HHS Public Access

Author manuscript

Neuron. Author manuscript; available in PMC 2017 August 03.

Published in final edited form as:

Neuron. 2016 August 3; 91(3): 652–665. doi:10.1016/j.neuron.2016.06.020.

Sublayer-specific coding dynamics during spatial navigation and learning in hippocampal area CA1

Nathan B. Danielson[1,2], **Jeffrey D. Zaremba**[1,2], **Patrick Kaifosh**[1,2], **John Bowler**[1,2], **Max Ladow**[1], and **Attila Losonczy**[1,3,4,*]

[1]Department of Neuroscience, Columbia University, New York, NY, 10032, USA

[2]Doctoral Program in Neurobiology and Behavior, Columbia University, New York, NY, 10032, USA

[3]Kavli Institute for Brain Science, Columbia University, New York, NY, 10032, USA

[4]Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, 10032, USA

Summary

The mammalian hippocampus is critical for spatial information processing and episodic memory. Its primary output cells, CA1 pyramidal cells (CA1 PCs), vary in genetics, morphology, connectivity, and electrophysiological properties. It is therefore possible that distinct CA1 PC subpopulations encode different features of the environment and differentially contribute to learning. To test this hypothesis, we optically monitored activity in deep and superficial CA1 PCs segregated along the radial axis of the mouse hippocampus and assessed the relationship between sublayer dynamics and learning. Superficial place maps were more stable than deep during head-fixed exploration. Deep maps, however, were preferentially stabilized during goal-oriented learning, and representation of the reward zone by deep cells predicted task performance. These findings demonstrate that superficial CA1 PCs provide a more stable map of an environment while their counterparts in deeper layers provide a more flexible representation that is shaped by learning about salient features in the environment.

Introduction

The mammalian hippocampus is critical for spatial information processing and episodic memory (Buzsáki and Moser, 2013; Eichenbaum and Cohen, 2014; O'Keefe and Nadel, 1978; Scoville and Milner, 1957; Squire et al., 1992). The canonical circuit motif supporting

*Corresponding author: Nathan B. Danielson (nbd2111@cumc.columbia.edu), or Attila Losonczy (al2856@cumc.columbia.edu).

Author Contribution

N.B.D. and A.L. conceived the project. N.B.D. performed experiments. N.B.D., J.D.Z., P.K., and A.L. designed the experiments, analyzed data, and wrote the manuscript. J.B. and M.L. assisted with experiments and provided technical infrastructure and input to the manuscript.

these functions is the trisynaptic loop, proceeding from the dentate gyrus input node, to area CA3, and finally to area CA1, where pyramidal cells (CA1 PCs) form the major output of the hippocampus. The extent to which individual CA1 PCs vary in their contribution to navigation and memory-related information processing remains largely unknown. Consequently, CA1 PCs are traditionally conceptualized as a relatively homogeneous population of processing units (Kesner and Rolls, 2015; Marr, 1971).

In fact, CA1 PCs vary greatly in their genetics, morphology, connectivity, and electrophysiological characteristics. Moreover, many of these differences are not randomly distributed throughout the hippocampus but rather are organized along its principal axes: dorsal-ventral (Fanselow and Dong, 2010; Jung et al., 1994; Kjelstrup et al., 2008, 2002; Maurer et al., 2005; Moser et al., 1993; Thompson et al., 2008), proximal-distal (Graves et al., 2012; Hartzell et al., 2013; Henriksen et al., 2010; Jarsky et al., 2008), and superficial-deep (Lee et al., 2014; Maroso et al., 2016; Mizuseki et al., 2011; Slomianka et al., 2011; Valero et al., 2015). The observed spatial organization of CA1 PC diversity is suggestive of regional specialization, whereby CA1 PCs along different axes could be biased in the information they process and in their contribution to different behaviors and forms of learning. This framework is not without precedent, as functional gradients are increasingly appreciated along the dorsal-ventral and proximal-distal axes in terms of spatial processing, episodic memory, and emotional responses (reviewed in Igarashi et al., 2014; Strange et al., 2014). Whether this framework extends to the third axis of the hippocampus, the superficial-deep (radial) axis, remains unknown.

A subdivision among deep (closer to *stratum oriens*) and superficial (closer to *stratum radiatum*) CA1 PCs has long been recognized (Lorente De Nó, 1934). Similar to the neocortex, deep and superficial cells are born during distinct neurogenic windows and differ both genetically and neurochemically (Angevine, 1965; Bayer, 1980; Cembrowski et al., 2016; Dong et al., 2009; Schlessinger et al., 1978; Slomianka, 2011). Such differences suggest the existence of radially organized microcircuits, and this notion is supported by recent work demonstrating prominent electrophysiological differences between deep and superficial CA1 PCs. Deep cells are more active, more likely to burst and to exhibit spatially tuned firing than superficial cells (Mizuseki et al., 2011), and the sublayers differentially participate in network synchrony events (Stark et al., 2014; Valero et al., 2015). Sublayer-specific perisomatic inhibition has been implicated in the organization of this microcircuit (Lee et al., 2014; Valero et al., 2015), and the observed functional differences are accompanied by radially-organized intrahipocamapal afferent and efferent connectivity (Groenewegen et al., 1987; Kohara et al., 2014; Lee et al., 2014; McGeorge and Faull, 1989). Previous functional studies of the radial axis have all been conducted under either *in vitro* or *in vivo* recording conditions in which cells were not monitored chronically during a learning task. Consequently, sublayer-specific dynamics have never been related to the dynamics of learning, which can evolve over several days. Thus, despite previous findings, the functional role of the radial microcircuit in guiding behavior remains unknown.

The most prominent behavioral correlate of CA1 PC activity is place coding, as individual “place cells” fire in restricted regions of the environment, their corresponding “place fields” (O’Keefe and Dostrovsky, 1971). As an ensemble, place cells provide an allocentric

representation of space, which is thought to support hippocampal spatial mnemonic function (Buzsáki and Moser, 2013; O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978). Consistent with this view, place fields remap in response to contextual manipulations (Colgin et al., 2008; Karlsson and Frank, 2008; Leutgeb et al., 2005, 2004; Muller and Kubie, 1987; Wilson and McNaughton, 1994), incorporate a wide range of non-spatial information, including salient sensory features of the environment, and are modulated by the internal behavioral state of the animal (Frank et al., 2000; Kobayashi et al., 1997; Moita et al., 2004; Pastalkova et al., 2008; Wood et al., 2000). Place maps are also modulated by proximal cues in an environment, which can serve as local landmarks to support allothetic navigation (Deshmukh and Knierim, 2013; Knierim and Hamilton, 2011; Knierim and Rao, 2003). Importantly, it has been repeatedly demonstrated that place maps incorporate goal-related information (Breese et al., 1989; Dupret et al., 2010; Fyhn et al., 2002; Gothard et al., 1996; Hok et al., 2007; Hollup et al., 2001; Hölscher et al., 2003; Kobayashi et al., 2003, 1997), and relatedly, that place coding is altered during behaviors requiring increased attention (Kentros et al., 2004; Markus et al., 1995; Muzzio et al., 2009). Again, however, while numerous functional differences have been described between deep and superficial CA1 PCs, it is unknown whether these differences extend to place coding dynamics and whether these subpopulations are differentially associated with these behavioural demands.

To address these questions, we performed two-photon Ca^{2+} imaging of identified deep and superficial CA1 PCs simultaneously in mice performing different head-fixed behaviors: random foraging (RF) and goal-oriented learning (GOL). While place coding during behavior has classically been studied using extracellular techniques, two-photon imaging allows for simultaneous recoding of large populations of CA1 PCs with unambiguously identified sublayer position and is also capable of assessing coding dynamics over long time scales with subcellular resolution (Dombeck et al., 2010; Kaifosh et al., 2013; Lovett-Barron et al., 2014; Sheffield and Dombeck, 2014). We found that deep cells remapped to a greater degree than superficial in the same context, as superficial place maps were more stable upon re-exposure to an identical context. Deep place maps, on the other hand, were stabilized during GOL, and activity in the deep sublayer predicted task performance to a significantly greater degree than superficial. These findings extend our understanding of the radial organization of CA1 PC dynamics during spatial navigation and learning and provide a mechanism through which the hippocampus can simultaneously convey both a stable map of space and behaviorally relevant task-related information.

Results

Two-photon calcium imaging of deep and superficial CA1 PCs *in vivo*

To longitudinally track the activity of deep and superficial CA1 PCs *in vivo*, we performed two-photon Ca^{2+} imaging (Fig 1a) in head-restrained mice as the animals engaged in two different behaviors, random foraging or goal-oriented learning, on a treadmill in a linear environment (Fig 1b, Danielson et al., 2016; Kaifosh et al., 2013; Lovett-Barron et al., 2014). Mice were stereotactically injected with rAAV1/2(*CaMKII-GCaMP6f*) to express the GCaMP6f Ca^{2+} indicator in CA1 PCs in the dorso-intermediate hippocampus, and implantation of a chronic window provided the optical access necessary for imaging

(Supplemental Movie 1, Dombeck et al., 2010; Kaifosh et al., 2013; Lovett-Barron et al., 2014). By coupling our image acquisition control to a piezoelectric crystal, we could rapidly toggle between z-coordinates to near-simultaneously image both sublayers with a resonant scanner (Fig 1a).

In total, we imaged 1,052 unique deep and 3,222 unique superficial CA1 PCs longitudinally as the animals engaged in RF (n = 8 mice) and GOL (n = 6 mice) behaviors. Motion-correction of Ca^{2+} imaging movies was performed using the SIMA package (Kaifosh et al., 2014, see Methods). Segmentation and tracking of CA1 PC somata across multiple imaging sessions was performed in ImageJ and ROI Buddy (Fig S1a), and signal extraction was performed in SIMA. Significant Ca^{2+} transients ($p < 0.05$, Fig 1c, S1b) were identified as previously described (Danielson et al., 2016; Dombeck et al., 2007; Lovett-Barron et al., 2014).

Deep CA1 PCs are a highly active place coding population

Consistent with previous work reporting sublayer-specific differences in activity (Mizuseki et al., 2011), we found that deep cells were significantly more active than superficial (Fig 1e, Supplemental Movie 2), as indicated by a higher frequency, larger amplitude, and longer duration of Ca^{2+} transients (Fig S2a). These results held independently in both the RF and GOL tasks.

In order to compare the degree of spatial tuning between the two populations, we first identified the place coding population within each sublayer according to two different criteria: spatial information (Skaggs et al., 1993) and tuning specificity (Danielson et al., 2016). To calculate tuning specificity, we represented each Ca^{2+} transient as a vector whose orientation indicated the mouse's position at the time of onset, and whose magnitude was inversely weighted by occupancy. The tuning vector was defined as the complex sum of transient vectors, and its magnitude and direction defined the tuning specificity and tuning direction, respectively (Danielson et al., 2016, Fig 1c, see Methods). A shuffle analysis (Fig S1b) allowed us to identify cells with significant spatial information and/or tuning specificity. Under both definitions, we found that the deep sublayer contained a significantly higher fraction of place cells (Fig 1d, f-g, S1c). The observed sublayer difference was consistent with previous work, but because two-photon imaging allows for the inclusion of silent cells, our unbiased estimates were slightly lower than previously reported (Mizuseki et al., 2011). Again these results held independently for the RF and GOL tasks.

While a higher proportion of deep than superficial CA1 PCs were identified as place cells, we found that deep cells fired transients more diffusely, resulting in a modest but consistent decrease in tuning specificity (Fig 1h). Deep CA1 PCs fired transients more consistently in their place fields lap-to-lap (higher sensitivity), though also more outside their place field (lower specificity) (Fig S2b). We did not detect any difference in place field width between the sublayers (Fig S2b). It is important to note that although the population means were not dramatically different for these metrics, the pairwise nature of our recordings allowed us to control for mouse-to-mouse variability and detect a consistent difference across the sublayers.

Superficial place maps are more stable than deep

In addition to representing place, CA1 PCs incorporate non-spatial information to form a representation of environmental context. Hippocampal spatial representations undergo a variable degree of transformation ("remapping") in response to contextual manipulation (Colgin et al., 2008; Karlsson and Frank, 2008; Knierim and Rao, 2003; Lee et al., 2015; Leutgeb et al., 2005, 2004; Muller and Kubie, 1987; Wilson and McNaughton, 1994). In addition, recent work using electrophysiological and single-photon population imaging approaches has revealed that even in familiar environments, hippocampal place cell activity is dynamic over timescales ranging from minutes to weeks (Mankin et al., 2015, 2012; Rubin et al., 2015; Ziv et al., 2013). As these methods are not well suited to resolve the precise radial position of recorded cells, we sought to assess the remapping dynamics and longitudinal stability of the sublayers using two-photon imaging under head-fixed conditions (Fig 2, S5). Mice were trained to perform a random foraging task, in which they ran on the treadmill for three randomly administered water rewards per lap. Following training (see Methods), mice underwent three 12 min imaging sessions separated by 60–90 minutes in distinct, novel, multisensory contexts A and B (Danielson et al., 2016, Fig 2a-b, S3a, S4a). These contexts shared the same treadmill belt but differed in the olfactory, auditory, visual, and tactile cues adhered to the belt. By exposing the animals to contexts A-B-B sequentially, we could compare the similarity of tuning in the A-B condition with that of the B-B condition. Each mouse underwent this protocol 1–3 times, yielding 16 fields-of-view (FOVs) across 7 mice. One mouse with a particularly low place cell fraction (resulting in little overlap between place coding populations) was excluded from the analysis.

We first examined the population vector (PV) correlation (Fig 2c) between deep and superficial place maps in the novel (A-B) and null (B-B) conditions. Population vectors were constructed from the rate maps (see Methods) of the unfiltered deep and superficial populations in each FOV. While place maps in both sublayers were more stable in the B-B condition than in A-B, the B-B stability of superficial maps exceeded that of deep, and the resulting Δstability was therefore greater for superficial than for deep. In addition to this population-level analysis, we assessed the similarity of place fields on a per-cell basis by calculating the tuning curve correlation, defined as the 1D correlation between the occupancy normalized transient rate maps in each session (Fig 2d). This additional analysis yielded similar results, with superficial cells displaying a higher mean correlation in B-B than deep, resulting in a greater Δstability (Fig 2d).

To investigate whether the remapping we observed was due to place cells anchoring their firing fields to the cues on the belt, we examined the distribution of distances to cues on belt B among cells that were cue-associated on belt A. We found only chance levels of anchoring to cues, indicating that the remapping we observed in the A-B condition was unlikely to be primarily cue-driven (Fig S4). We also found that the distributions of centroid shifts in the A-B condition were nearly uniform within FOVs (data not shown), indicating a global remapping rather than a coherent shift of the population.

In a separate cohort of animals (n = 5 mice), we examined whether the difference in stability we observed in the B-B condition extended to longer timescales over repeated exposures to the same context (Fig S5a). Following an identical training protocol, mice ran for water

rewards in context A or B (randomized across mice) once per day for one week. This allowed us to assess stability as a function of days elapsed. Again superficial place maps were more stable than deep, and after nearly a week, deep cells exhibited only chance levels of stability (Fig S5b-c). Put another way, deep cells remapped to a greater degree over time than superficial upon repeated exposure to an identical context.

Head-fixed goal-oriented learning

We next asked whether sublayer-specific spatial coding dynamics during RF extended to spatial goal-oriented learning (GOL). To this end we developed a GOL paradigm with a behavioral readout of task performance that could be implemented under head-fixed conditions. Drawing on previous freely-moving goal-oriented learning designs (Breese et al., 1989; Dupret et al., 2010; Hollup et al., 2001; Kobayashi et al., 2003), we developed a task in which mice had to learn the location of a hidden 10 cm reward zone on the treadmill. During pre-training, reward zones were shuffled each lap, which had the effect of training the mice to run and lick simultaneously, as the animals had no way of predicting the rewards (see Methods). Because of differences in pre-training regimens, we did not perform the RF and GOL experiments in the same animals.

Following pre-training, the GOL mice (n = 6) were placed in a novel context, and a single, fixed, uncued 10-cm reward zone was defined. Over the course of three daily 10-minute imaging sessions for 2–3 days, mice learned the position of the hidden reward zone (Condition I). The reward zone was then moved to a new location (Condition II), and the experiment was repeated (Fig 3a). Total licking and running was similar across days and sessions of the experiment (Fig S3b-c). By monitoring the spatial distribution of licks (Fig 3b), we could assess task performance. As they learned the task, the mice suppressed their licking outside the reward zone and developed an anticipatory lick signal prior to reward zone entry (Fig 3c-d). This resulted in a progressive increase in the fraction of licks occurring in the reward zone, providing a measurable behavioral expression of learning (Fig 3e).

Modulation of place coding dynamics during goal-oriented learning

Because task-related demands and attentional state affect place field stability (Kentros et al., 2004; Markus et al., 1995), we asked how CA1 PC stability during GOL compared with stability during RF (Fig 4). For each sublayer we compared session-to-session stability during the two behavioral states. We defined the task modulation for each sublayer as the difference in stability between RF and GOL.

We found significant differences in the stability of place maps between RF and GOL for the deep but not the superficial sublayer. Both the PV correlation (Fig 4a) and the tuning curve correlation (4c) were significantly higher during GOL than during RF for deep cells. Superficial cells were similarly stable in both conditions. Given that FOV-level measures, in which single-cell measures are averaged across large populations, exhibited correlations well above 0.5 (Fig. 2c-iii, d-iii), it is unlikely that ceiling effects would emerge at the mouse level in our GOL results, though it is difficult to eliminate this possibility with certainty.

In order to directly compare the magnitude of the task-related modulation of stability, we aggregated our stability measures by mouse and performed two separate shuffling analyses. In the first, we shuffled the superficial/deep cell identity within each experiment; in the second, the RF/GOL identities of the experiments were shuffled. Both shuffles for both metrics (Fig 4b,d) suggested that the magnitude of the task-related modulation from RF to GOL was greater for deep than for superficial CA1 PCs. This analysis suggests that the behavioral demands of the GOL task had a greater effect on the stability of deep than of superficial place maps.

Deep CA1 PCs are more strongly modulated by the reward zone

Given the behavioral significance of the reward zone, we next investigated how this region affected place coding within the sublayers (Fig 5). We compared the 24-hour stability of place fields near the reward zone with those away from it. As assessed by both the centroid shift and the tuning curve correlation, we found that deep place fields near the reward zone were significantly more stable than those away from it (Fig 5a-b). Superficial cells did not exhibit this relationship, and a direct, paired comparison between the sublayers revealed that deep CA1 PCs were stabilized by the reward to a greater degree than superficial. Although the stability of deep cells was selectively modulated by the reward zone, the stability of deep goal cells was still comparable to that of an average superficial cell. We also found that place fields near the reward were significantly narrower than those away from it (Fig 5c). Both sublayers exhibited this relationship, but the effect was again more pronounced in the deep sublayer. These results indicate that the behavioral significance of the reward location more strongly modulated deep place coding, and it suggests that the plasticity of deep cells is more sensitive to the behavioral salience of a location.

Deep CA1 PCs predict performance on the goal-oriented learning task

We next sought to analyze the relationship between reward-zone representation and performance in our task and to characterize sublayer-specific dynamics. To this end, we calculated for each place cell the angle between its centroid and the reward zone (Fig 6a). This allowed us for each recording session to associate the mean distance to the goal in deep and superficial sublayers (D_D and D_S, respectively) with the animal's performance in that session (P, the fraction of licks occurring in the reward zone) (Fig 6b). While the mean distance to reward was significantly below the chance distance of $\pi/2$ for both sublayers at the end of learning, we did not detect a significant difference in the magnitude of goal-zone over-representation between the sublayers (Fig 6c). We also asked whether CA1 PCs firing near the reward in Condition I also did so in Condition II (Fig 6d). While there were some cells that fired near both reward locations, we also found cells that fired near only one of the rewards, so the trend was not significant.

The work of Dupret et al., 2010, demonstrated that reward-zone representation was predictive of task performance during goal-oriented learning. This led us to ask whether one sublayer preferentially contributed to this phenomenon, and to this end we compared the relationship between D (distance to reward) and P (performance) between the sublayers. We modeled the mouse's task performance P as a linear combination of D_D, D_S, and t (D_D: mean distance to reward in the deep sublayer; D_S: mean distance to reward in the superficial

sublayer; *t*: day of the experiment). A simple regression analysis indicated highly significant relationships between D_D and t with P, and no relationship between D_S and P (Fig S6). This analysis indicates that over-representation of the reward in the deep, but not the superficial, sublayer is significantly related to task performance. To complement this analysis, we assessed the significance of the difference in relationships between D and P (after all, that one sublayer is predictive and the other is not does not guarantee that one sublayer is *more* predictive of performance than the other). We therefore calculated the difference between the sublayers in the Pearson's correlation between D and P and estimated how likely it was that we would have obtained the observed difference by chance (see Methods). The observed difference well exceeded the shuffle distribution. This model-independent approach supports the differential relationship to performance suggested by the previous analysis and demonstrates that reward-zone representation in the deep sublayer is significantly more predictive of performance than reward representation in the superficial sublayer (Fig 6f).

Discussion

By demonstrating cellular and population-level differences in activity between deep and superficial CA1 PCs, our work provides further insight into the functional architecture supporting spatial processing and learning in the hippocampus. We found that superficial maps were more stable than deep over multiple timescales, while deep place maps tended to remap over time. Deep place maps were preferentially stabilized during goal-oriented learning, and representation of the reward zone by deep cells predicted task performance to a greater degree than superficial. These results extend our understanding of the functional organization of the hippocampus and provide a behavioral link to the superficial-deep subdivision in CA1. Broadly interpreted, these findings demonstrate that superficial CA1 PCs provide a more stable map of an environment while their counterparts in deeper layers provide a more flexible representation that is shaped by learning about salient features in the environment.

The longitudinal and cellular-resolution access provided by two-photon imaging was critical in assessing differences between deep and superficial CA1 PCs during learning. Chronic imaging of CA1 PCs has been performed previously using single photon imaging technology (Rubin et al., 2015; Ziv et al., 2013), but our study is the first to track the activity of radial subpopulations of principal cells in the hippocampus over time, and importantly, it is the first to track CA1 PC activity as it evolves over the course of a multi-day spatial learning paradigm. Our findings extend our knowledge of the cellular dynamics underlying spatial goal oriented learning.

One possible explanation for the observed differences in stability between the sublayers is biased excitatory and inhibitory intrahippocampal afferent connectivity. Deep CA1 PCs are preferentially targeted by area CA2 (Kohara et al., 2014), which was recently shown to exhibit similarly unstable coding of context and space (Kay et al., 2016; Lee et al., 2015; Lu et al., 2015; Mankin et al., 2015). In addition, the work by Mizuseki et al., 2011, suggested deep CA1 PCs are more strongly driven by entorhinal input, and it is possible that cortical and subcortical (thalamic or neuromodulatory) inputs – which are known to be crucial for the stability and remapping of hippocampal representations (Brandon et al., 2014; Ito et al.,

2015; Lu et al., 2013; Miao et al., 2015) – also contributed to the observed differences. In addition, it was recently shown that dendritic activity in CA1 PCs is an important determinant of place field stability (Sheffield and Dombeck, 2014), and differences in compartment specific inhibitory inputs (Lee et al., 2014; Valero et al., 2015) are well suited for the regulation of dendritic activity.

It is also possible that radial subpopulations of CA1 PCs also respond differently to similar afferent inputs because of differences in intrinsic excitability or of differences in the coupling of neurotransmitters to various channels or signaling pathways. For example, in addition to the lamina-specific distribution of calbindin and zinc (Slomianka et al., 2011) a recent study has demonstrated transcriptional gradients in ~70 genes along the radial axis (Cembrowski et al., 2016), which could manifest in cell-intrinsic differences in protein products and electrophysiological properties. Consistent with this notion, studies have demonstrated significant electrophysiological differences along the radial axis. Superficial CA1 PCs exhibit more depolarized resting membrane potentials and larger h-current-mediated sag potentials in response to hyperpolarizing current pulses (Jarsky et al., 2008, Lee et al., 2014). Moreover, as recently reported (Maroso et al., 2016), superficial but not deep CA1 PCs exhibit a cannabinoid type-1 receptor (CB1R)-mediated regulation of h-currents.

In order to assess how the sublayers responded to a change in task demands and associated attentional state, we developed a novel head-fixed GOL task that allowed us to monitor subpopulation dynamics over the course of a multi-day learning paradigm. In freely moving animals it has been reported the attentional demands of a task strongly affect the selectivity and stability of place maps (Kentros et al., 2004; Markus et al., 1995). Our results confirm these findings and expand on them by showing that the magnitude of the effect was greater among deep CA1 PCs. Interestingly, the stability of deep place cells was strongly influenced by the reward zone. As the hippocampus receives neuromodulatory inputs related to both the behavioral state of the animal and the reward, it is possible that a differential response to these inputs, such as acetylcholine (Hasselmo, 2006, Brandon et al., 2014) or dopamine (Kentros et al., 2004; McNamara et al., 2014), may underlie sublayer-specific modulation observed here. That the enhanced overall stability of deep CA1 PCs during GOL was not entirely attributable to goal-zone cells suggests the involvement of both attention- and reward-related neuromodulation (Atherton et al., 2015).

It is also possible that the deep vs. superficial subdivision is a reflection of dual input streams. A number of prominent functional dissociations have been established between the medial (MEC) and lateral (LEC) entorhinal cortices, providing dual processing streams into the hippocampus (Knierim et al., 2006). In particular, the MEC is proposed to be primarily involved in path integration based on global frames of reference, while the LEC is proposed to primarily process information related to individual items and location based on a local frame of reference (Hargreaves et al., 2005; Knierim et al., 2014; Neunuebel et al., 2013). Importantly, these differences are mirrored by a gradient in place coding properties along the proximo-distal axis of hippocampal area CA1 (Henriksen et al., 2010). While the exact pattern of afferent connectivity of radial CA1 PC subpopulations with MEC and LEC remains to be determined, it is tempting to speculate that similar to the transverse axis, radial

subpopulations might also be differently innervated by the MEC and LEC. In this framework, our data would be consistent with a preferential innervation of superficial CA1 PCs by LEC, which may explain the increased stability of superficial place maps in our experiments, where proximal cues and landmarks are more relevant. Future experimental and theoretical work will help to delineate how these differences in neural response patterns emerge from differences in cellular properties or connectivity along the radial axis and how the functional dissociation along the radial axis interacts with other parallel information streams in the hippocampus (Knierim et al., 2006).

The remapping we observed during our GOL task is in line with previous reports of goal-zone over-representation (Breese et al., 1989; Dupret et al., 2010; Hollup et al., 2001; Kobayashi et al., 1997). While superficial place fields throughout the belt were similarly stable, deep cells near the reward were significantly more stable than those away from it. This may indicate that the plasticity of deep place fields is at least partly determined by the behavioral salience of a location. Importantly, population-level reward representation in the deep sublayer was significantly more predictive of task performance than was reward representation in the superficial sublayer. These results thus recapitulate the work of Dupret et al., 2010, and expand on it by suggesting that the observed effect is predominantly attributable to activity in the deep sublayer. Dupret et al. also found that reactivation of goal-related assembly patterns during sharp wave ripples (SPW-Rs) predicted memory performance. Recent extracellular (Stark et al., 2014) and intracellular (Valero et al., 2015) *in vivo* electrophysiological recordings demonstrated that superficial CA1 PCs preferentially participate in SPW-Rs in both anaesthetized and awake conditions. As these experiments were not performed during learning, it is possible that recruitment of deep cells into goal-related assemblies during SPW-Rs may be preferentially enhanced during learning, and deep cells may therefore provide a plastic substrate for the neuronal representation of new spatial memories related to remembered goals or specific routes (Dupret et al., 2010; Pfeiffer and Foster, 2013; Singer and Frank, 2009).

In analyzing the relationship between remapping and performance on the GOL task, we considered aggregate population level statistics, rather than focusing on the activity of individual cells, to compare the sublayers. It is possible that a small subset of cells drove the observed differences, and our data do not rule out this interpretation. Future experiments performing repeated manipulation of the reward position should be performed to more reliably assess the possible existence of a dedicated reward prediction-signaling population. Relatedly, it has recently been demonstrated that information is selectively routed to downstream targets via projection-specific subpopulations of CA1 PCs in the ventral hippocampus (Ciocchi et al., 2015). Variability in efferent connectivity has been described between the sublayers (Insausti and Muñoz, 2001; Lee et al., 2014; McGeorge and Faull, 1989), and sublayer-specific signaling to distinct downstream targets could be related to the differential relationship with performance we observed. Analyzing the relationship between task performance and place coding in projection-defined subpopulations should be an area of future investigation.

This study expands our understanding of the network dynamics underlying spatial goal-oriented learning, and it assigns behavioral relevance to the superficial-deep subdivision in

area CA1. Previous studies have described sublayer-specific differences in gene expression, neurochemical content, connectivity, and spike dynamics. By tracking sublayer activity over the course of a multi-day learning paradigm, our study extends these previous findings to include differences in place map stability and in dynamics during spatial reward learning. While further theoretical studies of hippocampal mnemonic functions may help to understand the computational consequences of functional segregation among CA1 PCs, it is worth considering that the radial subdivision may provide a means through which the hippocampus can simultaneously provide both a stable map of the animal's environment and behaviorally relevant task-related information. A more refined understanding of the functional organization of the hippocampus at its final output node will provide valuable insights into the fundamental mechanisms supporting diverse mnemonic demands.

Experimental Procedures[1]

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).

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

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). 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).

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. 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 (1.4×105 dB, Bruker) was used to amplify signals

[1]See Supplemental Information for complete Experimental Procedures.

prior to digitization. All experiments were performed at 2x digital zoom, covering 150 μm×150 μm in each imaging plane. For paired recordings of deep and superficial CA1 PCs, images were acquired at 7 Hz with 512×512 pixels per plane. Planes were separated by 25 μm.

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.

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 Supplemental Methods), 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 automatically warped to the remaining imaging sessions (Fig. S1a) using the SIMA project's ROI Buddy graphical user interface (Kaifosh et al., 2014).

Signal extraction—Dynamic GCaMP6f fluorescence signals were extracted using SIMA according to the previously described formulation (Kaifosh et al., 2014). We computed the relative fluorescence changes (ΔF/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 (Dombeck et al., 2007; Lovett-Barron et al., 2014).

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 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). The complete implementations of the tuning specificity and spatial information analyses are included in the Supplemental Experimental Procedures.

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. 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, it 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. The specific implementations of the three remapping metrics are provided in the Supplemental Experimental Procedures.

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 in the Supplemental Experimental Procedures.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Drs. Matthew Lovett-Barron and Andres Grosmark for helpful comments on the manuscript. We thank Dr. Boris V. Zemelman for the rAAV(*CaMKII-GCaMP6f*). N.B.D. is supported by NINDS F30NS090819. J.D.Z. is supported by NIMH 1F31MH105169. P.K. is a Howard Hughes Medical Institute Predoctoral Fellow. A.L. is supported by NIMH 1R01MH100631, 1U01NS090583, 1R01NS094668, the Searle Scholars Program, the Human Frontier Science Program, and the McKnight Memory and Cognitive Disorders Award.

References

Angevine JJ. Time of neuron origin in the hippocampal region: an autoradiographic study in the mouse. Exp. Neurol. 1965; 2:1–70.

Atherton LA, Dupret D, Mellor JR. Memory trace replay: The shaping of memory consolidation by neuromodulation. Trends Neurosci. 2015; 38:560–570. [PubMed: 26275935]

Bayer S. Development of the hippocampal region in the rat. I. Neurogenesis examined with 3H-thymidine autoradiography. J. Comp. Neurol. 1980; 190:87–114. [PubMed: 7381056]

Brandon MP, Koenig J, Leutgeb JK, Leutgeb S. New and Distinct Hippocampal Place Codes Are Generated in a New Environment during Septal Inactivation. Neuron. 2014; 82:789–796. [PubMed: 24853939]

Breese CR, Hampson RE, Deadwyler SA. Hippocampal place cells: stereotypy and plasticity. J. Neurosci. 1989; 9:1097–1111. [PubMed: 2703869]

Buzsáki G, Moser EI. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nat. Neurosci. 2013; 16:130–138. [PubMed: 23354386]

Cembrowski MS, Bachman JL, Wang L, Sugino K, Shields BC, Spruston N. Spatial Gene-Expression Gradients Underlie Prominent Heterogeneity of CA1 Pyramidal Neurons. Neuron. 2016; 89:351–368. [PubMed: 26777276]

Ciocchi S, Passecker J, Malagon-Vina H, Mikus N, Klausberger T. Selective information routing by ventral hippocampal CA1 projection neurons. Science. 2015; 348:560–563. [PubMed: 25931556]

Colgin LL, Moser EI, Moser M-B. Understanding memory through hippocampal remapping. Trends Neurosci. 2008; 31:469–477. [PubMed: 18687478]

Danielson NB, Kaifosh P, Zaremba JD, Lovett-Barron M, Tsai J, Denny CA, Balough EM, Goldberg AR, Drew LJ, Hen R, Losonczy A, Kheirbek MA. Distinct contribution of adult-born hippocampal granule cells to context encoding. Neuron. 2016; 90:1–12. [PubMed: 27054611]

Deshmukh SS, Knierim JJ. Influence of local objects on hippocampal representations: Landmark vectors and memory. Hippocampus. 2013; 23:253–267. [PubMed: 23447419]

Dombeck DA, Khabbaz AN, Collman F, Adelman TL, Tank DW. Imaging large-scale neural activity with cellular resolution in awake, mobile mice. Neuron. 2007; 56:43–57. [PubMed: 17920014]

Dombeck D, Harvey C, Tian L, Looger L, Tank D. Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. Nat. Neurosci. 2010; 13:1433–1440. [PubMed: 20890294]

Dong H-W, Swanson LW, Chen L, Fanselow MS, Toga AW. Genomic-anatomic evidence for distinct functional domains in hippocampal field CA1. Proc. Natl. Acad. Sci. U.S.A. 2009; 106:11794–11799. [PubMed: 19561297]

Dupret D, O'Neill J, Pleydell-Bouverie B, Csicsvari J. The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat. Neurosci. 2010; 13:995–1002. [PubMed: 20639874]

Eichenbaum H, Cohen NJ. Can we reconcile the declarative memory and spatial navigation views on hippocampal function? Neuron. 2014; 83:764–770. [PubMed: 25144874]

Fanselow MS, Dong H-W. Are the dorsal and ventral hippocampus functionally distinct structures? Neuron. 2010; 65:7–19. [PubMed: 20152109]

Frank LM, Brown EN, Wilson M. Trajectory encoding in the hippocampus and entorhinal cortex. Neuron. 2000; 27:169–178. [PubMed: 10939340]

Fyhn M, Molden S, Hollup S, Moser MB, Moser EI. Hippocampal neurons responding to first-time dislocation of a target object. Neuron. 2002; 35:555–566. [PubMed: 12165476]

Gothard KM, Skaggs WE, McNaughton BL. Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. J. Neurosci. 1996; 16:8027–8040. [PubMed: 8987829]

Graves AR, Moore SJ, Bloss EB, Mensh BD, Kath WL, Spruston N. Hippocampal Pyramidal Neurons Comprise Two Distinct Cell Types that Are Countermodulated by Metabotropic Receptors. Neuron. 2012; 76:776–789. [PubMed: 23177962]

Groenewegen HJ, Van der Zee EV, Kortschot A, Witter MP. Organization of the projections from the subiculum to the ventral striatum in the rat. A study using anterograde transport of Phaseolus vulgaris leucoagglutinin. Neuroscience. 1987; 23:103–120. [PubMed: 3683859]

Hargreaves EL, Rao G, Lee I, Knierim JJ. Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. Science. 2005; 308:1792–1794. [PubMed: 15961670]

Hartzell AL, Burke SN, Hoang LT, Lister JP, Rodriguez CN, Barnes CA. Transcription of the immediate-early gene Arc in CA1 of the hippocampus reveals activity differences along the proximodistal axis that are attenuated by advanced age. J. Neurosci. 2013; 33:3424–3433. [PubMed: 23426670]

Henriksen EJ, Colgin LL, Barnes CA, Witter MP, Moser M-B, Moser EI. Spatial representation along the proximodistal axis of CA1. Neuron. 2010; 68:127–137. [PubMed: 20920796]

Hok V, Lenck-Santini P-P, Roux S, Save E, Muller RU, Poucet B. Goal-Related Activity in Hippocampal Place Cells. J. Neurosci. 2007; 27:472–482. [PubMed: 17234580]

Hollup SA, Molden S, Donnett JG, Moser M, Moser EI. Accumulation of Hippocampal Place Fields at the Goal Location in an Annular Watermaze Task. J. Neurosci. 2001; 21:1635–1644. [PubMed: 11222654]

Hölscher C, Jacob W, Mallot HA. Reward modulates neuronal activity in the hippocampus of the rat. Behav. Brain Res. 2003; 142:181–191. [PubMed: 12798280]

Igarashi KM, Ito HT, Moser EI, Moser M-B. Functional diversity along the transverse axis of hippocampal area CA1. FEBS Lett. 2014; 588:2470–2476. [PubMed: 24911200]

Insausti R, Muñoz M. Cortical projections of the non-entorhinal hippocampal formation in the cynomolgus monkey (Macaca fascicularis). Eur. J. Neurosci. 2001; 14:435–451. [PubMed: 11553294]

Ito HT, Zhang S, Witter MP, Moser EI, Moser M. A prefrontal–thalamo– hippocampal circuit for goal-directed spatial navigation. Nature. 2015; 522:50–55. [PubMed: 26017312]

Jarsky T, Mady R, Kennedy B, Spruston N. Distribution of bursting neurons in the CA1 region and the subiculum of the rat hippocampus. J. Comp. Neurol. 2008; 506:535–547. [PubMed: 18067146]

Jia H, Rochefort NL, Chen X, Konnerth A. In vivo two-photon imaging of sensory-evoked dendritic calcium signals in cortical neurons. Nat. Protoc. 2011; 6:28–35. [PubMed: 21212780]

Jung M, Wiener S, McNaughton B. Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. J. Neurosci. 1994; 14:7347–7356. [PubMed: 7996180]

Kaifosh P, Lovett-Barron M, Turi GF, Reardon TR, Losonczy A. Septo-hippocampal GABAergic signaling across multiple modalities in awake mice. Nat. Neurosci. 2013; 16:1182–1184. [PubMed: 23912949]

Kaifosh P, Zaremba J, Danielson NB. SIMA: Python software for analysis of dynamic fluorescence imaging data. Front. Neuroinform. 2014; 8:1–10. [PubMed: 24501593]

Karlsson MP, Frank LM. Network dynamics underlying the formation of sparse, informative representations in the hippocampus. J. Neurosci. 2008; 28:14271–14281. [PubMed: 19109508]

Kay K, Sosa M, Chung J, Karlsson M, Larkin M, Grossrubatscher I, Frank L. A hippocampal network for spatial coding during immobility. Nature. 2016; 531:185–190. [PubMed: 26934224]

Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER. Increased attention to spatial context increases both place field stability and spatial memory. Neuron. 2004; 42:283–295. [PubMed: 15091343]

Kesner RP, Rolls ET. A computational theory of hippocampal function, and tests of the theory: new developments. Neurosci. Biobehav. Rev. 2015; 48:92–147. [PubMed: 25446947]

Kjelstrup KB, Solstad T, Brun VH, Hafting T, Leutgeb S, Witter MP, Moser EI, Moser M-B. Finite scale of spatial representation in the hippocampus. Science. 2008; 321:140–143. [PubMed: 18599792]

Kjelstrup KG, Tuvnes FA, Steffenach H-A, Murison R, Moser EI, Moser M-B. Reduced fear expression after lesions of the ventral hippocampus. Proc. Natl. Acad. Sci. U. S. A. 2002; 99:10825–10830. [PubMed: 12149439]

Knierim J, Lee I, Hargreaves EL. Hippocampal Place Cells: Parallel Input Streams, Subregional Processing, and Implications for Episodic Memory. Hippocampus. 2006; 16:755–764. [PubMed: 16883558]

Knierim JJ, Hamilton DA. Framing spatial cognition: neural representations of proximal and distal frames of reference and their roles in navigation. Physiol. Rev. 2011; 91:1245–1279. [PubMed: 22013211]

Knierim JJ, Neunuebel JP, Deshmukh SS. Functional correlates of the lateral and medial entorhinal cortex: Objects , path integration and local – global reference frames. Philos. Trans. R. Soc. London. 2014

Knierim JJ, Rao G. Distal landmarks and hippocampal place cells: Effects of relative translation versus rotation. Hippocampus. 2003; 13:604–617. [PubMed: 12921350]

Kobayashi T, Nishijo H, Fukuda M, Bures J, Ono T. Task-dependent representations in rat hippocampal place neurons. J. Neurophysiol. 1997; 78:597–613. [PubMed: 9307098]

Kobayashi T, Tran AH, Nishijo H, Ono T, Matsumoto G. Contribution of hippocampal place cell activity to learning and formation of goal-directed navigation in rats. Neuroscience. 2003; 117:1025–1035. [PubMed: 12654354]

Kohara K, Pignatelli M, Rivest AJ, Jung H-Y, Kitamura T, Suh J, Frank D, Kajikawa K, Mise N, Obata Y, Wickersham IR, Tonegawa S. Cell type-specific genetic and optogenetic tools reveal hippocampal CA2 circuits. Nat. Neurosci. 2014; 17:269–279. [PubMed: 24336151]

Lee H, Wang C, Deshmukh SS, Knierim JJ. Neural Population Evidence of Functional Heterogeneity along the CA3 Transverse Axis: Pattern Completion versus Pattern Separation. Neuron. 2015; 87:1093–1105. [PubMed: 26298276]

Lee S, Marchionni I, Bezaire M, Varga C, Danielson NB, Lovett-Barron M, Losonczy A, Soltesz I. Parvalbumin-positive basket cells differentiate among hippocampal pyramidal cells. Neuron. 2014; 82:1129–1144. [PubMed: 24836505]

Leutgeb S, Leutgeb JK, Barnes C, Moser E, McNaughton B, Moser M-B. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science. 2005; 309:619–623. [PubMed: 16040709]

Leutgeb S, Leutgeb JK, Treves A, Moser M-B, Moser EI. Distinct ensemble codes in hippocampal areas CA3 and CA1. Science. 2004; 305:1295–1298. [PubMed: 15272123]

Lorente De Nó R. Studies on the structure of the cerebral cortex. II. Continuation of the study of the ammonic system. J. für Psychol. und Neurol. 1934; 46:113–177.

Lovett-Barron M, Kaifosh P, Kheirbek MA, Danielson NB, Zaremba JD, Reardon TR, Turi GF, Hen R, Zemelman BV, Losonczy A. Dendritic inhibition in the hippocampus supports fear learning. Science. 2014; 343:857–863. [PubMed: 24558155]

Lu L, Igarashi KM, Witter MP, Moser EI, Moser MB. Topography of place maps along the CA3-to-CA2 axis of the hippocampus. Neuron. 2015; 87:1078–1092. [PubMed: 26298277]

Lu L, Leutgeb JK, Tsao A, Henriksen EJ, Leutgeb S, Barnes CA, Witter MP, Moser M-B, Moser EI. Impaired hippocampal rate coding after lesions of the lateral entorhinal cortex. Nat. Neurosci. 2013; 16:1085–1093. [PubMed: 23852116]

Mankin EA, Diehl GW, Sparks FT, Leutgeb S, Leutgeb JK. Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts. Neuron. 2015; 85:190–202. [PubMed: 25569350]

Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK. Neuronal code for extended time in the hippocampus. Proc. Natl. Acad. Sci. U.S.A. 2012; 109:19462–19467. [PubMed: 23132944]

Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J. Neurosci. 1995; 15:7079–7094. [PubMed: 7472463]

Maroso M, Szabo G, Kim HK, Alexander A, Bui A, Lee SH, Lutz B, Soltesz I. Cannabinoid control of learning and memory through HCN channels. Neuron. 2016; 89:1059–1073. [PubMed: 26898775]

Marr D. Simple memory: a theory for archicortex. Philos. Trans. R. Soc. Lond. B. Biol. Sci. 1971; 262:23–81. [PubMed: 4399412]

Maurer AP, Van Rhoads SR, Sutherland GR, Lipa P, McNaughton BL. Self-motion and the origin of differential spatial scaling along the septo-temporal axis of the hippocampus. Hippocampus. 2005; 15:841–852. [PubMed: 16145692]

McGeorge A, Faull R. The organization of the projection from the cerebral cortex to the striatum in the rat. Neuroscience. 1989; 29:503–537. [PubMed: 2472578]

McNamara CG, Tejero-Cantero Á, Trouche S, Campo-Urriza N, Dupret D. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. Nat. Neurosci. 2014; 17:1658–1660. [PubMed: 25326690]

Miao C, Cao Q, Ito HT, Yamahachi H, Witter MP, Moser MB, Moser EI. Hippocampal remapping after partial inactivation of the medial entorhinal cortex. Neuron. 2015; 88:590–603. [PubMed: 26539894]

Mizuseki K, Diba K, Pastalkova E, Buzsáki G. Hippocampal CA1 pyramidal cells form functionally distinct sublayers. Nat. Neurosci. 2011; 14:1174–1181. [PubMed: 21822270]

Moita M, Rosis S, Zhou Y, LeDoux JE, Blair HT. Putting fear in its place: remapping of hippocampal place cells during fear conditioning. J. Neurosci. 2004; 24:7015–7023. [PubMed: 15295037]

Moser E, Moser MB, Andersen P. Spatial learning impairment parallels the magnitude of dorsal hippocampal lesions, but is hardly present following ventral lesions. J. Neurosci. 1993; 13:3916–3925. [PubMed: 8366351]

Muller RU, Kubie JL. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci. 1987; 7:1951–1968. [PubMed: 3612226]

Muzzio IA, Kentros C, Kandel E. What is remembered? Role of attention on the encoding and retrieval of hippocampal representations. J. Physiol. 2009; 587:2837–2854. [PubMed: 19525568]

Neunuebel JP, Yoganarasimha D, Rao G, Knierim JJ. Conflicts between local and global spatial frameworks dissociate neural representations of the lateral and medial entorhinal cortex. J. Neurosci. 2013; 33:9246–9258. [PubMed: 23719794]

O'Keefe J, Dostrovsky J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. 1971; 34:171–175. [PubMed: 5124915]

O'Keefe, J.; Nadel, L. The Hippocampus as a cognitive map. Oxford, UK: Clarendon; 1978.

Pastalkova E, Itskov V, Amarasingham A, Buzsáki G. Internally Generated Cell Assembly Sequences in the Rat Hippocampus. Science. 2008; 321:1322–1327. [PubMed: 18772431]

Pfeiffer BE, Foster DJ. Hippocampal place-cell sequences depict future paths to remembered goals. Nature. 2013; 497:74–79. [PubMed: 23594744]

Rubin A, Geva N, Sheintuch L, Ziv Y. Hippocampal ensemble dynamics timestamp events in long-term memory. Elife. 2015; 4:1–16.

Schlessinger A, Cowan W, Swanson L. The time of origin of neurons in Ammon's horn and the associated retrohippocampal fields. Anat. Embryol. (Berl). 1978; 154:153–173. [PubMed: 686395]

Scoville WB, Milner B. Loss of recent memory after bilateral hippocampal lesions. J. Neurol. Neurosurg. Psychiat. 1957; 20:11–21. [PubMed: 13406589]

Sheffield MEJ, Dombeck DA. Calcium transient prevalence across the dendritic arbour predicts place field properties. Nature. 2014; 517:200–204. [PubMed: 25363782]

Singer AC, Frank LM. Rewarded Outcomes Enhance Reactivation of Experience in the Hippocampus. Neuron. 2009; 64:910–921. [PubMed: 20064396]

Skaggs, WE.; Mcnaughton, BL.; Markus, EJ.; Gothard, KM. An Information- Theoretic Approach to Deciphering the Hippocampal Code. In: Hanson, S.; Cowan, J.; Giles, C., editors. Advances in Neural Information Process Systems (NIPS). Vol. 5. 1993. p. 1030-1037.

Slomianka L, Amrein I, Knuesel I, Sørensen JC, Wolfer DP. Hippocampal pyramidal cells: the reemergence of cortical lamination. Brain Struct. Funct. 2011; 216:301–317. [PubMed: 21597968]

Squire LR, Kosslyn S, Zola-Morgan S, Haist F, Musen G. Memory and the Hippocampus: A Synthesis From Findings With Rats, Monkeys, and Humans. Psychol. Rev. 1992; 99:195–231. [PubMed: 1594723]

Stark E, Roux L, Eichler R, Senzai Y, Royer S, Buzsáki G. Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. Neuron. 2014; 83:467–480. [PubMed: 25033186]

Strange BA, Witter MP, Lein ES, Moser EI. Functional organization of the hippocampal longitudinal axis. Nat. Rev. Neurosci. 2014; 15:655–669. [PubMed: 25234264]

Thompson CL, Pathak SD, Jeromin A, Ng LL, MacPherson CR, Mortrud MT, Cusick A, Riley ZL, Sunkin SM, Bernard A, Puchalski RB, Gage FH, Jones AR, Bajic VB, Hawrylycz MJ, Lein ES. Genomic anatomy of the hippocampus. Neuron. 2008; 60:1010–1021. [PubMed: 19109908]

Valero M, Cid E, Averkin RG, Aguilar J, Sanchez-Aguilera A, Viney TJ, Gomez-Dominguez D, Bellistri E, de la Prida LM. Determinants of different deep and superficial CA1 pyramidal cell dynamics during sharp-wave ripples. Nat. Neurosci. 2015; 18:1281–1290. [PubMed: 26214372]

Wilson MA, McNaughton BL. Reactivation of hippocampal ensemble memories during sleep. Science. 1994; 265:676–679. [PubMed: 8036517]

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron. 2000; 27:623–633. [PubMed: 11055443]

Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, El Gamal A, Schnitzer MJ. Long-term dynamics of CA1 hippocampal place codes. Nat. Neurosci. 2013; 16:264–266. [PubMed: 23396101]

Highlights

- Ca^{2+} imaging demonstrates sublayer-specific place coding dynamics
- Superficial place maps were more stable than deep over multiple timescales
- Goal-oriented learning modulates deep cells to a greater degree than superficial
- Representation of reward by deep cells predicts task performance.

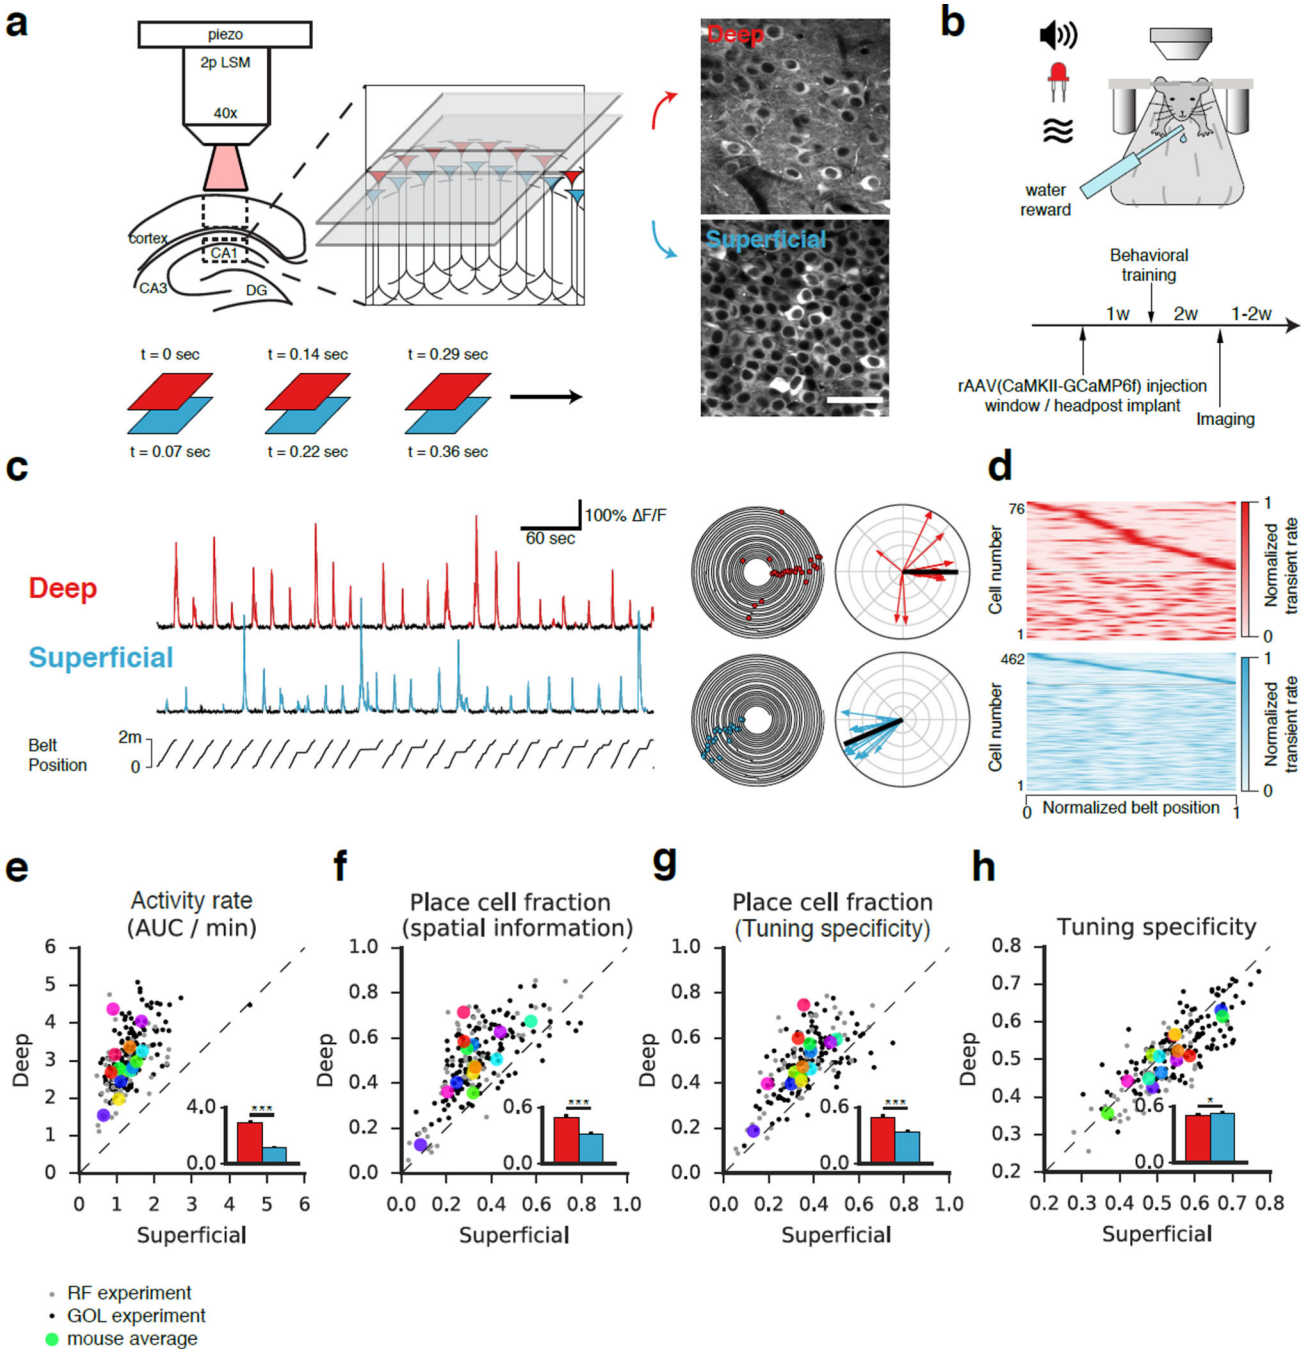

Figure 1. Optical dissection of the radial axis of hippocampal area CA1

(a) *Left:* Experimental schematic of simultaneous two-photon imaging of GCaMP6f in deep and superficial sublayers of dorsal CA1 using a piezoelectric crystal. *Right:* Time-averaged image sequences from a representative recording session. Planes were separated by 25 μm. *Bottom:* Timing of frame acquisition. **(b)** Schematic of the behavioral apparatus. Head-fixed mice ran on a cue-rich 2 m treadmill while tone, light, odor, and reward were presented by a microcircuit controller. The experimental timeline is shown below. **(c)** *Left:* Representative fluorescence traces from spatially-tuned deep and superficial CA1 PCs. Belt position is

plotted at bottom. Detected Ca^{2+} transients ($p<0.05$) are indicated in color. *Middle:* Polar trajectory plots from the same cells. Position and time are represented as the angular and radial coordinates, respectively. Significant running-related transients are indicated with dots. *Right:* Spatial tuning plots. Each running-related transient is represented as a vector (direction = position at onset, magnitude = inverse of occupancy). The complex sum of these vectors forms the tuning vector (black). **(d)** Spatial tuning heatmaps from a single recording session performed at low zoom. Each row represents a cell, and the x-axis represents the treadmill position. Heatmap intensity reflects the occupancy-normalized transient rate, and each cell is normalized to peak. Identified place cells are shown at top sorted by peak. **(e - f)** Firing properties and spatial coding features of deep and superficial CA1 PCs. Within-session averages of paired deep and superficial recordings (either during RF or GOL) are indicated by small closed dots. Within-mouse averages are indicated by the large colored dots. Means across mice are shown in the insets (mean +/− s.e.m.). **(e)** Activity rate (area under curve of significant transients per unit time, AUC/min) is defined as the area under significant Ca^{2+} transients divided by recording duration. Deep cells were significantly more active than superficial (n=14 mice, $p<0.001$, paired T-Test). **(f)** Place cell fraction as determined by spatial information. A higher fraction of deep CA1 PCs are identified as spatially tuned (n=14 mice, $p<0.001$, paired T-Test). **(g)** Place cell fraction as determined by tuning specificity. A higher fraction of deep CA1 PCs are identified as spatially tuned using this definition as well (n=14 mice, $p<0.001$, paired T-Test). **(h)** Tuning specificity was modestly but consistently higher in superficial than deep CA1 PCs (n=14 mice, $p<0.05$, paired T-Test).

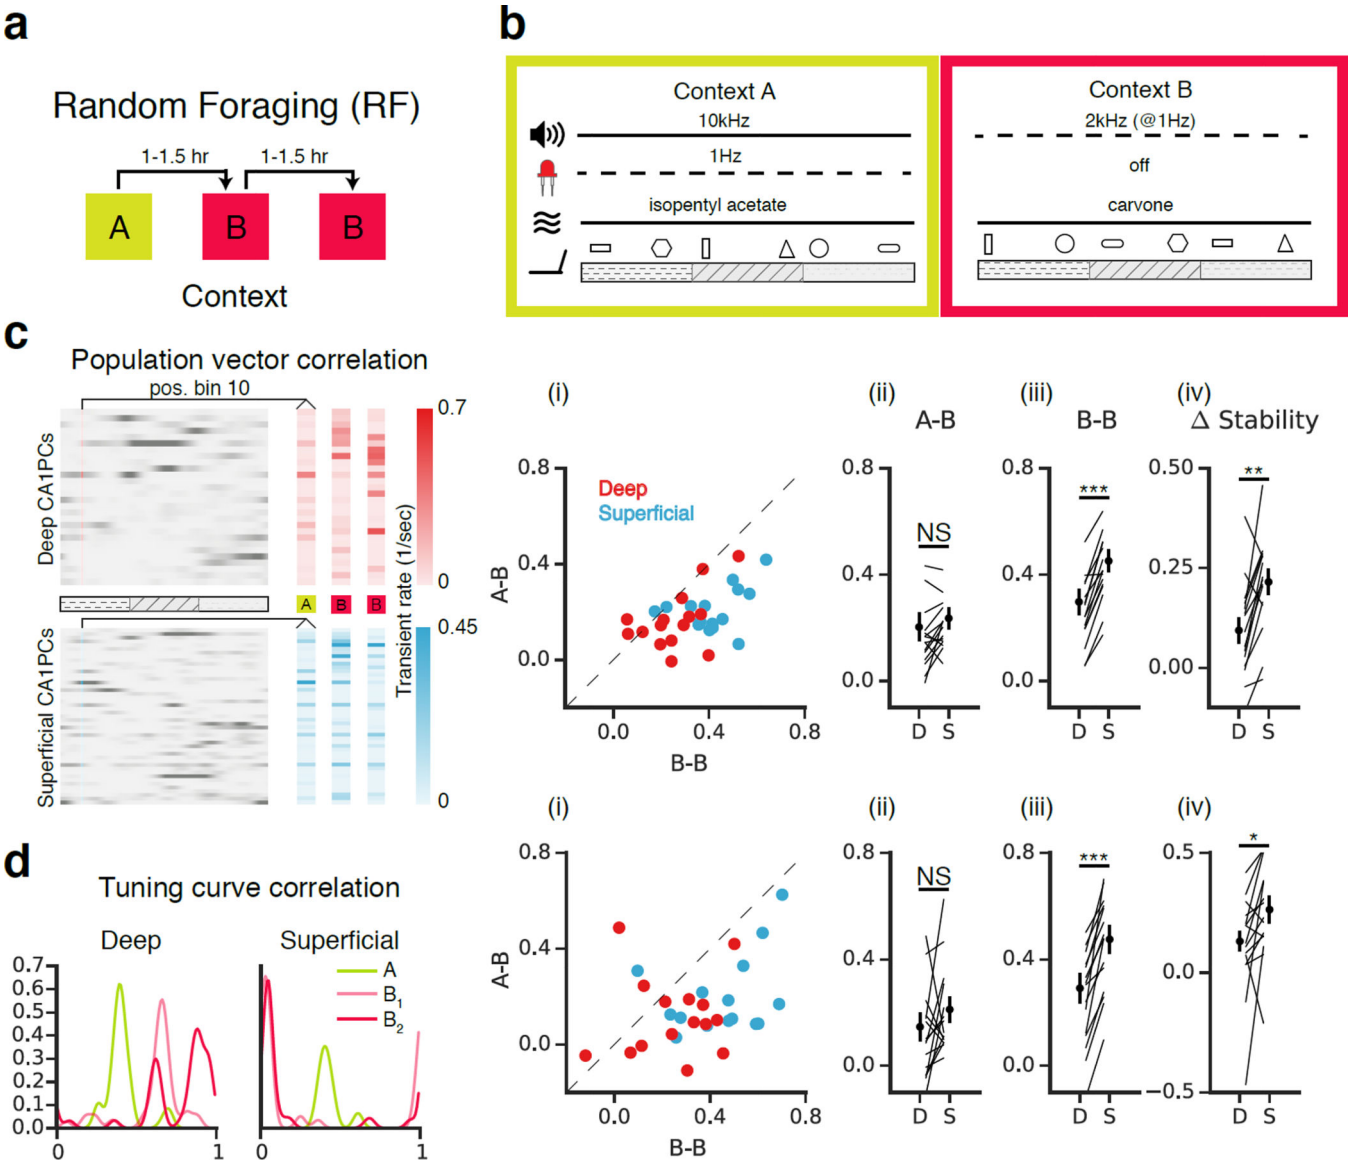

Figure 2. Remapping of superficial and deep place fields during random foraging (RF)

(a) Experimental schematic. Mice ran on the treadmill for randomly administered water rewards (3/lap) during three 12-minute sessions in contexts A and B. Sessions were separated by 60–90 min. Five mice underwent this protocol for 1–3 days. **(b)** Multisensory contexts A and B differed in auditory, visual, olfactory, and tactile cues. **(c)** Population vector analysis schematic. Rate maps for individual cells (rows) were combined to form population vectors (columns) as a function of position. Vectors corresponding to the same position were compared in the A-B and B-B conditions. *(i)* Scatterplots of the mean PV correlation across position bins in each FOV in the B-B vs. A-B conditions (n=15 FOVs). Statistical tests were performed by averaging across FOVs to obtain one measure per mouse (n=7). *(ii-iv)* Individual FOVs are represented as lines, and circles indicate means across mice (mean +/− s.e.m.). *(ii)* A-B: Deep and superficial cells remapped similarly (n=7 mice, p=0.31, paired T-Test). *(iii)* B-B: PV correlations were greater in superficial than in deep

(n=7 mice, p=<0.001, paired T-Test). *(iv)* [(A-B) – (B-B)]: Both sublayers were more stable in the B-B than A-B conditions (n=7 mice; Deep: p<0.05; Superficial: p<0.001; 1-Sample T-Test vs. 0), though the magnitude of the Δstability, the difference in stability between conditions, was greater among superficial cells (n =7 mice, p<0.01, paired T-Test). **(d)** Spatial tuning curves shown for example deep and superficial cells. The tuning curve correlation is the 1D correlation between spatial tuning curves. *(i)* Scatterplots of tuning curve correlations in A-B and B-B for deep and superficial FOVs. *(ii-iv)* Data presented as in (c). *(ii)* A-B: tuning curve correlations were similar between deep and superficial (n=7 mice, p=0.19, paired T-Test). *(iii)* B-B: tuning curve correlations were higher in superficial than deep (n=7 mice, p<0.001, paired T-Test). *(iv)* [(B-B) – (A-B)]: Both sublayers were more stable in the B-B than A-B conditions (n=7 mice; Deep: p<0.05; Superficial: p<0.01; 1-Sample T-Test vs. 0); the magnitude of the difference was higher in superficial than deep CA1 PCs (n=7 mice, p<0.05, paired T-Test).

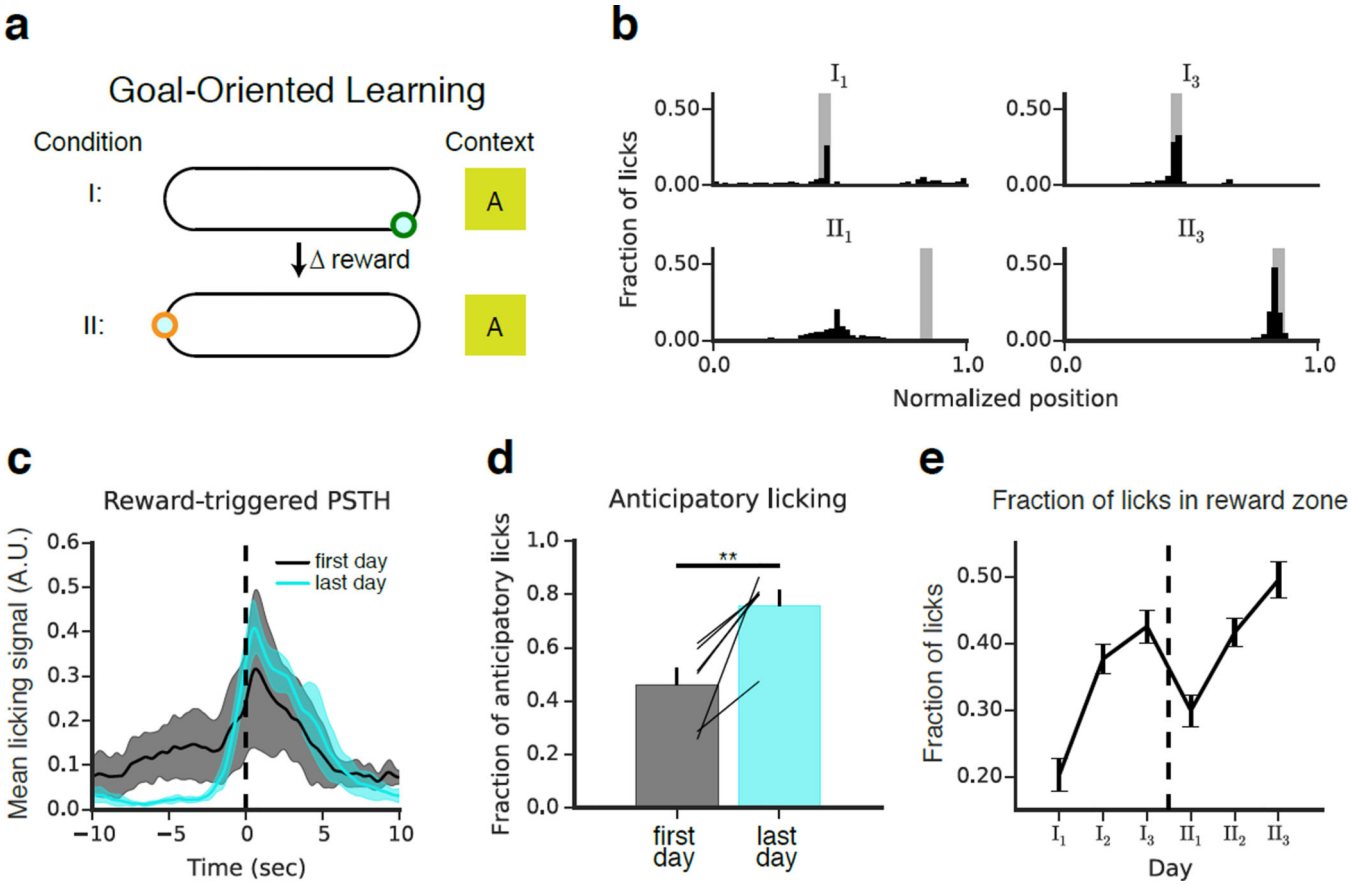

Figure 3. Goal-oriented learning task for head-fixed imaging

(a) Schematic of the goal-oriented learning (GOL) task. Mice (n=6) searched for an unmarked reward zone, and water rewards were administered only when the mice licked within the fixed 10 cm goal. At the end of condition I, the reward was moved to a new location of the belt, and the experiment was repeated. The same context (A) was maintained throughout. **(b)** Representative licking data from four individual experiments from one mouse performing the task. The fraction of total licks is plotted as a function of position on the belt (50 bins, 4 cm per bin). The reward zone is shaded gray. On the first day of the experiment, the mouse licked diffusely throughout the belt as it searched for the hidden reward zone. By the last day of condition I, the mouse licked selectively at the reward location. After the reward was moved, the mouse continued to lick at the original reward location. It eventually reverted to an exploratory licking state, and by the end of condition II the mouse selectively licked at the new reward location. **(c)** Peri-stimulus time-histogram (PSTH) of licking rate triggered on reward zone entry for the first (black) and last (cyan) days of each condition. PSTHs were calculated for each mouse and smoothed with a one-second Hamming filter. Shaded regions indicate mean +/− st. dev. across mice. Licking was initially diffuse. By the last day, licking outside the reward zone was largely suppressed and rose sharply prior to reward zone entry, reflecting expectation of reward. The effect was highly consistent across mice. **(d)** Anticipatory licking (fraction of non-reward zone licks in the 10 cm preceding the reward) increased significantly by the end of learning (n=6 mice, $p<0.01$, paired T-Test). Error bars indicate mean +/− s.e.m. across mice. **(e)** The fraction of

licks occurring within the reward zone aggregated by recording session and plotted by day (mean +/− s.e.m. across mice). Over time, licking became more selective for the reward zone.

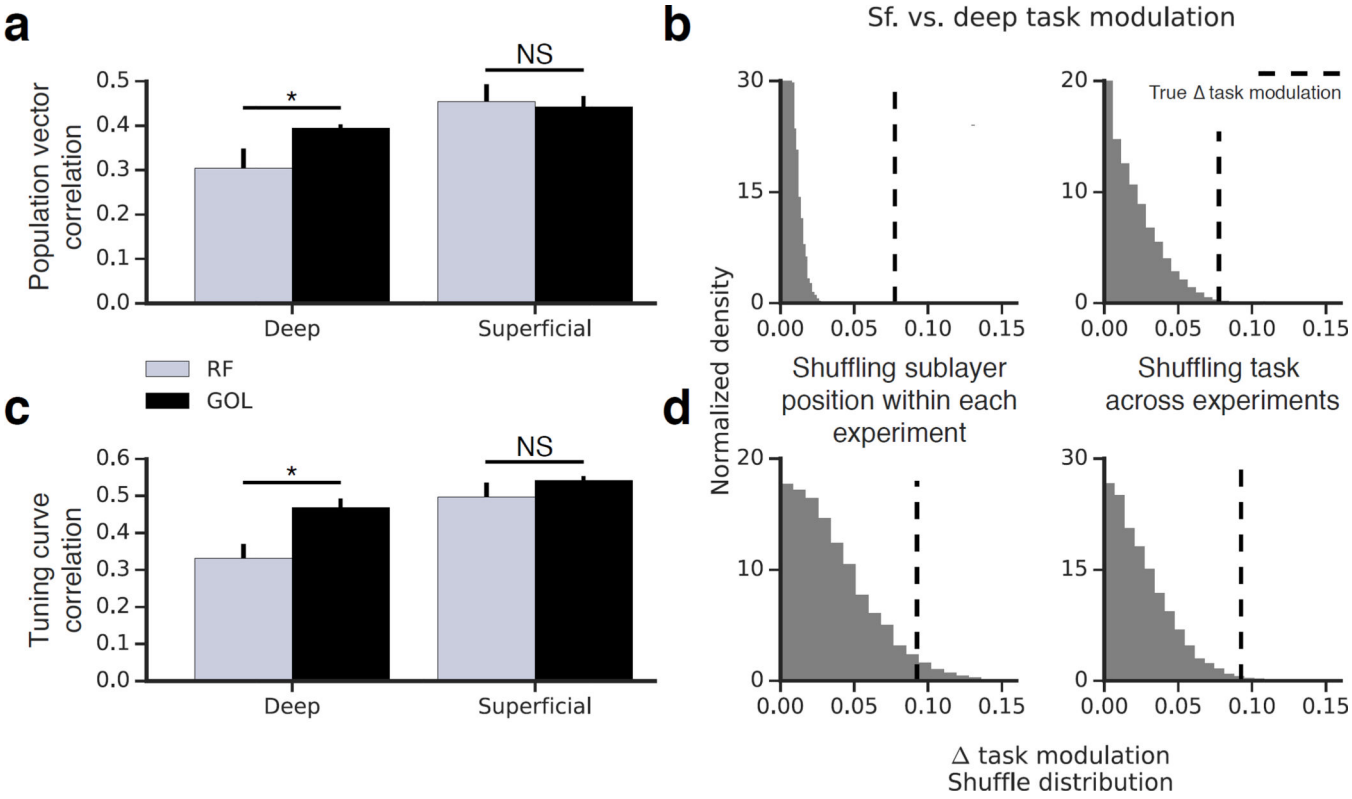

Figure 4. Sublayer-specific modulation of activity by the GOL task

(a) Stability in the B-B condition of RF compared to session-to-session stability in GOL (similar elapsed time of ~90 min.). Deep but not superficial CA1 PCs showed a significant increase in PV correlation in the GOL task as compared to RF (n=7 RF mice, 6 GOL mice; deep: $p<0.05$; superficial: $p=0.36$; Mann-Whitney U Test). Error bars indicate mean +/− s.e.m. across animals. A two-way ANOVA analysis is included in Supplemental Table 1. **(b)** The magnitude of the task modulation was compared across sublayers by performing two shuffling procedures: randomizing cell identity and randomizing experiment identity. Both comparisons suggested the magnitude of task modulation was greater for deep than for superficial cells ($p<0.001$, $p<0.01$). **(c)** The same analysis was performed with tuning curve correlation. Deep but not superficial CA1 PCs showed a significant increase in tuning curve correlation in the GOL task as compared to RF (n=7 RF mice, 6 GOL mice; deep: $p<0.05$; superficial: $p=0.11$; Mann-Whitney U Test). Error bars indicate mean +/− s.e.m. across animals. A two-way ANOVA analysis is included in Supplemental Table 1. **(d)** Both shuffles showed the magnitude of the task modulation was greater for deep than superficial ($p<0.01$, $p<0.01$).

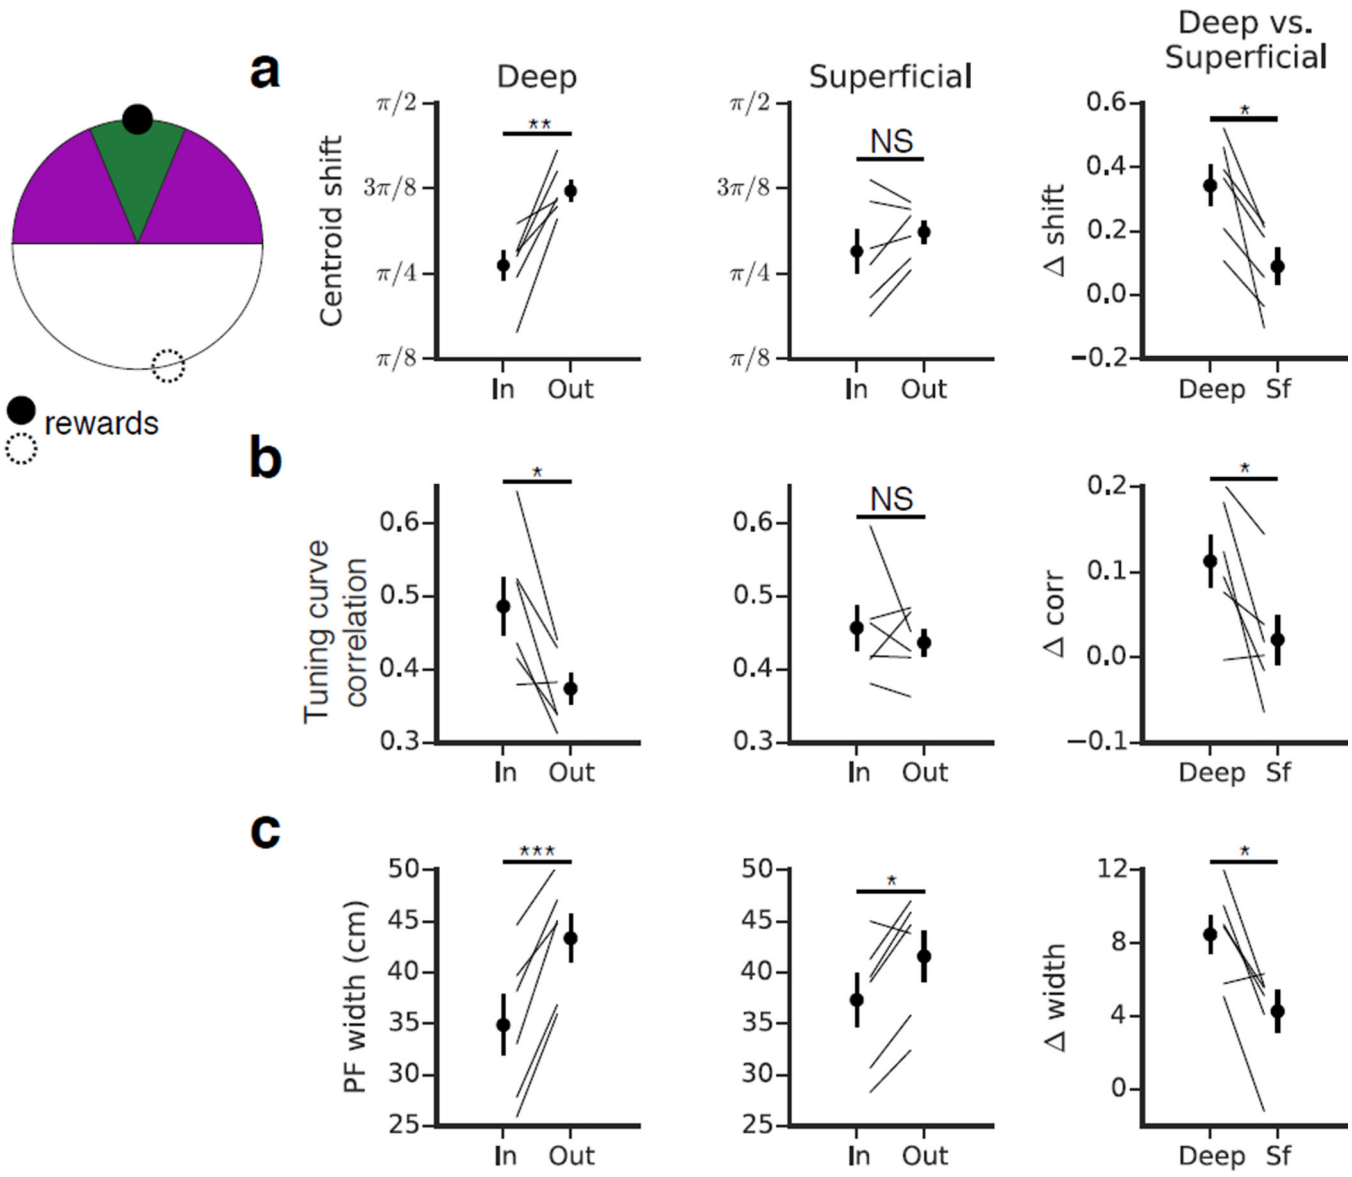

Figure 5. Sublayer-specific modulation of activity in the reward zone

(a - c) For each sublayer (first and second columns) we compared place field properties near the reward (green, within π/8 radians, ~12 cm on either side of the center of the reward zone) versus away from the reward (purple, π/8 to π/2 radians, ~12–50 cm). We performed a paired comparison of the in vs. out differences to directly compare reward-related modulation between the sublayers (third column). Error bars indicate mean +/− s.e.m. across animals. **(a)** Day-to-day centroid shift. Deep place fields near the reward were significantly more stable at 24 hours than those away from the reward (n=6 mice, $p<0.01$, paired T-Test). Superficial CA1 PCs did not show this trend (n=6 mice, $p=0.17$, paired T-Test). The direct, paired comparison of the in versus out stability [$(shift_{out} - shift_{in})_D$ vs. $(shift_{out} - shift_{in})_S$] showed that deep cells were stabilized by the reward to a greater degree than superficial cells (n=6 mice, $p<0.05$, paired T-Test). **(b)** Day-to-day tuning curve correlation. As in (a), deep place fields near the reward were more stable than those away from it (n=6 mice, $p<0.05$,

paired T-Test); this relationship was absent in the superficial sublayer (n=6 mice, p=0.51, paired T-Test). The direct comparison of the sublayers showed that deep cells were stabilized by the reward to a greater degree than superficial (n=6 mice, $p<0.05$, paired T-Test). **(c)** Place field width. In both sublayers, place fields were narrower near the reward than they were away from it (n=6 mice, deep: $p<0.001$, superficial: $p<0.05$, paired T-Test). The magnitude of the effect was greater for deep than for superficial (n=6 mice, $p<0.05$, paired T-Test).

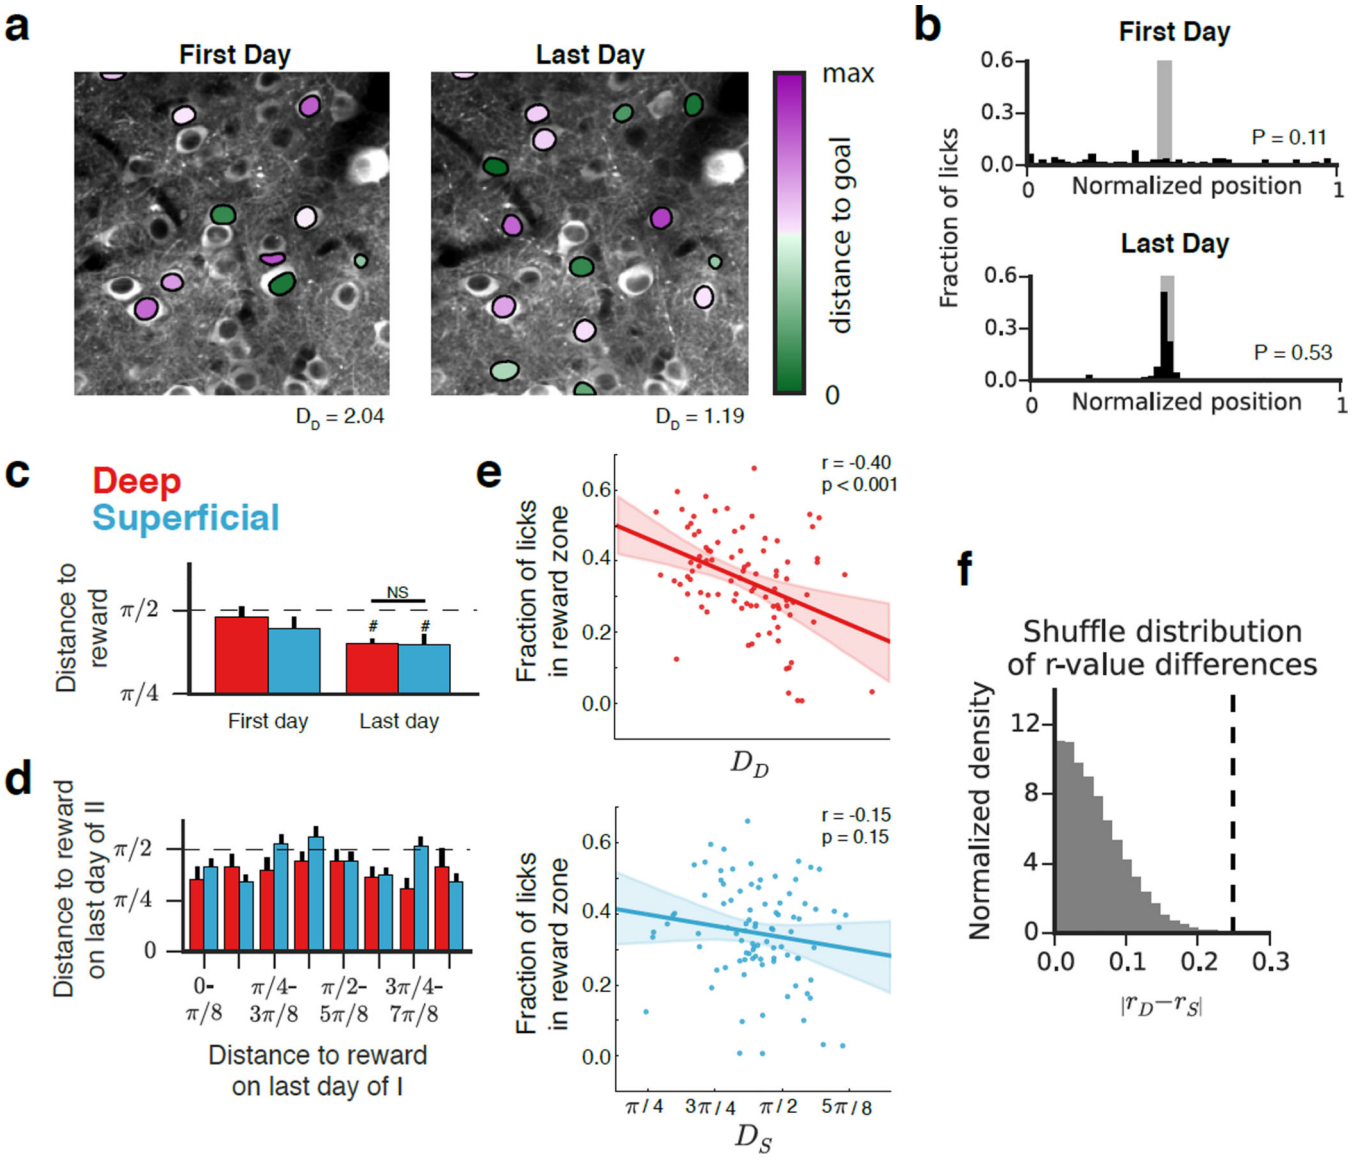

Figure 6. Reward zone representation versus performance on the GOL task

(a) Time-averaged images of the deep sublayer from two different recording sessions are shown in grayscale. Identified place cells from each session are overlaid and colored according to the distance of their centroid to the reward (green = near, purple = away). The mean distance to the reward zone is indicated (radians). **(b)** As in (4b), licking distributions from the two experiments corresponding to (a). **(c)** Mean distance of each place cell centroid to the reward on first and last days of the experiment. On the last day, the mean distance to reward was significantly different from the chance level of π/2 for both sublayers (dashed line; n=6 mice, one sample T-Test, deep: $p<0.001$, superficial: $p<0.05$), but was not significantly different between sublayers (n=6 mice, $p=0.95$, paired T-Test). Error bars indicate mean +/− s.e.m. across animals **(d)** We did not detect a relationship between distance to reward at the end of condition II with distance at the end of I, sublayer, or with the interaction [Type II ANOVA, n=172 deep, 336 superficial place cells, F(Distance end of

I) = 0.10, p(Distance end of I) = 0.76; F(layer) = 0.54, p(layer) = 0.46; F(interaction) = 0.11, p = 0.74]. The dashed line represents the mean distance expected in the case of a uniform place field distribution, and error bars indicate mean +/− s.e.m. across cells. **(e)** Fraction of licks in reward zone (*P*) plotted against D_D (top) and D_S (bottom). Individual points represent single recording sessions. The dashed line indicates the linear fit with the 95% confidence interval shaded. We observed a significant relationship between D_D and P (n = 91 sessions, r= −0.40, p<0.001, Pearson's r) but not between D_S and P (n=91 sessions, r= −0.15, p = 0.15, Pearson's r). **(f)** In order to directly compare the sublayers' relationships to performance on the GOL task, we compared the magnitude of the difference in correlation coefficients (0.25, dashed line) relative to a null distribution (see Methods). The true difference (dashed line) fell outside the shuffle distribution (p<0.001).