A learned map for places and concepts in the human MTL

Nora A. Herweg[1], Lukas Kunz[2], Armin Brandt[2], Paul A. Wanda[1], Ashwini D. Sharan[3], Michael R. Sperling[4], Andreas Schulze-Bonhage[2], Michael J. Kahana[1]

[1] Computational Memory Lab, Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

[2] Epilepsy Center, Medical Center – University of Freiburg, Faculty of Medicine, Freiburg im Breisgau, Germany

[3] Department of Neurosurgery, Thomas Jefferson University Hospital, Philadelphia, PA, USA

[4] Department of Neurology, Thomas Jefferson University, Philadelphia, PA, USA

* Correspondence:
Dr. Nora A Herweg
nherweg@sas.upenn.edu

or

Dr. Michael J Kahana
kahana@psych.upenn.edu

Keywords: cognitive map, spatial navigation, hippocampus, time, local field potential

Abstract

We examined the information coded by the human medial temporal lobe (MTL) during navigation of a complex virtual environment. Using multivariate decoding, we show that local field potentials in the MTL represent places, their semantic meaning, and the temporal order in which they are visited. These representations strengthen with experience, suggesting that the MTL constructs a learned map of physical and semantic spaces that can be used to predict future trajectories.

Cell populations in the medial temporal lobe (MTL) encode the spatial [1,2], temporal [3,4], and semantic [5,6] features of our perceptual and mnemonic experience. These features have rarely been investigated simultaneously, which is particularly important given their correlational structure: navigation through space takes time, and objects that we encounter in temporal or spatial proximity are often semantically related to one another. To investigate the roles of these variables in shaping neural activity, we trained a multivariate decoder on MTL micro-wire recordings from 20 patients with epilepsy as they navigated a virtual city. Using the decoder's predictions, we simultaneously assessed the representation of a subject's virtual spatial location, the semantics associated with that location, and the temporal sequence of visited locations.

We further examined how these representations interact with the degree of experience that subjects have as they learn the layout of the virtual environment. Repeated or extended exposure to a particular spatial environment, temporal sequence, or conceptual space routinely leads to improvements in behavioral performance on tasks that require knowledge of these regularities. But it is unclear how these changes in behavior relate to changes in the underlying representations of space, time and semantics in the human MTL. In our study, subjects navigated a virtual city to deliver objects to a series of target stores. Over time, subjects became more efficient at navigating to their target stores, as indicated by a reduction in excess path length over trials ($\chi^2_{(1)} = 7.77$, $p = 0.005$). Using our decoding approach, we directly assessed how this improvement in behavioral performance is reflected in a change of representational structure in the MTL.

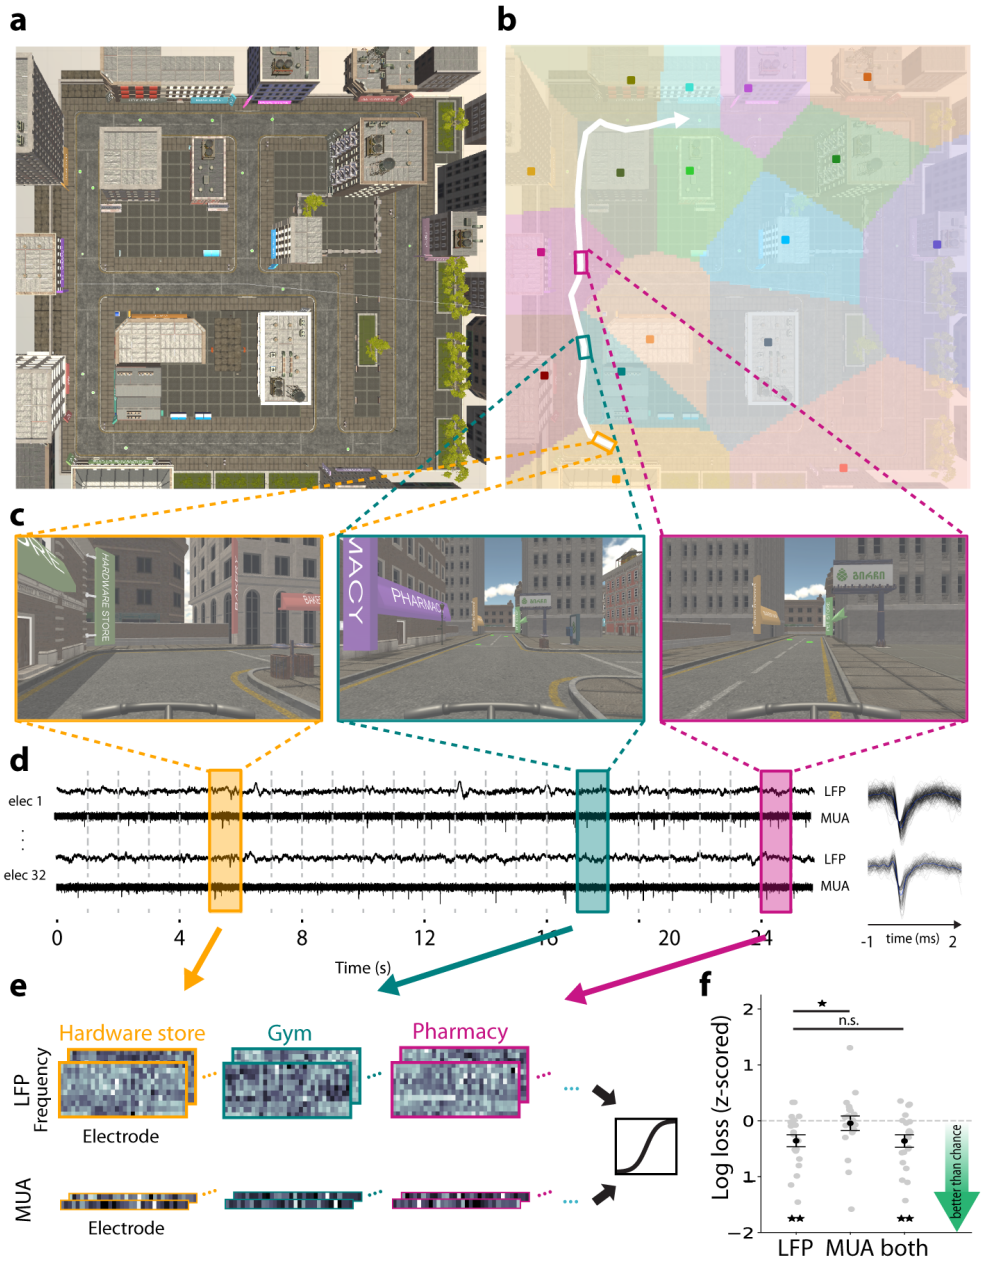

Figure 1. Multivariate decoding of location. a) Subjects navigated to a series of target stores in a virtual city. b) The target stores served as meaningful labels for different locations in the environment. Specifically, we assigned each coordinate in the environment to the store with the shortest Euclidean distance to that coordinate. The resulting borders are depicted along with a subject's example path from the hardware store to the bike shop (white arrow). c) Snapshots taken during navigation for the three exemplary locations highlighted with colored rectangles, which were labelled as hardware store (yellow), gym (teal; note that the pharmacy is in view despite being further away) and pharmacy (pink). d) Continuous micro-wire recordings were preprocessed to obtain local field potential (LFP) and multi-unit activity (MUA) traces, which were segmented in 1 s epochs (dashed gray lines). The waveforms of all neuronal spikes recorded on the depicted electrodes during the depicted time window are overlaid on the right. e) For each electrode, we computed MUA firing rate as well as LFP power for 10 approximately log-spaced frequencies between 3 and 200 Hz. The $N_{frequency}$ x $N_{electrode}$ LFP power matrices and the $N_{electrode}$ firing rate vectors were used as inputs to a multinomial logistic regression classifier predicting the subject's instantaneous location for each 1 s epoch. f) Across subjects, location decoding from LFP power, but not from MUA, was significantly better than chance (i.e. compared with classifiers trained on randomly shifted training labels). Error bars show standard error of the mean across subjects.

We trained a multinomial logistic regression model to decode a subject's location in the virtual city. To this end, we segmented the data into 1s epochs and labelled each epoch with the target store that was closest to the subject's instantaneous position (**Figure 1a-e**). This definition of location will later allow us to simultaneously assess the effects of spatial, temporal and semantic structure on neural activity patterns in the MTL. In line with recent animal work, the location of our human subjects could be decoded from local field potential (LFP) spectral power [7] (**Figure 1f**; $t_{(19)} = -3.33$, $p = 0.003$, Cohen's $d = -0.75$; 10 log-spaced frequencies between 3 and 200 Hz; significance assessed by comparing logarithmic loss against classifiers trained on shuffled training labels). Decoding from multi-unit firing rates, in turn, was no better than chance ($t_{(19)} = -0.33$, $p = 0.74$, Cohen's $d = -0.07$; see Supplement for discussion). There was no difference in decoding performance based on different frequency-bands or MTL sub-regions (hippocampus, N= 17 subjects; parahippocampal gyrus, N = 12 subjects; amygdala, N = 12 subjects; see Supplement). We therefore focused all following analyses on the LFP-based decoding model trained on all frequencies and MTL sub-regions.

We used our decoding model to understand whether the MTL encodes the spatial, temporal and semantic structure of a subject's experience. In a second step, we asked whether these representations strengthen over time. To answer these questions we analyzed the classifier's predictions, specifically those for all locations at which a subject was not currently located. The classifier predictions for a subject's true location determine the classifier's performance (which does not change over trials: $\chi^2_{(1)} = 0.01$, $p = 0.90$), but they do not tell us what kind of representation the classifier is detecting. Our classifier could, for instance, achieve high performance based on an internal representation of the stores that lacks any information on the spatial relations between stores. But if the MTL represents spatial relations, our classifier should assign high probabilities

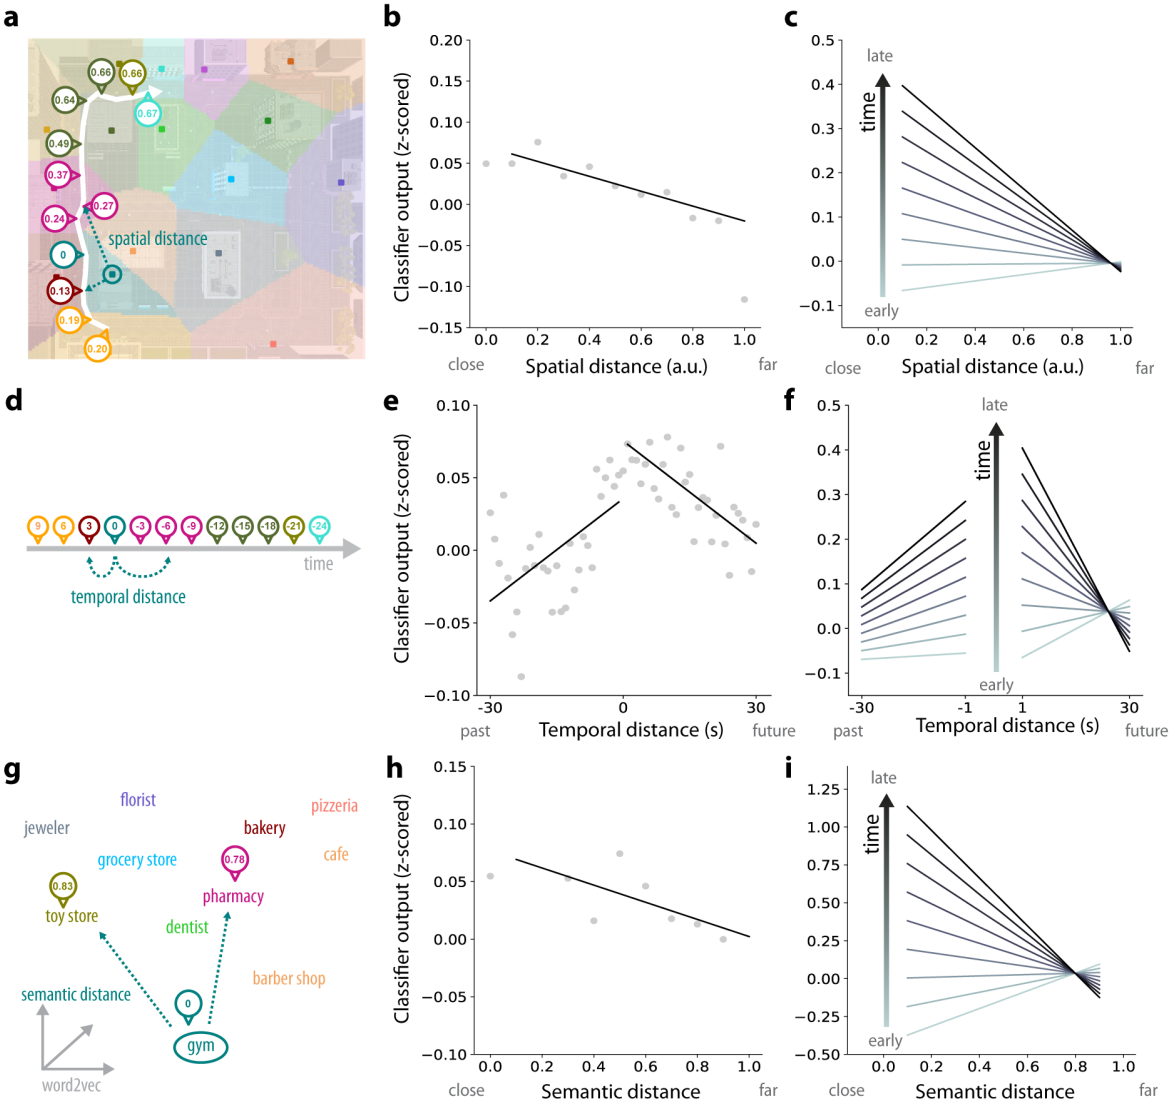

Figure 2. The representation of space, time and semantics strengthens with experience. a) A subject's spatial distance from an example store, the gym, changes as the subject navigates. If the MTL represents the stores' spatial locations, the classifier's output probabilities for a given store should decrease with the subject's distance to that store. b) We find that this is indeed the case: classifier output probabilities decrease with the spatial distance between a given store and the subject's true instantaneous position. All individual data points are averages over subjects and classes; all fitted lines are visualizing the parameter estimates of mixed effects models. c) This effect strengthens over time, suggesting that spatial knowledge manifests as increased representational similarity between nearby locations or an increased tendency for the MTL to represent proximate over distant spatial locations. Panels c, f and i show the effect of distance for eight equally spaced epoch numbers between 0 and 8000, corresponding to 133 min of navigation. d) Next, we asked how classifier output for a given stores relates to the time that has passed since the subject has visited that store (negative distance), or to the time that will pass until the subject arrives at that store (positive distance). e) Classifier output decreases with temporal distance, suggesting that subjects represent stores that they are about to visit or have just visited. This effect is asymmetric: classifier output is higher for stores that a subject is about to visit than for those that have been visited in the recent past. f) Like for spatial distance, the effect of temporal distance increases over time and this effect is stronger for stores on a subject's future path, meaning that with experience the MTL more strongly predicts a subject's proximate future trajectory. g) Subjects not only navigate physical but also semantic space as they travel from store to store. In this example, the subject is semantically closer to the gym when located at the pharmacy compared with the toy store. We assessed the effect of semantic distance in word2vec space on classifier

output for each store. h) Classifier output probabilities are high for stores that are semantically similar to the store at which the subject is currently located and low for stores that are semantically dissimilar to the subject's current location. i) Again, the effect of semantic distance strengthens over time, suggesting that the MTL increasingly represents the subject's location in semantic space.

to spatially proximate (but incorrect) locations and low probabilities to spatially distant locations. We find that this is the case: specifically, we observed a significant effect of the spatial distance between a given store and the subject's true location on classifier output probabilities for that respective store (**Figure 2a**, middle panel). To control for correlations between spatial, temporal and semantic distance, we first assessed the effects of individual predictors (e.g. spatial distance) in the "individual model" and then assessed the effects of all predictors (i.e. spatial, temporal and semantic distance) in a "joint model". The effect of spatial distance was significant in both models, meaning that spatial distance uniquely affects the classifier's predictions (individual model: $z = -5.22$, $\chi^2_{(1)} = 27.25$, $p < 0.001$; joint model: $z = -2.18$, $\chi^2_{(1)} = 4.76$, $p = 0.03$).

Furthermore, the effect of spatial distance on classifier output probabilities grew stronger over time, as indicated by a significant interaction of spatial distance and epoch number (individual: $z = -7.41$, $\chi^2_{(1)} = 54.95$, $p < 0.001$; joint: $z = -2.76$, $\chi^2_{(1)} = 7.62$, $p = 0.006$). This effect is illustrated in the right panel of **Figure 2a**, which shows the fitted effect of spatial distance on classifier output probabilities for eight equally spaced epoch numbers between 0 and 8000, corresponding to 133 min of navigation (this number is based on the maximum number of epochs across subjects). This finding provides evidence that spatial representations in the human MTL change over the course of learning: Specifically, spatial knowledge manifests as increased representational similarity between nearby locations or an increased tendency for the MTL to represent proximate over distant spatial locations.

Having shown that the MTL encodes the spatial relations between stores, we next asked whether its activity reflects the temporal sequence of visited locations. In other words, is there an

effect of the temporal distance between a given time point and the last or next time point at which a subject will be located at a particular store on the classifier's predictions. We find that this is the case when assessing temporal distance individually (**Figure 2b**, middle panel; $z = -5.13$, $\chi^2_{(1)} = 26.36$, $p < 0.001$) and we find a trend when controlling for the other variables ($z = -1.82$, $\chi^2_{(1)} = 3.30$, $p = 0.07$). This suggests that the MTL not only preferentially represents spatially proximate locations but specifically those that subjects have just visited or about to visit on their path through the environment. Furthermore, this effect is asymmetric: locations that will be visited in the near future are represented more strongly than those that have been visited in the recent past (individual: $z = 5.58$, $\chi^2_{(1)} = 31.16$, $p < 0.001$; joint: $z = 5.59$, $\chi^2_{(1)} = 31.28$, $p < 0.001$). And as for virtual space, the degree to which the MTL prefers temporally proximate over distant locations increases over time, especially for locations on a subject's future trajectory (**Figure 2b**, right panel; three-way interaction, individual: $z = -3.62$, $\chi^2_{(1)} = 13.11$, $p < 0.001$; joint: $z = -3.44$, $\chi^2_{(1)} = 11.81$, $p < 0.001$). The MTL thus predicts proximate upcoming locations more strongly, the more time subjects have spent navigating the environment.

Finally, we asked whether the MTL represents the semantic structure of the environment. To this end, we calculated the distances in word2vec [8] semantic space between all target stores and assessed their effect on the classifier's predictions. We find that the classifier assigns higher probabilities to semantically related stores, when assessed individually (**Figure 2c**, middle panel; $z = -2.09$, $\chi^2_{(1)} = 4.36$, $p = 0.04$) and we observe a similar trend when controlling for temporal and spatial distance (joint: $z = -1.96$, $\chi^2_{(1)} = 3.84$, $p = 0.05$). Again, the degree to which the MTL represents stores semantically related to the subject's current position (e.g. representing the pizzeria while the subject is located at the bakery) increases over time (**Figure 2c**, right panel; individual: $z = -16.60$, $\chi^2_{(1)} = 275.22$, $p < 0.001$; joint: $z = -16.15$, $\chi^2_{(1)} = 260.61$, $p < 0.001$). This

finding complements recent reports of concept coding in the human MTL [5,9]. Even though the semantic meaning of the store names has been learned prior to the experiment, activity in the MTL seems to reflect subjects' increasing knowledge of the environment's specific semantic subspace.

Our findings demonstrate that field potentials in the human MTL jointly represent a subject's virtual spatial location, the semantics associated with that location, and a subject's temporal trajectory. The representation of space, time and semantics becomes more pronounced with experience, indicating that the MTL stores a learned map of virtual and local semantic space that can be used to predict future trajectories.

Methods

Participants

20 patients with medication-resistant epilepsy undergoing clinical seizure monitoring at Thomas Jefferson University Hospital (Philadelphia, PA, USA) and the University Clinic Freiburg (Freiburg, GER) participated in the study. The study protocol was approved by the Institutional Review Board at each hospital and subjects gave written informed consent. Subjects were implanted with Behnke-Fried macro-micro depth electrodes (AdTech, Racine, WI) in the MTL. The location of these electrodes was determined based on clinical considerations.

Experimental design and task

Subjects played the role of a courier in a hybrid spatial-episodic memory task, riding a bicycle and delivering parcels to stores located within a virtual town (consisting of roads, stores,

and task-irrelevant buildings; **Figure 1**). Subjects completed a variable number of experimental sessions (min: 1, max: 4, mean = 2.45). Each experimental session consisted of a series of delivery days (i.e. trials; except for learning sessions described below), during which subjects navigate to deliver objects and, at the end of each trial, recall those objects. In the current study, we used data from the navigation phases only; data from the recall periods for an overlapping sample of subjects has been reported elsewhere [10,11]. Subjects completed slightly different versions of this paradigm, the details of which are described in the following paragraphs. We believe these task adaptations to be negligible in terms of the scope of the paper and report all results collapsed over all subjects/versions. The tasks were programmed and displayed to subjects using the Panda Experiment Programming Library [12], which is a Python based wrapper around the open source game engine Panda3d (with 3D models created using Autodesk Maya™) or the Unity Game Engine (Unity Technologies, San Francisco, CA).

Prior to starting the first delivery day, subjects viewed a static or rotating rendering of each store in front of a black background (later referred to as 'store familiarization'). Each store had a unique storefront and a sign that distinguished it from task-irrelevant buildings. Each delivery day trial consisted of a navigation phase (**Figure 1**) and a recall phase (for which no data is reported here). For the navigation phase, 13 stores were chosen at random out of the total number of 16 or 17 stores. Subjects were informed about their upcoming goal by on-screen instructions (e.g. "Please find store the hardware store.") and navigated to each store using the joystick or buttons on a game pad. Location-store mappings in the town were fixed for 4 subjects and random for 16 subjects (the layout was always fixed across experimental sessions; i.e. each subject experienced the same town layout across sessions). Upon arrival at the first 12 stores, subjects were presented with an audio of a voice naming the object or an image of the object they just delivered. Object

presentation was followed by the next on-screen navigation instruction. Upon arrival at the final store, where no item was presented, the screen went black and subjects heard a beep tone. After the beep, they had 30 or 90 s to recall as many objects as they could remember in any order. A final free recall phase followed on the last delivery day within each session.

To make sure that subjects did not spend too much time driving around the environment randomly when they did not know where a certain store was located, waypoints helped a subset of subjects (N = 4 subjects) navigate upon difficulties finding a store. Considering each intersection as a decision point, arrows pointing in the direction of the target store appeared on the street after three bad decisions (i.e. decisions that increased the distance to the target store). In a different task version (N = 5 subjects), subject had to complete a pointing task prior to navigation to each store. From their current location, they were asked to use the game pad to point an arrow in a straight-line path to where they thought the target store was located. This task served a similar purpose as the waypoints, since it provided subjects with feedback on the correct direction of the target store prior to navigation.

Most subjects (N = 16) completed an initial learning session prior to the first delivery day session, in which the store familiarization phase was followed by a 'town familiarization' phase. Here, subjects were instructed to navigate from store to store without delivering parcels (and recalling objects), visiting each store three times in pseudo-random order (each store was visited once, before the first store was visited the second time). LFP data during the learning session was available only for a subset of subjects (N = 4). For these subjects, data from the very first round of navigation was not used in the analyses. For some subjects (N = 4), a single town familiarization trial was repeated at the beginning of all following sessions prior to the first delivery day trial.

Behavioral analyses of navigation performance

Behavioral data were analyzed using Python version 3.6. As a measure of navigation performance, we computed the excess path length ratio. This is the actual length of a subject's path (i.e. from the initial instruction "please find the hardware store" to arrival at that store) divided by the shortest possible path between the start and arrival locations. A value of 1 indicates perfect performance, higher values indicate a suboptimal path. We used a mixed linear model with a random subject intercept to assess the effect of trial number on navigation performance. We performed a likelihood ratio test between a model including the fixed effect of trial number and an intercept-only model to test for a significant main effect of trial number.

Micro-wire data acquisition and preprocessing

LFPs were recorded from the inner micro-wire bundle of Behnke-Fried macro-micro depth electrodes (AdTech, Racine, WI) located in the MTL (including hippocampus, N = 17 subjects; parahippocampal gyrus, N = 12 subjects and amygdala, N = 12 subjects). Data were recorded at a sampling rate of 20,000 or 30,000 Hz using the NeuroPort ™ system (Blackrock, Salt Lake City, UT), a Neuralynx system (Neuralynx, Bozemann, MT) or an inomed system (inomed, Emmendingen, Germany). Data from the eight recording wires on each bundle were referenced online to a ninth designated reference wire. In cases where the reference wire was suspected to have low signal quality, one of the regular recording wires was used as an online reference instead. No re-referencing was performed offline. Coordinates of the radiodense macro-electrode contacts were derived from a post-implant CT or MRI scan and then registered with the pre-implant MRI scan in MNI space using SPM or Advanced Normalization Tools (ANTS)[13]. For 7 subjects, micro-

wire bundles were localized manually by a neurologist or a radiologist, for the other 13 subjects, manual annotations were not available and micro-wire bundles were instead localized by extrapolating 0.5 cm (the approximate length the micro-wires extend from the tip of the electrode) from the location of the most distal macro-electrode contact. Labels were derived using the probabilistic Harvard Oxford atlas with a threshold of 25%.

LFP data were analyzed using Python version 3.6 along with the Python Time Series Analyses (ptsa; https://github.com/pennmem/ptsa_new) and MNE [14] software packages. LFP data were aligned with behavioral data via pulses sent from the behavioral testing laptop to the recording system. A zero-phase notch filter was applied to filter line noise at 50 Hz (for data recorded in Germany) or 60 Hz (for data recorded in the US) and harmonics up to 240 Hz (stopband width: frequency/200). To extract multi-unit activity (MUA), data were additionally high-pass filtered between 300 and 3000 Hz using a zero-phase FIR filter. For each trial, a threshold of 3 times the interquartile range (IQR) above the 75th or below the 25th percentile was used to detect potential spikes. These were excluded from further analysis if their duration was too short (< 0.1 ms), too long (> 3 ms), or if they were surrounded by an artifact within a 100 ms window. Artifacts were defined as data points that were deviating from the 25th or 75th percentile either 20 times the IQR in the same direction as the spike in question or 10 times the IQR in the opposite direction as the spike in question. We then calculated the firing rate based on all valid spikes for each electrode and 1 s epoch. To extract LFP power, continuous data were down-sampled to 1000 Hz and convolved with complex Morlet wavelets (5 cycles) for 10 approximately log-spaced frequencies between 3 and 200 Hz. Some frequencies were slightly shifted to avoid overlap with line noise frequencies (3, 5, 8, 12, 19, 31, 40, 79, 130, 210 Hz). After convolution, data were log-transformed and averaged over 1 s non-overlapping epochs.

Multivariate classification

We used L2 penalized multinomial logistic regression to decode a subject's current location across these 1 s navigation epochs. As features we used the $N_{electrode}$ firing rate vectors, the $N_{electrode}$ * $N_{frequency}$ LFP power feature matrices, or both. All epochs were labeled using the stores in the environment (N classes = N stores). Specifically, for each of the 1 s epochs we determined the subject's current location by finding the store with the shortest Euclidean distance to the subject's instantaneous position. Since these labels may change over successive samples within a 1 s epoch, we assigned each epoch the most frequent label out of the set of labels for that epoch.

We assessed classifier performance using a nested leave-one-trial-out cross-validation procedure (each trial includes navigation to 13 stores). To make sure that all test trials were sufficiently separated in time from their respective training trials, we never tested the classifier on trials from the initial learning session. Because navigation in the learning session was continuous and not separated from one another by a recall phase or a short break, these trials were only used for classifier training. Nested cross validation was used to optimize the penalty parameter C using a grid search with 20 log spaced values between $\log_{10}(10^{-4})$ and $\log_{10}(200)$. To this end, the training data within each outer cross-validation fold was again divided into inner cross-validation folds. The classifier was trained on each inner training set for all 20 C values. The optimal C for a given outer fold was chosen as the C that maximized classifier performance across the inner test sets. For each inner and outer fold, all data were z-scored with respect to the mean and standard deviation of the respective training set. Samples were weighted inversely proportional to their class frequency to avoid bias towards more frequently occurring classes.

Performance was evaluated using logarithmic (or cross-entropy) loss (again weighting every sample inversely by its class frequency), which is defined as the negative logarithm of the probability assigned to the true class. Log loss therefore produces high values, when the assigned probability for the true class approaches zero (i.e. a confident incorrect classification) and values close to zero when the assigned probability for the true class approaches one (i.e. a confident correct classification). The value of log loss that corresponds to chance depends on the number of classes. The number of stores in the environment was variable across subjects (N_{stores} = 16 or 17) and this number was reduced for individual subjects with few trials due to insufficient sampling. Specifically, any classes that did not occur in every training set were removed to keep the number of classes constant across validation folds. We therefore used a permutation procedure ($N_{permutations}$ = 200) to z-score each subject's performance measure against their individual chance distribution. Specifically, we repeatedly trained the classifier on a shuffled version of the training data (again using the nested cross-validation scheme outlined above). Specifically, the true y-vector was flipped and rolled by a random number of elements in the vector. This procedure removes any true relation between features and to-be-predicted labels, while keeping the autocorrelation of the to-be-predicted labels unchanged. We then z-scored true classifier performance for each subject using the mean and standard deviation of that subject's random distribution. z-scores < 0 indicate above-chance performance and z-scores > 0 indicate below chance performance.

A one-sample t-test was used to compare z-scored classifier performance for individual conditions to chance. Paired t-tests were used to assess the effects of feature set (LFP, MUA, both). To evaluate the importance of activity from different MTL sub-regions we trained the classifier on a reduced feature set, which included only wire bundles in the hippocampus, parahippocampal gyrus

or amygdala. Likewise, we trained separate classifiers on LFP power in frequencies below (LF) or above (HFA) 30 Hz (see Supplement).

Analyses of classifier output probabilities

Classifier output probabilities from the test sets were z-scored with respect to their permutation distribution (see above). We then used a mixed linear model with a random intercept for class nested in subject to assess the effect of spatial distance, temporal distance, semantic distance, as well as their interactions with epoch number on classifier output probabilities. We excluded from this analysis all probabilities for the subject's true location, meaning that no statistical effect can be driven by a change in classifier performance per se. For each class (i.e. store) and each 1 s epoch, spatial distance was calculated as the Euclidean distance between the subject's current position and the location of the respective store. Temporal distance was calculated as the absolute temporal distance in seconds between the current epoch and the last or next epoch during which the subject was or will be located at the respective store. An additional binary regressor was included to model past vs. future time points. Because there was less data for long temporal distances, we only included data points with an absolute temporal distance < 30 s. Semantic distance was calculated as the Euclidean distance in word2vec space [8] between the name of a given store and the name of the store at which the subject is currently located. We assessed significance of fixed effects using likelihood ratio tests between a full model (all main effects or all main effects and all interaction terms) and a reduced model (main effect or interaction effect in question removed).

Data availability

Data that can be shared without compromising research participant privacy/consent will be available at http://memory.psych.upenn.edu/Electrophysiological_Data.

Acknowledgements

This work was supported by German Research Foundation (DFG) grant HE 8302/1-1 to NAH, National Science Foundation (NSF) grant BCS-1724243 to MJK, National Institutes of Health (NIH) grants MH061975 and NS113198 to MJK and Federal Ministry of Education and Research (BMBF) grant 01GQ1705A to ASB. We thank Alison Xu, Zeinab Helili, Katherine Hurley, Deb Levy, Logan O'Sullivan, Ada Aka and Allison Kadel for help with data acquisition and post-processing, Jonathan Miller and Ansh Johri for their contributions to the task design, Corey Novich and Ansh Patel for programming the Unity based experiment, as well as Joel Stein, Rick Gorniak and Sandy Das for electrode localization support. We are most grateful to all patients and their families who selflessly volunteered their time to participate in this research.

Author contributions

NAH and MJK designed research, NAH analyzed data, NAH and MJK wrote the paper, AB, LK, PAW and NAH collected data, AS-B, MRS, ADS recruited participants and performed clinical duties associated with data collection including neurosurgical procedures or patient monitoring, all authors reviewed and edited the manuscript.

References

1. Moser, E. I., Kropff, E. & Moser, M.-B. Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.* **31**, 69–89 (2008).
2. Herweg, N. A. & Kahana, M. J. Spatial representations in the human brain. *Front. Hum. Neurosci.* (2018).
3. Eichenbaum, H. Time cells in the hippocampus: A new dimension for mapping memories. *Nat. Rev. Neurosci.* **15**, 732–744 (2014).
4. Umbach, G. *et al.* Time cells in the human hippocampus and entorhinal cortex support episodic memory. *bioRxiv* (2020) doi:10.1101/2020.02.03.932749.
5. Quiroga, R. Q. Concept cells: the building blocks of declarative memory functions. *Nat. Rev. Neurosci.* **13**, 587–597 (2012).
6. Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-spatial dimension by the hippocampal/entorhinal circuit. *Nature* **543**, 719–722 (2017).
7. Agarwal, G. *et al.* Spatially Distributed Local Fields in the Hippocampus Encode Rat Position. *Science (80-.).* **344**, 626–630 (2014).
8. Mikolov, T., Corrado, G., Chen, K. & Dean, J. Efficient estimation of word representations in vector space. *arXiv* 1301.3781 (2013).
9. Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E. J. Organizing conceptual knowledge in humans with a gridlike code. *Science (80-.).* **352**, 1464–1467 (2016).
10. Miller, J. F. *et al.* Neural acitivity in human hippocampal formation reveals the spatial context of retrieved memories. *Science (80-.).* **365**, 1111–1114 (2013).
11. Herweg, N. A. *et al.* Reactivated spatial context guides episodic recall. *J. Neurosci.* **40**, 1640–19 (2020).
12. Solway, A., Miller, J. F. & Kahana, M. J. PandaEPL: A library for programming spatial navigation experiments. *Behav. Res. Methods* **45**, 1293–1312 (2013).
13. Avants, B. B., Epstein, C. L., Grossman, M. & Gee, J. C. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**, 26–41 (2008).
14. Gramfort, A. *et al.* MNE software for processing MEG and EEG data. *Neuroimage* **86**, 446–460 (2014).