

Parallel cognitive maps for multiple knowledge structures in the hippocampal formation

Xiaochen Y. Zheng^{1,*}, Martin N. Hebart^{2,3}, Filip Grill^{1,4}, Raymond J. Dolan^{5,6}, Christian F. Doeller^{2,7,8}, Roshan Cools^{1,9},
Mona M. Garvert^{2,10,11,12}

¹Donders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 EN, Nijmegen, the Netherlands,

²Max-Planck-Institute for Human Cognitive and Brain Sciences, 04103, Leipzig, Germany,

³Department of Medicine, Justus Liebig University, 35390, Giessen, Germany,

⁴Radboud University Medical Center, Department of Neurology, 6525 GA, Nijmegen, the Netherlands,

⁵Wellcome Centre for Human Neuroimaging, University College London, London WC1N 3AR, United Kingdom,

⁶Max Planck University College London Centre for Computational Psychiatry and Ageing Research, University College London, London WC1B 5EH, United Kingdom,

⁷Kavli Institute for Systems Neuroscience, Centre for Neural Computation, The Egil and Pauline Braathen and Fred Kavli Centre for Cortical Microcircuits, Jebsen Centre for Alzheimer's Disease, NTNU, 7491, Trondheim, Norway,

⁸Wilhelm Wundt Institute of Psychology, Leipzig University, 04109, Leipzig, Germany,

⁹Radboud University Medical Center, Department of Psychiatry, 6525 GA, Nijmegen, the Netherlands,

¹⁰Max Planck Research Group NeuroCode, Max Planck Institute for Human Development, 14195, Berlin, Germany,

¹¹Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Berlin, Germany,

¹²Faculty of Human Sciences, Julius-Maximilians-Universität Würzburg, Würzburg, Germany

*Corresponding author: Donders Centre for Cognitive Neuroimaging, Kapittelweg 29, 6525 EN, Nijmegen, The Netherlands. Email: xiaochen.zheng@donders.ru.nl

The hippocampal-entorhinal system uses cognitive maps to represent spatial knowledge and other types of relational information. However, objects can often be characterized by different types of relations simultaneously. How does the hippocampal formation handle the embedding of stimuli in multiple relational structures that differ vastly in their mode and timescale of acquisition? Does the hippocampal formation integrate different stimulus dimensions into one conjunctive map or is each dimension represented in a parallel map? Here, we reanalyzed human functional magnetic resonance imaging data from Garvert et al. (2017) that had previously revealed a map in the hippocampal formation coding for a newly learnt transition structure. Using functional magnetic resonance imaging adaptation analysis, we found that the degree of representational similarity in the bilateral hippocampus also decreased as a function of the semantic distance between presented objects. Importantly, while both map-like structures localized to the hippocampal formation, the semantic map was located in more posterior regions of the hippocampal formation than the transition structure and thus anatomically distinct. This finding supports the idea that the hippocampal-entorhinal system forms parallel cognitive maps that reflect the embedding of objects in diverse relational structures.

Key words: fMRI adaptation; hippocampus; relational knowledge; semantic representation.

Introduction

The hippocampal-entorhinal system builds rich models of the world, called cognitive maps, that account for the relationships between locations, events, and experiences (e.g. Tolman 1948; O'Keefe and Nadel 1978; Moser et al. 2008; Eichenbaum and Cohen 2014; Behrens et al. 2018). These maps capture the similarity between symmetric, high-dimensional relationships in a cognitive space, satisfying geometric constraints such as betweenness and equidistance (Gärdenfors 2004; Bellmund et al. 2018). Abstracting and organizing relational information in this way facilitates flexible behavior, enabling generalization and inference. Beyond classical findings on the importance of cognitive maps for spatial navigation (e.g. Burgess et al. 2002; Ekstrom and Ranganath 2018; O'Keefe and Nadel 1978), they are also thought to organize the relationships between objects (Constantinescu et al. 2016; Garvert et al. 2017; Garvert et al. 2023; Morton et al. 2020; Theves et al. 2019, 2020; Viganò et al. 2021), to represent temporal

distances (Bellmund et al. 2019; Bellmund et al. 2022; Burgess et al. 2002; Schapiro et al. 2012; Solomon et al. 2019), and to structure knowledge in the context of social cognition (Park et al. 2020; Son et al. 2021; Tavares et al. 2015). While cognitive mapping is thus proposed to be a universal, domain-unspecific coding principle to systematically organize knowledge (Stachenfeld et al. 2017; Behrens et al. 2018; Bellmund et al. 2018), it is unclear how the brain handles stimuli that are simultaneously embedded in multiple relational structures that are very distinct in terms of their mode and timescale of acquisition. Does the hippocampal-entorhinal system form one conjunctive map that integrates similarities along the different stimulus dimensions or does it form anatomically separable maps for each stimulus dimension?

In the study by Garvert et al. (2017), participants acquired new relational knowledge about everyday objects which were already linked by semantic connections. Here, participants were exposed to object sequences following a pseudo-random walk along a

Received: July 28, 2023. Revised: November 27, 2023. Accepted: November 30, 2023

© The Author(s) 2024. Published by Oxford University Press. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

graph. Within the hippocampal-entorhinal system, the similarity of neural object representations reflected the link distance between them on the graph (Garvert et al. 2017). In this situation, besides the newly learned transition structure between objects, participants can be assumed to have explicit knowledge about the semantic relationships between the same objects (e.g. rabbit and dog are both animals). Previous research has provided evidence that semantic relationships are represented in the hippocampus (Romero et al. 2019; Solomon et al. 2019; Pacheco Estefan et al. 2021) but also across various cortical regions (Charest et al. 2014; Clarke and Tyler 2014; Bracci et al. 2015; Price et al. 2015; Huth et al. 2016; Frisby et al. 2023).

Here, we ask whether prior semantic knowledge about objects would be simultaneously mapped in the same hippocampal system which also represents knowledge about transition structure. We reanalyzed the functional magnetic resonance imaging (fMRI) data from Garvert et al. (2017). Specifically, we constructed a model of object similarity that isolates the semantic relationships reflecting high-level conceptual knowledge acquired from experience from the low-level perceptual attributes of specific objects (Tversky 1977; Rosch and Lloyd 1978). To this end, we matched the stimuli used in Garvert et al. (2017) with photographs of the same objects and asked a separate participant population to assess their similarity using a triplet odd-one-out task (Hebart et al. 2020). Using fMRI adaptation analysis, we found evidence consistent with a map of semantic relationships between objects that was precisely localized in the hippocampus, alongside the previously identified map which coded for the newly learnt transition structure. Notably, although both map-like structures were represented in the hippocampal-entorhinal system, the semantic map was localized in more posterior regions than the transition structure. By showing that the hippocampal formation represents distinctive types of relationships simultaneously in parallel maps, our results thus demonstrate that the hippocampal formation does not construct conjunctive maps that integrate similarities across distinct stimulus dimensions. Instead, different stimulus dimensions are organized in anatomically separable maps, at least in situations where the mode and timescale of acquisition are very distinct.

Materials and methods

Experimental design

fMRI study

We reanalyzed the data from the fMRI study by Garvert et al. (2017), where 23 human participants (15 male, 8 female, $mean_{age} = 23.5$, $SD_{age} = 3.7$, age range 18–31) were tested.

On the day of training (day 1), participants were exposed to object sequences in an implicit learning task. The object transitions followed a pseudo-random walk along a graph (Fig. 1A) that was unknown to the participants. This means that each object could only be followed by an immediate neighbor in the graph structure. Participants performed a behavioral cover task, in which they learned to associate a random stimulus orientation with a specific button press. For example, the left-facing motorcycle was linked to button F, while the right-facing motorcycle corresponded to button J. The graph structure was the same for all participants. The link distance between any pair of objects in the graph is defined as the minimum number of links between this pair of objects (e.g. in the example displayed in Fig. 1A, the link distance between the rabbit and the leaf is two), which ranges from one to four. For each participant, a subset of 12 objects was selected from a total of 31 objects used in the study, and randomly

assigned to the 12 nodes on the graph. The objects covered a wide range of semantic categories (e.g. furniture, plants, body parts, animals; see Fig. 2B, top rows for the full set of objects used). Only one object within a semantic category was assigned to a participant (e.g. either banana or strawberry, but not both) and each participant was assigned a unique set of objects. Participants were trained for 12 blocks, with 132 transitions in each block.

In the scanning session (day 2), 7 out of the 12 training objects were used and presented in randomized order to reduce the total number of stimulus–stimulus transitions and thereby increase statistical power for the fMRI adaptation analysis. The transitions no longer followed the graph structure, but were pseudo-randomized in such a way that each possible stimulus–stimulus transition occurred exactly ten times per block (no stimulus repetitions). To reduce the motor responses in the scanner, a different behavioral cover task was employed that was orthogonal to the imaging analysis of interest: In 10% of the fMRI trials, participants performed an unrelated cover task, reporting whether a gray patch had been present on the preceding object (Fig. 1B). This means that participants were not required to pay active attention to the object identity. The fMRI session consisted of three blocks, with 420 transitions per block. Stimuli were presented for 1 s, with a jittered inter-trial interval generated from a truncated Poisson distribution with a mean of 2 s.

The study was in line with the Declaration of Helsinki and was approved by the University College London Hospitals Ethics Committee. All participants provided written consent.

Behavioral experiment of object similarity

Participants

A separate group of 128 human participants from the online crowdsourcing platform Amazon Mechanical Turk took part in a triplet odd-one-out task (55 male, 73 female, $mean_{age} = 42.6$, $SD_{age} = 11.9$, age range 20–70). All participants were located in the United States. The online research was approved by the Office of Human Research Subject Protection and all participants provided written informed consent. The participants were compensated financially for their time.

Stimuli and task

We used a triplet odd-one-out task on the 31 objects in the original study (Garvert et al. 2017, Fig. 2A) to measure the object similarity. By repeatedly varying the third object for a pair of target objects, their similarity could be assessed in a wide range of different contexts (Hebart et al. 2020). All images depict colored and shaded objects and were selected from the “Snodgrass and Vanderwart” database (Rossion and Pourtois 2004).

The task was carried out in sets of 20 trials. Participants could choose how many sets they would like to take part in. Participants engaged in a variable number of trials, ranging from a minimum of 20 to a maximum of 1460 (median = 50, 25th percentile = 20, 75th percentile = 145), with a median RT of 2221 ms (Supplementary Material S1). On each trial, participants were shown three object images side by side and were asked to select the image that was the least similar to the other two. Each object triplet and the order of stimuli were chosen randomly, but such that after collection of the entire dataset each cell in the 31×31 similarity matrix had been sampled at least once. The object similarity was defined as the probability $p(i,j)$ of the participants choosing objects i and j to belong together, irrespective of context imposed by the third object (Hebart et al. 2020).

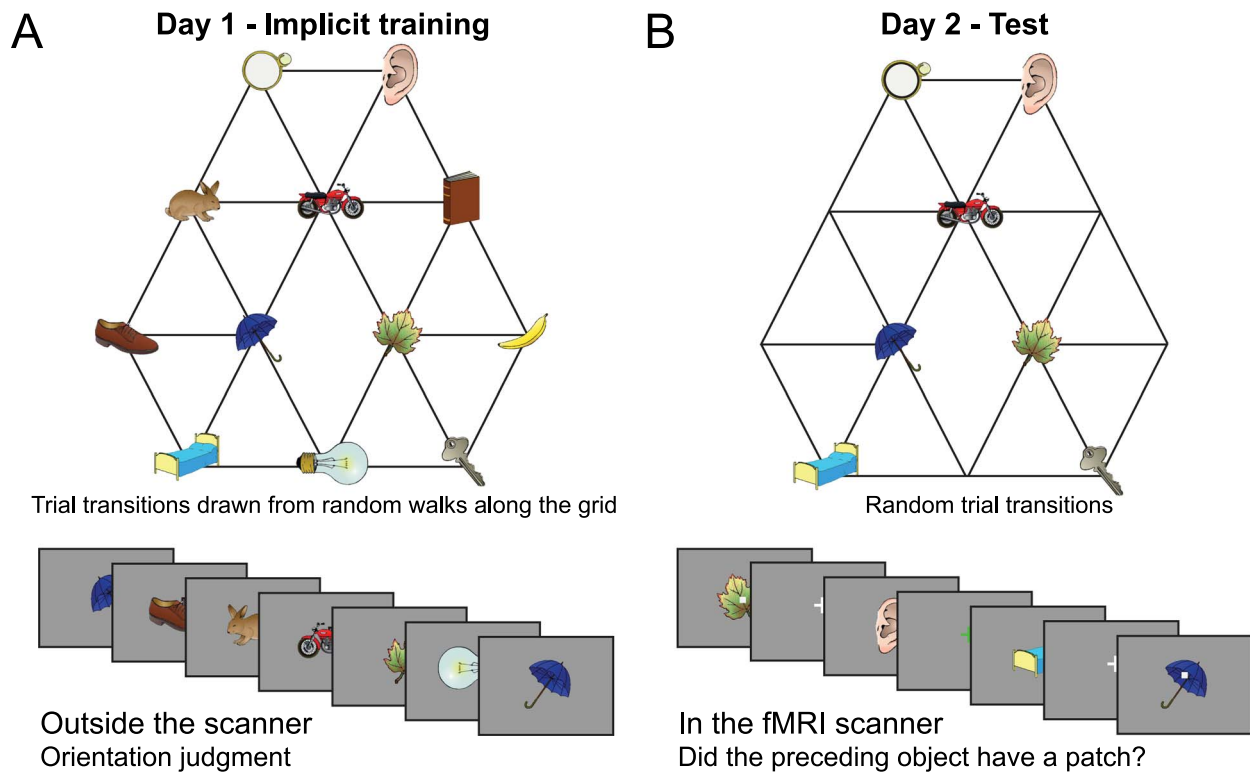


Fig. 1. Experimental design. (A) Graph structure used to generate stimulus sequences on day 1. Trial transitions were drawn from random walks along the graph. (B) Objects on reduced graph presented to participants in the scanner on day 2. Trial transitions were random. In both sessions, participants performed simple behavioral cover tasks. Figure adapted from Garvert et al. (2017).

Statistical analysis

Computation of the parametric regressors

For each participant, we computed a link distance matrix and two matrices describing object relations (i.e. semantic distance and residual distance). The semantic distance matrix and the residual distance matrix are derived from the object similarity matrices, with the latter computed directly from the triplet odd-one-out task described above (explained below).

Whereas the link distance matrix (values range from 1 to 4) was identical for all participants, the object similarity matrices were unique for each participant given that each participant received a different set of object stimuli in the training. To turn object similarities into a distance measure, we computed object dissimilarities by subtracting the similarity values from 1. Therefore, a number close to 1 means that two objects are dissimilar to each other, whereas a number close to 0 means the objects are very similar to each other.

To isolate semantic relationships, we made use of an independent similarity rating from a different dataset (Hebart et al. 2020). The rationale is that the part of variance that can be explained by an independent rating reflects semantic knowledge that is independent of the precise visual display of a particular object and instead reflects more abstract, semantic knowledge about stimulus relationships. The independent rating is based on 1854 images from the THINGS database (Hebart et al. 2019) which depicts photographs of objects embedded in a natural background, rated by a total of 5301 participants using the same triplet odd-one-out task. From the 1854 images, we selected 31 pictures that depict the same objects as the 31 objects in the original study of Garvert et al. (2017) and computed a sub-matrix of these 31 objects. Since each fMRI participant was trained on 12 out of the total 31 stimuli, we linearly regressed for each participant the 12×12 dissimilarity

matrix based on object images from the THING database (X in the regression below, Hebart et al. 2020) onto the 12×12 dissimilarity matrix based on object images from the original fMRI study (Y in the regression, Garvert et al. 2017). Both matrices were z-scored; therefore, no intercept was included in the regression:

$$Y \sim \beta * X + \text{residual}$$

We consider $\beta * X$ to reflect the variance of object dissimilarity that could be explained by the independent rating, likely to capture mostly semantic relationships. In contrast, the residual values reflect the variance that is not shared across different visual object stimuli, including specific perceptual features.

To visualize the object relatedness acquired from the triplet odd-one-out task, we performed multidimensional scaling (MDS) on the semantic and residual distance matrices. In the output MDS, objects are arranged in a two-dimensional space, where the Euclidean distances reflect the dissimilarities between objects as well as possible. Note that MDS can only be performed on matrices with positive entries. We therefore subtracted the minimum value of the matrix and added 1. The addition of constants does not affect the resulting visualization of distances.

fMRI adaptation analysis

Following the approach adopted in the original study, we exploited fMRI adaptation (Grill-Spector et al. 2006; Barron et al. 2016) to investigate the neural representations for transition structure and semantic similarities. fMRI adaptation relies on the observation that the repeated activation of the same population of neurons leads to a suppressed response. In this way, the amount of suppression can serve as a proxy for the similarity of the underlying neural representations.

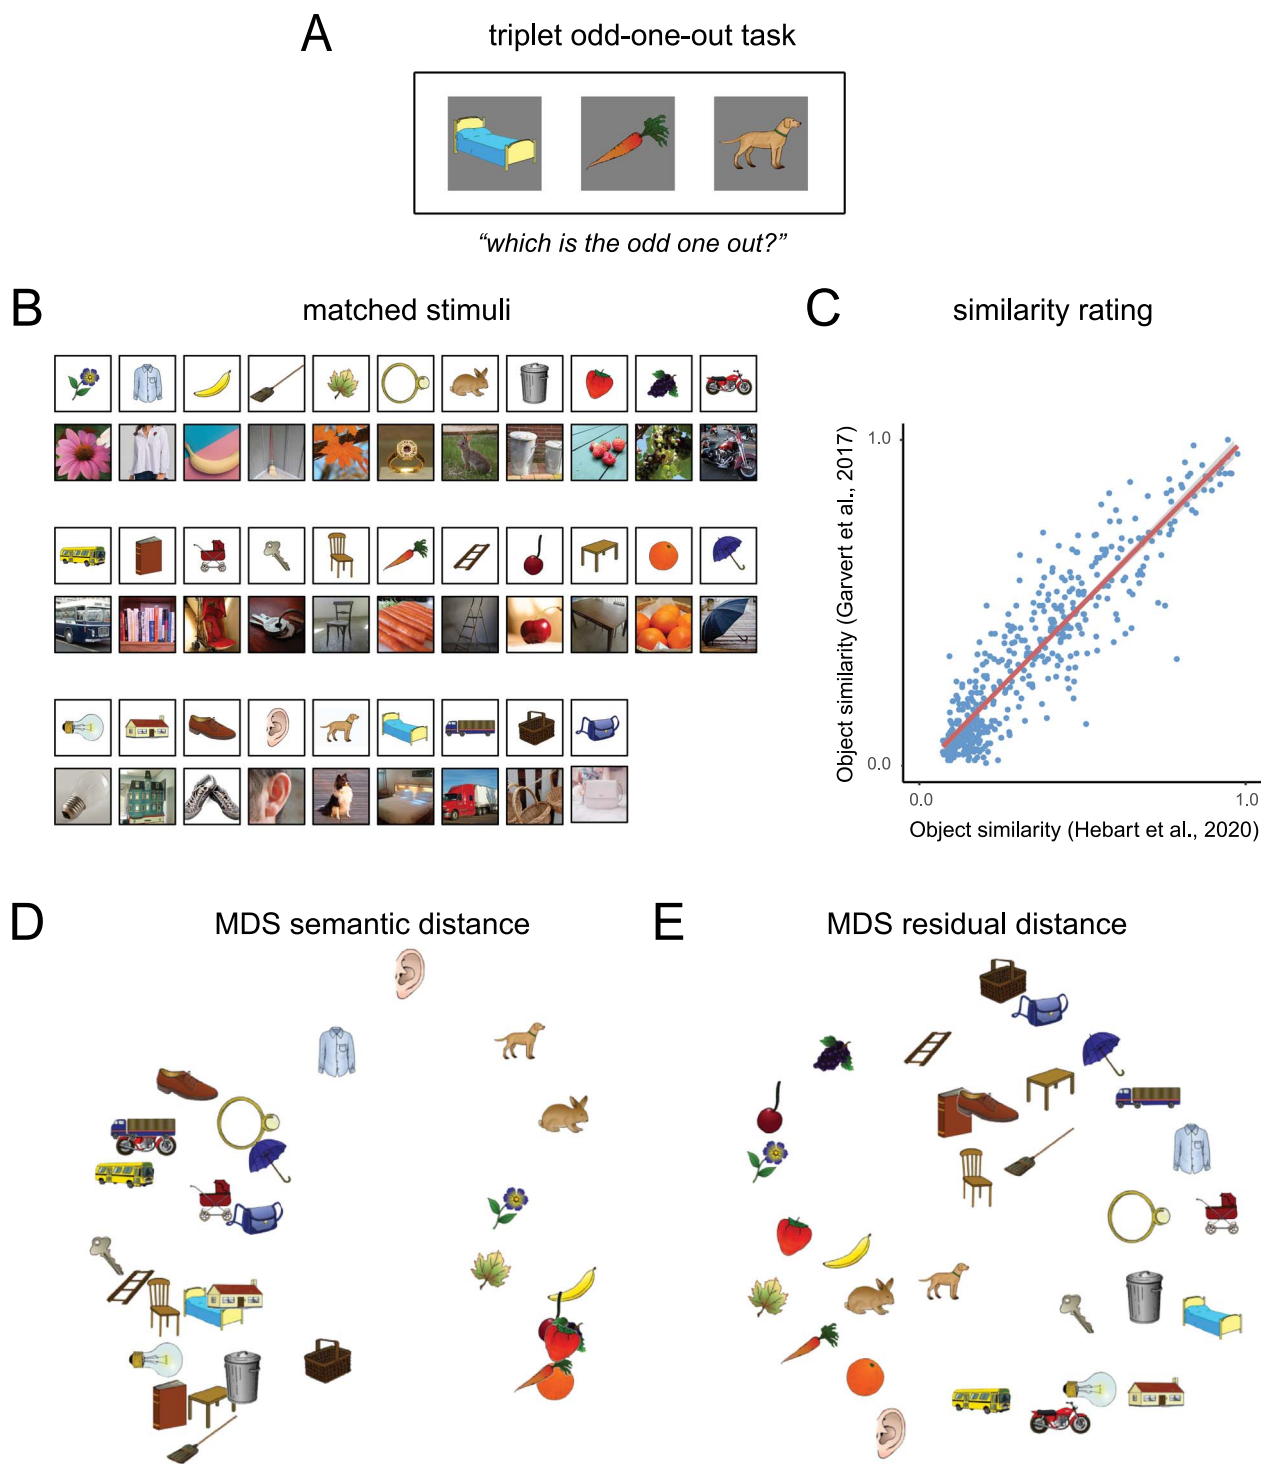


Fig. 2. Semantic distance constructed using the triplet odd-one-out task. (A) An example trial of the triplet odd-one-out task. The task measures object similarity as the probability of participants choosing two objects together, irrespective of the context imposed by the third object (Hebart et al. 2020). (B) Stimuli used in the odd-one-out task. Top rows: all 31 stimuli from the original study; bottom rows: a subset of stimuli from the THINGS database, matched with the 31 object stimuli used in the original study. The rating of the matched objects is done in the context of a total of 1854 objects (Hebart et al. 2020). (C) Correlation between similarity ratings based on our own stimuli and ratings based on the corresponding stimuli from the THINGS database (Spearman’s $\rho = 0.70$, $P < 0.001$). (D) Visualization of the 31 objects’ semantic distance in a two-dimensional space according to MDS. (E) 2D MDS visualization of the 31 objects’ residual distance.

We used an event-related generalized linear model (GLMs) to analyze the fMRI data. We included separate onset regressors for each of the seven objects with a patch and without a patch. Each onset regressor was accompanied by three parametric regressors describing the link distance, the semantic distance, and the

residual distance in relation to its previous object, respectively. All regressors were standardized in the GLM. No orthogonalization was applied. By analogy to the original analysis, we included a button press regressor as a regressor of no interest in the GLM. Trials associated with a button press and the two subsequent

trials were not included in the main regressors in order to avoid button press-related artifacts. All regressors were convolved with a canonical hemodynamic response function. In addition, we included the same 6 motion regressors and the 17 physiological regressors (10 for cardiac phase, 6 for respiratory phase, and 1 for respiratory volume) used in the original analysis were also included in the current GLM. Blocks were modeled separately within the GLM. Only the non-patch trials were included for our contrasts of interest.

The contrast images of all participants from the first level were analyzed as a second-level random effects analysis. We expected both the semantic information and the transition structure to be mapped in the hippocampal formation. Therefore, we focused our analysis on this region. The anatomical mask is created using FreeSurfer (Fischl 2012) segmentation in MNI space, combining bilateral hippocampus and bilateral entorhinal cortex (Supplementary Material S2). We consider our results significant if they survived family-wise error (FWE) correction at the peak-level of $P < 0.05$ within this anatomically defined mask (small volume correction, SVC). To explore the cortical semantic representation, we performed additional SVC using two anatomically defined masks: the left anterior temporal lobe and the left angular gyrus, two regions previously reported to be important in semantic processing (Visser et al. 2010; Humphreys et al. 2021). Both masks are defined using the Harvard-Oxford cortical structural atlas with a probabilistic threshold of 30% (Supplementary Material S2). Activations in other brain regions were only considered if they survived whole-brain peak-level FWE correction at $P < 0.05$. All statistical parametric maps visualized in the manuscript are thresholded at $P < 0.01$ uncorrected and unmasked for illustration.

To illustrate the non-overlapping clusters for the link distance effect and the semantic distance effect, we defined two functional regions of interest (ROIs) based on the two parametric estimations from non-patch trials in the main GLM. The link distance effect revealed a cluster in bilateral entorhinal cortex, which we used to define the entorhinal ROI; the semantic distance effect revealed a cluster in bilateral hippocampus, which we used to define the hippocampal ROI. For both ROIs, we included all the voxels exceeding a t -value of 2.5, corresponding to $P < 0.01$. From the two ROIs, we then extracted parameter estimates for each of the 23 participants for the two effects. Due to the statistical dependence between the data and the ROI definition, no statistical inference was made regarding the interaction. In order to demonstrate both effects as response gradients along the hippocampal long axis, we defined nine voxel-size ROIs along the left and right hippocampus, respectively. Group-level t -stats were extracted and plotting for these ROIs for link and semantic distances.

The statistical analysis was done using SPM12 (Wellcome Trust Centre for Neuroimaging, <http://www.fil.ion.ucl.ac.uk/spm>). Visualization was done using FSLeyes (Wellcome Centre for Integrative Neuroimaging, <https://git.fmrib.ox.ac.uk/fsl/fsleyes/fsleyes/>) and Connectome Workbench (<http://www.humanconnectome.org/software/connectome-workbench>).

Results

Human participants were trained on object sequences whose transition probabilities followed a discrete, non-spatial graph (Fig. 1A, Garvert et al. 2017). Without being consciously aware of the hidden graph structure, participants' neural activity in the hippocampus and entorhinal cortex reflected the transitional relationships they had experienced between the objects on a

subsequent day (Figs 1B, 3A). However, the brain may not only represent the newly learned transition structure, but also the semantic relationships between the objects. Thus, we asked whether this information is also mapped in the same system.

To address this question, we measured the semantic relationship between these objects using a triplet odd-one out task, where participants were shown three objects on each trial and asked to select the image that was the least similar to the other two (Fig. 2A; Hebart et al. 2020). To separate the semantic relationships between objects from the perceptual similarities, we matched the similarity rating of our objects used in the original study (Fig. 2B, top rows) to a separate rating of corresponding real-world photographs of the same objects in the THINGS database (Fig. 2B, bottom rows; Hebart et al. 2019). The THINGS database shows photographs of objects embedded in a natural environment as opposed to simple line drawings of stereotypical objects used in our fMRI study. In this way, the shared perceptual similarity between objects in these two data sets should be reduced (see Supplementary Material S3 for additional analysis on object similarity). Although rated by different groups of participants and in different contexts, the similarity ratings obtained for our objects and the corresponding objects in the THINGS database were highly correlated (Spearman's $\rho = 0.89$, $P < 0.001$, Fig. 2C, Supplementary Material S1). This suggests that semantic relations are preserved across datasets.

We regressed the matched similarity matrix (x-axis, Fig. 2C) onto our original similarity matrix (y-axis, Fig. 2C). By doing this, we were able to separate the variance into two parts: (i) the part that could be explained by an independent measure of object similarity obtained from a matched set of stimuli, and (ii) the part that could not be explained by the independent measure. Although semantic and perceptual features are unlikely to be fully separable in this way, we consider the first part to primarily reflect the semantic relationships between our objects that are preserved across different ways of visualizing objects; while the second, residual part reflected a combination of features that are not accounted for in terms of semantics, including perceptual similarities. The MDS of the semantic distance reveals that the similarity ratings led to the emergence of object category clusters (e.g. fruit, animals, man-made objects) and replicates well-known distinctions between “animate - inanimate” and “natural - man-made” (Hebart et al. 2020, Fig. 2D). The residual distance continued to express differences between man-made and natural objects, however the overall arrangement was less structured (Fig. 2E). Neither the semantic distance nor the residual distance was correlated with link distance (semantic: Spearman's ρ mean = 0.03, SD = 0.12, range = -0.25 - 0.30, $t_{22} = 1.04$, $P = 0.31$; residual: Spearman's ρ mean = -0.03, SD = 0.09, range = -0.22 - 0.15, $t_{22} = -1.52$, $P = 0.14$).

Following the approach adopted in the original study, we exploited fMRI adaptation to investigate the representational similarity for different objects. In line with the decrease in fMRI adaptation as a function of link distance observed in the hippocampal formation (Garvert et al. 2017) and previous work measuring semantic distance using fMRI (Bedny et al. 2008; Kim et al. 2009; Yee et al. 2010; Conca et al. 2021), we reasoned that in areas representing object relationships, objects that are closer to each other in the corresponding representational space should have a more similar representation and therefore fMRI adaptation should scale with the corresponding semantic distance measure. The results replicated our original link distance effect after accounting for the semantic distance and the residual distance. The fMRI adaptation analysis showed a cluster

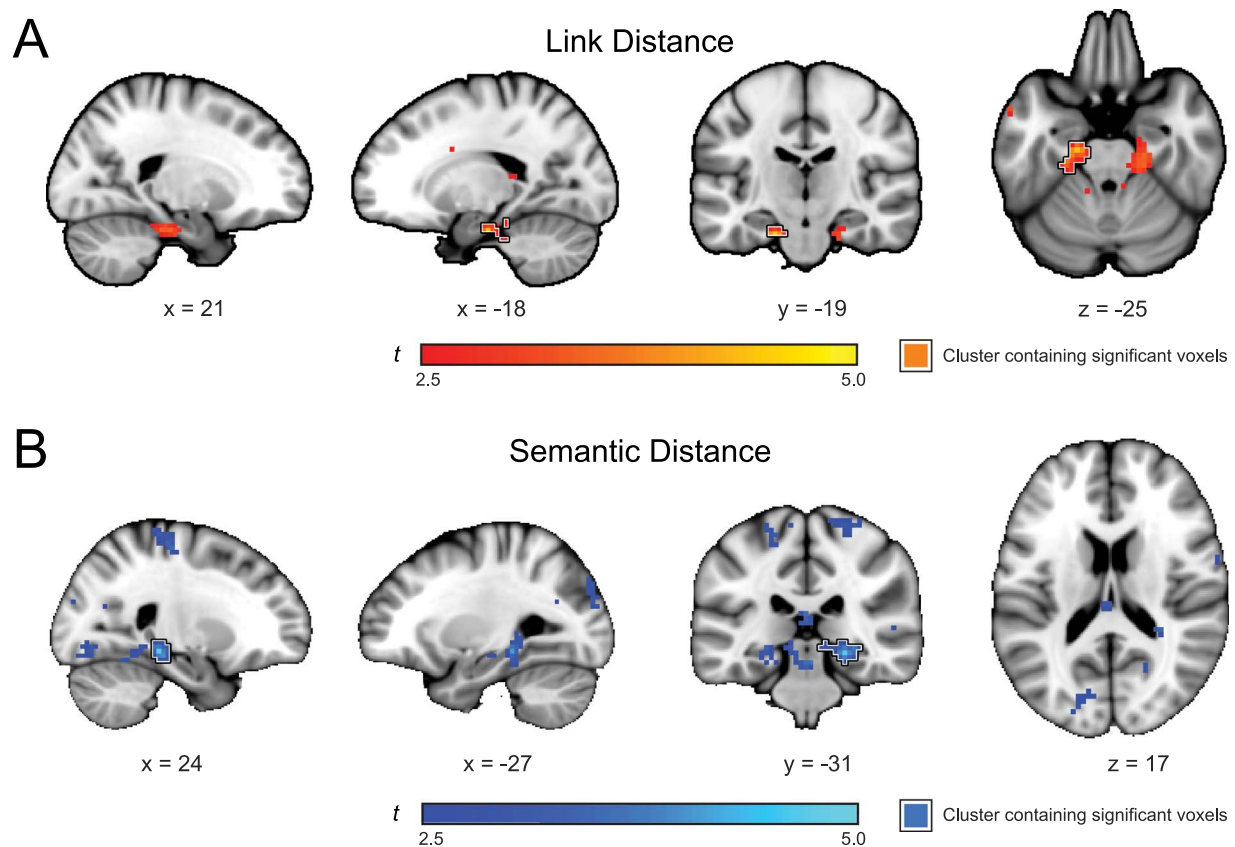


Fig. 3. Transition structure and semantic similarities are represented in the hippocampal-entorhinal system. (A) Whole-brain analysis showing a decrease in fMRI adaptation with link distance in the hippocampal formation, when link distance, semantic distance and residual distance are included in the model. (B) Whole-brain analysis showing a decrease in fMRI adaptation with semantic distance in the hippocampal formation, when link distance, semantic distance and residual distance are included in the model. Both (A) and (B) are thresholded at $P < 0.01$, uncorrected for visualization. The clusters containing voxels surviving correction for multiple comparisons (FWE, $P < 0.05$) are highlighted in solid black lines.

bilaterally in the entorhinal cortex (Fig. 3A; FWE corrected at peak level, peak $t_{22} = 4.44$, $P = 0.042$, $[-18, -19, -25]$). Critically, we also observed a semantic distance effect in the bilateral hippocampus (Fig. 3B; peak $t_{22} = 4.69$, $P = 0.028$, $[24, -31, -10]$). No other regions showed fMRI adaptation effects as a function of semantic distance (all P s > 0.73), including the left anterior temporal lobe (see Visser et al. 2010, for a meta-analysis) and the left angular gyrus (see Humphreys et al. 2021, for a recent review), two ROIs that have previously been associated with semantic processing. This suggests that the reported semantic effect is specific to the hippocampal formation. No brain region showed fMRI adaptation effects covarying with residual distance (all P s > 0.99 , no suprathreshold cluster using SVC).

While the link distance and the semantic distance are both represented in the hippocampal formation, they were located in two non-overlapping clusters (Fig. 4A). To investigate this at a more fine-grained level, we defined two functional ROIs: (i) a bilateral ROI in the entorhinal cortex defined by the link distance effect (EC ROI), and (ii) a bilateral ROI in the hippocampus defined by the semantic distance effect (HC ROI). Figure 4(B) shows parametric estimates of the semantic and the link distance effects extracted from the ROIs defined by the opposite contrast, respectively. The individual semantic distance effect extracted from the EC ROI and the link distance effect extracted from the HC ROI were not significantly different from 0 (semantic: $t_{22} = -0.25$, $P = 0.81$; link: $t_{22} = 0.54$, $P = 0.60$). To better understand the relative localization of the link distance effect and the semantic distance effect, we visualized the difference between the two contrasts

in the hippocampal formation (Fig. 4C) and extracted t-values from spherical ROIs located along the hippocampal axis (Fig. 4D). These analyses demonstrate that the semantic similarity effect is localized in more posterior regions of the hippocampal formation, whereas the transition structure effect resides in more anterior regions. This difference, found in both hemispheres, suggests the existence of a posterior–anterior gradient along the hippocampal long axis (Poppenk et al. 2013). This effect is particularly pronounced in the right hemisphere where peaks do not overlap. This suggests that the brain does not integrate the two relational structures into one conjunctive map, but instead forms separate relational structures.

Together, our results suggest that both recently learned transition structure of which participants have no conscious awareness, as well as semantic relationships that are explicitly accessible and acquired over the course of a lifetime, are represented in the hippocampal formation simultaneously, albeit in different subregions: While transition structures are represented in more anterior hippocampal regions, semantic relationships are represented in more posterior regions.

Discussion

The brain forms cognitive maps of the relationships between landmarks that help an animal navigate their physical environment (Tolman 1948; O'Keefe and Nadel 1978; Burgess et al. 2002; Ekstrom and Ranganath 2018). Previous studies have shown that the same organizing principle also applies to other




Fig. 4. Anatomical localization of transition structure and semantic similarities. (A) Left: link distance (red) and semantic distance (blue) are represented in non-overlapping clusters (thresholded at $P < 0.01$, uncorrected). Two ROIs were defined based on the link distance effect (in red) and the semantic distance effect (in blue) and included voxels exceeding a cluster-defining threshold of $P < 0.01$, uncorrected (both ROIs highlighted in solid lines). Right: boxplot of the parameter estimates for the link distance and semantic distance effects extracted from these two ROIs. The thick horizontal line inside the box indicates the median, and the bottom and top of the box indicate the first and third quartiles of each condition. Each dot represents one participant. The plot is for visualization only, since the contrast used for defining the ROIs is not independent from the interaction effect of interest here. (B) Anatomical location where the link distance is represented more strongly (red) versus where the semantic distance is represented more strongly (blue). The analysis is restricted to the hippocampal formation (incl. hippocampus and entorhinal cortex). (C) Visualization of response gradient along the hippocampal long axis. In both the left and the right hippocampus, the semantic distance peaks at more posterior locations compared to the link distance.

non-spatial types of relational information (Constantinescu et al. 2016; Garvert et al. 2017; Theves et al. 2019, 2020; Morton et al. 2020; Viganò et al. 2021; Garvert et al. 2023). For example, when participants acquire new knowledge about the relationships between objects by being exposed to experimentally generated object sequences, the hippocampal formation extracts the associated transition structure and stores it as map-like structural representations (Garvert et al. 2017). However, participants also already have existing knowledge about the semantic relationships between these objects acquired over an entire lifetime. Here we show that this prior semantic knowledge is also simultaneously mapped in the same neural system that codes for newly learned structural information. Specifically, we observed that repetition suppression of signals in the hippocampus scales with semantic distance. This representation aligns with the defining features of a cognitive map: relationships can be quantified in terms of a metric; this metric is symmetric and it adheres to geometric norms (Gärdenfors 2004; Gärdenfors and Zenker 2015; Bellmund et al. 2018). Not only are both knowledge structures mapped in the hippocampal-entorhinal system, they also both adhere to geometric coding principles whereby similar states are represented more similarly. This suggests that different types of relational knowledge, regardless of whether that knowledge was gathered over short durations or over a lifetime, might be structured within a similar cognitive mapping framework in the hippocampus.

Cognitive maps have been proposed to be an organizing principle that underlies our ability to generalize and make inferences (Behrens et al. 2018). The representation of both graph structure and semantic relationships in the same system is remarkable, given their very different timescales and modes of acquisition. However, while the two relational structures were represented in the same neural system, they were only partially represented in overlapping voxels. This suggests that the brain extracts separable relational structures in parallel rather than integrating them into one compositional map (Spiers 2020). Parallel representations of separable maps likely facilitate generalization and inference in an ever-changing environment where the relevance of stimulus dimensions can shift rapidly. When different stimulus dimensions become relevant at different times, the parallel coding of multiple knowledge structures allows for flexible selection of relevant information. Such cognitive computations enable the hippocampus to adaptively generalize based on task demands (Garvert et al. 2023) and to guide goal-directed behavior in novel situations (Whittington et al. 2020). Additionally, attention can be selectively allocated to relevant state representations (Radulescu et al. 2021) and multiple relational structures can be flexibly combined into more complex compositional structures for generalization (Saanum et al. 2021).

Our finding that object relationships are represented in the hippocampus is also consistent with previous findings that the hippocampus is involved in the retrieval of semantic memory, particularly for relational knowledge between concepts, and that hippocampal activity reflects distances in semantic spaces (Pacheco Estefan et al. 2021; Romero et al. 2019; Solomon et al. 2019). The hippocampus thus seems to support domain-general processing of semantic knowledge (Staresina et al. 2011; Ranganath and Ritchey 2012; Morton et al. 2021). By decoding the semantic map without explicitly manipulating the semantic similarity of the stimuli or demanding this information in the task, we demonstrate how prevalent the representation of relational information is in the hippocampus.

However, previous research on semantic representations often shows the additional involvement of a broader set of brain areas,

including cortical regions such as the anterior temporal lobe, the angular gyrus, the inferior frontal gyrus and the fusiform gyrus (Bracci et al. 2015; Charest et al. 2014; Clarke and Tyler 2014; Huth et al. 2016; Price et al. 2015; Tucciarelli et al. 2019; see Frisby et al. 2023 for a recent review). These regions are thought to be involved in various aspects of semantic processing, such as semantic categorization, semantic retrieval, and the integration of semantic and perceptual information (Binder et al. 2009; Bookheimer 2002; Lambon Ralph et al. 2017; Visser et al. 2010). Importantly, none of these functions were task-relevant in our study. Participants were not even required to pay attention to the objects, as they only had to attend to the presence of a gray patch on the screen. It is likely that the other regions that are involved in processing of semantic knowledge only become involved in situations where semantic knowledge is more relevant to task performance (see also Martin et al. 2018). It is also worth noting that object-specific semantic representations have been identified previously in the perirhinal cortex (Clarke and Tyler 2014), a cortical region close to the hippocampal formation. However, due to fMRI signal drop-out in this region, we could not examine whether our effects of interest are also represented there.

Notably, we found an anatomical gradient along the anterior-posterior axis of the hippocampus (Poppenk et al. 2013; Strange et al. 2014), with the graph structure represented in more anterior parts of the hippocampal formation and the semantic map in more posterior parts. Alternatively, this could also be viewed as two functionally dissociated clusters, with the cluster residing in the entorhinal cortex encoding statistical information about transition structures and the cluster in hippocampus encoding semantic similarities between specific objects. Distinct functional clusters would suggest more specialized processing within the hippocampus, suggesting that different types of knowledge are more rigidly localized, perhaps facilitating categorization of information for more systematic retrieval. A gradient on the other hand suggests a more integrated and potentially overlapping functionality within the hippocampus, perhaps facilitating processing in ambiguous situations and retrieval of information in context-rich situations. Due to spatial correlations inherent to fMRI data, it is not possible to completely disentangle a gradient from two separable clusters. Future studies, potentially employing higher-resolution fMRI or intracranial recordings, can provide more definitive answers.

In either case, the anatomical segregation of the two maps may reflect differences in the nature of the underlying knowledge structures (Peer et al. 2021). The semantic relationships may reflect taxonomic knowledge derived from shared features and properties between objects that participants formed over their lifetimes. The transition structure on the other hand could stem from recent associative learnings. Our findings are thus consistent with the idea that structural information about statistical regularities is encoded separately from semantic similarities between specific objects (Whittington et al. 2020). Several previous investigations should also be mentioned here. Leshinskaya and Thompson-Schill (2020) suggested that perceptual features, newly acquired associations as well as generalizable relational knowledge manifest in neighboring regions of the lateral temporal areas. However, in contrast to our own observations, the authors did not find any evidence of associative coding in medial temporal lobes or the hippocampus. In addition, Mirman et al. (2017) report a neural dissociation between taxonomic and thematic semantics across a set of studies (e.g. Schwartz et al. 2011; Davey et al. 2016; Kalénine and Buxbaum 2016). These studies suggest that anterior temporal lobes (ATL) predominantly encode taxonomic

semantic knowledge and the temporo-parietal cortex (TPC) encodes thematic semantic processing. However, the literature on this neural observation is by no means conclusive, and many studies, including our own, do not echo this ATL-TPC dissociation.

Our findings suggest a more integrative role for the hippocampus, accommodating various types of relational knowledge, both taxonomic (semantic) and associative/temporal (transition structure) (Peer et al. 2021), underscoring the dynamic and flexible nature of hippocampal codes. This is further supported by the anatomical gradient reminiscent of the gradient observable in the scale of hippocampal spatial codes, where anterior parts of the hippocampus display coarser spatial codes than posterior parts of the hippocampus (Poppenk et al. 2013; Strange et al. 2014; Brunec et al. 2018). This hints at broader organizational principles within the hippocampus.

The anatomical separability we report could also be attributable to the temporal disparity in the acquisition and consolidation of semantic relationships versus newly learned relations. Semantic relationships, built and reinforced over a lifetime, have undergone extensive consolidation processes, perhaps resulting in more stable and distinct neural representations within the hippocampus. In contrast, relationships acquired over a short duration, such as those from a single training session, might still be in the early phases of consolidation (Walker and Stickgold 2004; Squire et al. 2015). In short, several features differ between the two relational structures in our study, including the recency of learning, the nature of the type of relational knowledge, and the degree to which knowledge is explicit or implicit. The observed spatial segregation in the hippocampus is likely driven by a combination of these features, potentially reflecting the nature of the encoded information.

It is also worth noting that there are other plausible measures that might better characterize the neural representation of the transition structure, which are discussed comprehensively in Garvert et al. (2017). In the domain of reinforcement learning, the utility of a cognitive map is greatly enhanced when the representation of a state not only embodies the present but is also predictive, encompassing a spectrum of probable future states. This concept is encapsulated in the successor representation (Dayan 1993; Momennejad et al. 2017; Russek et al. 2017), which is suggested to be encoded by hippocampal place cells (Stachenfeld et al. 2014, 2017). From this perspective, hippocampal place cells are posited to encode not the immediate location of an animal, but a predictive array of forthcoming locations. Such a representation is advantageous for reinforcement learning, as it amalgamates predictive insights of future states with reward information, thereby facilitating the swift computation of potential navigational paths (Baram et al. 2018; Dayan 1993; Momennejad et al. 2017; Russek et al. 2017). Analogous to the successor representation, graph theory introduces the matrix resolvent as a means to quantify “communicability” or the closeness between nodes. Similarly, the matrix exponential, another graph theory measure, computes a weighted summation over future states and exhibits versatility across various dimensions and contexts (Estrada and Hatano 2008). Both the successor representation and these graph-theoretic measures explain the fMRI adaptation effects observed by Garvert et al. (2017). Nonetheless, disentangling their unique neural contributions presents a challenge, primarily due to the high intercorrelations among these distinct distance metrics.

In sum, our study shows that the hippocampal-entorhinal system encodes diverse relational structures in which a stimulus is embedded. Both the semantic relationships and transition

structures are represented simultaneously but with distinct spatial organization, even when neither structure is task relevant. This enables flexible selection of relevant knowledge in order to guide goal-directed behavior in novel situations (Behrens et al. 2018; Spiers 2020; Whittington et al. 2020).

CRediT statement

Xiaochen Zheng (Conceptualization, Data curation, Formal analysis, Investigation, Project administration, Visualization, Writing—original draft), Martin Hebart (Formal analysis, Investigation, Methodology, Visualization, Writing—review and editing), Filip Grill (Formal analysis, Writing—review and editing), Raymond Dolan (Funding acquisition, Writing—review and editing), Christian Doeller (Funding acquisition, Writing—review and editing), Roshan Cools (Conceptualization, Funding acquisition, Supervision, Writing—review and editing), Mona Garvert (Conceptualization, Formal analysis, Investigation, Methodology, Visualization, Writing—original draft).

Supplementary material

Supplementary material is available at *Cerebral Cortex* online.

Funding

This study was supported by the Dutch Research Council NWO (Gravitation grant to the Language in Interaction Consortium, grant number: 024.001.006) and the Max Planck Society (including a Max Planck Research Group Grant to MNH, grant number: M.TN.A.NEPF0009).

Conflict of interest statement: None declared.

Data and code accessibility

Data and codes of the original study are available on datadryad (<https://doi.org/10.5061/dryad.nk08s>).

Additional data and codes for the current follow-up study are available at the Donders Repository (<https://doi.org/10.34973/8m6q-qj39>).

References

- Baram AB, Muller TH, Whittington JC, Behrens TE. Intuitive planning: global navigation through cognitive maps based on grid-like codes. *BioRxiv*. 2018:421461. <https://doi.org/10.1101/421461>.
- Barron HC, Garvert MM, Behrens TE. Repetition suppression: a means to index neural representations using BOLD? *Philos Trans R Soc B: Biol Sci*. 2016;371(1705):20150355. <https://doi.org/10.1098/rstb.2015.0355>.
- Bedny M, McGill M, Thompson-Schill SL. Semantic adaptation and competition during word comprehension. *Cereb Cortex*. 2008; 18(11):2574–2585. <https://doi.org/10.1093/cercor/bhn018>.
- Behrens TE, Muller TH, Whittington JC, Mark S, Baram AB, Stachenfeld KL, Kurth-Nelson Z. What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*. 2018;100(2): 490–509. <https://doi.org/10.1016/j.neuron.2018.10.002>.
- Bellmund JL, Gärdenfors P, Moser EI, Doeller CF. Navigating cognition: spatial codes for human thinking. *Science*. 2018; 362(6415):eaat6766. <https://doi.org/10.1126/science.aat6766>.
- Bellmund JL, Deuker L, Doeller CF. Mapping sequence structure in the human lateral entorhinal cortex. *Elife*. 2019;8:e45333. <https://doi.org/10.7554/eLife.45333>.

- Bellmund JL, Deuker L, Montijn ND, Doeller CF. Mnemonic construction and representation of temporal structure in the hippocampal formation. *Nat Commun*. 2022;13(1):3395. <https://doi.org/10.1038/s41467-022-30984-3>.
- Binder JR, Desai RH, Graves WW, Conant LL. Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cereb Cortex*. 2009;19(12):2767–2796. <https://doi.org/10.1093/cercor/bhp055>.
- Bookheimer S. Functional MRI of language: new approaches to understanding the cortical organization of semantic processing. *Annu Rev Neurosci*. 2002;25(1):151–188. <https://doi.org/10.1146/annurev.neuro.25.112701.142946>.
- Bracci S, Caramazza A, Peelen MV. Representational similarity of body parts in human occipitotemporal cortex. *J Neurosci*. 2015;35(38):12977–12985. <https://doi.org/10.1523/JNEUROSCI.4698-14.2015>.
- Brunec IK, Bellana B, Ozubko JD, Man V, Robin J, Liu ZX, Grady C, Rosenbaum RS, Winocur G, Barense MD, et al. Multiple scales of representation along the hippocampal anteroposterior axis in humans. *Curr Biol*. 2018;28(13):2129–2135. <https://doi.org/10.1016/j.cub.2018.05.016>.
- Burgess N, Maguire EA, O'Keefe J. The human hippocampus and spatial and episodic memory. *Neuron*. 2002;35(4):625–641. [https://doi.org/10.1016/S0896-6273\(02\)00830-9](https://doi.org/10.1016/S0896-6273(02)00830-9).
- Charest I, Kievit RA, Schmitz TW, Deca D, Kriegeskorte N. Unique semantic space in the brain of each beholder predicts perceived similarity. *Proc Natl Acad Sci*. 2014;111(40):14565–14570. <https://doi.org/10.1073/pnas.1402594111>.
- Clarke A, Tyler LK. Object-specific semantic coding in human perirhinal cortex. *J Neurosci*. 2014;34(14):4766–4775. <https://doi.org/10.1523/JNEUROSCI.2828-13.2014>.
- Conca F, Catricalà E, Canini M, Petrini A, Vigliocco G, Cappa SF, Della Rosa PA. In search of different categories of abstract concepts: a fMRI adaptation study. *Sci Rep*. 2021;11(1):22587. <https://doi.org/10.1038/s41598-021-02013-8>.
- Constantinescu AO, O'Reilly JX, Behrens TE. Organizing conceptual knowledge in humans with a gridlike code. *Science*. 2016;352(6292):1464–1468. <https://doi.org/10.1126/science.aaf0941>.
- Davey J, Thompson HE, Hallam G, Karapanagiotidis T, Murphy C, De Caso I, Krieger-Redwood K, Bernhardt BC, Smallwood J, Jefferies E. Exploring the role of the posterior middle temporal gyrus in semantic cognition: integration of anterior temporal lobe with executive processes. *NeuroImage*. 2016;137:165–177. <https://doi.org/10.1016/j.neuroimage.2016.05.051>.
- Dayan P. Improving generalization for temporal difference learning: the successor representation. *Neural Comput*. 1993;5(4):613–624. <https://doi.org/10.1162/neco.1993.5.4.613>.
- Eichenbaum H, Cohen NJ. Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron*. 2014;83(4):764–770. <http://dx.doi.org/10.1016/j.neuron.2014.07.032>.
- Ekstrom AD, Ranganath C. Space, time, and episodic memory: the hippocampus is all over the cognitive map. *Hippocampus*. 2018;28(9):680–687. <https://doi.org/10.1002/hipo.22750>.
- Estrada E, Hatano N. Communicability in complex networks. *Phys Rev E*. 2008;77(3):036111: 1–12. <https://doi.org/10.1103/PhysRevE.77.036111>.
- Fischl B. FreeSurfer. *NeuroImage*. 2012;62(2):774–781. <https://doi.org/10.1016/j.neuroimage.2012.01.021>.
- Frisby SL, Halai AD, Cox CR, Ralph MA, Rogers TT. Decoding semantic representations in mind and brain. *Trends Cogn Sci*. 2023;27(3):258–281. <https://doi.org/10.1016/j.tics.2022.12.006>.
- Gärdenfors P. Conceptual spaces: the geometry of thought. Cambridge, MA: MIT Press; 2004. <https://doi.org/10.7551/mitpress/2076.001.0001>.
- Gärdenfors P, Zenker F. (Ed.) *Applications of conceptual spaces: the case for geometric knowledge representation*. Cham: Springer Verlag; 2015. <https://doi.org/10.1007/978-3-319-15021-5>.
- Garvert MM, Dolan RJ, Behrens TE. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *Elife*. 2017;6:e17086. <https://doi.org/10.7554/eLife.17086>.
- Garvert MM, Saanum T, Schulz E, Schuck NW, Doeller CF. Hippocampal spatio-predictive cognitive maps adaptively guide reward generalization. *Nat Neurosci*. 2023;26(4):615–626. <https://doi.org/10.1038/s41593-023-01283-x>.
- Grill-Spector K, Henson R, Martin A. Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci*. 2006;10(1):14–23. <https://doi.org/10.1016/j.tics.2005.11.006>.
- Hebart MN, Dickter AH, Kidder A, Kwok WY, Corriveau A, Van Wicklin C, Baker CI. THINGS: a database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLoS One*. 2019;14(10):e0223792. <https://doi.org/10.1371/journal.pone.0223792>.
- Hebart MN, Zheng CY, Pereira F, Baker CI. Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nat Hum Behav*. 2020;4(11):1173–1185. <https://doi.org/10.1038/s41562-020-00951-3>.
- Humphreys GF, Ralph MA, Simons JS. A unifying account of angular gyrus contributions to episodic and semantic cognition. *Trends Neurosci*. 2021;44(6):452–463. <https://doi.org/10.1016/j.tins.2021.01.006>.
- Huth AG, De Heer WA, Griffiths TL, Theunissen FE, Gallant JL. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*. 2016;532(7600):453–458. <https://doi.org/10.1038/nature17637>.
- Kalénine S, Buxbaum LJ. Thematic knowledge, artifact concepts, and the left posterior temporal lobe: where action and object semantics converge. *Cortex*. 2016;82:164–178. <https://doi.org/10.1016/j.cortex.2016.06.008>.
- Kim JG, Biederman I, Lescroart MD, Hayworth KJ. Adaptation to objects in the lateral occipital complex (LOC): shape or semantics? *Vis Res*. 2009;49(18):2297–2305. <https://doi.org/10.1016/j.visres.2009.06.020>.
- Lambon Ralph MA, Jefferies E, Patterson K, Rogers TT. The neural and computational bases of semantic cognition. *Nat Rev Neurosci*. 2017;18(1):42–55. <https://doi.org/10.1038/nrn.2016.150>.
- Leshinskaya A, Thompson-Schill SL. Transformation of event representations along middle temporal gyrus. *Cereb Cortex*. 2020;30(5):3148–3166. <https://doi.org/10.1167/19.10.91a>.
- Martin CB, Douglas D, Newsome RN, Man LL, Barense MD. Integrative and distinctive coding of visual and conceptual object features in the ventral visual stream. *elife*. 2018;7:e31873. <https://doi.org/10.7554/eLife.31873>.
- Mirman D, Landrigan JF, Britt AE. Taxonomic and thematic semantic systems. *Psychol Bull*. 2017;143(5):499–520. <https://doi.org/10.1037/bul0000092>.
- Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw ND, Gershman SJ. The successor representation in human reinforcement learning. *Nat Hum Behav*. 2017;1(9):680–692. <https://doi.org/10.1038/s41562-017-0180-8>.
- Morton NW, Schlichting ML, Preston AR. Representations of common event structure in medial temporal lobe and frontoparietal cortex support efficient inference. *Proc Natl Acad Sci*. 2020;117(47):29338–29345.

- Morton NW, Zippi EL, Noh SM, Preston AR. Semantic knowledge of famous people and places is represented in hippocampus and distinct cortical networks. *J Neurosci*. 2021;41(12):2762–2779. <https://doi.org/10.1073/pnas.1912338117>.
- Moser EI, Kropff E, Moser MB. Place cells, grid cells, and the brain's spatial representation system. *Annual Review Neuroscience*. 2008;31(1):69–89. <https://doi.org/10.1146/annurev.neuro.31.061307.090723>.
- O'Keefe J, Nadel L. The hippocampus as a cognitive map. Oxford, UK: Clarendon Press; 1978:27:263–267. <https://doi.org/10.5840/philstudies19802725>.
- Pacheco Estefan D, Zucca R, Arsiwalla X, Principe A, Zhang H, Rocamora R, Axmacher N, Verschure PF. Volitional learning promotes theta phase coding in the human hippocampus. *Proc Natl Acad Sci*. 2021;118(10):e2021238118. <https://doi.org/10.1073/pnas.2021238118>.
- Park SA, Miller DS, Nili H, Ranganath C, Boorman ED. Map making: constructing, combining, and inferring on abstract cognitive maps. *Neuron*. 2020;107(6):1226–1238. <https://doi.org/10.1016/j.neuron.2020.06.030>.
- Peer M, Brunec IK, Newcombe NS, Epstein RA. Structuring knowledge with cognitive maps and cognitive graphs. *Trends Cogn Sci*. 2021;25(1):37–54. <https://doi.org/10.1016/j.tics.2020.10.004>.
- Poppenk J, Evensmoen HR, Moscovitch M, Nadel L. Long-axis specialization of the human hippocampus. *Trends Cogn Sci*. 2013;17(5):230–240. <https://doi.org/10.1016/j.tics.2013.03.005>.
- Price AR, Bonner MF, Peelle JE, Grossman M. Converging evidence for the neuroanatomic basis of combinatorial semantics in the angular gyrus. *J Neurosci*. 2015;35(7):3276–3284. <https://doi.org/10.1523/JNEUROSCI.3446-14.2015>.
- Radulescu A, Shin YS, Niv Y. Human representation learning. *Annu Rev Neurosci*. 2021;44(1):253–273. <https://doi.org/10.1146/annurev-neuro-092920-120559>.
- Ranganath C, Ritchey M. Two cortical systems for memory-guided behaviour. *Nat Rev Neurosci*. 2012;13(10):713–726. <https://doi.org/10.1038/nrn3338>.
- Romero K, Barense MD, Moscovitch M. Coherence and congruency mediate medial temporal and medial prefrontal activity during event construction. *NeuroImage*. 2019;188:710–721. <https://doi.org/10.1016/j.neuroimage.2018.12.047>.
- Rosch E, Lloyd BB. *Cognition and categorization*. New Jersey: John Wiley & Sons; 1978.
- Rossion B, Pourtois G. Revisiting Snodgrass and Vanderwart's object pictorial set: the role of surface detail in basic-level object recognition. *Perception*. 2004;33(2):217–236. <https://doi.org/10.1068/p5117>.
- Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND. Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Comput Biol*. 2017;13(9):e1005768. <https://doi.org/10.1371/journal.pcbi.1005768>.
- Saanum T, Schulz E, Speekenbrink M. Compositional generalization in multi-armed bandits. *Proceedings of the Annual Meeting of the Cognitive Science Society* 2021:43. <https://escholarship.org/uc/item/5nn9q6zc>.
- Schapiro AC, Kustner LV, Turk-Browne NB. Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Curr Biol*. 2012;22(17):1622–1627. <https://doi.org/10.1016/j.cub.2012.06.056>.
- Schwartz MF, Kimberg DY, Walker GM, Brecher A, Faseyitan OK, Dell GS, Mirman D, Coslett HB. Neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. *Proc Natl Acad Sci*. 2011;108(20):8520–8524. <https://doi.org/10.1073/pnas.1014935108>.
- Solomon EA, Lega BC, Sperling MR, Kahana MJ. Hippocampal theta codes for distances in semantic and temporal spaces. *Proc Natl Acad Sci*. 2019;116(48):24343–24352. <https://doi.org/10.1073/pnas.1906729116>.
- Son JY, Bhandari A, Feldman HO. Cognitive maps of social features enable flexible inference in social networks. *Proc Natl Acad Sci*. 2021;118(39):e2021699118. <https://doi.org/10.1073/pnas.2021699118>.
- Spiers HJ. The hippocampal cognitive map: one space or many? *Trends Cogn Sci*. 2020;24(3):168–170. <https://doi.org/10.1016/j.tics.2019.12.013>.
- Squire LR, Genzel L, Wixted JT, Morris RG. Memory consolidation. *Cold Spring Harb Perspect Biol*. 2015;7(8):a021766. <https://doi.org/10.1101/cshperspect.a021766>.
- Stachenfeld KL, Botvinick M, Gershman SJ. Design principles of the hippocampal cognitive map. *Adv Neural Inf Proces Syst*. 2014:27.
- Stachenfeld KL, Botvinick MM, Gershman SJ. The hippocampus as a predictive map. *Nat Neurosci*. 2017;20(11):1643–1653. <https://doi.org/10.1038/nn.4650>.
- Staresina BP, Duncan KD, Davachi L. Perirhinal and parahippocampal cortices differentially contribute to later recollection of object-and scene-related event details. *J Neurosci*. 2011;31(24):8739–8747. <https://doi.org/10.1523/JNEUROSCI.4978-10.2011>.
- Strange BA, Witter MP, Lein ES, Moser EI. Functional organization of the hippocampal longitudinal axis. *Nat Rev Neurosci*. 2014;15(10):655–669. <https://doi.org/10.1038/nrn3785>.
- Tavares RM, Mendelsohn A, Grossman Y, Williams CH, Shapiro M, Trope Y, Schiller D. A map for social navigation in the human brain. *Neuron*. 2015;87(1):231–243. <https://doi.org/10.1016/j.neuron.2015.06.011>.
- Theves S, Fernandez G, Doeller CF. The hippocampus encodes distances in multidimensional feature space. *Curr Biol*. 2019;29(7):1226–1231. <https://doi.org/10.1016/j.cub.2019.02.035>.
- Theves S, Fernández G, Doeller CF. The hippocampus maps concept space, not feature space. *J Neurosci*. 2020;40(38):7318–7325. <https://doi.org/10.1523/JNEUROSCI.0494-20.2020>.
- Tolman EC. Cognitive maps in rats and men. *Psychol Rev*. 1948;55(4):189–208. <https://doi.org/10.1037/h0061626>.
- Tucciarelli R, Wurm M, Baccolo E, Lingnau A. The representational space of observed actions. *Elife*. 2019;8:e47686. <https://doi.org/10.7554/eLife.47686>.
- Tversky A. Features of similarity. *Psychol Rev*. 1977;84(4):327–352. <https://doi.org/10.1037/0033-295X.84.4.327>.
- Viganò S, Rubino V, Di Soccio A, Buiatti M, Piazza M. Grid-like and distance codes for representing word meaning in the human brain. *NeuroImage*. 2021;232:117876. <https://doi.org/10.1016/j.neuroimage.2021.117876>.
- Visser M, Jefferies E, Lambon Ralph MA. Semantic processing in the anterior temporal lobes: a meta-analysis of the functional neuroimaging literature. *J Cogn Neurosci*. 2010;22(6):1083–1094. <https://doi.org/10.1162/jocn.2009.21309>.
- Walker MP, Stickgold R. Sleep-dependent learning and memory consolidation. *Neuron*. 2004;44(1):121–133. <https://doi.org/10.1016/j.neuron.2004.08.031>.
- Whittington JC, Muller TH, Mark S, Chen G, Barry C, Burgess N, Behrens TE. The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell*. 2020;183(5):1249–1263. <https://doi.org/10.1016/j.cell.2020.10.024>.
- Yee E, Drucker DM, Thompson-Schill SL. fMRI-adaptation evidence of overlapping neural representations for objects related in function or manipulation. *NeuroImage*. 2010;50(2):753–763. <https://doi.org/10.1016/j.neuroimage.2009.12.036>.