SUPPLEMENTAL INFORMATION

Figure S1
Imaging analysis workflow and identification of spatially tuned cells (related to Fig. 1).

Figure S2
Transient statistics and spatial coding metrics of deep and superficial CA1 PCs (related to Fig. 1).

Figure S3
Behavior in the random foraging (RF) and goal-oriented learning (GOL) tasks (Figs. 1-6).

Figure S4
Assessment of cue-associated firing (related to Fig. 2).

Figure S5
Long-term remapping dynamics of deep and superficial CA1 PCs during random foraging (related to Fig. 2).

Figure S6
Regression analysis of performance vs. goal-zone representation (related to Fig. 6).

Movie S1.
Two-photon z-stack of the CA1 pyramidal layer (related to Fig. 1).

Movie S2.
Functional imaging of deep and superficial CA1 PCs (related to Fig. 1).

Table S1.
Summary of statistical tests performed (related to all figures).

Supplemental Experimental Procedures
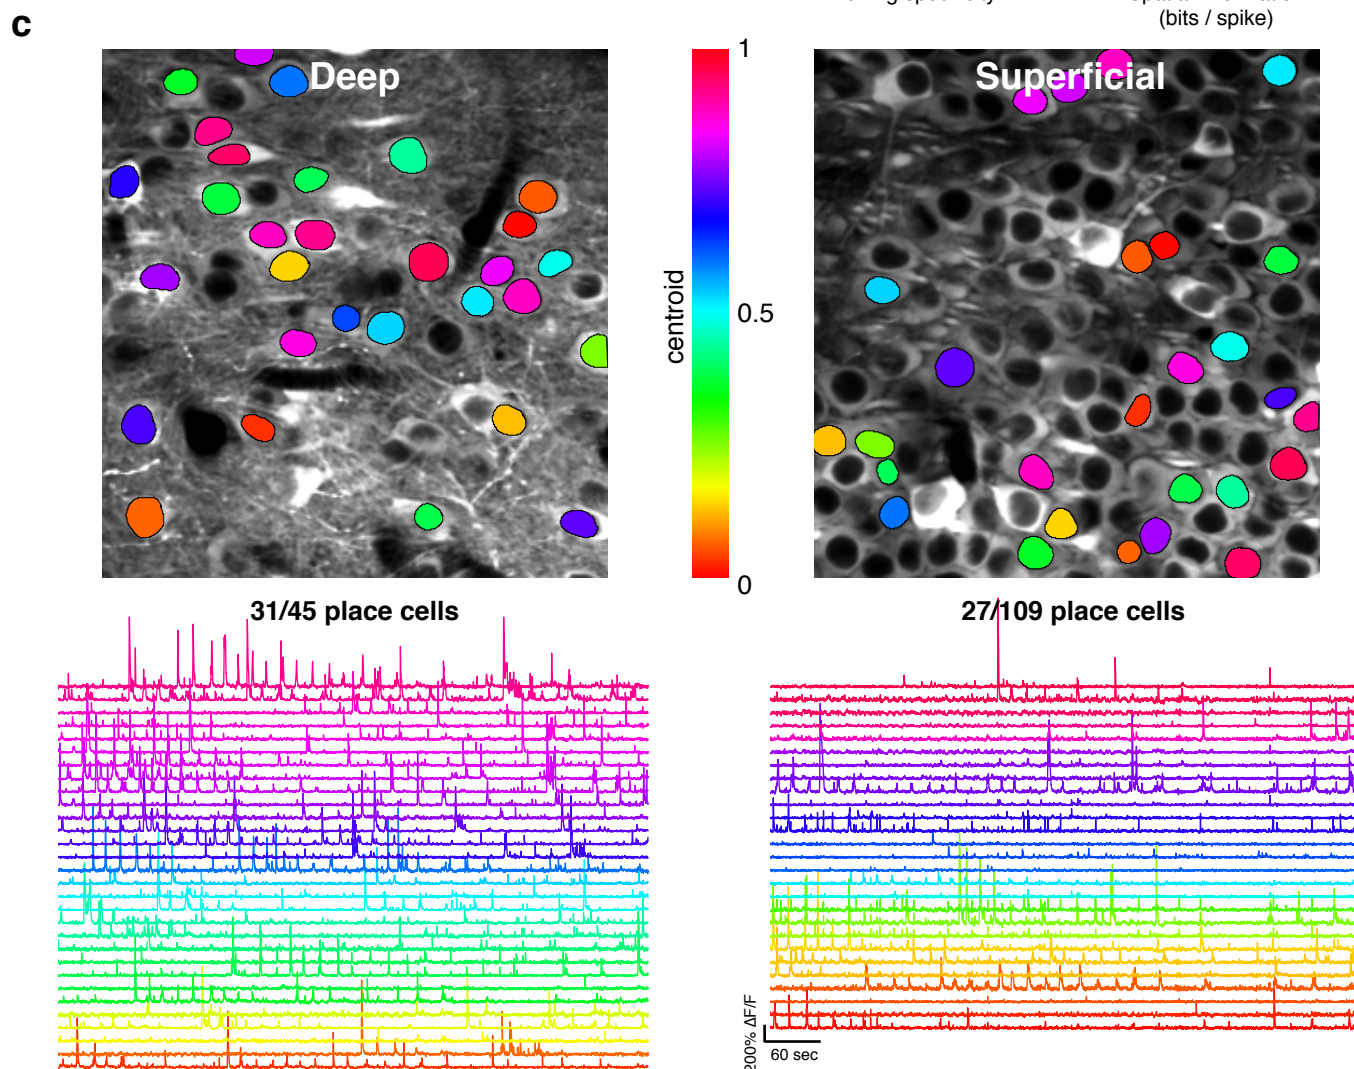
Supplemental References

Supplemental Figure 1

a

register images and warp ROIs in SIMA   manually draw ROIs in ImageJ   register images and warp ROIs in SIMA

Session 1   Session 2   Session 3

50μm

b

200% ΔF/F   60 sec

Belt position   2m   0

compute true tuning specificity / spatial information

↓

shuffle transients in time

↓

x100,000   recompute tuning specificity / spatial information

Shuffle distributions

true value

Normalized density

Tuning specificity

Spatial information (bits / spike)

c

Deep   Superficial

centroid   1   0.5   0

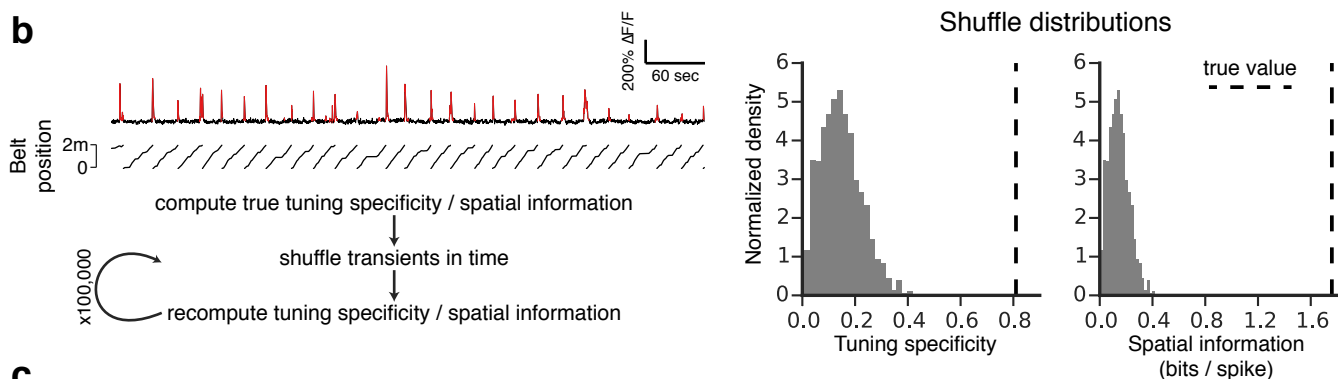31/45 place cells   27/109 place cells

200% ΔF/F   60 sec

Figure S1. Imaging analysis workflow and identification of spatially tuned cells (related to Fig. 1). (a) Longitudinal tracking of cells in the same FOV across imaging sessions. The time-averaged images from a deep FOV in the RF dataset is shown for three imaging sessions separated by 90 min. Regions of interest (ROIs) were drawn manually for one session in ImageJ over CA1 PC somata. These polygons were then imported into SIMA (Kaifosh et al., 2014) and assigned unique labels. Using the ROI Buddy software, an affine registration was calculated between the time-averaged imaged images, and the resulting transformation was applied to the manually drawn ROIs, aligning them with the target imaging session. Minor adjustments were made to the alignments post-hoc. **(b)** Identification of place cells based on both spatial information and tuning specificity. A representative fluorescence trace from a spatially-tuned CA1 PC is shown with detected Ca^{2+} transients ($p < 0.05$) in red. The mouse's position on the treadmill is plotted below. For each cell we calculated the tuning specificity and spatial information (see Methods). We then shuffled the running-related transients in time (restricted to running intervals) and recomputed the tuning specificity and spatial information. This procedure was performed 100,000 times, and the resulting shuffle distributions for both metrics are shown at right. The true tuning specificity and spatial information are indicated by the dashed lines. The p-values for this cell were highly significant ($p < 0.001$) relative to both distributions, and therefore this cell was classified as a place cell according to both definitions. This procedure was performed for each cell in every recording in order to identify the place cell populations. **(c)** The same FOV as in (a) is shown on the left, and the corresponding superficial FOV is shown on the right. Identified place cells are overlaid and are colored according to the centroid of their place field (for cells with multiple fields, the field with the highest in-field transient rate was used to calculate the centroid). The ΔF/F traces for each of the identified place cells (as assessed by spatial information) are shown below colored by centroid.
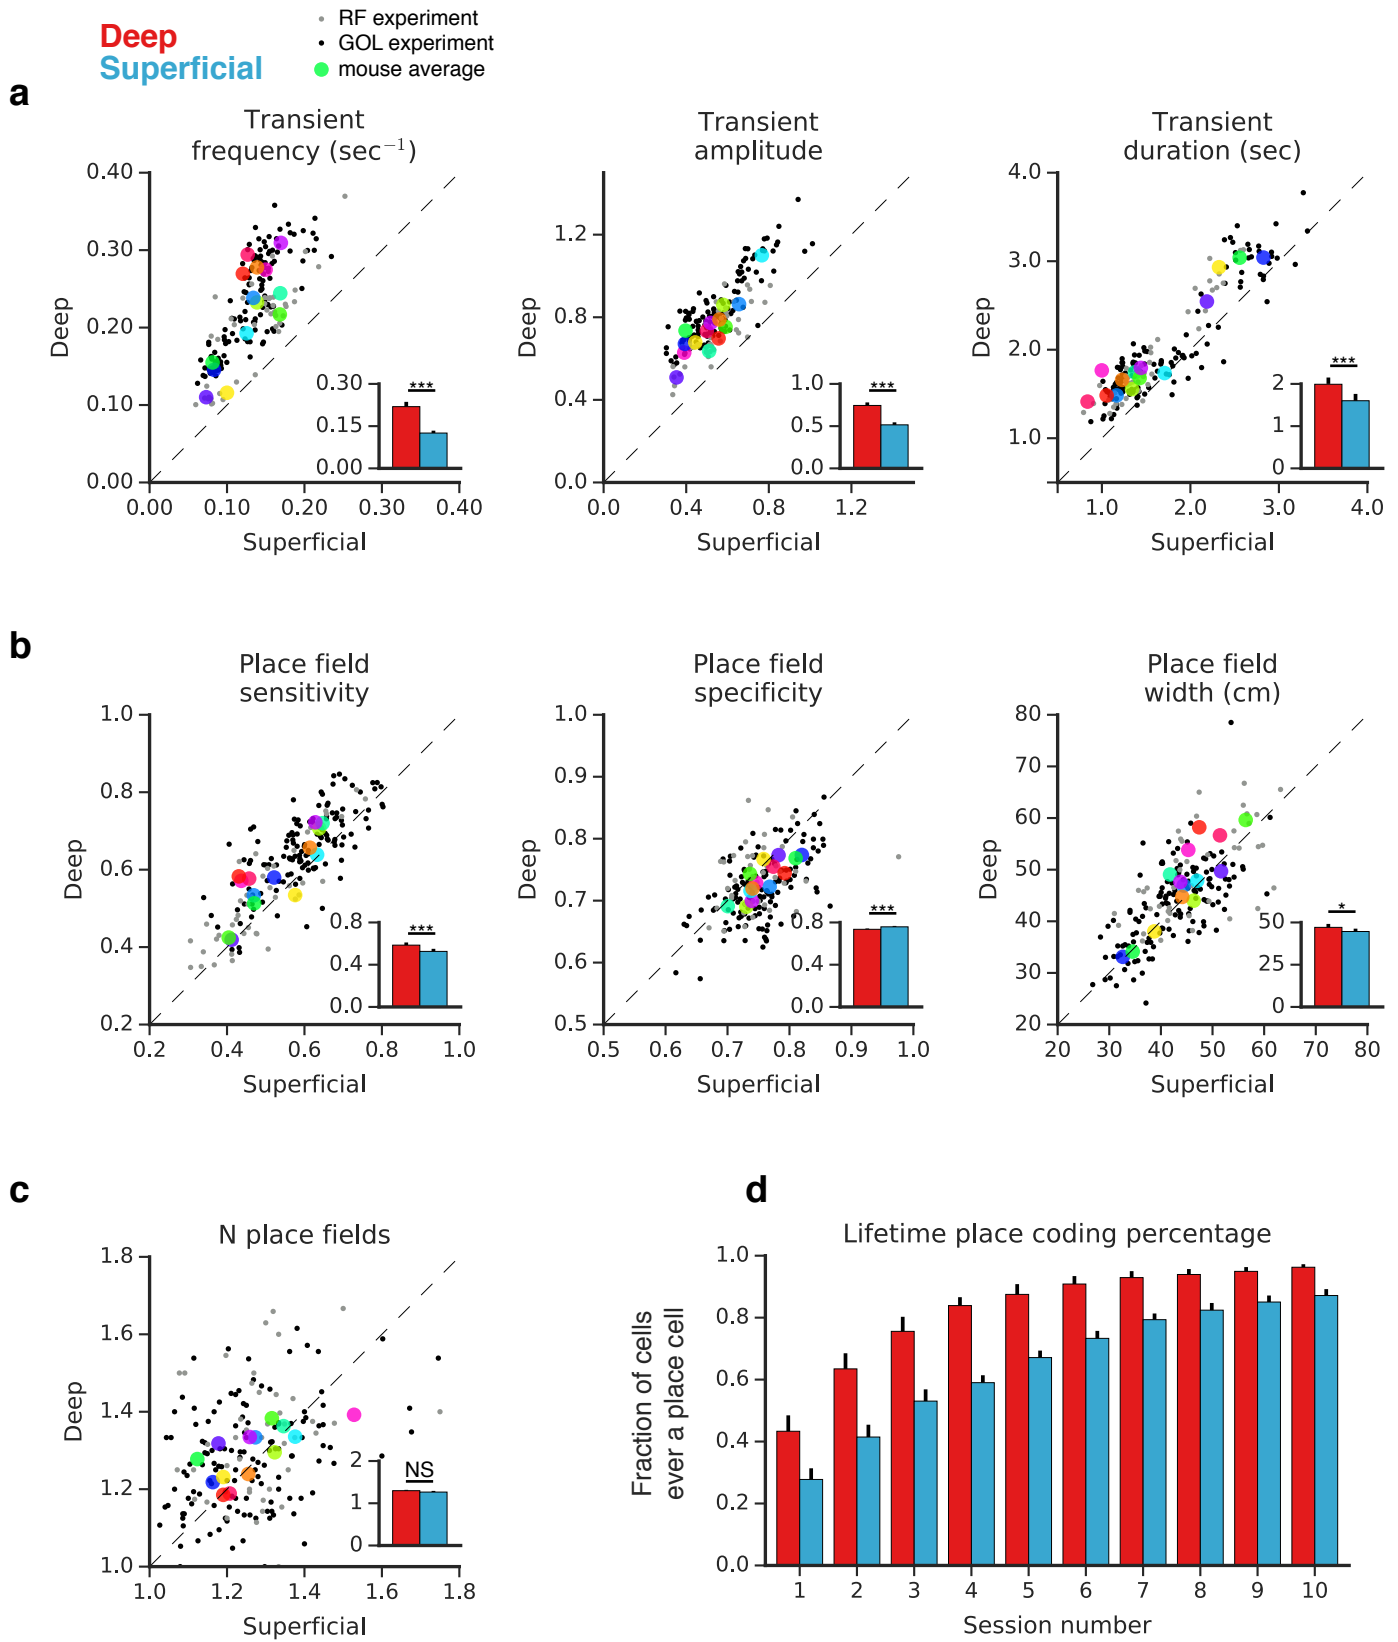
Supplemental Figure 2

Figure S2. Transient statistics and spatial coding metrics of deep and superficial CA1 PCs (related to Fig. 1).
(a - c) Within-session averages of paired deep and superficial recordings (either during RF or GOL) are indicated by small closed dots. Mouse averages are indicated by the large colored dots. Means across animals are shown in the inset bar plots as mean +/- s.e.m. **(a)** Transient statistics. *(i)* The transient frequency, defined as the total number of transients per unit time, was significantly higher for deep than superficial CA1 PCs (n = 14 mice, $p < 0.001$, paired T-Test). *(ii)* The mean transient amplitude was significantly larger in deep than in superficial CA1 PCs (n = 14 mice, $p < 0.001$, paired T-Test), as was *(iii)* the mean transient duration (n = 14 mice, $p < 0.001$, paired T-Test). **(b)** Spatial coding metrics. *(i)* The place field sensitivity (defined as the fraction of complete forward passes through the place field associated with a significant Ca^{2+} transient) was significantly higher in deep than in superficial (n = 14 mice, p = 0.001, paired T-Test), but the *(ii)* place cell specificity (defined as the fraction of running-related transients occurring within the place field), was significantly higher in superficial than in deep (n = 14 mice, $p < 0.001$, paired T-Test). Deep fields were slightly larger than superficial (n = 14 mice, $p < 0.05$, paired T-Test). **(c)** The mean number of place fields was similar between the sublayers (n = 14 mice, p = 0.15). **(d)** The lifetime place coding percentage was defined as the fraction of cells that had ever been identified as a place cell after the nth recording. For all levels of n, the lifetime place coding percentage was higher in deep than superficial, and consequently the deep population converged faster to 100%. A two-way ANOVA was performed for fraction as a function of days and sublayer (n = 168 mice X days). There was a significant effect of layer (F = 76.0, $p < 0.001$) and days (F = 39.7, $p < 0.001$) but not of the interaction (F = 0.70, p = 0.12).

Supplemental Figure 3

Running   Licking

a

Random foraging

b
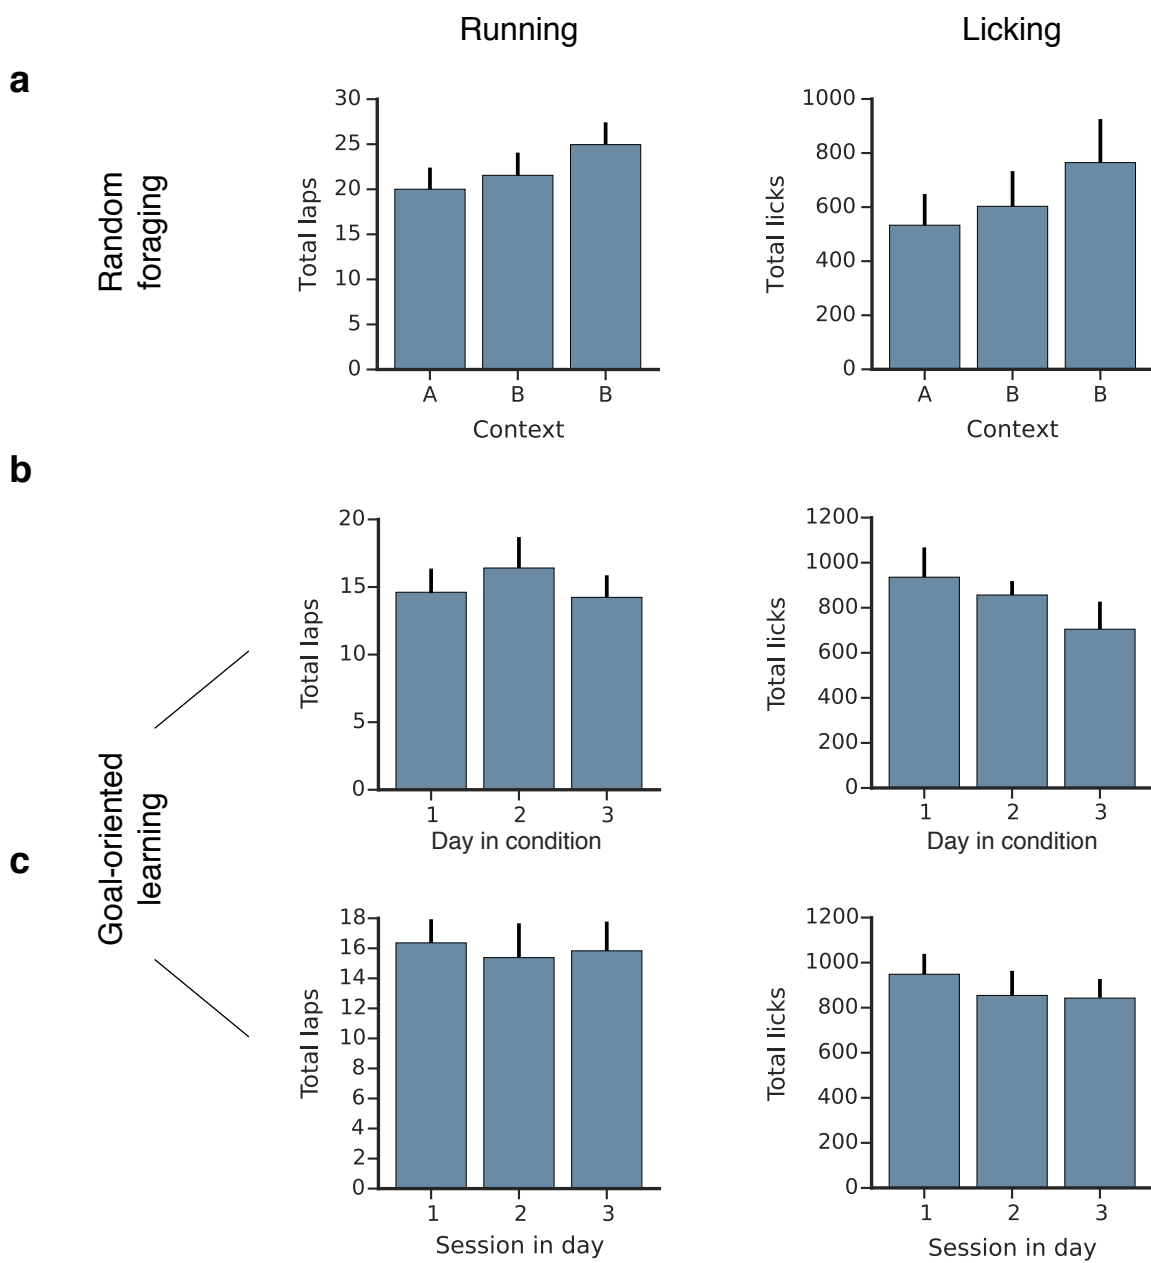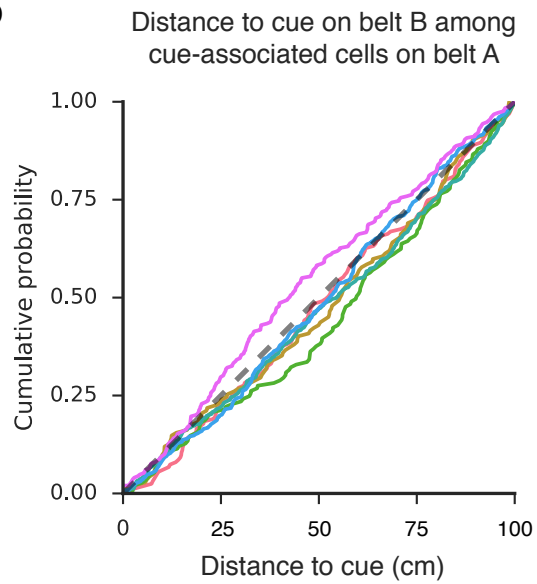
Goal-oriented learning

c

Figure S3. Behavior in the random foraging (RF) and goal-oriented learning (GOL) tasks (related to Figs. 1-6). (a-c) Error bars indicate mean +/- s.e.m. across animals **(a)** Licking and running across sessions of RF experiments. The total number of laps run in the RF experiments was not significantly different across sessions (n = 24 mice X sessions, F(session) = 1.84, p = 0.19, Type II ANOVA), nor was the total number of licks (n = 24 mice X sessions, F(session) = 1.31, p = 0.27, Type II ANOVA). **(b)** Licking and running across days of the GOL experiment. The total number of laps run per session of the GOL experiment was not significantly different across days (n = 16 mice X days, F(day) = 0.28, p = 0.76, Type II ANOVA), nor was the total number of licks (n = 16 mice X days, F(day) = 0.82, p = 0.46, Type II ANOVA). **(c)** Licking and running across sessions of the GOL experiment aggregated across days. The total number of laps run per session was not significantly different across sessions within each day (n = 18 mice X sessions, F(session) = 0.05, p = 0.95, Type II ANOVA), nor was the total number of licks (n = 18 mice X sessions, F(session) = 0.31, p = 0.74, Type II ANOVA).

Supplemental Figure 4

a

Belt A

Cue 1　Cue 2
Cue 3　Cue 4
Cue 5　Cue 6

66 cm

Belt B

Cue 3　Cue 6
Cue 2　Cue 5
Cue 1　Cue 4

66 cm

b

Distance to cue on belt B among
cue-associated cells on belt A

Cumulative probability

Distance to cue (cm)

c

Fraction of belt A cue-associated cells
near cue on belt B

Fraction near reward
(<10 cm)

Cue

Figure S4. Assessment of cue-associated firing (related to Fig. 2). (a) Images of the two belts (A and B) used in the RF experiments. The order of the fabrics was the same for the two belts, but the six cues (1: silver glitter masking tape; 2: green pom poms; 3: velcrow; 4: glue gun spikes; 5: pink foam strips; 6: jewels) were affixed in a different order. **(b)** In order to look for evidence of cells anchoring their place fields to specific cues in the A-B condition of the RF experiments, we first identified cue-associated cells on belt A (defined as having a tuning vector within 10 cm of the center of the cue). For the population of cells associated with each cue, we then calculated the distance from their tuning vectors on belt B to the associated cue. The distribution for each cue is shown in a unique color. A population of cells strongly anchoring to the cue should be indicated by an increased fraction of cells at small distances. **(c)** The true fraction of cells falling within 10 cm of the cue on belt B is quantified for each cue. For none of the cues was the fraction of associated cells higher than would be expected by chance (10%, dashed line), and there was no difference detected across cues ($F(cue) = 0.01$, $p = 0.99$, 1-way ANOVA). Error bars indicate mean +/- s.e.m. across animals.
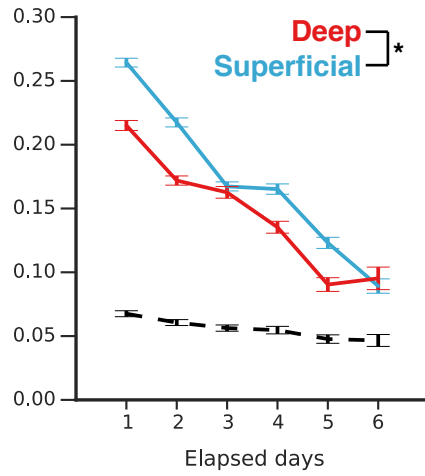
Supplemental Figure 5

a

Random foraging (RF) day-to-day stability

b

PV correlation

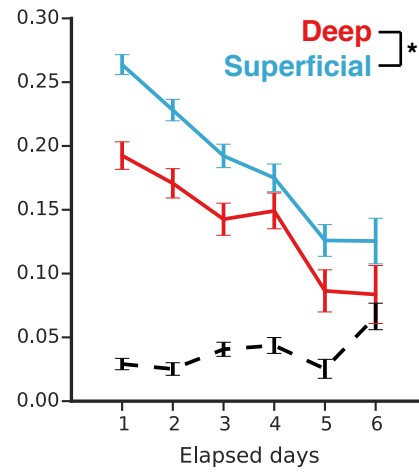

c

Tuning curve correlation

Figure S5. Long-term remapping dynamics of deep and superficial CA1 PCs during random foraging (related to Fig. 2). (a) Experimental schematic. Five animals ran in a single context for seven consecutive days (one 12-minute session per day) for randomly administered water rewards (three per lap). The context was novel on the first day of the experiment. **(b)** The population vector correlation (calculated across all cells) is plotted as a function of elapsed days and aggregated by mouse. Superficial maps were more stable than deep (n = 60 mice X layer X days, $F(layer) = 5.87$, $p < 0.05$, Type II ANOVA). The shuffle distribution is shown by the dashed line and was obtained by randomly pairing cell identities. The effect of days elapsed was also significant ($F(days) = 60.1$, $p < 0.001$), but the interaction was not ($F(layer \ X \ days) = 1.47$, $p = 0.23$). Error bars indicate mean +/- s.e.m. **(c)** The tuning curve correlation is plotted as a function of elapsed days and aggregated by mouse. Superficial CA1 PC place maps were more stable than deep over days-long timescales (n = 60 mice X layer X days, $F(layer) = 4.12$, $p < 0.05$, Type II ANOVA). The shuffle distribution is shown by the dashed line and was obtained by randomly pairing cell identities. The effect of time was also highly significant ($F(days) = 32.4$, $p < 0.001$), but the interaction between layer and elapsed days was not significant ($F(layer \ X \ days) = 1.24$, $p = 0.27$). Error bars indicate mean +/- s.e.m. across animals.

Supplemental Figure 6

a

Partial regression plots

b

Component-component-plus-residual plots

c

$$P = \beta_1 D_D + \beta_2 D_S + \beta_3 t + y$$

No. Observations:	91	R-squared:	0.585
Df Residuals:	87	F-statistic:	40.83
Df Model:	3	Prob (F-statistic):	1.45e-16

	coef	std err	t	P> ǀtǀ	[95.0% Conf. Int.]	
Intercept	0.5505	0.066	8.341	0.000	0.419	0.682
D_D	-0.2052	0.048	-4.258	0.000	-0.301	-0.109
D_S	-0.0026	0.045	-0.058	0.954	-0.092	0.087
day (t)	0.1147	0.012	9.407	0.000	0.090	0.139

Omnibus:	0.695	Durbin-Watson:	1.531
Prob(Omnibus):	0.707	Jarque-Bera (JB):	0.814
Skew:	-0.128	Prob(JB):	0.666
Kurtosis:	2.613	Cond. No.	19.3

Figure S6. Regression analysis of performance vs. goal-zone representation (related to Fig. 6). (a-c) Performance on the goal-oriented learning task (P, the fraction of licks in the reward zone) was modeled as a linear combination of D_D, D_S, and t (the day of the experiment). **(a)** Partial regression plots illustrate the strength of the linear relationship between the covariates and performance. **(b)** Component-component-plus-regression (CCPR) plots show the change in performance attributable to each covariate after accounting for the effects of the remaining two variables. **(c)** Summary of the regression analysis results. The simple linear regression model is shown above. Parameters were estimated through ordinary least squares. There was a significant relationship between D_D and P and between t and P, but not between D_S and P. Omnibus and Jarque-Bera tests indicate that the residuals are normally distributed. The Durbin-Watson test indicates the absence autocorrelation in the residuals. The condition number (Cond No.) indicates that multicolinearity is not a significant issue in the model (Greene, 2003).

Movie S1. Two-photon z-stack of the CA1 pyramidal layer (related to Fig. 1). A two-photon z-stack was obtained with 2 μm spacing between planes *in vivo*. 250 volumes were acquired sequentially. The data was motion-corrected in SIMA (2D HMM) and averaged across time to obtain a high-resolution structural volume. The movie was created and edited in ImageJ and iMovie.

Movie S2. Functional imaging of deep and superficial CA1 PCs (related to Fig. 1). This movie shows a near-simultaneous recording of GCaMP6f activity in deep and superficial sublayers during a GOL experiment. Motion artifacts were removed as described (2D HMM). See Methods for details of *in vivo* two-photon imaging strategy. The movie was created and edited in ImageJ and iMovie.

Table S1. Summary of statistical tests performed (related to all figures). For all statistical tests performed in the manuscript we have included the full test parameters and results.

Supplemental Experimental Procedures

Mice and viruses:
All experiments were conducted in accordance with the US National Institutes of Health guidelines and with the approval of the Columbia University Institutional Animal Care and Use Committee. Experiments were performed with adult male and female wild-type mice on a C57/Bl6 background (Jackson Laboratory). Animals were kept in a vivarium on a 12 hr light/dark cycle and were housed with 3-5 littermates per cage.

Recombinant adeno-associated virus (rAAV) expressing GCaMP6f under the CaMKII promoter, rAAV1/2(*CaMKII-GCaMP6f*), was used for broad expression of GCaMP6f in the CA1 *str. pyramidale*.

Viral injection and hippocampal window/headpost implant:

Viral delivery to hippocampal area CA1 was performed via stereotactic viral injection with a Nanoject syringe, as previously described (Kaifosh et al., 2013; Lovett-Barron et al., 2014). 60 nL (6 x 10nL pulses) of rAAV stocks were injected at each of three dorsoventral locations (-2.3 mm AP, -1.5 mm ML, -0.9, -1.05, -1.2 mm DV relative to cortical surface).

Mice were surgically implanted with an imaging window (diameter: 3.0-mm; height: 1.5-mm) over the left dorso-intermediate hippocampus and with a stainless-steel headpost for head fixation during imaging. Imaging cannulas were constructed by adhering (Narland optical adhesive) a 3-mm glass coverslip (64-0720, Warner) to a cylindrical steel cannula. We used the same headpost as in previous publications, and the surgical procedure was performed as described previously (Kaifosh et al., 2013; Lovett-Barron et al., 2014). Briefly, following induction of anesthesia (Isoflurane: 3% induction, 1.5-2.0% maintenance; 1.0 L/min O_2) and administration of analgesia (buprenorphine 0.05-0.1mg/kg), the scalp was removed, and a 3 mm diameter craniotomy was performed with a fine-tipped dental drill (V00033, Henry-Schein). The dura was removed, and the underlying cortex was aspirated (B30-50, B27-50, SAI) until the medial-lateral fibers of the external capsule were clearly visible. The cannula was gently fit into the craniotomy until the glass was in contact with the capsular fibers. The headpost was affixed to the skull using dental cement (675572, Dentsply). Mice were active within 15 minutes of surgery, and analgesia was continued for three days post-operatively.

In vivo two-photon imaging:

All imaging was conducted using a two-photon 8 kHz resonant scanner (Bruker). A piezoelectric crystal was coupled to the objective (Nikon 40X NIR water-immersion, 0.8 NA, 3.5 mm working distance), allowing for rapid displacement of the imaging plane in the z-dimension. In order to align the CA1 pyramidal layer with the imaging plane, we adjusted the angle of the mouse's head using two goniometers (±10° range, Edmund Optics). For excitation we used a 920nm laser (50-100 mW, Coherent). Red (tdTomato) and green (GCaMP6f) channels were separated by an emission cube set (green, HQ525/70m-2p; red, HQ607/45m-2p; 575dcxr, Chroma Technology), and fluorescence signals were collected with photomultiplier tubes (green GCaMP fluorescence, GaAsP PMT, Hamamatsu Model 7422P-40; red tdTomato fluorescence, multi-alkali PMT, Hamamatsu R3896). A custom dual stage preamp (1.4 x 105 dB, Bruker) was used to amplify signals prior to digitization. All experiments were performed at 2x digital zoom, covering 150 μm x 150 μm in each imaging plane. For paired recordings of deep and superficial CA1 PCs, images were 512 x 512 pixels per plane, and we allowed the piezo to settle for 40 ms after switching planes to maintain the flatness of the planes (jump-settle-acquire-jump-settle-acquire…), resulting in a frame rate of 7 Hz. Deep and superficial planes were separated by 25 μm. The superficial imaging plane was defined as the first complete layer of CA1 PC somata above the *stratum pyramidale/radiatum* border (typically 10 - 15μm from the border), and the deep imaging plane was 25μm above (**Supplemental Movie 1**).

Behavioral training

Random foraging (RF): Mice were water-restricted (>90% pre-deprivation weight) and trained to run on a cue-deplete burlap treadmill belt for water rewards over the course of 1-2 weeks. We applied a progressively restrictive water reward schedule, with mice initially receiving 40 randomly placed rewards per lap and ultimately receiving 3 randomly placed rewards per lap. Mice were trained for 20 minutes daily until they regularly ran at least one lap per minute. Mice were habituated to the optical instrumentation (presence of objective, laser, shutter sounds) prior to imaging experiments.

Goal-oriented learning (GOL): Mice were water-restricted (>90% pre-deprivation weight) and trained to run and lick simultaneously for water rewards on a cue-deplete burlap treadmill belt over the course of 1-2 weeks. Mice were first trained to obtain water operantly through the water port (water delivered *ad libitum* in response to tongue contact with a capacitance-sensing cannula). Once mice reliably licked for rewards, they were required to advance the treadmill for reward eligibility (e.g. 5 steps of forward progress resulted in 5 sec of reward availability). Reward zones were randomized each lap. Similarly to the RF training protocol, we implemented a progressive training schedule in which the number of reward zones per lap, as well as the spatial extent and time-limit on each zone, was gradually decreased until mice ran for a single randomly place 10 cm reward zone per lap with a 3 sec time limit. Because the reward zone shifted from lap-to-lap, this protocol trained mice to run and lick simultaneously as they searched for rewards. Mice were habituated to the optical instrumentation prior to imaging experiments.

Contexts:

Similar to our previous work, each context (A and B) consisted of the same treadmill belt (3 joined fabric ribbons), but distinct in their visual, auditory, tactile, and olfactory stimuli (**Fig 2b**, Danielson et al., 2016; Lovett-Barron et al., 2014). To allow for comparison of deep and superficial CA1 PC activity between similar contexts, the belts were made of the same three fabrics in the same order, but the locations of all of the six tactile cues (1: silver glitter masking tape; 2: green pom poms; 3: velcrow; 4: glue gun spikes; 5: pink foam strips; 6: jewels) were shuffled between the two belts.

Stimulus presentation and behavioral readout:

Visual, auditory, and olfactory stimuli were presented and behavioral data were recorded as described previously (Danielson et al., 2016; Kaifosh et al., 2013; Lovett-Barron et al., 2014). In order to reliably track the position of the treadmill belt, we defined registration anchors at known positions along the belts and interpolated between them using a quadrature encoded movement signal tied to the rotation of the treadmill wheels. Registration anchors were marked by radio-frequency identification (RFID) buttons (16mm, 125kHz, SparkFun Electronics) at evenly spaced positions along the belt, and were read off as they passed over a fixed RFID reader (ID-12LA, SparkFun). The rotational quadrature signal was produced by marking treadmill wheels with offset tick marks, and this signal was encoded by a pair of photodiodes (SEN-0024, SparkFun) aligned to the wheels (< 0.5cm resolution).

Ca^{2+} *data processing*:

Motion correction: All imaging data were analyzed using the SIMA software package (Kaifosh et al., 2014). Motion correction was performed using a modified 2D Hidden Markov Model (Dombeck et al., 2007; Kaifosh et al., 2013), in which the model was re-initialized each plane in order to account for the 40ms settling time of the piezo (see Methods: *In vivo two-photon imaging*), resulting in discontinuous displacements across planes. This modified algorithm has been made freely available in version 1.3 of the SIMA package. In cases where motion artifacts were not adequately corrected, the affected data was discarded from further analysis.

Segmentation of CA1 PC somata: For each field-of-view, segmentation was performed manually in ImageJ (http://imagej.nih.gov/ij/) by conservatively outlining putative CA1 PC somata in each plane of the time-averaged images of motion-corrected movies. This was performed for one imaging session of each FOV, and the resulting polygons were imported into SIMA and assigned unique labels. Using the SIMA project's ROI Buddy graphical user interface (Kaifosh et al., 2014), an affine registration was calculated between the time-averaged images of the segmented session with each unsegmented session, and the resulting transformation was applied to the manually drawn ROIs, aligning them with the target imaging session (**Fig S1a**). Minor adjustments were made to the alignment post-hoc in ROI Buddy.

Signal extraction: Dynamic GCaMP6f fluorescence signals were extracted from ROIs using SIMA according to the previously described formulation. We computed the relative fluorescence changes ($\Delta F/F$) as described (Jia et al., 2011), with uniform smoothing window $t_1 = 3$ sec. and baseline size $t_2 = 60$ sec. We detected statistically significant Ca^{2+} transients as described previously (Danielson et al., 2016; Dombeck et al., 2007; Lovett-Barron et al., 2014). In order to improve our sensitivity, we then recalculated the baseline of the raw signal after masking frames identified previously as occurring during a significant transient. $\Delta F/F$ was then recalculated, and transients re-estimated. Transients less than one second were removed to reduce false positives. This iterative procedure was repeated three times and effectively removed the transient contamination from the calculated baseline.

Identification of spatially-tuned cells:

In order to identify spatially-tuned cells (**Fig S1b**), we implemented two approaches: one based on tuning specificity and one based on spatial information. We restricted our analysis to running epochs, defined as consecutive frames of forward locomotion (imaging frames in which at least one forward pair of beam breaks occurred) at least 1 sec in duration and with a minimum peak speed of 5 cm/sec. Consecutive epochs separated by < 0.5 seconds were merged. Running-related transients were defined as those that were initiated during a running-related epoch. Transient start was defined as the first imaging frame with mean fluorescence >= 2σ, with σ equal to the standard deviation of the baseline frames. Offset was defined as the first frame with mean fluorescence <= 0.5σ (Dombeck et al., 2007).

Tuning specificity: The spatial tuning vector (Danielson et al., 2016) was calculated as $\sum_j \frac{e^{i\theta_j}}{o(\theta_j)}$, where θ_j is the position of the mouse at the onset time of the j-th running-related transient, and o_j is the fraction of running frames acquired at position θ_j. Only the animal's position at the time of transient onset was used for constructing the tuning vector and rate maps. The tuning specificity was defined as the magnitude of the spatial tuning vector (equivalently, tuning specificity = 1 – circular variance of occupancy-normalized transients). In order to assess the significance of the tuning specificity, for each cell we generated a null tuning distribution by shuffling the transient onset times (restricted to running frames) and repeatedly recomputing the tuning specificity. This process was repeated 100,000 times, and the p-value was defined as the fraction of this distribution that exceeded the cell's true tuning specificity.

Spatial information: For each cell we first computed the spatial information content (Skaggs, et al., 1993) as $I_N = \sum_{i=1}^{N} \lambda_i \ln \frac{\lambda_i}{\lambda} p_i$ where λ_i and p_i are the transient rate and fraction of time spent in the ith bin, λ is the overall firing rate, and N is the number of bins. We computed I_N for multiple values of $N = 2, 4, 5, 8, 10, 20, 25,$ and 100. We then created 100,000 random reassignments of the transient onset times within the running-related epochs and re-computed the values of I_N^s, where s is the index of the shuffle. In calculated spatial information using binned data, it is important to note that this measure is biased by both the number of bins chosen and by the number of events fired by the cell. Performing shuffles on a per-cell basis addresses the latter bias. To roughly correct from the bias associated with binning, we subtracted the mean of this null distribution from all estimates to obtain values $\hat{I}_N = I_N - \frac{1}{100,000} \sum_{s=1}^{100,000} I_N^s$. Finally, we computed a single estimate of the information content for the true transient onset times, $\hat{I} = \max_N \hat{I}_N$, and for the shuffles, $\hat{I}_s = \max_N \hat{I}_N^s$. Note that by maximizing $\hat{I}_N^s$ over bin sizes we have allowed each iteration of the shuffle to maximize its value across bin sizes (an alternate less conservative choice would have been to enforce that each shuffle be calculated using the same bin size as was optimal for $\hat{I}$). The spatial tuning p-value was taken as the fraction of values of s for which $\hat{I}$ exceeded $\hat{I}_s$.

Remapping analysis:

In order to compute the similarity between spatial maps in different sessions, we computed three metrics throughout the manuscript: population vector (PV) correlation, centroid shift, and tuning curve correlation. In the PV correlation analysis, all cells were included regardless of degree of tuning or of activity. To be included in the centroid shift analysis, a cell needed to be identified as a place cell on the basis of its tuning specificity in both sessions of the comparison. Similarly, for a cell to be included in the tuning curve correlation analysis, a cell needed to be identified as a place cell on the basis of its spatial information in both sessions of the comparison. We chose this approach because the centroid shift metric is most appropriately applied to singly peaked tuning profiles, and the tuning specificity place cell criterion selects for such cells.

Population vector correlation: The 1D rate maps of all deep or superficial cells in the experiment (rows) were stacked, resulting in a matrix whose columns represented population vectors of activity for that sublayer at different positions on the belt. The same position bin was compared across the two sessions of the comparison, and correlations were averaged across position bins to obtain one measure per sublayer.

Centroid shift: For each cell identified as a place cell on the basis of tuning specificity, the centroid shift was taken as the angle (in radians) between tuning vectors in the two sessions being compared.

Tuning curve correlation: Rate maps were formed by dividing the number of transients starting in each bin by the occupancy of that bin. We calculated rate maps with 100 position bins and smoothed with a Gaussian kernel (σ=3

bins). The tuning curve correlation for each cell was defined as the Pearson correlation coefficient between the tuning curves.

Shuffling analyses:

RF vs. GOL (Fig 4): To compare the difference in task modulation between sublayers (Δ task modulation), we performed two shuffles. First, we compared the true Δ task modulation with a null distribution generated by randomizing cell identity (with regard to superficial/deep designation) within each experiment and recomputing the Δ task modulation (100,000 shuffles). In the second shuffle, we compared the true Δ task modulation with a null distribution generated by shuffling the RF/GOL identity of the experiments while maintaining the sublayer identity of each cell (100,000 shuffles)

Reward representation vs. performance (Fig 6): To estimate the likelihood of obtained the observed difference in r-values by chance [$|r(D_D$ vs $P) - r(D_S$ vs $P)|$], we compared the observed difference with a null distribution. The null distribution of r-value differences was generated by randomizing cell identity within each experiment and recomputing the r-values for each sublayer (100,000 shuffles).

Statistics:

All statistical tests are described in the corresponding figure legends. All comparisons were two-sided, and all data was aggregated by mouse except where indicated. A paired two-sample T-Test was performed wherever possible to compare population means across animals, and the non-parametric Mann-Whitney U test was used otherwise. ANOVA was used to compare behavior across sessions and days of the experiment. Null distributions were generated as described above.

Supplemental References

Danielson, N.B., Kaifosh, P., Zaremba, J.D., Lovett-Barron, M., Tsai, J., Denny, C.A., Balough, E.M., Goldberg, A.R., Drew, L.J., Hen, R., Losonczy, A., Kheirbek, M.A. (2016). Distinct contribution of adult-born hippocampal granule cells to context encoding. Neuron *90*, 1–12.

Dombeck, D.A., Khabbaz, A.N., Collman, F., Adelman, T.L., Tank, D.W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. Neuron *56*, 43–57.

Greene, W.H. (2003). Econometric Analysis, 5th ed. Prentice Hall.

Jia, H., Rochefort, N.L., Chen, X., Konnerth, A. (2011). In vivo two-photon imaging of sensory-evoked dendritic calcium signals in cortical neurons. Nat. Protoc. *6*, 28–35.

Kaifosh, P., Lovett-Barron, M., Turi, G.F., Reardon, T.R., Losonczy, A. (2013). Septo-hippocampal GABAergic signaling across multiple modalities in awake mice. Nat. Neurosci. *16*, 1182–4.

Kaifosh, P., Zaremba, J.D., Danielson, N.B., Losonczy, A. (2014). SIMA: Python software for analysis of dynamic fluorescence imaging data. Front. Neuroinform. *8*, 1-10.

Lovett-Barron, M., Kaifosh, P., Kheirbek, M.A., Danielson, N.B., Zaremba, J.D., Reardon, T.R., Turi, G.F., Hen, R., Zemelman, B.V, Losonczy, A. (2014). Dendritic inhibition in the hippocampus supports fear learning. Science *343*, 857–63.

Skaggs, W.E., Mcnaughton, B.L., Markus, E.J., Gothard, K.M. (1993). An Information-Theoretic Approach to Deciphering the Hippocampal Code, in: Hanson, S., Cowan, J., Giles, C. (Eds.), Advances in Neural Information Process Systems (NIPS), 5. pp. 1030–1037.