

A Learned Map for Places and Concepts in the Human Medial Temporal Lobe

 Nora A. Herweg,^{1,2}  Lukas Kunz,^{3,4}  Daniel Schonhaut,¹  Armin Brandt,⁴  Paul A. Wanda,¹  Ashwini D. Sharan,⁵  Michael R. Sperling,⁶  Andreas Schulze-Bonhage,⁴ and  Michael J. Kahana¹

¹Computational Memory Lab, Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania, 19104, ²Department of Neuropsychology, Institute of Cognitive Neuroscience, Faculty of Psychology, Ruhr University Bochum, 44801 Bochum, Germany, ³Department of Biomedical Engineering, Columbia University, New York, New York 10027, ⁴Epilepsy Center, Medical Center, University of Freiburg, Faculty of Medicine, 79106 Freiburg, Germany, ⁵Departments of Neurosurgery, and ⁶Neurology, Thomas Jefferson University, Philadelphia, Pennsylvania 19107

Distinct lines of research in both humans and animals point to a specific role of the hippocampus in both spatial and episodic memory function. The discovery of concept cells in the hippocampus and surrounding medial temporal lobe (MTL) regions suggests that the MTL maps physical and semantic spaces with a similar neural architecture. Here, we studied the emergence of such maps using MTL microwire recordings from 20 patients (9 female, 11 male) navigating a virtual environment featuring salient landmarks with established semantic meaning. We present several key findings. The array of local field potentials in the MTL contains sufficient information for above-chance decoding of subjects' instantaneous location in the environment. Closer examination revealed that as subjects gain experience with the environment the field potentials come to represent both the subjects' locations in virtual space and in high-dimensional semantic space. Similarly, we observe a learning effect on temporal sequence coding. Over time, field potentials come to represent future locations, even after controlling for spatial proximity. This predictive coding of future states, more so than the strength of spatial representations per se, is linked to variability in subjects' navigation performance. Our results thus support the conceptualization of the MTL as a memory space, representing both spatial- and nonspatial information to plan future actions and predict their outcomes.

Key words: cognitive map; hippocampus; local field potential; spatial navigation; time

Significance Statement

Using rare microwire recordings, we studied the representation of spatial, semantic, and temporal information in the human MTL. Our findings demonstrate that subjects acquire a cognitive map that simultaneously represents the spatial and semantic relations between landmarks. We further show that the same learned representation is used to predict future states, implicating MTL cell assemblies as the building blocks of prospective memory functions.

Received Jan. 24, 2022; revised Mar. 10, 2023; accepted Mar. 15, 2023.

Author contributions: N.A.H. and M.J.K. designed research; N.A.H., L.K., D.S., A.B., P.A.W., A.D.S., M.R.S., and A.S.-B. performed research; N.A.H. analyzed data; N.A.H. and M.J.K. wrote the paper.

This work was supported by German Research Foundation Grants HE 8302/1-1 to N.A.H. and KU 4060/1-1 to L.K., National Science Foundation Grant BCS-1724243 to M.J.K., National Institutes of Health (NIH)–National Institute of Mental Health Grant MH061975, NIH–National Institute of Neurological Disorders and Stroke Grant NS113198 to M.J.K., and Federal Ministry of Education and Research Grant 01GQ1705A to A.S.-B. We thank Alison Xu, Zeinab Helili, Katherine Hurley, Deb Levy, Logan O'Sullivan, Ada Aka, and Allison Kadel for help with data acquisition and postprocessing; Jonathan Miller and Ansh Johri for contributions to the task design; Corey Novich and Ansh Patel for programming the Unity-based experiment; Joel Stein, Rick Gorniak, and Sandy Das for electrode localization support; and the patients and families of patients who volunteered to participate in this research.

The authors declare no competing financial interests.

Correspondence should be addressed to Nora A. Herweg at nherweg@sas.upenn.edu or Michael J. Kahana at kahana@psych.upenn.edu.

<https://doi.org/10.1523/JNEUROSCI.0181-22.2023>

Copyright © 2023 the authors

Introduction

Ever since Tolman (1948) introduced the idea of a cognitive map, researchers have debated its scope beyond the domain of physical space (Bellmund et al., 2018). If place cells in the rodent hippocampus act as pointers in a spatial cognitive map (O'Keefe, 1976; O'Keefe and Nadel, 1978), how do they relate to the episodic and semantic memory deficits seen in human patients with lesions or resections of the hippocampus and surrounding medial temporal lobe (MTL) cortex (Scoville and Milner, 1957)? O'Keefe and Nadel (1978) discussed the possibility that the human MTL maps semantic space like it maps physical space, but it was not until decades later that researchers discovered hippocampal concept cells (Quiroga, 2012; Reber et al., 2019) that fire for particular concepts in high-dimensional semantic space. Finally, time cells fire at particular times in a sequence of events (MacDonald et al., 2011; Eichenbaum, 2014; Umbach et al., 2020; Reddy et al., 2021). Together, these cells provide all

components for the formation and retrieval of episodic memories, by defining what happened, where, and when (Tulving, 2001; Miller et al., 2013; Kunz et al., 2021).

However, few studies have examined spatial, semantic, and temporal representations in combination. Moreover, it remains largely unclear how these different representations change during learning and how they relate to behavior. Here, we address this gap by analyzing microwire recordings from the human MTL. Patients who were undergoing clinical seizure monitoring learned to navigate a virtual environment while instructed to successively find different target stores. Using a multivariate decoding approach, we asked whether the MTL would independently represent spatial, temporal, and semantic information while subjects moved through the virtual city. In addition, we aimed to characterize learning-related changes to these representations. Whereas place fields in rodents navigating small vista spaces become established within a few minutes (Frank et al., 2004), we reasoned that representational changes in a complex environment with several distinct start and target locations may occur on a slower time scale. Finally, we asked how representations of space and temporal sequences relate to subjects' behavior. Navigation performance in a novel environment tends to be variable and improve over time (Manning et al., 2014). However, it remains unclear how variability in behavior relates to variability in neural activity in human MTL populations.

To address these questions, we first trained a location decoder on the spectral features of the local field potential (LFP). The micro LFP is particularly well suited for this approach. It can be recorded on a large number of MTL channels (i.e., eight microwires extending from the tip of each MTL depth electrode, even in the absence of large isolated spikes), while also providing high temporal stability, in particular across recording sessions. We trained our decoder to predict the subject's location, defined as the label of the target store closest to the subject's instantaneous position.

We chose to define location in this manner to capture both semantic and spatial aspects of the subject's position; each landmark is associated with a particular location in space and also has a predefined semantic meaning. We can thus derive a measure of representational content from the classifier output probabilities. If the MTL represents spatial location, the instantaneous probability for a particular store should depend on the spatial distance between the subject's current location and that store; if the MTL represents semantic space, the probability should depend on the semantic distance between the two locations. Similarly, we assessed temporal information. If the MTL tracks the sequence of past events (MacDonald et al., 2011; Hsieh et al., 2014), the probability assigned to a store should depend on the time elapsed since the subject passed that store. We assumed that if subjects learned the associations between the stores they would also predict upcoming locations along their path. If these predictions are represented in the MTL, the probability should further depend on the time that will pass before a subject reaches a store.

Materials and Methods

Participants

Twenty patients (9 female, 11 male) with medication-resistant epilepsy undergoing clinical seizure monitoring at Thomas Jefferson University Hospital in Philadelphia and the University Clinic in Freiburg, Germany, 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) in the MTL. The location of these electrodes was determined based on clinical considerations.

Experimental design and statistical analyses

Behavioral task. Subjects played the role of a bicycle courier in a spatial memory task, delivering parcels to stores located within a virtual town (consisting of roads, stores, and task-irrelevant buildings; Fig. 1). Subjects completed a variable number of delivery days (mean = 9.8; minimum, 2; maximum, 22) across one or multiple experimental sessions (mean = 2.45; minimum, 1; maximum, 4) of variable duration (mean = 56.3 min; minimum, 14.2; maximum, 124.9). On each delivery day, subjects were first instructed to deliver a series of objects to specific stores in the city (see below). Then, after navigating to the last store, subjects were tested on their memory for these 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 previously (Miller et al., 2013; Herweg et al., 2020). 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 article and report all results collapsed over all subjects/versions. One of the differences was the spatial layout and visual implementation of the city. Figure 1 shows the newest version of the task ($N = 5$); Miller et al. (2013; $N = 11$) and Herweg et al. (2020; $N = 4$) provide a depiction of the older versions. The tasks were programmed and displayed to subjects using the Panda Experiment Programming Library (Solway et al., 2013), 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.

Before 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 consisted of a navigation phase (Fig. 1) and a recall phase (data not shown). For the navigation phase, 13 stores were chosen pseudorandomly from a total number of 16 or 17 stores. Subjects were informed about their upcoming goal by on-screen instructions (e.g., "Please find the hardware store.") and navigated to each store using the joystick or buttons on a game pad. The mapping between store names and store locations/visual appearance were random for most subjects (15 unique mappings were used across 20 subjects). The layout, however, was always fixed across experimental sessions (i.e., each subject experienced the same city layout across sessions). On 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. On 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 ensure that subjects did not spend extensive time searching for a store, waypoints helped a subset of subjects ($N = 4$ subjects) navigate. 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 version of the task ($N = 5$ subjects), subjects had to complete a pointing task before 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 because it provided subjects with feedback on the correct direction of the target store before navigation.

Most subjects ($N = 16$) completed an initial learning session before 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 or later recalling objects, visiting each store three times in pseudorandom order (each store was visited once before any repeat visit). 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

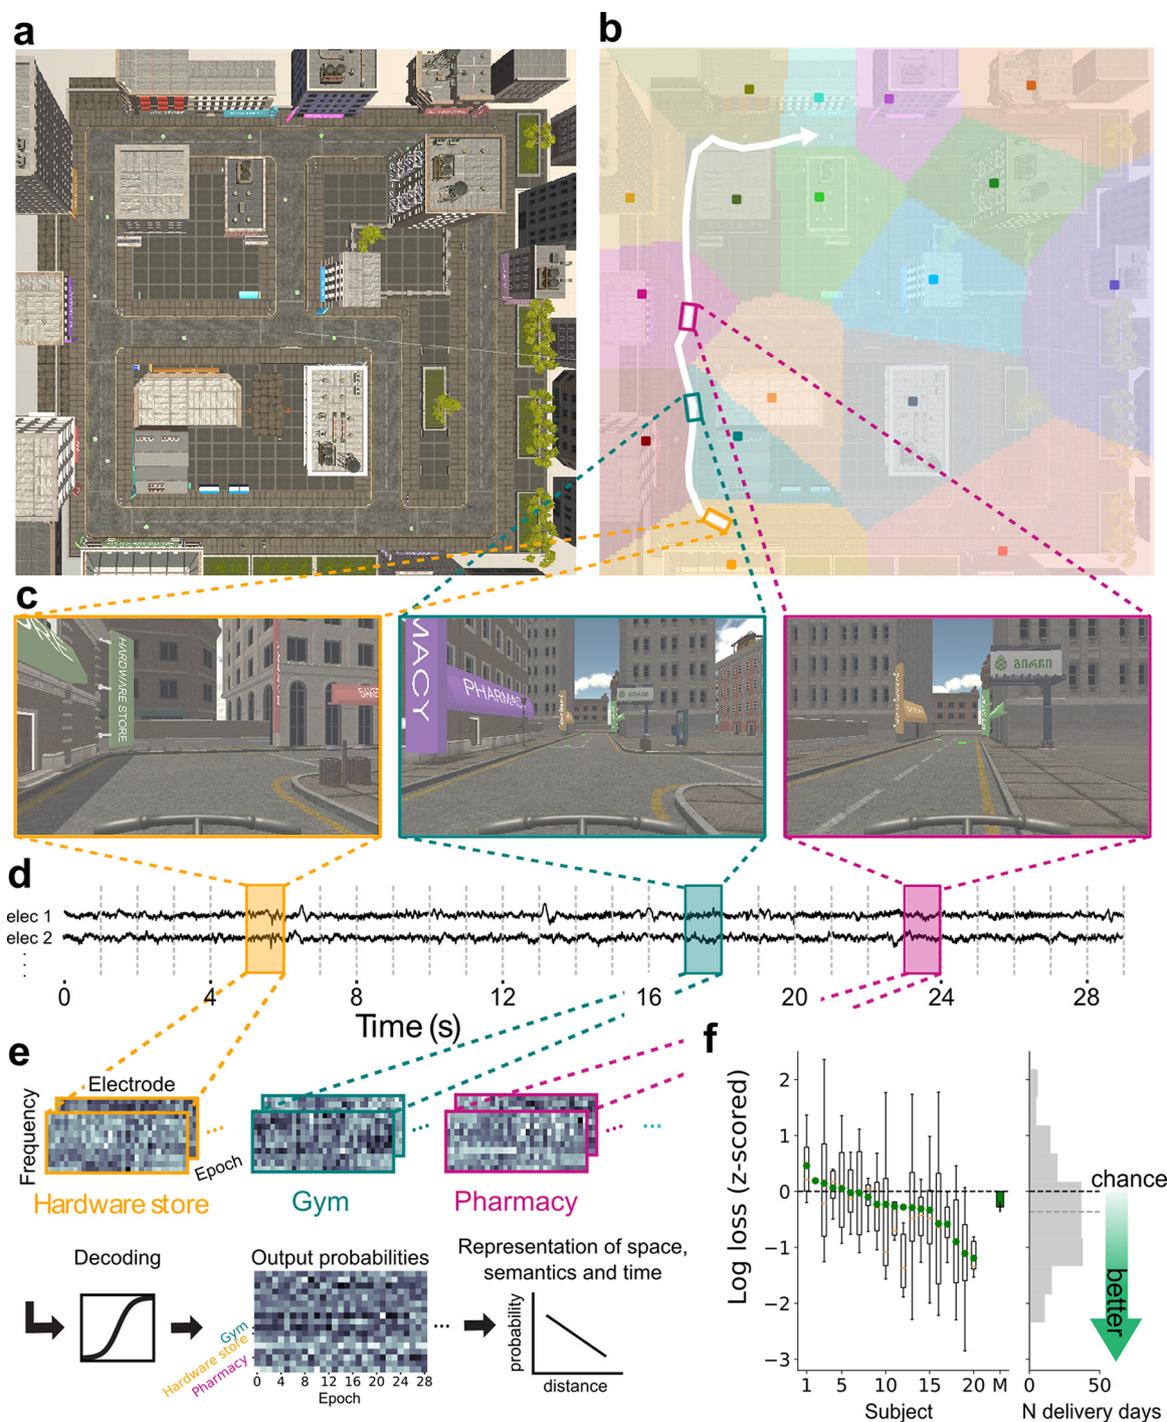


Figure 1. Multivariate decoding of location. **a**, Subjects navigated to a series of target stores in a virtual city. **b**, The stores served as meaningful labels for different locations, and we thus assigned each coordinate in the environment to the closest store (based on Euclidean distance). The resulting borders along with a subject's example path from the hardware store to the bike shop (white arrow) are illustrated in **b**. **c**, Snapshots taken during navigation for the three exemplary locations outlined with colored rectangles, which were labeled as hardware store (yellow), gym (teal; Note that the pharmacy is in view despite being farther away.), and pharmacy (pink). **d**, We extracted spectral power in 10 log-spaced frequencies between 3 and 200 Hz for each 1 s epoch of field potential recordings. **e**, The $N_{\text{frequency}} \times N_{\text{electrode}}$ LFP power matrices served as inputs to a multinomial logistic regression classifier predicting the subject's instantaneous location for each 1 s epoch. In a second step, we used the classifier output probabilities to infer whether the MTL represents virtual space, semantic space, and the temporal sequence of visited locations. To do this, we analyzed the relation between the predicted probabilities for each store and the spatial, semantic, or temporal distance between that store and the subject's true location (see Figure 2). **f**, We used logarithmic loss to assess decoding performance. A log loss of zero indicates perfect classification performance with no uncertainty (i.e., a consistent classifier output of 100% for the true class). We z-scored classification performance relative to a chance distribution derived from classifiers trained on randomly shifted labels. A z-score of zero (dashed black line) indicates chance performance, lower values indicate performance exceeding chance. Left, The distribution of z-scored performance across delivery days (outliers removed for visualization) for each subject, along with the subject-specific average (green points). Average z-scored performance across subjects (green bar \pm SEM) was significantly better than chance. Right, The aggregate distribution and average (gray dashed line) across all delivery days.

analyses. For some subjects ($N = 4$), a single town familiarization trial was repeated at the beginning of all following sessions before the first delivery day.

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 one 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 delivery day number on navigation performance. We performed a likelihood ratio test between a model including the fixed effect of delivery day number and an intercept-only model to test for a significant main effect.

Microwire data acquisition and preprocessing. LFPs were recorded from the inner microwire bundle of Behnke-Fried Macro-Micro Depth Electrodes (AdTech) located in the MTL (including hippocampus, $N = 17$ subjects; parahippocampal gyrus, $N = 12$ subjects; amygdala, $N = 12$ subjects). Data were recorded at a sampling rate of 20,000 or 30,000 Hz using the NeuroPort system (Blackrock Neurotech), a Neuralynx system, or an Inomed system. Data from the eight recording wires on each bundle were referenced online to a ninth designated reference wire. If the reference wire had low signal quality, one of the regular recording wires was used as an online reference instead. No rereferencing was performed off-line. Coordinates of the radiodense macroelectrode contacts were derived from a postimplant CT or MRI scan and then registered with the preimplant MRI scan in MNI space using SPM or Advanced Normalization Tools (ANTs; Avants et al., 2008). For seven subjects, microwire bundles were localized manually by a neurologist or a radiologist. For the other 13 subjects, microwire bundles were localized by extrapolating 0.5 cm (the approximate length the microwires extend from the tip of the electrode) from the location of the most distal macroelectrode 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 Analysis (https://github.com/pennmem/pts_a_new) and MNE (Gramfort et al., 2014) 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 United States) and harmonics up to 240 Hz (stop-band width, frequency/200). Based on the finding that data cleaning can decrease statistical power for multivariate decoding of memory states from intracranial EEG data (Meisler et al., 2019), we performed no additional data cleaning. Instead, we performed a set of additional analysis to assess whether the presence of artifacts affects any of our results (see below, Artifacts). 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. When using exact log spacing (i.e., 3, 5, 8, 12, 19, 31, 50, 79, 125, and 200 Hz), some frequencies would have been close to line noise frequencies or their harmonics at 50, 60, 100, 120, 150, 180, and 200 Hz. We therefore shifted these frequencies to ensure a 10 Hz distance to the closest line noise frequency. The extracted frequencies were 3, 5, 8, 12, 19, 31, 40, 79, 130, and 210 Hz. After convolution, data were log transformed and averaged over 1 s nonoverlapping 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 $N_{\text{electrode}} * N_{\text{frequency}}$ LFP power feature matrices. 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. As these labels may change over successive samples within a 1 s epoch, we assigned each epoch the most frequent label of the set of labels for that epoch. The average number of epochs per class (i.e., store) was 182.5 (minimum, 16; maximum, 1033). The average ratio between the most and least frequently occurring classes per subject was 10:1 (minimum, 3:1; maximum, 41:1). Class imbalance is a concern in classification

analyses because a classifier may perform well if it exhibits a simple bias toward predicting the majority class. We address this issue in two ways. When computing classifier performance, we weighted samples inversely proportional to their class frequency. In addition, we z -scored performance against a permutation distribution that exhibits the same class imbalance as the original data (see below). If a classifier performs better on the original than on the permuted data, this difference cannot be explained by class imbalance.

We assessed classifier performance using a nested leave-one-delivery-day-out cross-validation procedure (each delivery day includes navigation to 13 stores and lasted several minutes). To make sure that test and training data were sufficiently separated in time, we never tested the classifier on data 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 data 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 leave-one-delivery-day-out 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 SD of the respective training set. Samples were weighted inversely proportional to their class frequency to avoid bias toward more frequently occurring classes.

We evaluated performance 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, that is, stores. The number of stores in the environment was variable across subjects ($N_{\text{stores}} = 16$ or 17). We therefore used a permutation procedure ($N_{\text{permutations}} = 200$) to z -score the log loss for each delivery day against its individual chance distribution. Specifically, we repeatedly trained the classifier on a shuffled version of the training data. On each permutation, the true y vector (i.e., location labels) was flipped and circularly shifted by a random number of elements in the vector before the entire nested cross-validation scheme was repeated (Valente et al., 2021). 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 delivery day using the mean and SD of the random distribution of that delivery day. A z -score < 0 indicates above-chance performance, and a z -score > 0 indicates below-chance performance.

A one-sample t test across subjects was used to compare average z -scored classifier performance to chance. To evaluate the importance of activity from different MTL subregions, we additionally 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 or above 30 Hz. We used a t test across subjects to assess the effect of frequency band and a mixed linear model with a random subject intercept to assess the effect of subregion on classifier performance. A mixed-effects model is better suited for the latter effect because not all subjects had electrodes in all brain regions of interest. In the mixed model, we assessed significance of the single fixed effect (brain region) using a likelihood ratio test between the full model (main effect included) and an intercept-only model. All tests were two sided.

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 (i.e., unique locations) nested in subject to assess the effect of spatial distance, temporal distance, and semantic distance, as well as their interactions with epoch

number or navigation performance on classifier output probabilities. We excluded from this analysis all probabilities for the subject's true location, meaning that no statistical effect could be driven by a change in classifier accuracy. 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 versus future time points. Because there was less data for long temporal distances, we only included data points with an absolute temporal distance <30 s, meaning distances for stores that the subject has visited within the previous 30 s or is about to visit in the next 30 s. Semantic distance was calculated as the Euclidean distance in word2vec space (Mikolov et al., 2013) 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).

Artifacts. We detected interictal epileptiform discharges (IEDs) and other artifacts based on extreme voltage values. We defined extreme values using the interquartile range (IQR) as a robust measure of dispersion (Tukey, 1977). IEDs are very large in amplitude compared with the ongoing EEG. Consistent with prior work (Herweg et al., 2020), we therefore used a conservative threshold of 5 times the IQR above the 75th percentile or 5 times the IQR below the 25th percentile to mark individual time points as artifactual. We then computed an index reflecting the percentage of artifactual data points for each epoch, averaged over microwire channels.

To assess whether the presence of artifacts had an impact on our ability to decode location, we averaged this artifact index over epochs, resulting in a single number per subject, which we correlated with decoding performance. We also asked whether artifacts may inflate the effect of spatial, semantic, or temporal distance on classifier output probabilities. This may, for instance, be the case if artifacts are temporally clustered and result in increased similarity in temporally proximate epochs. To rule out any such influence, we repeated the analyses assessing the effects of spatial/semantic/temporal distance on classifier output while adding the epochwise artifact index and its interaction with the respective distance measure as a predictor. If IEDs or other artifacts are inflating the effect of temporal/semantic/spatial distance, we would expect to see an interaction between the respective distance measure and the artifact index such that a higher artifact index is linked to stronger distance modulation.

Data availability

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

Results

We analyzed MTL microwire recordings from 20 patients undergoing clinical seizure monitoring, who navigated a virtual city to deliver objects to different target stores (Fig. 1*a–c*). Subjects completed a variable number of delivery days (mean = 9.8; minimum, 2; maximum, 22) across one or multiple experimental sessions (mean = 2.45; minimum, 1; maximum 4) of variable duration (mean = 56.3 min; minimum = 14.2 maximum, 124.9). On each delivery day, subjects first navigated to a series of 13 stores that were pseudorandomly selected from the total number of stores (16–17) without replacement. On-screen instructions informed subjects of their upcoming goal (e.g., “Please find the hardware store.”). Subjects navigated to each store using the joystick or buttons on a game pad. On arrival at each of the first 12 stores, they were presented with the object they just delivered.

Then, after navigating to the final 13th store, subjects were tested on their memory for these objects. Here, we report data from the navigation phase only; retrieval data from an overlapping set of subjects have been published previously (Miller et al., 2013; Herweg et al., 2020). In total, subjects contributed an average of 62 min of navigation data (minimum, 24; maximum, 146). We quantified subjects' navigation performance using a ratio between their path length for each delivery and the shortest available path for that delivery. Over time (i.e., delivery days), subjects became more efficient at navigating to their target stores [$\chi^2(1) = 7.77, p = 0.005$], confirming that they acquired knowledge for the spatial relations between the target stores.

We used a multinomial logistic regression classifier to decode location from LFP spectral power. To train our decoder, we segmented the LFP into 1 s epochs and labeled each epoch with the target store that was closest to the subject's instantaneous position (Fig. 1*b–e*). We assessed significance of our model using a permutation procedure in which we trained chance classifiers on shuffled training labels (see above, Materials and Methods).

LFP spectral power provides sufficient information to decode location

We find that the LFP contains sufficient information to predict the virtual location of human subjects with above-chance accuracy (Fig. 1*f*; $t_{(19)} = -2.91, p = 0.009$, Cohen's $d = -0.65$; 10 log-spaced frequencies between 3 and 200 Hz). To determine whether decoding performance varied by MTL subregion or frequency band, we trained separate classifiers on microwire bundles located in the hippocampus ($N = 17$ subjects), parahippocampal gyrus ($N = 12$ subjects), or amygdala ($N = 12$ subjects). Similarly, we trained separate classifiers on low-frequency (<30 Hz) or high-frequency (>30 Hz) features only. When compared with chance, decoding was significant in the hippocampus ($t_{(16)} = -2.86, p = 0.01$, Cohen's $d = -0.69$; parahippocampal gyrus, $t_{(11)} = -1.90, p = 0.08$, Cohen's $d = -0.55$; amygdala, $t_{(11)} = -0.97, p = 0.35$, Cohen's $d = -0.28$) and for high-frequency features ($t_{(19)} = -3.50, p = 0.002$, Cohen's $d = -0.78$; low frequency, $t_{(19)} = -0.85, p = 0.40$, Cohen's $d = -0.19$). However, when directly assessing the effects of brain region and frequency band, there was no effect of brain region [$\chi^2(2) = 3.29, p = 0.19$; likelihood ratio test comparing mixed-effects models; see above, Materials and Methods] and no effect of frequency band ($t_{(19)} = 0.80, p = 0.43$, Cohen's $d = 0.25$) on classifier performance. We therefore focused all following analyses on the LFP-based decoding model trained on all frequencies and MTL subregions.

Having shown that we can predict the subject's location from LFP spectral power, we next asked whether the MTL represents locations in physical or in semantic space. We define the former using the locations of the landmarks (i.e., the target stores) in the virtual environment and the latter using their locations in word2vec space (Mikolov et al., 2013). Because all target stores had established semantic meaning, our model could achieve high performance based on either type of representation; the LFP could be representing either the subject's spatial distance or their semantic distance to the stores [e.g., “at the landmark with coordinates (x_i, y_i) , which is spatially proximate to the landmark with coordinates (x_j, y_j) ” vs “at the pizzeria, which is semantically similar to the bakery”].

Representation of virtual space strengthens over time

First, we examined representations of virtual space, using linear mixed effects models (see Materials and Methods). Specifically,

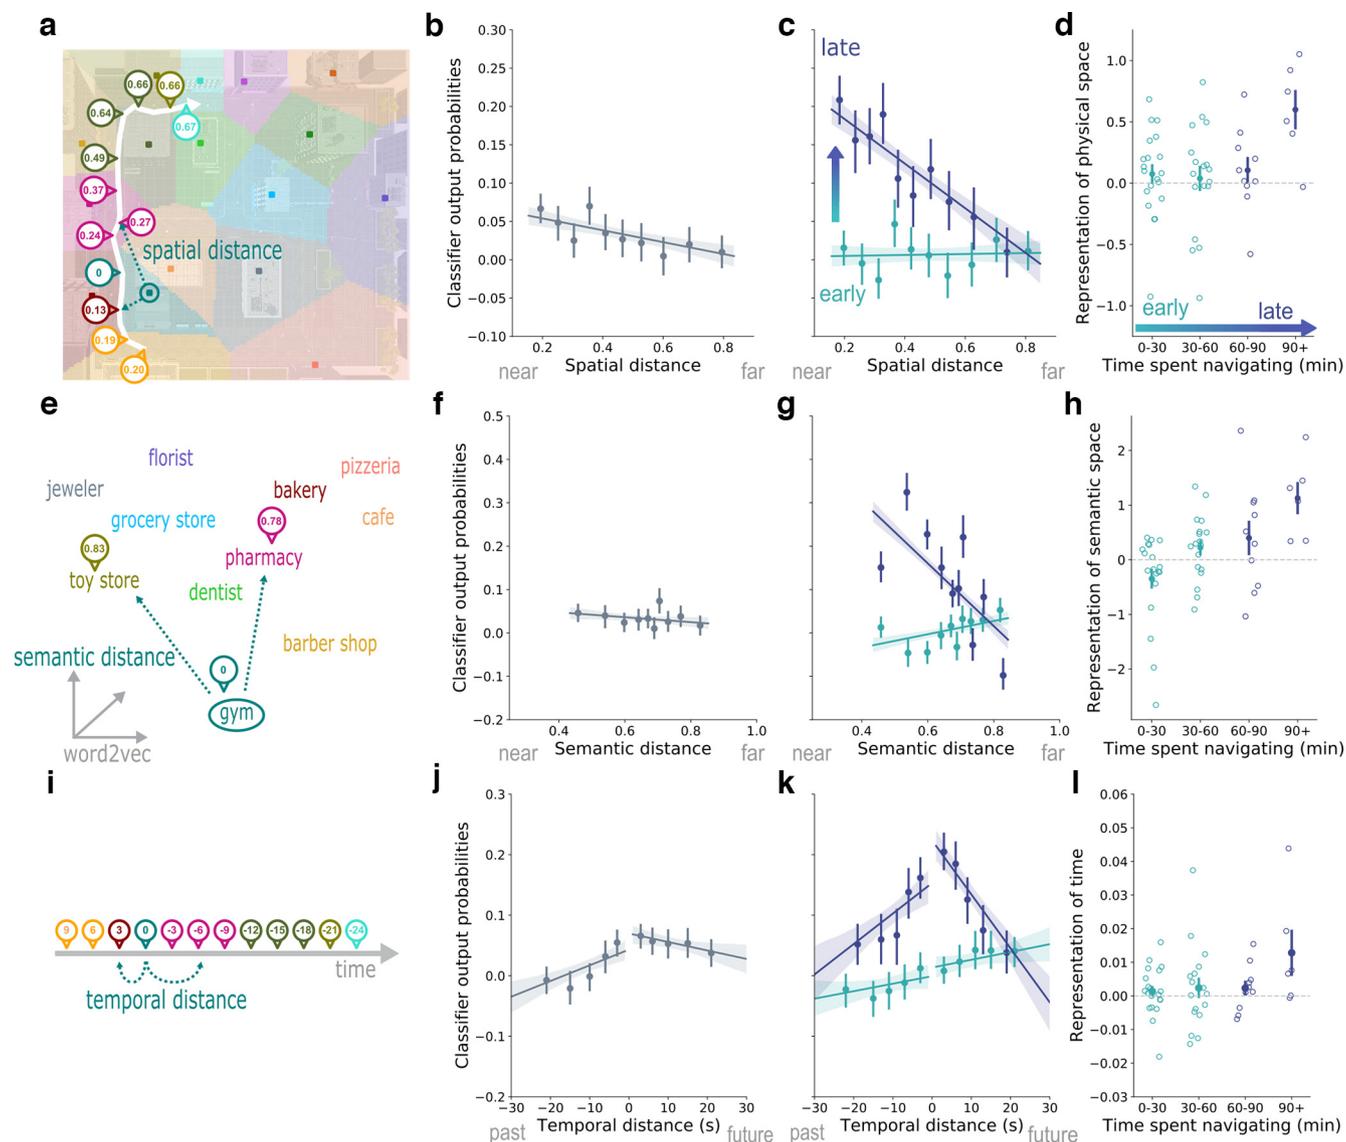


Figure 2. Representations of space, time, and semantics strengthen with experience. **a**, A subject's spatial distance from an example store, the gym, changes as the subject navigates. If the MTL represents spatial location, classifier output probabilities for a given store should decrease with the subject's distance to that store. **b–i**, We find this is the case; classifier output decreases with the spatial distance between a given store and the subject's true instantaneous position (**b**). Data are visualized after removing each subject's estimated random effect. Error bars (**b**, **c**, **f**, **g**, **j**, **k**) depict a 95% confidence interval. (**c**) The effect of spatial distance strengthens as subjects gain experience navigating. For visualization, we partition data into an early and a late phase at the 60 min mark. **d**, Showing the effect (i.e., slope) for four smaller time bins ($N = 20, 18, 10, 6$), reveals strengthened representations of space after 90 min. The graph shows inverted slopes estimated on the data from each bin. Positive values indicate higher classifier output for spatially proximate locations. The average slope for early and late time periods matches the slope of the regression line depicted in **c**. Error bars (**d**, **h**, **l**) depict SEM. **e**, Subjects move through semantic space as they navigate between stores. In this example, the subject is semantically closer (in word2vec space) to the gym when located at the pharmacy compared with the toy store. We assessed the effect of semantic distance on classifier output for each store. **f**, There was no overall effect of semantic distance on classifier output probability. **g**, **h**, The effect of semantic distance did, however, strengthen over time, suggesting that the MTL increasingly represents the subject's location in a task-relevant semantic space. **i**, Finally, we asked how classifier output for a given store relates to the elapsed time 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). **j**, Classifier output is higher for stores that a subject is about to visit than for those that have been visited in the recent past. **k**, **l**, As with spatial and semantic distance, the effect of temporal distance increases over time, especially for stores on a subject's future path, indicating that with experience the MTL more strongly predicts a subject's future trajectory.

we analyzed the distribution of classifier output probabilities for all stores as a function of the subject's instantaneous spatial distance to those stores (Note that this analysis included only incorrect stores, thus excluding the one store closest to the subject's instantaneous position.). If the MTL represents spatial relations, the classifier should assign high probabilities to spatially proximate locations and low probabilities to spatially distant locations (Fig. 2a; Note that this is independent of classifier performance, which is calculated based on the probabilities assigned to the correct store.). 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 [Fig. 2b; $z = -4.00$, $\chi^2(1) = 15.99$, $p < 0.001$].

In our task, spatial and temporal distance were correlated ($r_{\text{spatial,temporal}} = 0.71$), but semantic distance was uncorrelated with both of them ($r_{\text{semantic,spatial}} = 0.05$, $r_{\text{semantic,temporal}} = 0.03$). To control for the positive correlation between spatial and temporal distance, we ran a second model in addition to the individual model reported above, that included both spatial and temporal distance as a predictor. Although this approach slightly complicated the interpretation of null effects (they may be true

null effects or observed because of the shared variance in space and time), we reasoned that it provides a conservative means to assess the independent contributions of each predictor. If spatial distance explains variability in classifier output independent of time, then it should remain significant in the joint model. Indeed, the effect of spatial distance was significant in both models [joint model, $z = -2.75$, $\chi^2(1) = 7.54$, $p = 0.006$]. Together, these results indicate that the LFP represents the subject's spatial location within a broad spatial map of the environment. Moreover, this effect remained significant even after accounting for correlations between movement through time and space.

Next, we examined whether the neural representation of spatial relations exhibits a learning effect. Having shown that subjects navigate more efficiently as they acquire spatial knowledge of the environment, we expected this behavioral improvement to be reflected in a change of the representational structure in the MTL. We find that the effect of spatial distance on classifier output probabilities strengthens over time, as shown by a significant interaction of spatial distance and epoch number [individual, $z = -7.23$, $\chi^2(1) = 52.32$, $p < 0.001$; joint, $z = -3.61$, $\chi^2(1) = 13.04$, $p < 0.001$; Epoch number is a continuous index from start to end of the experiment.]. This effect is illustrated in Figure 2c, which shows classifier output probabilities as a function of spatial distance for an early and a late phase of navigation, split at the 60 min mark. Here, early and late refer to the total time spent navigating from the beginning to the end of the experiment. Figure 2d displays the slope of this regression for four smaller time bins (30 min increments), revealing the strongest increase after 90 min of navigation. 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.

Representation of semantic space strengthens over time

We next asked whether MTL activity reflects the semantic structure of the environment. To this end, we calculated the distances in word2vec (Mikolov et al., 2013) semantic space between all target stores and assessed their relationship with the classifier's predictions (Fig. 2e). We find no overall effect of semantic distance on classifier output probabilities [Fig. 2f; $z = -1.54$, $\chi^2(1) = 2.37$, $p = 0.12$].

As with spatial distance, we assessed whether the relation between classifier output and semantic distance changes over time. Whereas we had a strong expectation that subjects' spatial knowledge would increase over time, one may argue that semantic information (e.g., the concept of a bakery) should have been firmly established through prior experience and would thus not be expected to change over the course of the experiment. However, it is possible that MTL neural activity adapts to the specific semantic subspace spanned by our task, reflecting semantic similarity on fewer relevant dimensions. To evaluate these alternative predictions, we asked whether the degree to which the MTL represents the subject's semantic distance to

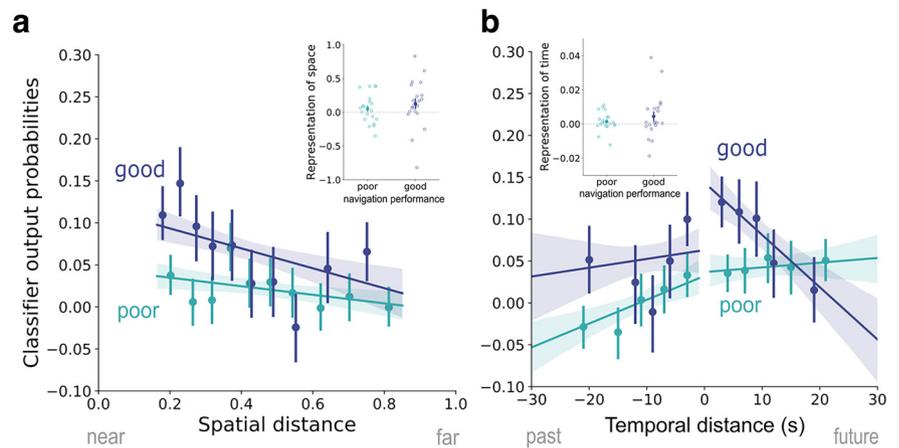


Figure 3. Representation of time linked to navigation performance. **a**, Representations of space (i.e., the effect of spatial distance on classifier output) may not only change over time but may also be linked to navigation performance (visualized here for good vs poor performance trials). We find that this is the case when assessing spatial distance individually but not when controlling for temporal distance. Data are visualized after removing the estimated random effect for each subject. Error bars indicate 95% confidence intervals. Inset, The average slopes across subjects \pm SEM. **b**, The interaction between navigation performance and temporal distance remains significant even after accounting for spatial distance, particularly for upcoming locations. The slope is steeper for good compared with poor performance trials, particularly for future locations (right). Although general knowledge of the spatial layout of the city enables spatial planning, it seems that the ability to predict upcoming locations on specific routes most closely relates to subjects' navigation performance. Inset, The average slopes across subjects \pm SEM.

each store changes over time (i.e., the distance between the nearest store and all other stores at each point in time). As with spatial information, we find that the link between semantic distance and classifier output strengthens over time [Fig. 2g,h; $z = -16.77$, $\chi^2(1) = 281.07$, $p < 0.001$]. This means that the MTL increasingly represents task-relevant semantic information the longer subjects have been navigating the virtual environment.

Representation of temporal sequence strengthens over time

Finally, we asked whether the LFP represents the temporal sequence of visited locations when controlling for the spatial distance between them. We expected that representations of stores would linger after they were visited and that the MTL would come to predict upcoming locations along a subject's future trajectory. To address this question, we assessed the link between classifier output for each store and the temporal distance between the current time point t_0 and the last time a subject had visited that store t_{0-j} (past) or the next time a subject was going to visit that store t_{0+i} (future; Fig. 2i). We find that temporal distance does affect classifier output, when assessed individually [Fig. 2j; $z = -3.01$, $\chi^2(1) = 9.06$, $p = 0.003$] but not when controlling for spatial distance [$z = 0.02$, $\chi^2(1) < 0.01$, $p = 0.99$]. However, we also find that the effect of time is asymmetric, in that future visits are more strongly represented than those in a subject's past [individual, $z = 5.77$, $\chi^2(1) = 33.24$, $p < 0.001$; joint, $z = 5.81$, $\chi^2(1) = 33.71$, $p < 0.001$]. So although there seems to be no fine-grained sequence information, the MTL broadly differentiates between upcoming and recently visited locations.

Although the MTL does not represent fine-grained sequence information when including all data, we asked whether sequence coding, and in particular the prediction of upcoming locations, gets stronger over time. Because subjects need to know the city's spatial layout before being able to plan routes and predict future states, we hypothesized that temporal sequence coding may emerge later in the experiment. We could confirm this prediction, in particular for locations on a subject's future trajectory, as

indicated by a three-way interaction of absolute temporal distance, epoch number, and future versus past visits [Fig. 2*k,l*; individual, $z = -4.20$, $\chi^2(1) = 17.63$, $p < 0.001$]. This effect holds when controlling for spatial distance [joint, $z = -4.67$, $\chi^2(1) = 21.78$, $p < 0.001$]. The MTL thus predicts upcoming locations more strongly, as subjects spend more time navigating the environment.

Effects of spatial, semantic, and temporal distance cannot be explained by artifacts

We conducted an additional set of analyses to ensure that artifacts, such as IEDs, did not affect our main results. We computed an artifact index, reflecting the percentage of artifactual data points for each epoch, averaged over microwire channels. This index was smaller than 1% for all subjects, ranging from 0.0003 to 0.8% (mean = 0.3%). We observed no correlation between the average artifact index for each subject and classifier performance ($r = 0.13$, $p = 0.59$), demonstrating that artifacts in the data did not systematically increase or decrease the classifier's ability to decode location.

We then asked whether artifacts may have inflated the effect of spatial, semantic, or temporal distance on classifier output probabilities. This may, for instance, be the case, if artifacts are temporally clustered and result in increased similarity in temporally proximate epochs. To rule out any such influence, we repeated the analyses assessing the effects of spatial, semantic, and temporal distance on classifier output while adding the epochwise artifact index and its interaction with the respective distance measure as a predictor. A significant interaction between artifact index and distance, with a higher artifact index being linked to stronger distance modulation, would suggest that artifacts contributed to the reported effects. We observed no such interaction for spatial [$z = 0.60$, $\chi^2(1) = 0.36$, $p = 0.55$], semantic [$z = -0.82$, $\chi^2(1) = 0.67$, $p = 0.41$], and temporal [$z = -1.36$, $\chi^2(1) = 1.85$, $p = 0.17$] distance, demonstrating that our findings cannot be explained by the presence of IEDs or other artifacts.

Temporal sequence coding is linked to navigation performance

Having shown that representations of time and space get sharper over time, we wondered whether variance in the sharpness of these representations may be linked to subjects' navigation performance. Do subjects navigate more efficiently when MTL activity reflects the spatial layout of the environment and the particular temporal sequence of visited locations? We find that this is the case. While controlling for time on task (i.e., epoch number), we observe a significant interaction between spatial distance and excess path ratio on classifier output probabilities [$z = 2.06$, $\chi^2(1) = 4.25$, $p = 0.04$], suggesting that the MTL represents the city's spatial layout more strongly at times when subjects navigate to their goal efficiently. Similarly, we find a significant interaction between temporal distance and excess path ratio on classifier output [$z = 4.17$, $\chi^2(1) = 17.34$, $p < 0.001$], suggesting that temporal sequence is represented more strongly during deliveries with high navigation performance. As expected, we observed no such effect for semantic representations, which were equally strong on trials with good and poor performance [$z = 0.37$, $\chi^2(1) = 0.13$, $p = 0.71$]. Again, we repeated the analyses for time and space, including the other factor, respectively, to see whether temporal and spatial distances explain independent variance. When controlling for temporal distance, the effect of spatial distance disappears [$z =$

-1.18 , $\chi^2(1) = 1.39$, $p = 0.24$]. The effect of temporal distance, however, withstands correction for spatial distance [$z = 3.82$, $\chi^2(1) = 14.59$, $p < 0.001$] and is stronger for locations that lie in the future compared with the past [three-way interaction, $z = 1.69$, $\chi^2(1) = 4.24$, $p = 0.04$]. Together, these results suggest that the MTL function that is most strongly linked to subjects' navigation performance is the simulation of specific future paths.

Discussion

Here, we used a decoding-based approach to study neural representations of places, concepts, and temporal sequences. As a first step, we trained a location decoder on the micro LFP while subjects navigated a virtual city. We find that the spectral features of the micro LFP provide sufficient information to decode subjects' instantaneous location in the virtual city with above-chance accuracy. As such, our results provide the first demonstration of location decoding from local field potentials in the human MTL, and they are in line with a prior study in rodents showing that hippocampal theta oscillations can predict rat position (Agarwal et al., 2014). Because research suggests that the human hippocampus does not exhibit theta oscillations in the same continuous way as that of rodents (Watrous et al., 2013; Aghajian et al., 2017), and since the number of micro-electrodes is limited by clinical considerations, we defined our feature set more broadly including multiple MTL subregions (hippocampus, parahippocampal gyrus, and amygdala) and frequency bands. We find that when working with subsets of the data classification did not differ between MTL regions or frequency bands. Similar findings have been obtained with human single-unit recordings, showing that place-selective neurons are not confined to the hippocampus but can also be observed in the parahippocampal gyrus and amygdala (Jacobs et al., 2013; Miller et al., 2013). It is, however, possible that the regional nonspecificity can partly be explained by the difficulty of precisely localizing microwire electrodes in human subjects.

We used our trained decoding model to characterize the representations that underlie successful decoding. As subjects move around the virtual environment, navigating from store to store, they change their virtual location, but because all stores have pre-established semantic meaning, they also change their semantic location (e.g., activating the concept of bakery). Our decoding model has been trained in a way that is blind to this distinction (i.e., using the stores as labels for different parts of the environment), meaning that it could achieve above-chance performance based on either type of representation, spatial or semantic. By analyzing the classifier output probabilities for each store at every time point, we found that the MTL represents locations in both virtual space (similar neural representations for spatially proximate stores) and semantic space (similar neural representations for semantically similar stores).

Considering virtual space, we show that the classifier output probabilities depend on the spatial distance between a subject's true location and each store. The classifier, on average, assigns a high probability to the store that is closest to the subject's true location. For other locations, the classifier assigns decreasing probabilities to stores that are farther from the subject's true location. The representational space spanned by MTL spectral features, hence, mirrors the virtual space spanned by the landmarks in the city. These results resemble similar findings obtained with representational similarity analysis on fMRI data from the human hippocampus (Deuker et al., 2016). Moreover,

we observed evidence that the spatial coding effect strengthens over time, providing a window into the learning-related changes of representational structure in the MTL. The more time subjects have spent navigating the virtual city, the more strongly the MTL represents the city's spatial layout. Whereas place fields in rats foraging in small open environments are known to stabilize within a few minutes of entering the environment (Wilson and McNaughton, 1993; Frank et al., 2004), the changes we observed here occurred over a duration of more than an hour. This time scale is plausible given the complexity of the environment and the fact that we observed concurrent increases in subjects' spatial navigation performance.

We further assessed whether the MTL additionally tracks subjects' movement through semantic space. Whereas we observed no overall effect of semantic distance, we show that task-relevant semantic information in the MTL increases over time. For instance, while the subject is near the bakery, the classifier increasingly tends to assign high probabilities to related stores such as the pizzeria. This finding complements previous reports of concept coding in the human MTL (Quiroga, 2012; Constantinescu et al., 2016). However, it may be surprising in the context of our study, given that the concepts we used were familiar to subjects at the outset (e.g., subjects did not have to newly learn the association between bakery and pizzeria). So why did semantic information increase over time? Our results are consistent with the notion that concept coding in the MTL is dynamic and task-dependent (Bottini and Doeller, 2020; Theves et al., 2020). Specifically, we interpret the observed effect as reflecting an increased focus on particular aspects in high-dimensional semantic space. Each target store may be invoking many semantically associated concepts at the beginning of the experiment. As time progresses, subjects learn the set of concepts that is relevant in the current context. They thus begin to focus on the attributes of each concept that make it similar or dissimilar from other concepts in the set (e.g., if all stores were selling food, that attribute would not be helpful in discriminating between them and may be disregarded). At the level of MTL neural activity, this may explain why similar concepts become more similar, and dissimilar concepts become even less similar (Bottini and Doeller, 2020).

Finally, we assessed whether the MTL represents the temporal sequence of visited locations. We find that the classifier assigns higher probabilities to stores in temporal proximity, suggesting that the MTL does represent the temporal sequence of events. Furthermore, probabilities are higher for stores in a subject's future than for those visited in the recent past. However, when we controlled for spatial distance, only this latter, broader signal remained significant. There are two explanations for the absence of an independent fine-grained effect of time. Spatial and temporal distance shared significant variance in our experiment, so controlling for space while assessing the effects of time and vice versa likely reduced the power to observe either effect. In addition, we show that the temporal effect gets stronger over time and thus likely depends on learning. The effect may therefore be small when considering all data. The learning effect was particularly strong in the forward direction, meaning that the prediction of future locations, more so than the lingering of locations in the recent past, is shaped by experience. The more knowledge subjects acquire about the spatial layout of the virtual city, the more they are capable of planning specific trajectories when looking for their target store, including the specific sequence of stores they will pass. And, intriguingly, it is this predictive signal that is linked to subjects' navigation

performance for individual deliveries, more so than it is a general representation of the spatial layout of the city. This finding implicates the MTL not only in the formation and storage of a cognitive map but also in accessing the map to predict future states.

Conclusions

We demonstrate that over time, field potentials in the human MTL come to jointly represent a subject's virtual spatial location, the semantics associated with that location, and a subject's temporal trajectory. Our findings indicate that the MTL holds a map-like representation of virtual and semantic space that is shaped by experience and can be used to predict future trajectories.

References

- Agarwal G, Stevenson IH, Berényi A, Mizuseki K, Buzsáki G, Sommer FT (2014) Spatially distributed local fields in the hippocampus encode rat position. *Science* 344:626–630.
- Aghajani ZM, Schuette P, Fields TA, Tran ME, Siddiqui SM, Hasulak NR, Tchong TK, Eliashiv D, Mankin EA, Stern J, Fried I, Suthana N (2017) Theta oscillations in the human medial temporal lobe during real-world ambulatory movement. *Curr Biol* 27:3743–3751.e3.
- Avants BB, Epstein CL, Grossman M, Gee JC (2008) Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med Image Anal* 12:26–41.
- Bellmund JLS, Gärdenfors P, Moser EI, Doeller CF (2018) Navigating cognition: spatial codes for human thinking. *Science* 362:eaat6766.
- Bottini R, Doeller CF (2020) Knowledge across reference frames: cognitive maps and image spaces. *Trends Cogn Sci* 24:606–619.
- Constantinescu AO, O'Reilly JX, Behrens TEJ (2016) Organizing conceptual knowledge in humans with a gridlike code. *Science* 352:1464–1468.
- Deuker L, Bellmund JL, Navarro Schröder T, Doeller CF (2016) An event map of memory space in the hippocampus. *Elife* 5:e16534.
- Eichenbaum H (2014) Time cells in the hippocampus: a new dimension for mapping memories. *Nat Rev Neurosci* 15:732–744.
- Frank LM, Stanley GB, Brown EN (2004) Hippocampal plasticity across multiple days of exposure to novel environments. *J Neurosci* 24:7681–7689.
- Gramfort A, Luessi M, Larson E, Engemann DA, Strohmeier D, Brodbeck C, Parkkonen L, Hämäläinen MS (2014) MNE software for processing MEG and EEG data. *Neuroimage* 86:446–460.
- Herweg NA, Sharan AD, Sperling MR, Brandt A, Schulze-Bonhage A, Kahana MJ (2020) Reactivated spatial context guides episodic recall. *J Neurosci* 40:2119–2128.
- Hsieh LT, Gruber MJ, Jenkins LJ, Ranganath C (2014) Hippocampal activity patterns carry information about objects in temporal context. *Neuron* 81:1165–1178.
- Jacobs J, Weidemann CT, Miller JF, Solway A, Burke JF, Wei X, Suthana N, Sperling MR, Sharan AD, Fried I, Kahana MJ (2013) Direct recordings of grid-like neuronal activity in human spatial navigation. *Nat Neurosci* 16:1188–1190.
- Kunz L, Brandt A, Reinacher PC, Staresina BP, Reifenschein ET, Weidemann CT, Herweg NA, Patel A, Tsitsiklis M, Kempter R, Kahana MJ, Schulze-Bonhage A, Jacobs J (2021) A neural code for egocentric spatial maps in the human medial temporal lobe. *Neuron* 109:2781–2796.e10.
- MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H (2011) Hippocampal "time cells" bridge the gap in memory for discontinuous events. *Neuron* 71:737–749.
- Manning JR, Lew TF, Li N, Sekuler R, Kahana MJ (2014) MAGELLAN: a cognitive map-based model of human wayfinding. *J Exp Psychol Gen* 143:1314–1330.
- Meisler SL, Kahana MJ, Ezzayat Y (2019) Does data cleaning improve brain state classification? *J Neurosci Methods* 328:108421.
- Mikolov T, Corrado G, Chen K, Dean J (2013) Efficient estimation of word representations in vector space. *arXiv:1301.3781*. <https://doi.org/10.48550/arXiv.1301.3781>.
- Miller JF, Neufang M, Solway A, Brandt A, Trippel M, Mader I, Hefft S, Merkwow MB, Polyn SM, Jacobs J, Kahana MJ, Schulze-Bonhage A (2013)

- Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science* 342:1111–1114.
- O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurol* 51:78–109.
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford: Clarendon.
- Quiroga RQ (2012) Concept cells: the building blocks of declarative memory functions. *Nat Rev Neurosci* 13:587–597.
- Reber TP, Bausch M, Mackay S, Boström J, Elger CE, Mormann F (2019) Representation of abstract semantic knowledge in populations of human single neurons in the medial temporal lobe. *PLoS Biol* 17:e3000290–17.
- Reddy L, Zoefel B, Pospel JK, Peters J, Dijksterhuis DE, Poncet M, van Straaten ECW, Baayen JC, Idema S, Self MW (2021) Human hippocampal neurons track moments in a sequence of events. *J Neurosci* 41:6714–6725.
- Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatry* 20:11–21.
- Solway A, Miller JF, Kahana MJ (2013) PandaEPL: A library for programming spatial navigation experiments. *Behav Res Methods* 45:1293–1312.
- Theves S, Fernández G, Doeller CF (2020) The hippocampus maps concept space, not feature space. *J Neurosci* 40:7318–7325.
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55:189–208.
- Tukey JW (1977) *Exploratory data analysis*. Reading, MA: Addison-Wesley.
- Tulving E (2001) Episodic memory and common sense: how far apart? *Philos Trans R Soc Lond B Biol Sci* 356:1505–1515.
- Umbach G, Kantak P, Jacobs J, Kahana M, Pfeiffer BE, Sperling M, Lega B (2020) Time cells in the human hippocampus and entorhinal cortex support episodic memory. *Proc Natl Acad Sci U S A* 117:28463–28474.
- Valente G, Castellanos AL, Hausfeld L, De Martino F, Formisano E (2021) Cross-validation and permutations in MVPA: validity of permutation strategies and power of cross-validation schemes. *Neuroimage* 238:118145.
- Watrous AJ, Lee DJ, Izadi A, Gurfkoff GG, Shahlaie K, Ekstrom AD (2013) A comparative study of human and rat hippocampal low-frequency oscillations during spatial navigation. *Hippocampus* 23:656–661.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.