The study of Drinnenberg et al. (2018) is a remarkable technical achievement. It combines advances in viral engineering (Deverman et al., 2016) with chemogenetics (Magnus et al., 2011) to suppress the light responses of horizontal cells selectively, completely, and reversibly. It introduces new AAV promoters to monitor and manipulate the activity of cones and ganglion cells without the need for transgenic intersection. It records spike trains of thousands of ganglion cells on large-scale CMOS arrays. The precision of their manipulation and the scale of their observations, enable the authors to provide a comprehensive account of the contributions of horizontal cells to visual processing in the retina.

Horizontal cells are best known for mediating lateral inhibition in the outer retina where their negative feedback reduces photoreceptor responses to large stimuli (Baylor et al., 1971). Lateral inhibition in the outer retina is thought to shape antagonistic receptive fields surrounds of ganglion cells in the inner retina (Thoreson and Mangel, 2012). Drinnenberg et al. (2018) confirm that horizontal cells are preferentially activated by large stimuli and mediate lateral inhibition in the outer retina. However, horizontal cell contributions to ganglion cell surrounds appear to be minor and uniform across cell types. Instead, Drinnenberg et al. (2018) discover that the dominant vertical consequences of horizontal cell function are cell-type-specific changes in response dynamics and response range of ganglion cells (Chaya et al., 2017; Ströh et al., 2018). Their model explains how this diversity arises through parallel processing of negative feedback.

Finally, the study of Drinnenberg et al. (2018) highlights how quickly interactions of canonical computational elements can produce complex and counterintuitive results, suggesting that studies of similar experimental precision and scale, and theoretical acumen, will be needed to understand the input-output transformations of circuits throughout the nervous system.

REFERENCES

Baylor, D.A., Fuortes, M.G., and O'Bryan, P.M. (1971). Receptive fields of cones in the retina of the turtle. J. Physiol. *214*, 265–294.

Chaya, T., Matsumoto, A., Sugita, Y., Watanabe, S., Kuwahara, R., Tachibana, M., and Furukawa, T. (2017). Versatile functional roles of horizontal cells in the retinal circuit. Sci. Rep. *7*, 5540.

Deverman, B.E., Pravdo, P.L., Simpson, B.P., Kumar, S.R., Chan, K.Y., Banerjee, A., Wu, W.L., Yang, B., Huber, N., Pasca, S.P., and Gradinaru, V. (2016). Cre-dependent selection yields AAV variants for widespread gene transfer to the adult brain. Nat. Biotechnol. *34*, 204–209.

Drinnenberg, A., Franke, F., Morikawa, R.K., Jüttner, J., Hillier, D., Hantz, P., Hierlemann, A., Azeredo da Silveira, R., and Roska, B. (2018). How diverse retinal functions arise from feedback at the first visual synapse. Neuron *99*, this issue, 117–134.

Euler, T., Haverkamp, S., Schubert, T., and Baden, T. (2014). Retinal bipolar cells: elementary building blocks of vision. Nat. Rev. Neurosci. *15*, 507–519.

Helmstaedter, M., Briggman, K.L., Turaga, S.C., Jain, V., Seung, H.S., and Denk, W. (2013). Connectomic reconstruction of the inner plexiform layer in the mouse retina. Nature *500*, 168–174.

Kemmler, R., Schultz, K., Dedek, K., Euler, T., and Schubert, T. (2014). Differential regulation of cone calcium signals by different horizontal cell feedback mechanisms in the mouse retina. J. Neurosci. *34*, 11826–11843.

Magnus, C.J., Lee, P.H., Atasoy, D., Su, H.H., Looger, L.L., and Sternson, S.M. (2011). Chemical and genetic engineering of selective ion channel-ligand interactions. Science *333*, 1292–1296.

Ströh, S., Puller, C., Swirski, S., Hölzel, M.B., van der Linde, L.I.S., Segelken, J., Schultz, K., Block, C., Monyer, H., Willecke, K., et al. (2018). Eliminating glutamatergic input onto horizontal cells changes the dynamic range and receptive field organization of mouse retinal ganglion cells. J. Neurosci. *38*, 2015–2028.

Thoreson, W.B., and Mangel, S.C. (2012). Lateral interactions in the outer retina. Prog. Retin. Eye Res. *31*, 407–441.

Are We There Yet? Identification of Reward-Selective Cells in the Hippocampus

Marielena Sosa[1] **and Loren M. Frank**[1,2,*]
[1]Neuroscience Graduate Program, Kavli Institute for Fundamental Neuroscience and Department of Physiology, University of California, San Francisco, San Francisco, CA 94158, USA
[2]Howard Hughes Medical Institute, San Francisco, CA, USA
*Correspondence: loren@phy.ucsf.edu
https://doi.org/10.1016/j.neuron.2018.06.037

Navigation to a previously visited reward site requires a reliable and accurate spatial memory. In this issue of *Neuron*, Gauthier and Tank (2018) use two-photon calcium imaging to uncover a discrete hippocampal subpopulation specialized for encoding reward location.

An animal's ability to navigate back to a previously discovered food or water source is critical to its survival. In order to locate such a reward successfully, the memory of the reward location must be precise, reliable, and updatable, particularly when the reward location or the surrounding environment changes. Furthermore, this memory must include detailed

spatial information if there are no reward-specific landmarks or cues. Remarkably, the neural underpinnings of this representation have remained elusive.

The hippocampus has long been known to be critical for both memory formation and spatial navigation. Neurons in the hippocampus known as place cells represent specific locations in space via their "place fields," regions of space where each neuron is active. Place cells often change their firing rate and preferred location when the environment changes, a phenomenon known as remapping (Colgin et al., 2008).

Remapping is useful for creating new representations that are different from previously stored ones, but if all cells remap, nothing would be preserved in the hippocampus from one experience to the next. Alternatively, there might be specialized populations of neurons that specifically encode important features such as reward. As the animal gets close to a reward site, this subpopulation could activate to alert the animal to the reward's proximity (Burgess and O'Keefe, 1996). While the utility of such a population is clear, prior to Gauthier and Tank's elegant study (this issue of *Neuron*), no such cells had been identified in the hippocampus.

Previous work had established that receipt of reward influences hippocampal activity on the population level. Place fields of hippocampal neurons tend to cluster around reward sites after learning, particularly in the absence of reward-site-specific cues (Dupret et al., 2010; Poucet and Hok, 2017). In addition, place cells with fields generally distributed in the environment can exhibit excess firing near reward locations (Poucet and Hok, 2017). At the network level, receipt of reward increases the occurrence of awake sharp-wave ripple (SWR) events, which coincide with bursts of hippocampal place cell sequences that recapitulate past or precede future trajectories (Singer and Frank, 2009). These reward-associated SWRs tend to engage reward-associated place cells in a manner correlated with behavioral performance (Dupret et al., 2010) and in general are more likely to recruit accurate place cell sequences (Singer and Frank, 2009). As a result, this amplification of SWR activity could support the association of reward locations with the spatial paths that lead to them.

Despite the evidence for reward-related modification of hippocampal patterns, the presence of a precise, reliable, and updatable code for reward location had not been established on the single-cell level. First, most previous studies introduced confounding sensory cues or predictable motor patterns at reward locations, such that the neural representation of the reward and the reward-associated cue or behavior could not be dissociated. Second, even with well-controlled behavior, it remained unclear whether some individual hippocampal neurons were specialized to fire consistently near reward locations, as opposed to any location where the animal stopped moving (Kay et al., 2016). Third, no neurons had been reported that could systematically remap across environments to maintain their reward signaling. Such regularity would actually be unexpected given previous evidence that the majority of hippocampal neurons seem to remap randomly between distinct environments (Colgin et al., 2008). Lastly, previous studies using extracellular recording techniques sampled relatively small populations of cells. This means that reward-specific cells might have been overlooked if they comprise a small fraction of the hippocampal population.

In this issue of *Neuron*, Gauthier and Tank (2018) used two-photon calcium imaging to reveal a small but remarkably reliable population of hippocampal neurons that code for reward location across contingencies and environments. They employed transgenic mice expressing GCaMP3 to track pyramidal cells in the output regions of the hippocampus, CA1 and subiculum. Mice were trained to run on a ball in a virtual reality linear track environment, which contained a continuously repeating set of distinct visual "wall" cues. To ask whether hippocampal neurons track rewards as they change location within the same environment, the authors provided a water reward to the mouse in two places relative to the wall cues: either toward the end of the track, or in the middle. The reward location was switched back and forth across blocks of trials. Critically, no additional cues were given at the reward location, and reward was delivered from a fixed water spout that was always in front of the animal's mouth.

The authors found that while most hippocampal place fields were uniformly distributed around the track, there was an excess density of fields around the reward location. This excess density could be explained by a reward-associated group of neurons that maintained their firing near the reward site across the contingency switch. Of the cells that continued to fire when the contingency changed, non-reward place cells either remained stable or remapped to random locations on the track. In contrast, the reward-associated cells (which we will refer to as "reward cells" for this Preview) reliably remapped to the new reward location on each switch. This remapping was rapid, occurring over the course of only a few traversals. These results indicate that the previously reported clustering of place fields near reward sites (Dupret et al., 2010) reflects the consistent activity of this small subset of reward cells. Moreover, this subset is specialized to remain anchored to reward locations in the same environment even when they change.

It remained unclear, however, whether reward cells are a distinct cell class that is specialized for reward sites in general. One possibility is that in any given environment, a random subset of cells gets selected from the network to be reward cells, and then the entire population remaps in a new environment, yielding a new subset. Alternatively, the same subpopulation of reward cells could be maintained, which might better support the transfer of learned reward search strategies across environments. To dissociate these two possibilities, the authors introduced a second, shorter, and visually distinct environment in which the reward was delivered similarly near the end of the track. Following this switch, most place cells with non-reward fields again remapped to random locations on the track, while the reward cells immediately remapped to the new reward location. Strikingly, all cells active on both environments maintained their identity, with no place cells becoming reward cells or vice versa. This strict categorical boundary between representation types has not been previously reported in the hippocampus. Overall, these results suggest that reward cells comprise a unique hippocampal cell class dedicated to representing reward location in any environment.

Gauthier and Tank (2018) also demonstrated a clear dissociation between reward cell activity and behaviors involved in reward approach and consumption. A series of controls allowed them to show that while reward cell activity preceded and was correlated with anticipatory behaviors (e.g., stopping and licking), it could not be explained by the behaviors themselves. Specifically, they found that about a third of reward cells became active prior to the reward site, while a third were activated consistently after (the remaining third changed their timing across contingencies). Those that were active prior they termed "reward-predictive." These reward-predictive cells were strongly correlated with slowing on approach to reward but did not fire above baseline levels when the mouse randomly stopped and "rested" at other non-reward locations. They were also significantly less active on error trials, when the mouse stopped at the previous contingency location, than on correct trials, despite comparable movement speeds. Since not all slowing-correlated cells were reward cells, this suggests that the reward cells are a distinct cell class from previously described hippocampal neurons active during immobility in general (Kay et al., 2016).

The current study by Gauthier and Tank (2018) fills a critical gap in our understanding of how the hippocampus encodes goal information. They have demonstrated that a unique cell class quickly, reliably, and flexibly represents reward location in a manner that predicts reward-seeking behaviors. This cell class is small, comprised of roughly 5% of CA1 or subiculum cells active in both contingencies or environments, and less than 1% of all cells recorded. While this is a substantial fraction of cells active at any given time, it is small enough to have been overlooked in previous studies due to sample size or the nature of the task.

What other features set this reward-specific cell class apart? Notably, in order to simultaneously target CA1 and subiculum, the study imaged a more posterior and intermediate region of the hippocampus, further ventral on the hippocampal dorsoventral axis than most electrophysiology studies. Increasing evidence points to a gradient of molecular expression, anatomical connectivity, and functional differences along the dorsoventral axis. In particular, reward-site-associated firing has been observed at more ventral sites, and intermediate hippocampus may be poised to integrate the spatial selectivity of dorsal hippocampus with the enhanced limbic connectivity and reward representation of ventral hippocampus (Strange et al., 2014).

In addition, Gauthier and Tank (2018) note that they largely imaged the most dorsal sublayer of CA1, which corresponds to the deepest part of stratum pyramidale. Cells in this deep sublayer have previously been shown to overrepresent reward locations and remap most stably during goal-driven learning, whereas the more superficial or ventral sublayer consists mostly of stable place cells encoding other regions of space (Danielson et al., 2016). It is possible that with further advances in optical imaging, more sublayer and intra-subregion differences will emerge as we are better able to dissect subpopulations of neighboring hippocampal neurons.

An intriguing possibility is that the reward cells described in Gauthier and Tank (2018) are preferentially connected with other brain areas. The authors discuss this issue with respect to unique outputs, such as to the nucleus accumbens to broadcast reward prediction information, or to the orbitofrontal cortex to broadcast reward confidence. Additionally, hippocampal reward cells may be innervated more strongly by certain inputs, such as dopaminergic or noradrenergic inputs from the ventral tegmental area or locus coeruleus. Perhaps the cells themselves are enriched in neuromodulatory receptors. Given the increasing availability of mouse lines and viruses engineered to tag cell subtypes, there is a unique opportunity to study how the diversity of connections to and from the hippocampus may affect physiological function.

The current work also inspires questions about the nature of the reward signal and how it might be used to guide behavior. First, if hippocampal reward cells reflect some quality of the reward itself, it would be interesting to know how manipulations of reward size or value might affect this subpopulation. Would the activity level change within the subpopulation, or would additional place cells get recruited to represent reward sites? Second, as Gauthier and Tank point out, it will be important to understand whether different types of rewards are represented similarly, and whether this type of reward representation is specific to certain types of tasks. The reward in their task is unmarked by visual cues, whereas a cued reward might not engage a hippocampal reward representation (Dupret et al., 2010). In addition, the study's well-controlled virtual reality task produces an extremely stereotyped approach behavior in which the animal's path cannot deviate. This might generate a highly regular reward prediction signal that depends primarily on distal spatial cues and distance traveled. By contrast, tasks in which the animal engages in flexible approach to reward might generate lower confidence prediction signals due to higher variability in sensory input. More flexible tasks may either mask the activity pattern of this small population or rely less on a hippocampal reward signal.

Finally, could the hippocampal network use reward cells to guide behavior outside of the immediate reward location? The authors demonstrated that reward cells are simultaneously active with place cells and fire in sequences, suggesting that reward cells may be incorporated into memory traces of entire trajectories. Reward cells therefore could be preferentially reactivated at select times when the hippocampal network represents locations far from the animal's location, as during awake replay events and "theta sequences" (Wikenheiser and Redish, 2015). As the authors suggest, this activation could be coordinated with representations of reward elsewhere in the brain to support trajectory planning or consolidation of rewarded paths.

REFERENCES

Burgess, N., and O'Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. Hippocampus *6*, 749–762.

Colgin, L.L., Moser, E.I., and Moser, M.B. (2008). Understanding memory through hippocampal remapping. Trends Neurosci. *31*, 469–477.

Danielson, N.B., Zaremba, J.D., Kaifosh, P., Bowler, J., Ladow, M., and Losonczy, A. (2016). Sublayer-specific coding dynamics during spatial navigation and learning in hippocampal area CA1. Neuron *91*, 652–665.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial

memory performance. Nat. Neurosci. *13*, 995–1002.

Gauthier, J.L., and Tank, D.W. (2018). A dedicated population for reward coding in the hippocampus. Neuron *99*, this issue, 179–193.

Kay, K., Sosa, M., Chung, J.E., Karlsson, M.P., Larkin, M.C., and Frank, L.M. (2016). A hippocampal network for spatial coding during immobility and sleep. Nature *531*, 185–190.

Poucet, B., and Hok, V. (2017). Remembering goal locations. Curr. Opin. Behav. Sci. *17*, 51–56.

Singer, A.C., and Frank, L.M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. Neuron *64*, 910–921.

Strange, B.A., Witter, M.P., Lein, E.S., and Moser, E.I. (2014). Functional organization of the hippocampal longitudinal axis. Nat. Rev. Neurosci. *15*, 655–669.

Wikenheiser, A.M., and Redish, A.D. (2015). Decoding the cognitive map: ensemble hippocampal sequences and decision making. Curr. Opin. Neurobiol. *32*, 8–15.