

Hippocampal theta sequences reflect current goals

Andrew M Wikenheiser¹ & A David Redish²

Hippocampal information processing is discretized by oscillations, and the ensemble activity of place cells is organized into temporal sequences bounded by theta cycles. Theta sequences represent time-compressed trajectories through space. Their forward-directed nature makes them an intuitive candidate mechanism for planning future trajectories, but their connection to goal-directed behavior remains unclear. As rats performed a value-guided decision-making task, the extent to which theta sequences projected ahead of the animal's current location varied on a moment-by-moment basis depending on the rat's goals. Look-ahead extended farther on journeys to distant goals than on journeys to more proximal goals and was predictive of the animal's destination. On arrival at goals, however, look-ahead was similar regardless of where the animal began its journey from. Together, these results provide evidence that hippocampal theta sequences contain information related to goals or intentions, pointing toward a potential spatial basis for planning.

In humans, the hippocampus is critical for the simulation of potential or imagined future possibilities, flexibly using past experiences to make predictions about the future^{1–5}. As such, the hippocampus has an important role in goal-directed decision-making^{6–9}. The rodent hippocampus, however, is most often associated with spatial processing^{10,11}. Performance of many spatial memory tasks depends on hippocampal function, and the location-specific firing of hippocampal place cells suggests that space may be a primary organizing principle around which hippocampal representations are constructed.

If the hippocampus mediates similar cognitive functions across species in the performance of goal-directed behavior, planning-related signals in the rodent hippocampus might occur in the framework of spatial representations. Hippocampal theta sequences, time-compressed, ensemble representations of trajectories through the environment^{12–16}, are a promising candidate neural mechanism for prospective planning². As implied by hippocampal phase precession^{17,18}, the average theta sequence representation begins slightly behind the rat's current position in space and projects forward a small distance beyond the rat^{12,15,16}. However, a previous study¹⁹ examined individual theta sequences in detail and found considerable variability in trajectory representations, with some sequences being confined to a narrow region around the rat's actual position and others extending beyond the rat's current position in a manner reminiscent of previously described forward-directed CA3 ensemble representations²⁰.

Although such modulation of theta sequence content suggests a means by which the hippocampus could use a spatial framework to support goal-directed decision-making, an alternative possibility is that forward-shifted theta representations reflect a sensory-cued recall process, where landmark-place associations drive predictive representations of upcoming locations on the maze¹⁹. If theta sequences contain computations of prospective plans, ensemble spiking should reflect currently active spatial goals. To test this idea, we examined

theta sequences as rats performed a value-guided decision-making task. Because behavior on the task was guided by rats' preferences rather than their attempts to determine and match reward contingencies set by the experimenter, this task offered a unique opportunity for testing how the hippocampus contributes to volitional navigation decisions akin to those made in natural settings.

RESULTS

We trained rats to perform a foraging task in which they chose whether or not to wait for food delivered after varying amounts of delay. Rats ran unidirectional laps around a circular track with three food pellet dispensers spaced evenly around the perimeter, each associated with a fixed-length delay. If subjects remained at a feeder site until the delay period passed, food pellets (the same type and quantity at each feeder) were dispensed. However, the rat was also free to move on to the next site. In either case, the feeder site became inactive until the rat approached it on the subsequent lap. In a session, the delay at each site remained fixed; however, across sessions, different sets of three delays were counterbalanced across the sites. A full account of behavior on this task has been presented previously²¹.

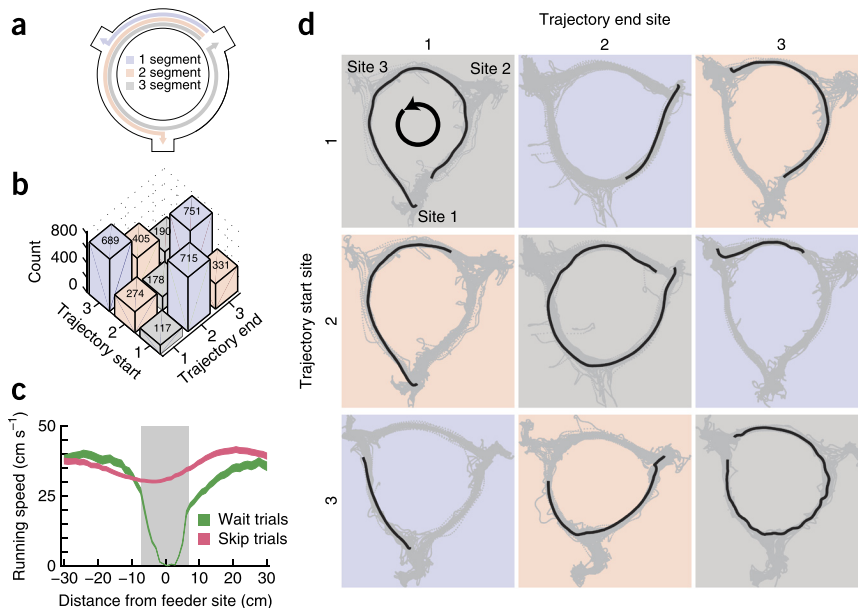
Rats ran different patterns of trajectories between sites, depending on the spatial arrangement of delays in the session and their willingness to wait for delayed reward (**Fig. 1**). Subjects could perform any of three types of trajectories: running to and stopping at the next site completed a one-segment trajectory, skipping the next site and running to the subsequent feeder constituted a two-segment trajectory, and skipping the next two feeder sites and making a complete lap around the track produced a three-segment trajectory.

Rats were implanted with tetrode arrays targeting dorsal CA1 hippocampus. Consistent with previous reports^{14,15,19}, place cell spiking was organized into theta sequences (**Fig. 2** and **Supplementary Fig. 1**). We observed that some theta sequences were related to rats' upcoming decisions on the task. For instance, **Figure 2a** shows two

¹Graduate Program in Neuroscience, University of Minnesota, Minneapolis, Minnesota, USA. ²Department of Neuroscience, University of Minnesota, Minneapolis, Minnesota, USA. Correspondence should be addressed to A.D.R. (redish@umn.edu).

Received 12 August 2014; accepted 25 November 2014; published online 5 January 2015; doi:10.1038/nn.3909

Figure 1 Behavioral task. (a) Rats allocated their time between three food delivery sites, each with unique delays that remained fixed within session, but varied across sessions. Rats ran unidirectional laps; thus, from any feeder site, subjects could run a one-segment trajectory, a two-segment trajectory or a three-segment trajectory. (b) The histogram indicates the frequency of trajectories beginning and ending at each physical feeder location (across all rats and sessions) and of trajectories of each type (color). (c) Rats' average running speed (\pm s.e.m.) on wait trials (in which the rat stopped at a feeder site for food delivery) and skip trials (in which rats bypassed a site in lieu of a different location) are plotted for the regions of space around the food delivery site. The shaded region denotes the feeder trigger zone (Online Methods). (d) Examples of trajectories beginning and ending at each feeder site. Tracking data for the entire session is plotted in gray; data for a single trajectory in each square is plotted in black.



sequences from different trials during the same recording session. In one trial, the rat ran to and stopped at the upper left feeder, and the theta sequence that occurred represented a trajectory up to, but not beyond, that site. Later in the session, the rat skipped the upper left site, and the theta sequence traced a trajectory past the site. In **Figure 2b**, theta sequences recorded at different positions along six

trajectories, each beginning at the upper right feeder and ending at the bottom center feeder (that is, two-segment trajectories skipping the upper left site) are shown. Many of these sequences contained spikes from cells ahead of the rat's location, near its goal destination.

To test whether theta sequences consistently related to upcoming behavior, we measured theta sequence look-ahead (the distance that theta sequences extended forward beyond the rat's current location¹⁹) as subjects performed the task. Data were divided into one-, two- and three-segment trajectory cases, and theta look-ahead was examined as rats traversed the first portion of each trajectory (**Fig. 3a**), where behavior was consistent, but rats' intended destination varied. Look-ahead distance was longest for three-segment trajectories and shortest for one-segment trajectories (**Fig. 3b–d**). We performed a one-way ANOVA to test whether goal destination influenced theta look-ahead and found a significant effect of trajectory type on look-ahead distance ($F_{2,20268} = 172.09$, $P < 0.001$). *Post hoc* comparison indicated that look-ahead was greatest for three-segment trajectories, intermediate for two-segment trajectories and shortest for one-segment trajectories ($P < 0.001$ for all pairwise comparisons; Tukey's HSD test). Theta sequence look-ahead showed a different pattern when aligned to arrival at a goal (**Fig. 4**). As rats approached their goal destination at the end of a trajectory, sequence look-ahead was not modulated by trajectory type ($F_{2,20789} = 0.56$, $P = 0.57$, one-way ANOVA). Thus, theta look-ahead depended on how far rats were from their goal destination, but not on how far they had traveled to arrive at their current location.

The rank ordering of mean theta look-ahead distance of trajectory types was consistent across rats (**Supplementary Fig. 2a**), a pattern that is unlikely to have arisen by chance ($P = 0.0046$, Online Methods).

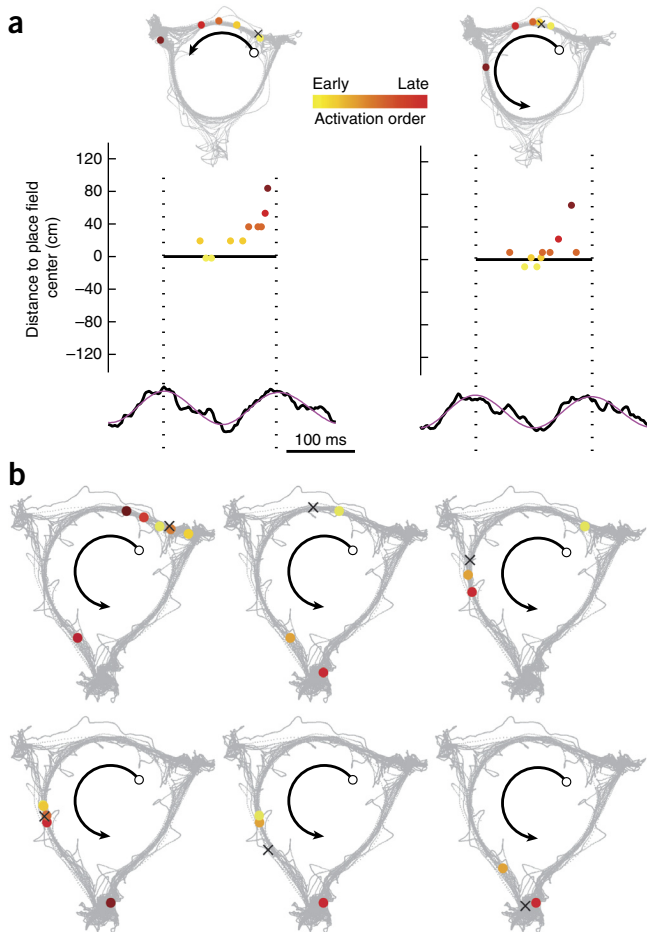


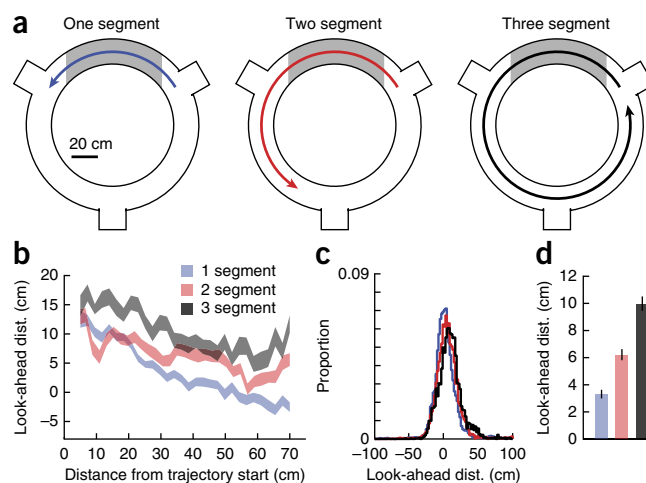
Figure 2 Theta sequences reflect future choices. (a) When the rat was planning to stop and wait for food delivery at the upper left feeder (left) a theta sequence represented a trajectory up to that feeder. Later in the session, when the rat was about to skip that feeder (right), the representation instead traced a trajectory past the feeder. (b) Each box displays the spatial sequence represented in single theta cycles recorded during six separate trajectories between the upper right and lower center feeders. Place cells near the goal destination were frequently active along with place cells near the rat's actual position. The black 'x' in spatial plots indicates the location of the rat.

Figure 3 Look-ahead distance varied with the length of planned trajectories. (a) Data were aligned to trajectory initiation, divided by how far the rat would run and examined over the initial limb of each trajectory (shaded region). (b–d) For each trajectory type, plots display the mean look-ahead distance across the initial trajectory segment (\pm s.e.m.; $n = 20,271$ theta cycles; b), distributions of look-ahead distance (c) and 95% confidence intervals (d).

Theta look-ahead distance did not depend on the physical portion of the track that trajectories occurred on (Supplementary Fig. 2b) or the threshold used to detect trajectories between feeders (Supplementary Fig. 2c). Alternate approaches to computing look-ahead distance produced similar results (Supplementary Fig. 3), and the firing rates of individual place cells did not vary substantially across trajectory types (Supplementary Fig. 4). Theta period was relatively constant for different look-ahead distances (Supplementary Fig. 5a), as was the asymmetry of the theta oscillation²² (Supplementary Fig. 5b) and the concentration of ensemble spiking in theta cycles (Supplementary Fig. 5c).

Rats' running speed and acceleration profiles differed slightly across trajectory types (speed: $F_{2,20268} = 310.69$, $P < 0.001$; acceleration: $F_{2,20268} = 62.73$, $P < 0.001$; Supplementary Fig. 6a,b). To ensure that these variables could not explain goal-dependent differences in look-ahead distance, we used a bootstrapping procedure to generate surrogate data sets assuming that running speed, acceleration or particular combinations of running speed and acceleration determined theta look-ahead¹⁹. None of these model data sets (Supplementary Fig. 6c–e) showed a significant effect of goal location on look-ahead distance (speed model: $F_{2,20268} = 0.67$, $P = 0.51$; acceleration model: $F_{2,20268} = 0.28$, $P = 0.75$; speed \times acceleration model: $F_{2,20268} = 0.19$, $P = 0.83$; one-way ANOVA). These results suggest that differences in theta sequence look-ahead were driven by representations related to rats' intended goal destination, rather than by sensory or motor variables.

The differences in theta look-ahead observed across trajectory types suggest that the hippocampus represents spatial locations up to rats' goal destinations as they perform the task. We computed the proportion of theta sequence representations in each session that ended in different spatial regions of the track as animals traversed the first limb of trajectories (Supplementary Fig. 7). On initiation of one-segment trajectories, most theta sequences ended in the region between the rat's position and the next feeder site, his intended destination. During two- and three-segment trajectories, however, the proportion



of sequences ending in this region decreased, and a greater fraction were directed to portions of the maze further beyond the rat's actual location. Over the final portion of each trajectory, as rats arrived at goal locations, there was no trajectory-dependent influence on the proportion of theta sequences that ended beyond the animal's target feeder site, suggesting that theta sequences represented paths that tended not to cross animals' goal destinations.

We next considered whether theta sequence look-ahead modulation carried sufficient information about trajectory type to allow prediction of rats' spatial goals. We used a multinomial logistic regression (MNL) approach to predict trajectory type from theta look-ahead distance. A sample of theta look-ahead values was drawn randomly and without replacement from the full population of theta cycles recorded over all animals and recording sessions. Half of the sample was used to train the regression model, whereas the other half of the sample was used to test the model (that is, predict trajectory type from look-ahead distance). This sampling and testing procedure was repeated 100 times, and for each repetition we also trained a model with shuffled data, allowing us to examine the predictive power each run of the model achieved by chance alone. Receiver operating characteristic (ROC) curves were constructed for each trajectory type on each run of the model, and the area under the ROC curves (AUC) was measured for actual and shuffled data sets.

Theta sequence look-ahead values recorded as subjects departed on trajectories predicted goal location at above-chance levels ($z_{299} = 17.03$, $P < 0.001$; sign test; Fig. 5a). On arrival at goal destinations, however, the area under ROC curves resulting from shuffled and actual data sets were not statistically different ($z_{299} = 0.28$, $P = 0.77$; sign test; Fig. 5b). As an additional control, we tested a model trained with running speed and acceleration as predictors of trajectory type. The prediction achieved by this model did not exceed chance levels ($z_{299} = 0.64$, $P = 0.53$; sign test; Fig. 5c). These analyses suggest that hippocampal representations encoded information relevant to currently active goals in theta cycles beyond any information incidental to speed or acceleration differences.

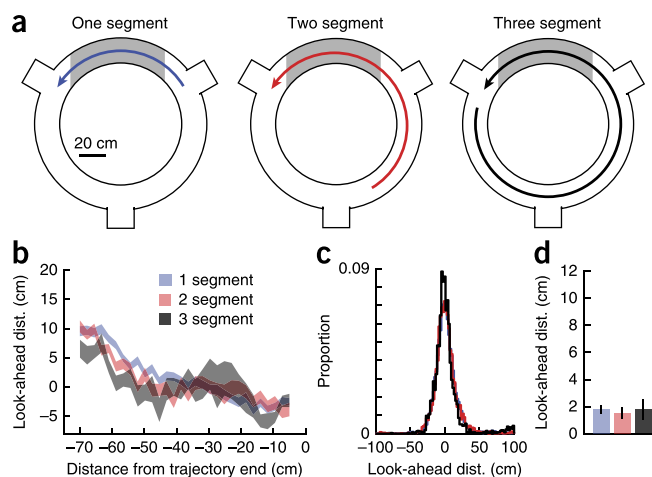
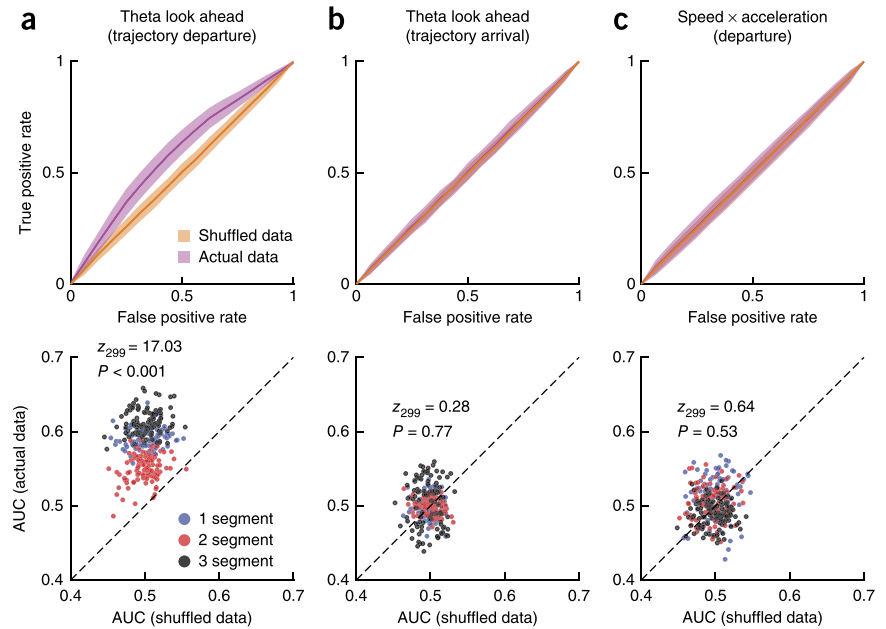


Figure 4 Look-ahead was constant on goal arrival. (a) Data were aligned to trajectory completion as rats arrived at their goal destination and examined over the final limb of each trajectory (shaded region). (b) Look-ahead distance did not differ over space. The shading indicates s.e.m. ($n = 20,792$ theta cycles). (c) Distributions of theta look-ahead distance were overlapping across trajectories. (d) Mean look-ahead distance (shown with 95% confidence intervals) did not differ across trajectories.

Figure 5 Theta look-ahead predicted rats' current goals. ROC curves (upper row) and AUC plots (lower row) were computed for regression models that estimated trajectory type from theta look-ahead distance or running speed and acceleration. Curves were computed for each of the three trajectory types using the output of models trained with actual data and shuffled data. ROC curves were averaged across trajectory types (shading indicates s.e.m.), and the area under ROC curves resulting from each run of the model was plotted separately for the three trajectory types. **(a)** On initiation of trajectories, prediction of trajectory type by theta look-ahead distance exceeded chance levels. **(b)** As rats arrived at goal destinations, however, theta look-ahead was no longer more predictive of trajectory type than shuffled data. **(c)** Classification of trajectory type using speed and acceleration as predictors was not better than chance.



On a single-cell level, increased theta look-ahead distance implies an earlier activation of neurons as subjects approach place fields, and consequently predicts that place fields located solely on longer trajectories should be larger than those on short trajectories. To test this prediction, we identified place fields situated such that rats passed through them almost exclusively during completion of only a single type of trajectory ($n = 258$ fields), and compared their sizes (**Fig. 6a**). A one-way ANOVA identified a significant effect of trajectory type on place field size ($F_{2,255} = 5.53, P = 0.0044$), and *post hoc* testing revealed that place fields on three-segment trajectories were significantly larger (on average ~20%, or 10 cm longer) than those on one-segment trajectories ($p_{1,2} = 0.11, p_{1,3} = 0.003, p_{2,3} = 0.25$; Tukey's HSD test).

Trajectory-dependent modulation of look-ahead distance also predicts that the place fields that rats traverse en route to multiple, different goal sites will vary in size from trial to trial, depending on the rats' current spatial goal destination on individual passes through the field. To quantify place field size on a trial-by-trial basis, we measured the distance between each place field's center and the rat's position when the first and last spike occurred for each pass through the field (**Fig. 6b**). Consistent with goal modulation of look-ahead distance, the location of the first place cell spike on each pass varied significantly depending on the rat's intended destination ($F_{2,10656} = 122.02, P < 0.001$, one-way ANOVA), with Tukey's HSD *post hoc* test showing that the initial spike on three-segment trajectory passes was shifted significantly toward the early portion of the place field relative to the initial spike on one-segment passes ($p_{1,2} < 0.001, p_{1,3} < 0.001, p_{2,3} = 0.02$). Notably, this result held when only mixed-case place fields (fields that rats traversed while completing more than one trajectory

type; $n = 104$ fields) were compared ($F_{2,2886} = 6.99, P < 0.001$, one-way ANOVA), strongly suggesting that the trajectory dependence of look-ahead distance that we observed is a result of modulation of place cell activation on single trials, and not of differences in the place fields associated with particular trajectory types. Although the position of the last spike on passes through place fields also varied with goal destination ($F_{2,10656} = 3.49, P = 0.03$, one-way ANOVA), the difference across trajectory types was smaller than the difference in initial spike location (**Fig. 6b**), suggesting that increased look-ahead distance primarily affected the early portion of place fields, leaving the ends of place fields largely intact^{23,24}.

Previous work has shown that many place cells exhibit asymmetric expansion over time^{25–27}, with place field center of mass (COM) shifting in a direction opposite of the animal's movement, resulting in a negative relationship between place field COM and the number of passes through the place field. To examine how asymmetric expansion interacted with look-ahead modulation, we computed the COM of place cell spiking for each pass through a place field and fit lines to the relationship between lap number and COM, separately for each trajectory type (**Fig. 7a**). Asymmetric expansion of place fields predicts that such lines would have a characteristic slope that is unaffected by current goals, whereas trial-to-trial look-ahead modulation predicts that line intercepts should vary with goal destination. A one-way ANOVA revealed no effect of trajectory type on line slope ($F_{2,488} = 0.10, P = 0.91$; **Fig. 7b**), but showed that the intercepts of lines fit for each goal destination varied significantly ($F_{2,488} = 18.72, P < 0.001$). *Post hoc* testing with Tukey's HSD test revealed that intercepts for three-segment trajectories were shifted significantly further toward the initial portion of the place field than one-segment trajectories

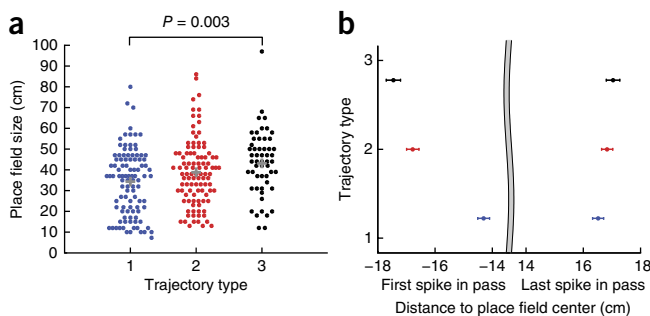


Figure 6 Place field size was modulated by trajectory type. **(a)** Consistent with trajectory-dependent modulation of theta look-ahead, place fields that rats traversed on three-segment trajectories were significantly larger than fields that rats passed through on one-segment trajectories. **(b)** Place field size varied similarly with goal location from trial to trial. The first spike of passes through place fields during three-segment trajectories occurred further in front of the place field center than on one-segment trajectory passes. The location of the last spike also varied, but to a lesser extent, across trajectory types. Error bars indicate s.e.m.

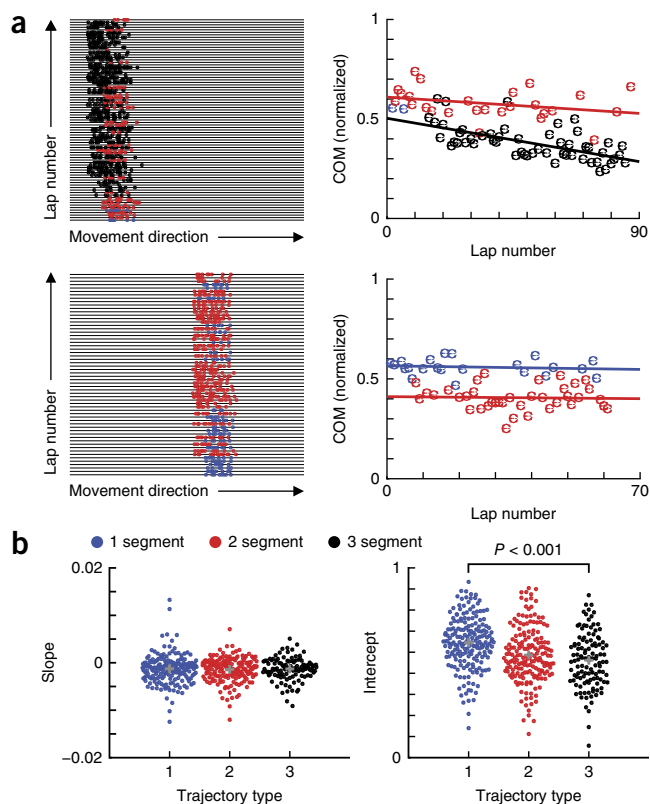


Figure 7 Look-ahead was modulated by goal location on single trials. (a) Scatter plots (left) show the spiking of two place cells that rats passed through on trajectories of different lengths throughout the course of the session. Circles mark the location of the rat when spikes occurred, and color indicates the trajectory type the rat was completing. Place field COM (normalized to span from 0 at the beginning of the field to 1 at the end of the field) is plotted for each lap. (b) We fit lines to the COM by lap relationship as plotted in a, separately for each trajectory type. The slopes of these lines were not modulated by trajectory type, whereas line intercepts were significantly forward-shifted for longer, three-segment trajectories, consistent with trial-by-trial modulation of look-ahead distance. Gray crosses indicate the means of distributions.

($p_{1,2} = 0.15$, $p_{1,3} < 0.001$, $p_{2,3} = 0.23$). These results held even when only mixed-case place fields (fields subjects passed through en route to more than one goal site) were compared (slope: $F_{2,267} = 0.10$, $P = 0.90$; intercept: $F_{2,267} = 7.23$, $P < 0.001$; one-way ANOVA), suggesting that look-ahead modulation occurred on a single-trial basis, in tandem with, but separable from, the place field expansion phenomena.

DISCUSSION

The look-ahead distance of hippocampal theta sequences increased when rats executed longer spatial trajectories. Consequently, in addition to representing animals' location as they performed the task, CA1 theta sequences also included information about current spatial goals. Beyond the sensory features of landmarks, covert cognitive factors such as the motivational importance that subjects attached to regions of space (that is, whether or not they would skip a feeder site) influenced the expression of theta sequences.

These findings strongly suggest that CA1 theta sequence modulation is unlikely to be a purely sensory-driven recall process. Sensory properties common to all feeder locations (for example, olfactory food cues) would have exerted a similar influence on theta sequences

regardless of rats' intentions to skip sites or wait for food delivery. We observed, however, that sequence look-ahead varied with rats' goals on a moment-to-moment basis in a way that reflected future choices; these data are incompatible with sequence expression being determined solely by proximity to potential food delivery sites. If sequence modulation were instead driven by the sensory features unique to each feeder location, patterns of look-ahead would have varied across feeder sites, which we did not observe (**Supplementary Fig. 2b**). These results extend previous work by showing that the influence of sensory and incentive properties of landmarks in the environment on theta sequence expression can be dissociable, at least under some conditions.

Previous studies have reported hippocampal representations related to goals or future behavior during quiescence, coincident with sharp-wave ripple (SWR) complexes^{28,29}, and disruption of hippocampal spiking during SWRs impairs performance of a hippocampally dependent, spatial memory task³⁰. Notably, our results here and in previous work indicate that hippocampal representations related to upcoming behaviors also occur during theta^{20,31–33}. The relationship between theta- and SWR-associated hippocampal planning signals remains untested³⁴, and an important question for future studies to address is whether and how decision-related representations during SWRs^{28–30} interact with goal-related theta sequences. That hippocampal theta sequences continue to look-ahead toward the target goal throughout a journey suggests that the rodent hippocampus may maintain plans online, as behavior is executed^{5,35}.

Consistent with a role for the hippocampus in actively coordinating ongoing behavior, we found that theta sequences reflected rats' future trajectories; it remains unclear, however, whether such representations are critical for task performance or simply a reflection of behavior. Previous studies^{36,37} used a pharmacological approach to show that disrupting the temporal organization of theta sequences (while minimally affecting other place cell properties) reversibly abolishes correct performance of a delayed-alternation T-maze task. Because rats in these studies had achieved asymptotic performance before theta sequences were manipulated, a role for theta sequences in decision-making (as opposed to facilitating or modulating learning-related processes) seems likely.

Although the functions of theta sequences and phase precession remain unclear, a growing body of evidence suggests that hippocampal ensemble spiking performs a moment-by-moment prediction of upcoming spatial positions across the theta cycle^{23,24}, coordinated by gamma frequency oscillations^{38–40}. Such representations have clear utility in decision-making situations. By passing information about imminent environmental features to other brain regions, goal-directed theta look-ahead could coordinate the computation of other representations relevant for decision-making. For instance, ramp cells in the ventral striatum (neurons whose firing rates encode proximity to reward) phase precess to hippocampal theta⁴¹, raising the possibility that they too participate in theta-paced sequential representations, imbuing the spatial paths encoded by the hippocampus with information about potential upcoming rewards⁴².

The hippocampus also interacts with the prefrontal cortex (PFC), a brain region associated with planning and other higher order cognitive functions. PFC neural activity shows consistent organization relative to hippocampal theta^{43,44}, coherence in the local field potential at theta frequency increases between these structures during decision-making⁴⁵, and theta entrainment of PFC spiking predicts accurate behavioral performance⁴⁶. Prospective hippocampal representations of upcoming spatial locations could facilitate decision-related information processing in these and other brain regions to adaptively

guide behavior^{47,48}. Pyramidal neurons in a range of brain regions are known to be exquisitely sensitive to temporally patterned dendritic inputs^{49,50}, suggesting a possible mechanism by which cells receiving hippocampal inputs could detect even subtle differences in the content of theta sequence representations.

METHODS

Methods and any associated references are available in the [online version of the paper](#).

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

The authors thank members of the Redish laboratory for discussion about this work. A.M.W. was supported by US National Institutes of Health grant T32-DA-07234 and a University of Minnesota Doctoral Dissertation Fellowship. A.D.R. was supported by US National Institutes of Health grant R01-MH-080318. This work was also supported by US National Institutes of Health grant R56-MH-080318.

AUTHOR CONTRIBUTIONS

A.M.W. and A.D.R. designed the experiments. A.M.W. collected the data. A.M.W. and A.D.R. analyzed the data. A.M.W. and A.D.R. wrote the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

- Cohen, N.J. & Eichenbaum, H. *Memory, Amnesia and the Hippocampal System* (MIT Press, 1993).
- Lisman, J. & Redish, A.D. Prediction, sequences and the hippocampus. *Phil. Trans. R. Soc. Lond. B* **364**, 1193–1201 (2009).
- Schacter, D.L. & Addis, D. On the nature of medial temporal lobe contributions to the constructive simulation of future events. *Phil. Trans. R. Soc. Lond. B* **364**, 1245 (2009).
- Barron, H.C., Dolan, R. & Behrens, T. Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat. Neurosci.* **16**, 1492–1498 (2013).
- Wikenheiser, A. & Redish, A. Decoding the cognitive map: ensemble hippocampal sequences and decision making. *Curr. Opin. Neurobiol.* published online, doi:10.1016/j.conb.2014.10.002 (17 October 2014).
- Peters, J. & Büchel, C. Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron* **66**, 138–148 (2010).
- Viard, A., Doeller, C.F., Hartley, T., Bird, C.M. & Burgess, N. Anterior hippocampus and goal-directed spatial decision making. *J. Neurosci.* **31**, 4613–4621 (2011).
- Bornstein, A.M. & Daw, N.D. Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. *PLoS Comput. Biol.* **9**, e1003387 (2013).
- Spiers, H.J. & Barry, C. Neural systems supporting navigation. *Curr. Opin. Behav. Sci.* **1**, 47–55 (2015).
- O'Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map* (Clarendon Press, Oxford, 1978).
- Redish, A.D. *Beyond the Cognitive Map: From Place Cells to Episodic Memory* (MIT Press, 1999).
- Skaggs, W.E., McNaughton, B.L., Wilson, M.A. & Barnes, C.A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–172 (1996).
- Tsodyks, M.V., Skaggs, W.E., Sejnowski, T.J. & McNaughton, B.L. Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus* **6**, 271–280 (1996).
- Dragoi, G. & Buzsáki, G. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* **50**, 145–157 (2006).
- Foster, D.J. & Wilson, M.A. Hippocampal theta sequences. *Hippocampus* **17**, 1093–1099 (2007).
- Maurer, A.P., Burke, S.N., Lipa, P., Skaggs, W.E. & Barnes, C.A. Greater running speeds result in altered hippocampal phase sequence dynamics. *Hippocampus* **22**, 737–747 (2012).
- O'Keefe, J. & Recce, M. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317–330 (1993).
- Maurer, A.P. & McNaughton, B.L. Network and intrinsic cellular mechanisms underlying theta phase precession of hippocampal neurons. *Trends Neurosci.* **30**, 325–333 (2007).
- Gupta, A.S., van der Meer, M., Touretzky, D. & Redish, A. Segmentation of spatial experience by hippocampal θ sequences. *Nat. Neurosci.* **15**, 1032–1039 (2012).
- Johnson, A. & Redish, A.D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* **27**, 12176–12189 (2007).
- Wikenheiser, A.M., Stephens, D. & Redish, A. Subjective costs drive overly patient foraging strategies in rats on an intertemporal foraging task. *Proc. Natl. Acad. Sci. USA* **110**, 8308–8313 (2013).
- Belluscio, M.A., Mizuseki, K., Schmidt, R., Kempter, R. & Buzsáki, G. Cross-frequency phase-phase coupling between theta and gamma oscillations in the hippocampus. *J. Neurosci.* **32**, 423–435 (2012).
- Jensen, O. & Lisman, J.E. Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learn. Mem.* **3**, 279–287 (1996).
- Lisman, J.E. & Jensen, O. The theta-gamma neural code. *Neuron* **77**, 1002–1016 (2013).
- Mehta, M.R. & McNaughton, B.L. Rapid changes in hippocampal population code during behavior: A case for Hebbian learning *in vivo*. *Soc. Neurosci. Abstr.* **758.3** (1996).
- Mehta, M.R., Barnes, C.A. & McNaughton, B.L. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA* **94**, 8918–8921 (1997).
- Mehta, M.R., Quirk, M.C. & Wilson, M.A. Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* **25**, 707–715 (2000).
- Pfeiffer, B.E. & Foster, D. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).
- Singer, A.C., Carr, M., Karlsson, M. & Frank, L. Hippocampal swr activity predicts correct decisions during the initial learning of an alternation task. *Neuron* **77**, 1163–1173 (2013).
- Jadhav, S.P., Kemere, C., German, P. & Frank, L. Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454–1458 (2012).
- Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).
- Takahashi, M., Lauwereyns, J., Sakurai, Y. & Tsukada, M. A code for spatial alternation during fixation in rat hippocampal CA1 neurons. *J. Neurophysiol.* **102**, 556–567 (2009).
- Takahashi, M., Nishida, H., Redish, A.D. & Lauwereyns, J. Theta phase shift in spike timing and modulation of gamma oscillation: a dynamic code for spatial alternation during fixation in rat hippocampal area ca1. *J. Neurophysiol.* **111**, 1601–1614 (2014).
- Schmidt, B. & Redish, A.D. Neuroscience: navigation with a cognitive map. *Nature* **497**, 42–43 (2013).
- Howard, L.R. *et al.* The hippocampus and entorhinal cortex encode the path and euclidean distances to goals during navigation. *Curr. Biol.* **24**, 1331–1340 (2014).
- Robbe, D. *et al.* Cannabinoids reveal importance of spike timing coordination in hippocampal function. *Nat. Neurosci.* **9**, 1526–1533 (2006).
- Robbe, D. & Buzsáki, G. Alteration of theta timescale dynamics of hippocampal place cells by a cannabinoid is associated with memory impairment. *J. Neurosci.* **29**, 12597–12605 (2009).
- Colgin, L.L. *et al.* Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**, 353–357 (2009).
- Bieri, K.W., Bobbitt, K.N. & Colgin, L.L. Slow and fast gamma rhythms coordinate different spatial coding modes in hippocampal place cells. *Neuron* **82**, 670–681 (2014).
- Cabral, H.O. *et al.* Oscillatory dynamics and place field maps reflect hippocampal ensemble processing of sequence and place memory under NMDA receptor control. *Neuron* **81**, 402–415 (2014).
- van der Meer, M.A.A. & Redish, A.D. Theta phase precession in rat ventral striatum links place and reward information. *J. Neurosci.* **31**, 2843–2854 (2011).
- Malhotra, S., Cross, R.W. & van der Meer, M.A. Theta phase precession beyond the hippocampus. *Rev. Neurosci.* **23**, 39–65 (2012).
- Siapas, A.G., Lubenov, E.V. & Wilson, M.A. Prefrontal phase locking to hippocampal theta oscillations. *Neuron* **46**, 141–151 (2005).
- Jones, M.W. & Wilson, M.A. Phase precession of medial prefrontal cortical activity relative to the hippocampal theta rhythm. *Hippocampus* **15**, 867–873 (2005).
- Benchenane, K. *et al.* Coherent theta oscillations and reorganization of spike timing in the hippocampal-prefrontal network upon learning. *Neuron* **66**, 921–936 (2010).
- Hyman, J.M., Zilli, E.A., Paley, A.M. & Hasselmo, M.E. Working memory performance correlates with prefrontal-hippocampal theta interactions but not with prefrontal neuron firing rates. *Front. Integr. Neurosci.* published online, doi:10.3389/fnint.2010.002.2010 (10 March 2010).
- Gupta, K., Erdem, U. & Hasselmo, M. Modeling of grid cell activity demonstrates *in vivo* entorhinal look-ahead properties. *Neuroscience* **247**, 395–411 (2013).
- Erdem, U.M. & Hasselmo, M.E. A biologically inspired hierarchical goal directed navigation model. *J. Physiol. Paris* **108**, 28–37 (2014).
- Mel, B.W. NMDA-based pattern discrimination in a modeled cortical neuron. *Neural Comput.* **4**, 502–517 (1992).
- Branco, T., Clark, B.A. & Häusser, M. Dendritic discrimination of temporal input sequences in cortical neurons. *Science* **329**, 1671–1675 (2010).

ONLINE METHODS

Subjects. Four male, Fisher-Brown Norway hybrid rats aged 6–14 months, (Harlan) were subjects for the experiment. A formal power analysis was not conducted, but this sample size is consistent with similar studies examining information processing in neural ensembles recorded from behaving rats. Before beginning behavioral training, rats were handled for 7 d and acclimated to eating the food pellets that would be delivered during the behavioral task (45-mg sucrose pellets, Test Diet). Rats were maintained on a 12-h light-dark cycle, and behavioral sessions occurred at the same time daily, during the light phase. Subjects were food restricted to maintain their weight at $\geq 80\%$ of their free-feeding weight; water was always available in the home cage. All experimental and animal care procedures complied with US National Institutes of Health guidelines for animal care and were approved by the Institutional Animal Care and Use Committee at the University of Minnesota.

Training and behavior. Rats performed the foraging task on an elevated, circular track (diameter = 80 cm) with three food pellet dispensers (Med Associates) positioned evenly around the perimeter. An overhead camera recorded subjects' position via a light-emitting diode affixed around the rat's body (before electrode implantation) or light-emitting diodes mounted to the headstage (after electrode implantation). Data were recorded with a Cheetah 160 acquisition system (Neuralynx). Custom Matlab software (MathWorks) controlled the task. Subjects first performed a training task in which they ran unidirectional laps around the track to earn food at each feeder. Pellets were delivered as soon as rats arrived at each feeder site. Attempts to run backwards were blocked by the experimenter during training sessions (during neural recording sessions, rats were well trained on the task and only ran forward). After rats ran 30 or more laps for three consecutive sessions, the training phase was considered complete and task performance began.

During each daily, 30-min session of the foraging task, rats could earn food pellets from the three feeder locations after a delay period. The delay began when the subject approached within 7 cm of a feeder site. Entry into this zone was signaled by a tone sequence (200-ms pulses, repeated once per second). The tone's frequency was proportional to the site's delay. Six sets of delays were used, defining six unique session types. Rats experienced session types in pseudorandom order (the same type was never repeated on consecutive days). Delays were counterbalanced across feeder sites to ensure that delay distributions at each location were equivalent across sessions. The sets of delays were: 1, 6, 12 s (session type 1); 4, 6, 12 s (session type 2); 4, 12, 20 s (session type 3); 6, 12, 20 s (session type 4); 6, 12, 26 s (session type 5); 6, 12, 32 s (session type 6). Subjects performed each session type four times.

Surgery and recordings. Following completion of the behavioral sequence (24 sessions) subjects were allowed *ad libitum* access to food for at least 24 h, and then implanted with tetrode arrays targeting the dorsal CA1 region of hippocampus in the right hemisphere (-3.8 mm anteroposterior, 2.0 mm lateral from bregma). Surgical procedures have been described elsewhere^{51,52}. Hyperdrives (Kopf) containing 12 tetrodes and two electrodes consisting of four tetrode wires electrically fused together (for recording LFPs) were used. Tetrodes were advanced slowly over approximately 1 week until estimates of electrode depth and electrophysiological signatures were consistent with the CA1 pyramidal layer. One LFP electrode was placed in the corpus callosum above hippocampus and the other was placed in the hippocampal fissure. Data were recorded by a 64 channel Neuralynx Cheetah system. The voltage on each tetrode channel was monitored at 32 kHz and filtered from 600–6,000 Hz. When the voltage on any channel exceeded a preset threshold, 1 ms of activity on each channel was saved to disk and time-stamped. Spikes were manually sorted into putative single units off-line using MClust 3.5 (ref. 53). Local field potentials were recorded from one channel per tetrode, sampled at 2 kHz and filtered from 1–425 Hz. Daily recording sessions continued for as long as large ensembles of well-isolated hippocampal neurons were recorded, with subjects cycling pseudorandomly through session types for the duration of recording sessions.

Statistics and general data analysis. All statistical tests were two-sided. Data were not formally tested for normality, but full data distributions are shown in the figures. We analyzed 1,263 cells recorded across 26 sessions. Ensemble sizes ranged from a minimum of 25 to 62 simultaneously recorded neurons.

Only well-isolated units were kept for further analysis; the median isolation distance of cells was 32.65, and the median L-ratio was 0.02 (ref. 54). Units were isolated blind to rats' behavior. All analyses were conducted using Matlab (MathWorks).

Detection of theta state. Because analyses focused on hippocampal representations during the theta network state, we were careful to exclude data from non-theta epochs (for example, during the large, irregular activity network state, when sharp-wave ripples are prominent⁵⁵). All analyses of hippocampal theta were based on the signal recorded from the hippocampal fissure, where theta amplitude is greatest. Sharp-wave/ripple analyses were based on recordings taken from the CA1 pyramidal cell layer. Recordings were first pre-processed to remove 0.5 s of data around any artifacts (instances of maximum/minimum voltage), and then bandpass filtered between 6–10 Hz to obtain the theta band signal, between 2–4 Hz to obtain the delta band signal, and between 140–220 Hz to obtain the sharp-wave/ripple band signal. Instantaneous amplitude (power) and phase were estimated via the Hilbert transform. Theta phase ranged from $-\pi$ to π , where a phase of zero corresponded to the trough of the oscillation recorded at the hippocampal fissure. Theta cycles were defined as the time between peaks of the 6–10 Hz band-pass filtered fissure LFP. The log-transformed ratio of theta to delta power was computed^{51,56}, and this quantity was averaged within each putative cycle of the theta rhythm. Only theta cycles with an average theta-delta ratio $> 1\sigma$ the session average theta-delta ratio were included for analyses. In addition, only theta cycles with a period corresponding to a 6–10-Hz oscillation were included. Spikes that occurred when ripple power was $> 4\sigma$ the session average were excluded from analyses.

Place fields. Place fields were detected as described previously^{52,57}. Units with a session mean firing rate < 0.05 Hz (non-task-responsive cells) or > 5 Hz (putative interneurons) were excluded from place field analyses. Spikes recorded when the animal's running speed was < 5 cm s^{-1} were excluded. The linearized maze was divided into approximately 2-cm bins and firing rate was computed within each spatial bin. Contiguous bins in which the firing rate was $\geq 5\%$ the cell's session maximum firing rate were considered place fields. Fields separated by ≤ 2 bins (≈ 4 cm) were merged. To avoid ambiguity about which positions neurons were representing, place cells with multiple firing fields (22 units) were excluded from further analyses.

Theta sequence look-ahead. To measure theta sequence look-ahead, ensemble place cell activity was rotated relative to the animal's current position (as in **Fig. 2** and **Supplementary Fig. 1**), and the look-ahead distance was taken as the average position of spikes that occurred in the final quarter of the theta cycle. In other words, the look-ahead distance of each theta cycle was the distance between the rat's location and the average of place field centers of cells active in the final quarter of that theta cycle, weighted by the number of spikes each cell fired. Theta cycles with no spikes in this region of the theta cycle (approximately 8% of cycles that otherwise met inclusion criteria) were excluded from analysis. Using variations on this approach (for example, taking the position of the final spike in the cycle; taking the position of the spike furthest forward of the animal, regardless of its timing within the theta cycle) produced similar estimates of look-ahead distance (**Supplementary Fig. 3**). In addition to the criteria described above for selection of theta cycles, look-ahead was computed only for theta cycles containing at least three spikes from a minimum of two cells. To compute the probability of mean theta look-ahead distance across trajectory types exhibiting the same rank ordering for all subjects by chance, we used the formula $P = (1/6)^4 * 6$, as there were six possible orderings of means for each of four rats.

Trajectories. Comparisons of look-ahead distance were made for data from journeys between feeders devoid of pauses or other behavioral irregularities. We isolated trajectories that occurred between rewarded visits to feeder sites (that is, cases where subjects left a site after receiving food and traveled to a site to await another food delivery). If running speed fell below 10 cm s^{-1} at any point after departure from the origin feeder site and before arrival at the destination feeder site the trajectory was excluded. **Figure 1** shows examples of one-, two- and three-segment trajectories initiated from each feeder site, and the relative frequencies of each trajectory type. We also tested trajectories detected in the same manner, but using a faster minimum running speed threshold of 25 cm s^{-1} (**Supplementary Fig. 2c**), and observed similar results (**Figs. 3d** and **4d**).

Movement controls. To rule out subtle variations in speed or acceleration driving the goal-dependent modulation of look-ahead distance, we used a bootstrapping analysis to construct model data sets that assume movement parameters solely determined look-ahead (**Supplementary Fig. 6**). Thus, every theta cycle used to construct **Figure 3** was assigned a surrogate look-ahead distance, drawn randomly from a distribution of all look-ahead values that occurred when the animal was running at a similar speed or accelerating at a similar rate. We also constructed a surrogate data set that resampled look-ahead values from theta cycles matched for both running speed and acceleration.

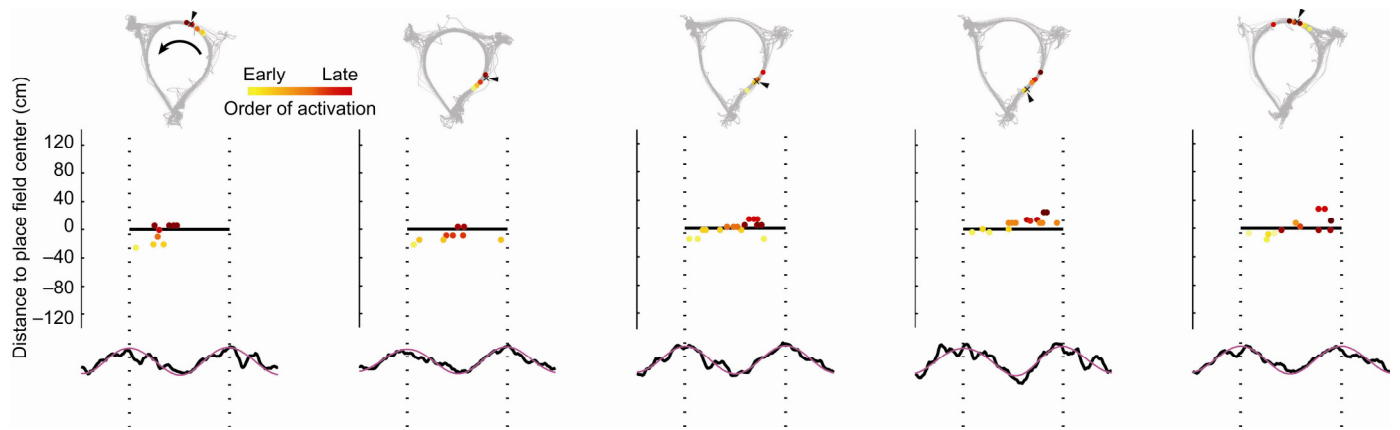
Predicting trajectory type from look-ahead distance. We used multinomial logistic regression (a generalization of binomial logistic regression to an arbitrary number of groups) to classify theta look-ahead values as occurring during one-, two- or three-segment trajectories⁵⁸. In this analysis, theta look-ahead distance was the independent variable (or predictor) and trajectory type was the dependent (outcome) variable. Theta look-ahead values (600 theta cycles each from one-, two- and three-segment trajectories) were sampled from the population of theta cycles recorded across all animals that met the inclusion criteria outlined above. One half of this sample was used to train the regression model. The remaining half of the sample was used as input to the model, which returned the probability of each sample being an instance of a one-, two- or three-segment trajectory look-ahead value. We then shuffled the trajectory type labels of the training samples and used this shuffled data set to construct a new regression model, which was tested on the same sample of look-ahead values used to test the unshuffled model. This sampling and testing procedure was repeated 100 times. In addition to using this procedure to test the predictive power of theta look-ahead values on trajectory arrival and departure, we also used this approach to test the model with both running speed and acceleration as predictors. ROC analysis was performed separately for each trajectory type, and ROC curves for trajectory types one, two, and three were averaged to produce the plots in **Figure 5**. The area under each ROC curve was computed, and values for shuffled and unshuffled data sets were plotted against one another (**Fig. 5**).

Place field size analyses. To test whether goal destination affected place field size (**Fig. 6a**), we identified fields that rats traveled through primarily on one of the three trajectory types. For every spike a place cell fired, we determined which

trajectory type the rat was completing at that moment; place cells with $\geq 90\%$ of spikes occurring on a single trajectory type were included in this analysis. To compute place field size on a trial-by-trial basis (**Fig. 6b**), we measured the rat's location when the first and last spike occurred on each pass through the place field, relative to that place field's center, and determined the trajectory type that rats were completing. Mixed trajectory fields were defined as place cells with $\leq 60\%$ of spikes accounted for by a single trajectory type, thus ensuring that a single mixed field contributed sufficient measurements to at least two trajectory types.

Place field expansion analyses. To measure backward expansion of place fields, we computed the average position of all the spikes a cell fired on a given pass through its place field. Spike positions were normalized to fractional portion of the field, such that the beginning of the field corresponded to zero and the terminus of the field corresponded to one. For the analysis of the COM–lap number relationship (**Fig. 7**), lines were fit only for trajectory types with ≥ 10 passes, to ensure reliable parameter estimates.

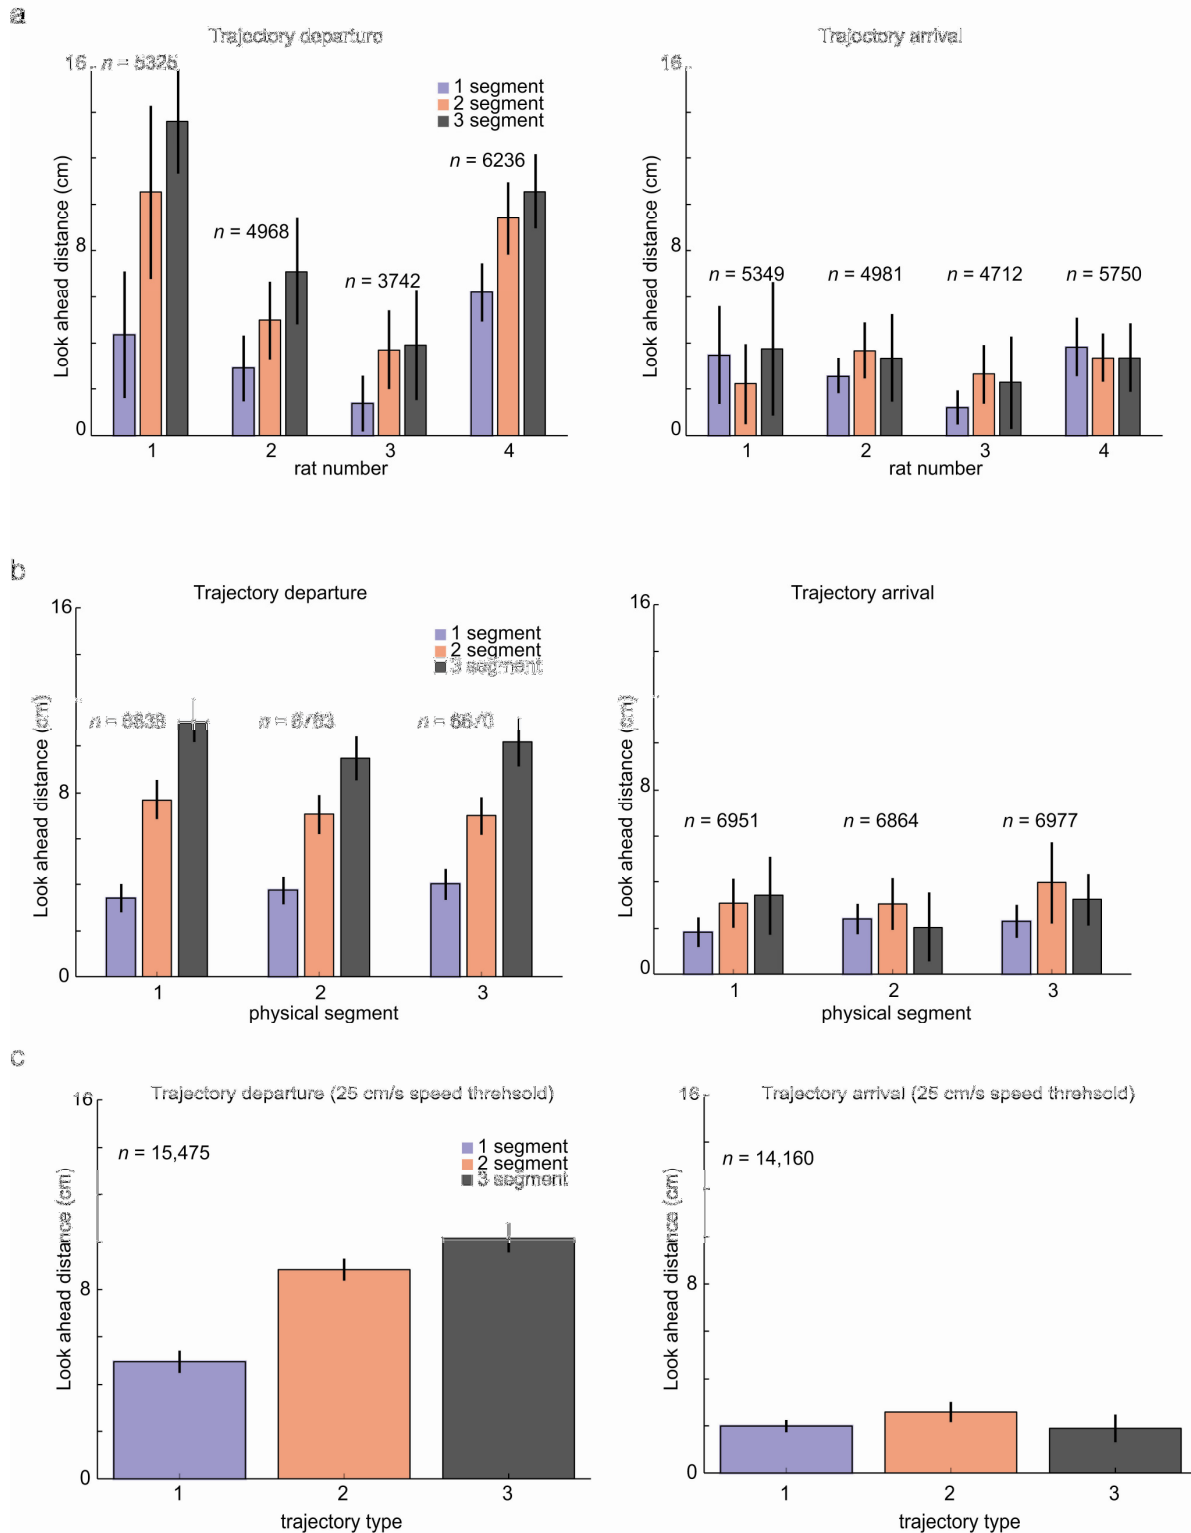
51. Jackson, J.C., Johnson, A. & Redish, A.D. Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. *J. Neurosci.* **26**, 12415–12426 (2006).
52. Wikenheiser, A.M. & Redish, A.D. Changes in reward contingency modulate the trial to trial variability of hippocampal place cells. *J. Neurophysiol.* **106**, 589–598 (2011).
53. Redish, A.D. & Schmitzer-Torbert, N.C. MCLUST spike sorting toolbox, version 3.0. <<http://redishlab.neuroscience.umn.edu/>> (2008).
54. Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K. & Redish, A. Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* **131**, 1–11 (2005).
55. Carr, M.F., Jadhav, S.P. & Frank, L.M. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* **14**, 147–153 (2011).
56. Csicsvari, J., Hirase, H., Czurkó, A. & Buzsáki, G. Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *J. Neurosci.* **19**, 274–287 (1999).
57. Gupta, A.S., van der Meer, M.A.A., Touretzky, D.S. & Redish, A.D. Hippocampal replay is not a simple function of experience. *Neuron* **65**, 695–705 (2010).
58. Hosmer, D.W. Jr., Lemeshow, S. & Sturdivant, R.X. *Applied Logistic Regression* (John Wiley & Sons, 2013).



Supplementary Figure 1

Theta sequences

Rasters of place cell spiking are plotted for five example theta cycles. Cells are sorted on the y (space) axis according to the location of their place field center relative to the rat's instantaneous location (positive numbers indicate space ahead of the rat). The LFP (unfiltered, and filtered from 6–10 Hz) is plotted below each raster. In the 2D spatial plots, dots denote the place field center of each cell that was active during the theta cycle, and a black x indicates the rat's position. In both plots, spikes are color coded to indicate the order in which cells became active.

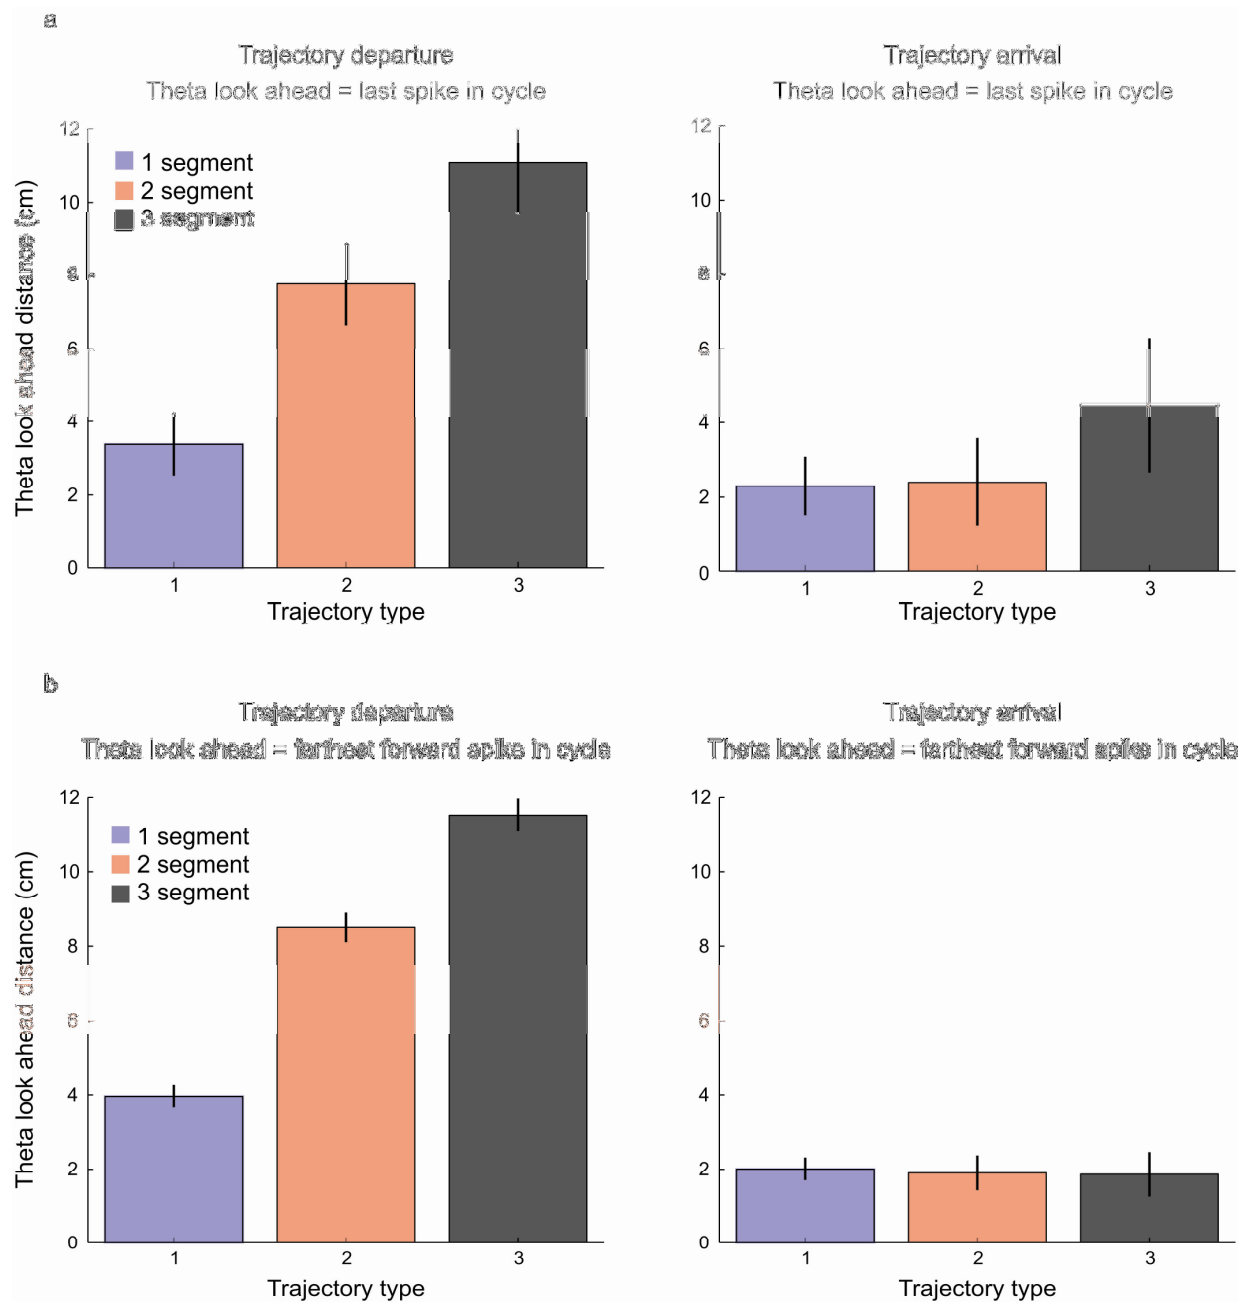


Supplementary Figure 2

Theta look-ahead was consistent within subject, across space, and for different methods of identifying trajectories.

(a) Mean look ahead distance ($\pm 95\%$ confidence interval) is plotted for individual rats on initiation of trajectories (left panel, cf. fig. 3 in the main text) and on completion of trajectories (right panel, cf. fig. 4 in the main text). (b) Look ahead distance was computed

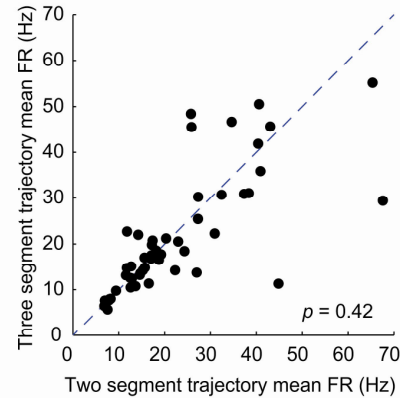
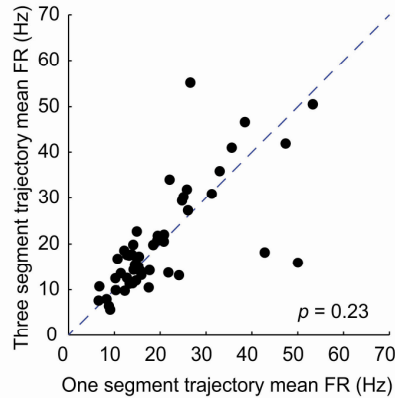
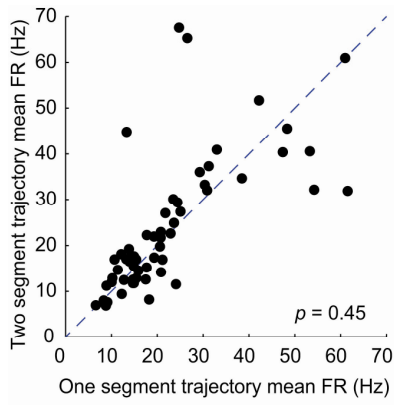
separately for one, two, and three segment trajectories beginning or ending at each of the three physical feeder locations. Patterns were similar in all cases. Error bars denote 95 % confidence intervals. (c) Using a more stringent running speed threshold for identifying trajectories (25 cm/s rather than 10 cm/s; see Methods) did not alter the pattern of theta sequence look ahead observed on trajectory departure or arrival at goal destinations. Error bars denote 95% confidence intervals.



Supplementary Figure 3

Alternative methods of computing theta look-ahead distance did not influence results.

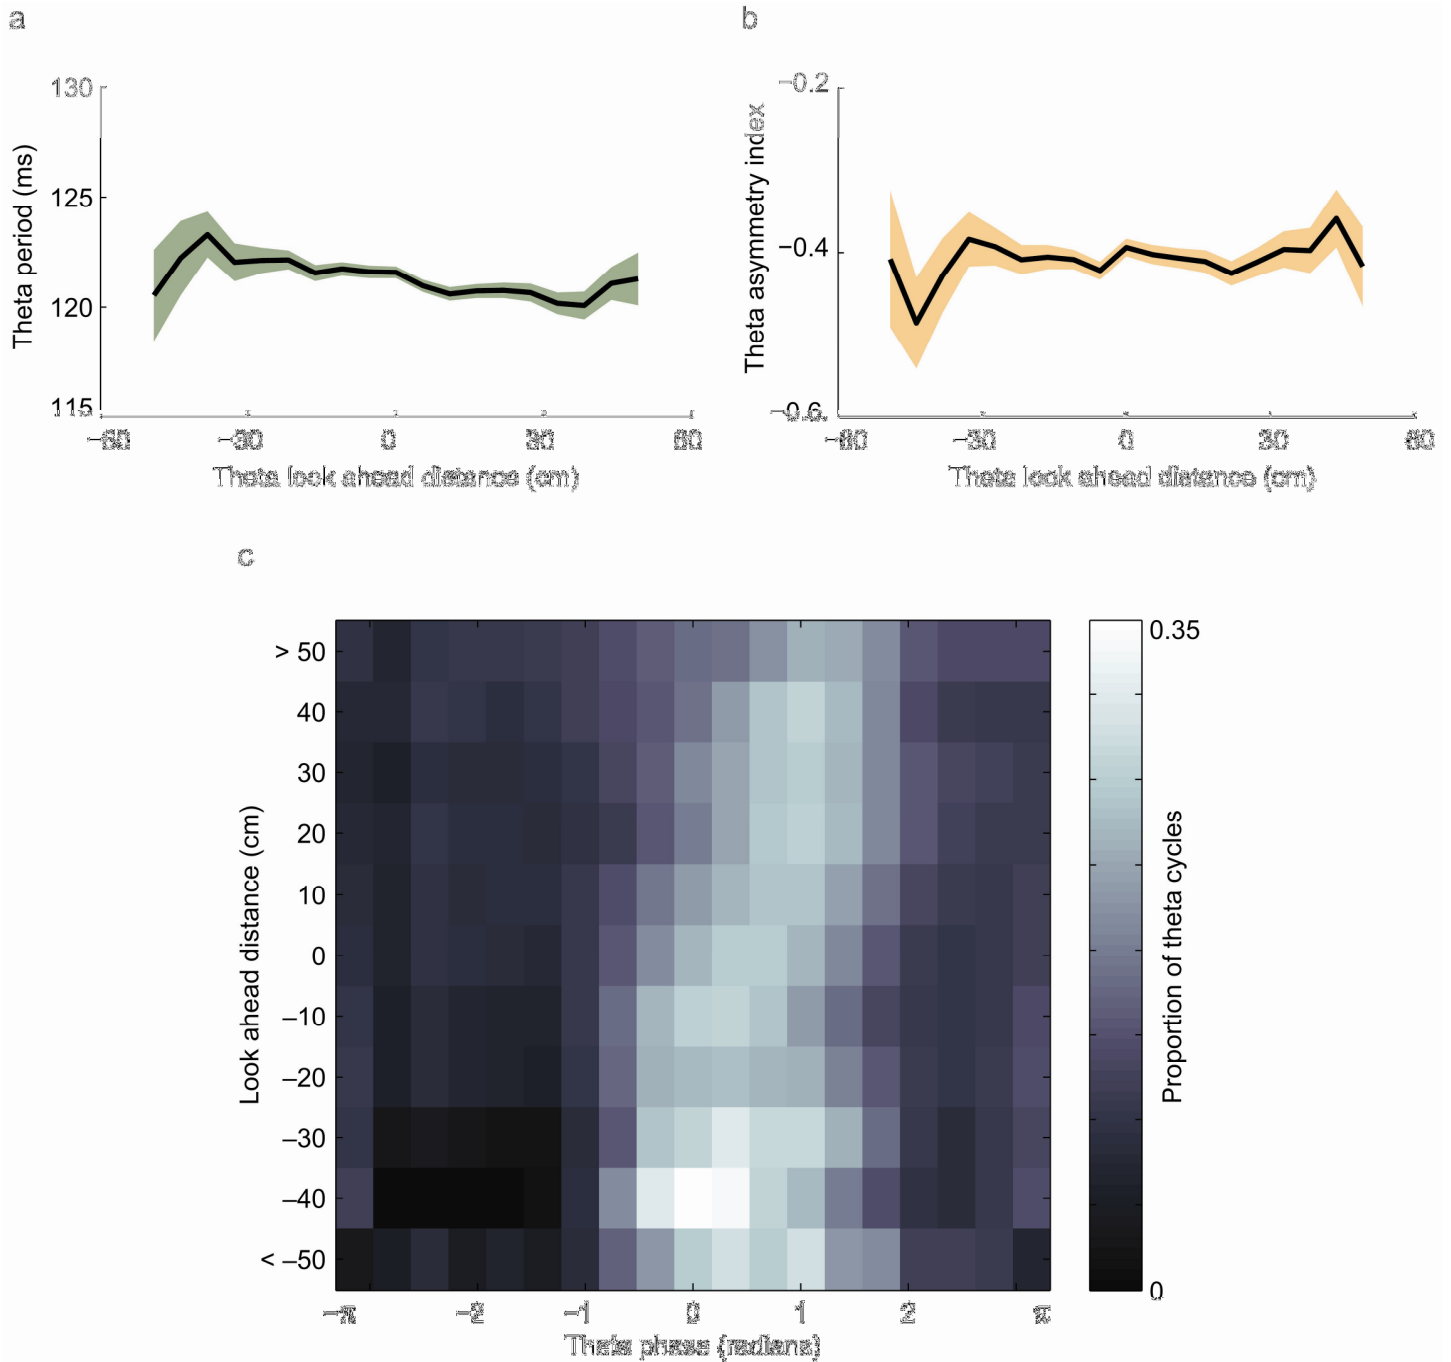
Taking the theta look ahead distance as either (a) the distance between the rat's current location and the place field center of the last spike to occur within each theta sequence, or (b) the distance between the rat's current location and the place field center of the cell farthest in front of rats' current location that was active in the cycle (regardless of when the cell fired in the theta cycle) produced similar patterns of results across trajectories.



Supplementary Figure 4

Place cell firing rates were similar across trajectories.

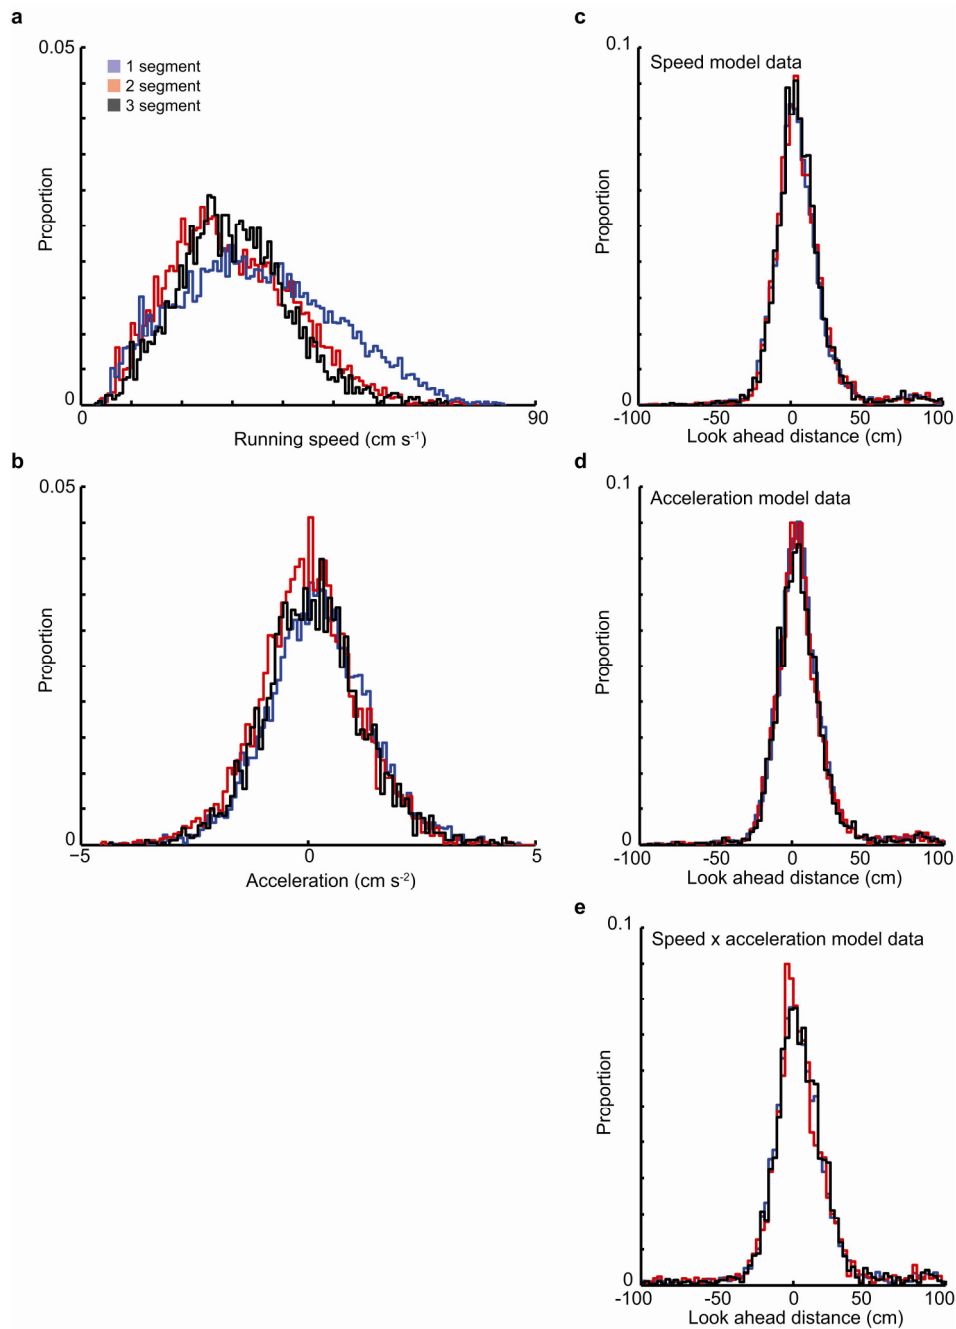
We computed the average firing rates of place cells separately for different trajectory types. Differences in rate across trajectories were not significant. The p values plotted in each figure are for sign tests.



Supplementary Figure 5

Properties of theta cycles.

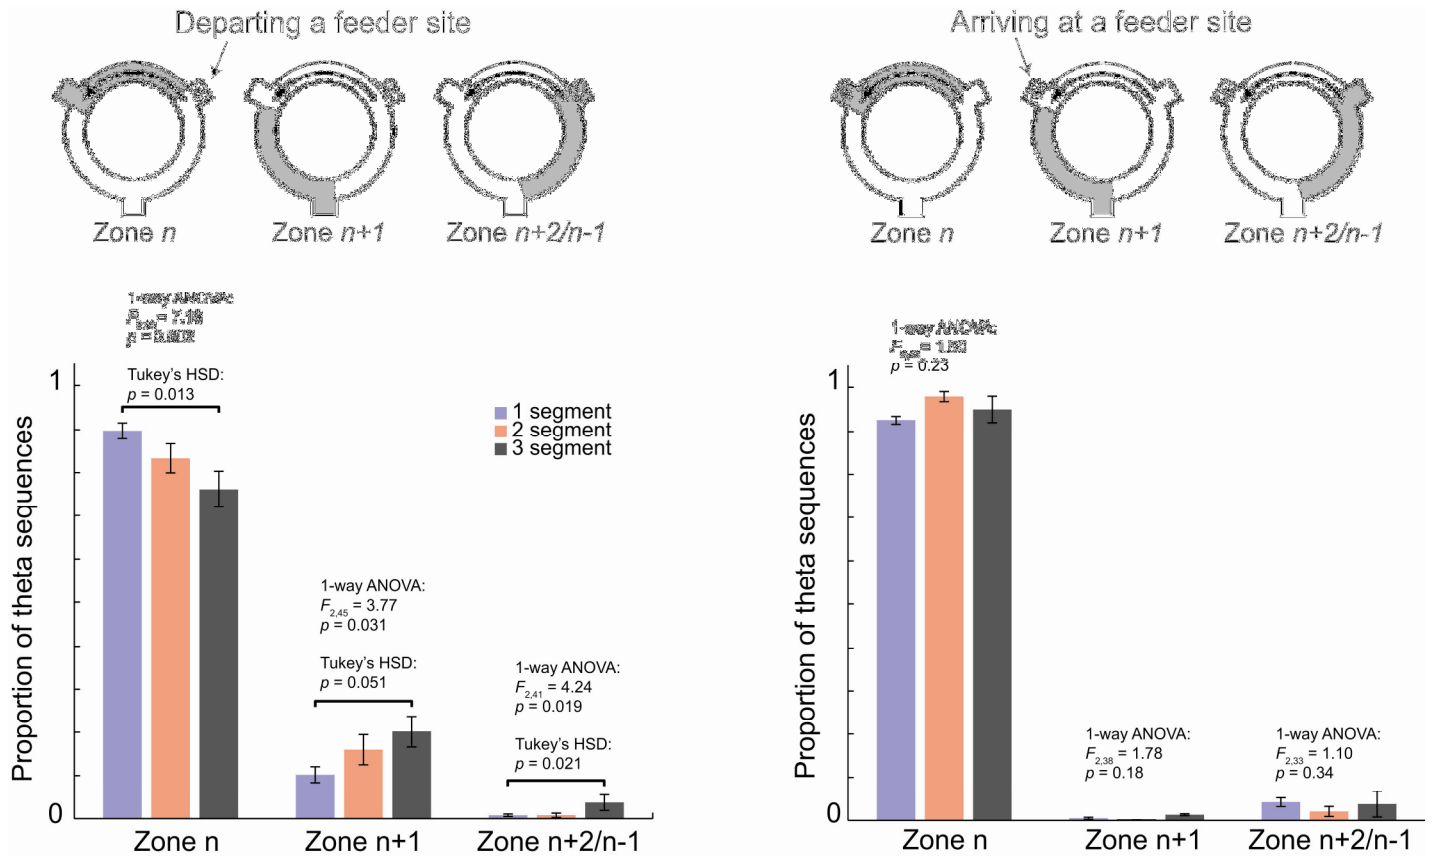
(a) The period of theta cycles was not consistently related to theta look ahead distance. Shaded regions indicate the standard error of the mean. (b) The symmetry of theta cycles did not vary substantially across different values of theta look ahead distance. Shaded regions indicated the standard error of the mean. (c) We constructed histograms of ensemble place cell spiking across the theta cycle for different look ahead distances. A theta phase of 0 corresponds to the trough of the theta oscillation as recorded from the hippocampal fissure. Histograms are normalized within row to display the proportion of place cell spikes at different phases of theta. Spiking was generally concentrated in the same portion of the theta cycle at different values of look ahead distance, although a slight shift to later phases is evident for theta cycles where look ahead was farthest in front of the animal.



Supplementary Figure 6

Movement parameters did not account for differences in look-ahead on trajectory departure.

Histograms display running speed (a) and acceleration (b) for the theta sequences analyzed in fig. 3 of the main text. Because there were differences between the three trajectory types, we constructed surrogate theta look ahead data sets (matched to the data used to construct fig. 3; $n = 20,269$ theta cycles) assuming that speed (c), acceleration (d), or particular combinations of speed and acceleration (e) determined theta look ahead. No goal-dependent effects were detected in these bootstrapped data sets.



Supplementary Figure 7

Forward-directed theta representations are more common on longer trajectories.

We computed the proportion of theta sequence trajectory representations that ended in different regions of the track as animals traversed the first limb of trajectories. On trajectory initiation (left panel), the proportion of sequences that ended in the space between the rat's start site and the next feeder on the maze (zone_n) was highest during one segment trajectories and decreased on two and three segment trajectories. The proportion of sequences that ended in zone_{n+1} was lowest for one segment trajectories and greatest for three segment trials. The proportion of sequences that extended to the zone farthest from the departure site zone_{n+2/n-1} was specifically increased on three segment trajectories. In contrast, as rats traversed the final limb of trajectories on arrival at goal destinations (right panel), no trajectory-dependent differences in the proportions of sequences ending in particular spatial zones were detected. In both panels, error bars indicate the standard error of the mean across recording sessions.