistical tests.

Population coding using logistic regression and feed-forward neural networks

We replicated the random forest decoder (RFD) analysis using two different algorithms, namely a logistic regression with L_2 regularization, as implemented in Scikit-learn (Pedregosa et al., 2011), and a fully connected feed-forward neural network (Figure S2). For the latter, we adapted the implementation published by (Glaser et al., 2017) based on the Keras library (https://keras.io). The network consisted of two hidden layers, each containing 50 neurons. Hidden layers had rectified linear unit (ReLU) activations (Glorot et al., 2011) and 30% dropout (Srivastava et al., 2014; dropout is a common method to reduce overfitting by randomly ignoring a subset of units during training). The output layer uses a Softmax activation function (Bishop, 2006). The model was trained with categorical cross-entropy loss and with the Adam optimizer (Kingma and Ba, 2014) for a maximum of 30 epochs. For the neural network, the test data was further split in equal-sized test and validation sets. The validation set was used for early stopping with delta 10^{-6} and patience of 10 epochs. This means that improvements in accuracy on the validation set smaller than 10^{-6} were considered absence of improvement, and training was terminated after 10 epochs with no improvement. To ensure sufficient test and validation data, we employed 5-fold cross-validation (reported accuracy was averaged over 5 repetitions of the cross-validation).

Receiver operating characteristic analysis

For the Receiver Operating characteristic (ROC) analysis, the spikes of each unit were first binned in 200 ms bins. We selected all spatial bins in which the animal spent at least 30 s in both the Front and the Mid Task. For each neuron we further restricted our analysis to the spatial bins in which at least 10 spikes were fired. Then, for each spatial bin, we constructed an ROC curve, which indicates how well the spiking activity of the neuron at that bin discriminates between the Front and Mid Task, and computed the area under the curve (AUC) and its 95% confidence interval using the method of DeLong et al. (1988) as implemented in the pROC package (Robin et al., 2011). The AUC value for a given neuron at a given bin was considered significant if its confidence interval did not contain 0.5. To obtain a single AUC value per neuron we took the maximum AUC values across all bins, and we refer to neurons as having a significant AUC when they have a significant AUC value in at least one spatial bin. The confidence intervals for the sensitivities (Figure 5A) were computed using bootstraps (Robin et al., 2011). Results obtained by different temporal binning, and by requiring a significant AUC in more than one spatial bin, are shown in Figure S5. Results with different requirements for the minimum number of seconds spent at a single location are shown in Figure S5A. Results obtained by different temporal binning, and by requiring a significant AUC in more than one spatial bin, are shown in Figures S5B–S5F. AUC values determined by taking the average across all spatial bins instead of the maximum were qualitatively equivalent, as shown throughout Figure S5.

Conditional mutual information

Spike trains of each neuron were binned in 200 ms bins, considering data from the full Task Phase, and excluding neurons which did not fire any spikes throughout this phase ($N = 3$). The positions of the animal and robot used to compute information theoretic

quantities were obtained by linearly interpolating the output of the position tracking at the times of the temporal bin centers. For information about Front versus Mid Task, we specified animal position in spatial bins of 3×3 cm and computed discrete mutual information (MI) $I(X, Y)$ between the neuron's binned spike count and the binary variable indicating task type, and the discrete conditional mutual information (cMI) $I(X, Y|Z)$ between the neuron's spike count and task type, conditional on the position of the rat. MI captures both linear and non-linear correlations between spiking activity and task type, while cMI additionally removes the information already provided by rat position. For information about robot position, we replaced the binary task type variable with the position of the robot and employed the first Kraskov-Stoegbauer-Grassberger (KSG) estimator (Kraskov et al., 2004; Lizier, 2014). Importantly, the estimation of information theoretic quantities with finite sample sizes may suffer from a bias, which causes non-zero estimated information, even if no information is actually shared between the variables. A common way to address this problem is to test whether the estimated information quantity is statically different from zero by comparing it to an empirical null distribution, generated with an appropriate resampling method (Lizier, 2014). For a given neuron, we generated 1000 surrogate datasets by randomly permuting the spike count vector $\boldsymbol{d}_i$ (importantly, the target variable and rat position vectors were not shuffled). We computed MI and cMI for each shuffled dataset to obtain the null distributions, then computed a p value equal to the fraction of surrogate datasets for which the values of MI and cMI was higher than the observed MI and cMI. This process was repeated for all neurons, and p values were adjusted with Bonferroni correction over all neurons. The reported scores are debiased by subtracting from the observed MI and cMI of each neuron the average MI and cMI of all its shuffled surrogates.

QUANTIFICATION AND STATISTICAL ANALYSIS

The dataset comprises 3 rats, with 15, 6 and 5 recording sessions per rat, yielding 330, 145 and 146 units, respectively. Details of statistical comparisons and the definition of center, dispersion, and precision measures are contained in the figure legends. In general, we applied non-parametric statistical tests: Wilcoxon's signed-rank test (paired data), Mann-Whitney's U test (unpaired data), Kruskal-Wallis ANOVA, and Friedman's test (with Tukey-Kramer's post hoc test). We tested at the $\alpha = 0.05$ level, and, whenever appropriate, corrected for multiple comparisons using False Discovery Rate (FDR) and Bonferroni correction. The cross-validation strategy for the decoding, the significance of AUC values in the ROC analysis, and the shuffling procedure used to debias information measures are detailed in the corresponding sections. Relevant software for the statistical analysis includes the pROC package (Robin et al., 2011) used to compute AUC values and their confidence intervals, and the Seaborn library (https://github.com/mwaskom/seaborn), used to obtain bootstrap confidence intervals for the bar plots (Figures 6C–6E).

DATA AND CODE AVAILABILITY

The dataset and code supporting the current study are available from the corresponding author upon request.