bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



# Predictive Representations in Hippocampal and Prefrontal Hierarchies

Iva K. Brunec[1], Ida Momennejad[2,3]

1 Department of Psychology, University of Toronto

2 Department of Biomedical Engineering, Columbia University

3 Microsoft Research, New York City

## Abstract

As we navigate the world, we use learned representations of relational structures to explore and reach goals. Studies of how relational knowledge enables inference and planning are typically conducted in controlled small-scale settings. It remains unclear, however, how people use stored knowledge in continuously unfolding navigation, e.g., walking long distances in a city. We hypothesized that predictive representations, organized at multiple scales along posterior-anterior prefrontal and hippocampal hierarchies, guide naturalistic navigation. We conducted model-based representational similarity analyses of neuroimaging data measured during navigation of realistically long paths in virtual reality. We tested the pattern similarity of each point–along each path–to a weighted sum of its successor points within different predictive horizons. We found that anterior PFC showed the largest predictive horizons, posterior hippocampus the smallest, with the anterior hippocampus and orbitofrontal regions in between. Our findings offer novel insights into how cognitive maps support hierarchical planning at multiple scales.

## Acknowledgements

The authors gratefully acknowledge Morris Moscovitch and Jason Ozubko for designing the original experiment, collecting and sharing the data, and sincerely thank Ken Norman, Buddhika Bellana, and Morgan Barense for helpful discussions. This work was supported by Canadian Institute of Health Research grant (#MOP49566 and #MOP125958) to Morris Moscovitch, a doctoral award from the Alzheimer Society of Canada (I.K.B.), grants from the James S. McDonnell Foundation and Natural Sciences and Engineering Research Council granted to Morgan Barense, NIMH Grant R01-MH104606 granted to Joshua Jacobs, and the John Templeton Foundation (I.M.), NIBIB R01EB022864.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



# Introduction

As we navigate the world, our brains construct representations in memory. Relational structures are updated using these representations during generalization and inference. This relational knowledge is later retrieved to make decisions, plan, and guide behavior (Behrens et al., 2018; Momennejad, 2020). This idea has been captured by computational models that account for planning in human behavior (Momennejad et al., 2017), relational knowledge in human fMRI (Garvert et al., 2017), and place and grid fields in rodent electrophysiology (Stachenfeld et al., 2017). This converging body of evidence suggests that the brain encodes predictive maps of relational structures, which can be used for fast and flexible planning. It has been suggested that these predictive representations are organized in a multi-scale fashion, each scale of representation corresponding to different gradients in the neural representational hierarchy, e.g., in the hippocampus (Momennejad & Howard, 2018; Stachenfeld et al., 2017) and the prefrontal cortex (Christoff & Gabrieli, 2000; Koechlin & Hyafil, 2007; Momennejad & Haynes, 2013). Here, we tested the hypothesis that predictive representations are organized along prefrontal and hippocampal hierarchies at multiple scales, consistent with predictions made by our computational models. We predicted that representations at multiple scales would be active simultaneously, and supported by different brain regions. Such hierarchical structure in the representations of states and trajectories could also enable the extraction of generalized schemas, or structured relationships at higher levels of abstraction that can be unfolded at lower levels when necessary (Figure 1A).

It is typically assumed that testing computational models is only possible in small-scale, highly controlled experiments. Here, we used model-based analysis of fMRI data to test predictions of a reinforcement learning (RL) model (Momennejad & Howard, 2018) on brain signals collected during virtual navigation of realistic distances. We show that even though participants had learned these paths in their real lives, representations in their prefrontal and hippocampal hierarchies followed predicted representations learned by our hypothesized RL model. Previously we had shown that the representations learned by these models capture human behavior in planning tasks (Momennejad et al. 2017). Here we test the hypothesis that simple principles of representation learning from our computational models may predict the representations in human brains for navigation in real life and everyday settings. We think this contribution opens up a new window into building and testing neurally plausible models of everyday cognition. Such models need to capture both behavioral and neural responses of humans performing the given task. Our approach offers a theory-rich perspective on testing models of multi-scale predictive representations in the hippocampus and the PFC using fMRI representational similarity analysis.

Multi-scale representations of space are supported by electrophysiological and neuroimaging evidence from rodents and humans. Hippocampal place cells fire within spatial fields of different sizes, and entorhinal grid fields tile the space at various levels of granularity (Brun et al., 2008;

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



Kjelstrup et al., 2008; Poppenk et al., 2013; Strange et al., 2014). Evidence from rodent electrophysiology suggests that the average place field size increases along the dorsoventral axis of the rodent hippocampus, with more ventral regions encoding space at a larger spatial scale and in a more overlapping manner (Contreras et al., 2018; Jung et al., 1994; Strange et al., 2014). Furthermore, human fMRI evidence suggests that the hippocampal posterior-anterior axis (homologous to the rodent dorsal-ventral axis) is also involved in finer- to coarser-grained spatial representations (Evensmoen et al., 2013), representing memories from shorter to longer time-scales (Nielson, Smith, Sreekumar, Dennis, & Sederberg, 2015), and inference from lower to higher levels of abstraction (Collin et al., 2015).

The larger scale representations in the anterior hippocampus are proposed to support goal-directed search (Ruediger et al., 2012), the integration of spatial and non-spatial states that are further apart (Collin et al., 2015), and longer time horizons (Brunec, Bellana et al., 2018; Nielson, Smith, Sreekumar, Dennis, & Sederberg, 2015). The representations in the posterior hippocampus are more myopic and might support smaller predictive scales, such as smaller place fields in spatial navigation (Strange et al., 2014) and more pattern separation in memory studies (Duncan & Schlichting, 2018; Leutgeb et al., 2007; Lohnas et al., 2018; Schlichting et al., 2015). Finally, recent computational models provide further support for multi-scale predictive cognitive maps. These models account for why place fields are skewed toward goal locations (Stachenfeld et al., 2017) and show that multi-scale predictive maps can capture distance to goal and reconstruct predicted sequences (Momennejad & Howard, 2018). Taking these results into account, our first hypothesis was that in virtual navigation with realistically long distances, the anterior hippocampus would display representational similarity to farther predictive horizons compared to the posterior hippocampus.

Another candidate region for processing hierarchical representations during planning is the prefrontal cortex (Badre & D'Esposito, 2007). Broadly, it has been proposed that the prefrontal cortex (PFC) is involved in navigation when it is *active* and requires planning (Behrens et al., 2018; Epstein et al., 2017; Spiers & Gilbert, 2015), in consideration of the number of paths to goal and alternative paths (Javadi et al., 2017), in reversal and detours (Kringelbach & Rolls, 2004; Spiers & Gilbert, 2015), and in retrospective revaluation mediated by offline replay (Momennejad, Otto, Daw, & Norman, 2018). Neuroimaging evidence suggests a prefrontal hierarchy in which more anterior PFC regions support relational reasoning (Christoff & Gabrieli, 2000; Christoff, Keramatian, Gordon, Smith, & Mädler, 2009), abstraction (Bunge, Kahn, Wallis, Miller, & Wagner, 2003; Christoff et al., 2001), and prospective memory (Gilbert, 2011; Haynes & Rees, 2006; Momennejad & Haynes, 2012, 2013). Bringing these findings together, our second hypothesis was that anterior prefrontal cortex encodes predictive maps with farther predictive horizons, i.e., information about states further away, while more posterior prefrontal regions maintain predictive maps that display representational similarity to more myopic predictive horizons.

Importantly, we expected the scales of anterior prefrontal cortex to exceed the highest predictive horizons of the hippocampus (Figure 1B). This is because PFC cells broadly have longer delays

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



enabling information to linger across longer time frames–as in working memory–and slower learning rates. The anterior PFC is the largest cytoarchitectonic region of the human prefrontal cortex, and a region in which we have the most difference with evolutionary ancestors (Ramnani & Owen, 2004). In contrast, the hippocampus has been suggested to be involved in rapid statistical learning at a faster rate (Schapiro et al., 2017) and is less heterogeneous across mammal species (Strange et al., 2014). It is important to note that the predictive horizon discussed here need not be merely temporal, or spatial, but can involve states that are far apart from one another in graph structures (Javadi et al., 2017) of relational knowledge acquired by statistical learning (Schapiro et al., 2013) and conceptual state spaces (Constantinescu et al., 2016). Furthermore, our predictions do not rely on the claim that supporting longer horizons is the only function of the PFC and remain agnostic to the generative mechanisms that may reinstate these representations.

Here we tested the hypothesis that hippocampal-prefrontal hierarchies simultaneously maintain multiple representations of the same underlying relational structures at different predictive horizons. We used representational similarity analysis of an existing dataset with functional magnetic resonance imaging (fMRI) data. We reanalyzed fMRI data from a previously published study of realistic virtual navigation of Goal-directed and novel routes (Brunec, Bellana et al., 2018). In this paradigm, participants underwent functional neuroimaging while they navigated Goal-directed and GPS-guided routes in a virtual version of the city they lived in (Toronto). Virtual Toronto was built using images from Google Street View. In the Goal-directed condition, participants navigated routes that they regularly traversed in their everyday lives. In the GPS-guided (GPS) condition, they were guided along unfamiliar routes by following a dynamic arrow.

This virtual navigation setup had important advantages for our purposes. First, it rendered the participants' experience as realistic as possible within the constraints of fMRI scanning. More importantly, it allowed us to compare pattern similarity for long horizons with realistically long distances in daily navigation (at the scale of kilometers). Finally, the experimental design benefited from participants' real world familiarity with certain paths. This allowed us to compare the scales or predictive horizons of well-learnt long routes vs. novel routes. In order to navigate GPS-guided routes, participants used the same control buttons as they did along Goal-directed routes. However, in the GPS condition they did not know the goal they were navigating towards (Figure 1C).

To test our hypotheses about multi-scale predictive representations along hippocampal and PFC hierarchies, we used two main representational similarity analyses (Figure 1D). To maximally benefit from the temporal resolution afforded by fMRI, paths were discretized into steps: each step corresponded to a TR, or repetition time, during which an entire brain volume was measured. In the first analysis, we computed the correlation between every given step (TR) and the average of all future steps (TRs) within a particular horizon (e.g., mean of future 10 TRs following the current TR). In the second analysis, following the equations for predictive or successor representations (Dayan, 1993; Momennejad et al., 2017), we computed the

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



correlation between every given step and the weighted sum of future steps within a horizon. The pattern across voxels at each future TR was weighed exponentially using a discount parameter (i.e., gamma value, ɣ) between 0 and 1, and the value of the discount parameter corresponded to the scale of abstraction, corresponding to different levels of a representational hierarchy (Momennejad & Howard, 2018). Consistent with our prediction, we found that on Goal-directed, compared to GPS paths, the anterior hippocampus and anterior prefrontal regions maintained predictive maps with longer horizons (i.e. displayed similarity to more distal states), while the posterior hippocampus cached predictive maps with smaller scales (i.e. displayed similarity to more proximal states).

# Methods

## Subjects

Twenty-two healthy right-handed volunteers were recruited. One participant was excluded because of excessive difficulty with the task (i.e., repeatedly getting lost). Two additional participants were excluded due to incomplete data or technical issues. Exclusions resulted in 19 participants who completed the study (9 males; mean age 22.58 years, range 19-30 years). All participants had lived in Toronto for at least 2 years (M = 10.45, SE = 1.81). All participants were free of psychiatric and neurological conditions. All participants had normal or corrected-to-normal vision and otherwise met the criteria for participation in fMRI studies. Informed consent was obtained from all participants in accordance with Rotman Research Institute at Baycrest's ethical guidelines. Participants received monetary compensation upon completion of the study.

## Experimental design and paradigm

We used a realistic navigation software drawing on 360° panoramic images from Google Street View. This allowed participants to walk through a virtual Toronto from a first-person, street-level perspective. The navigation software was written in MATLAB v7.5.0.342. Navigation was controlled using three buttons: left, right, and forward. A "done" button allowed participants to indicate that they had completed a route. The task was projected on a screen in the bore of the scanner viewed by the participants through a mirror mounted inside of the head coil. Participants navigated in 4 conditions, and navigated 16 routes in total (4 in each condition, in a randomized order). The details of the experimental design have been reported in a previously published study (Brunec, Bellana, et al., 2018).

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



Data from two conditions of interest were analyzed in the present manuscript: Goal-directed and GPS/arrow-following routes. The routes were constructed prior to the day of scanning: participants built routes with researcher assistance, using a computer program which showed overhead maps of Toronto. Additionally, sets of routes in areas of Toronto with which participants were generally unfamiliar were created. Four of these routes were randomly assigned to each participant to be used in the baseline (GPS) condition. In the scanner, participants were provided with Goal-directed route destinations and asked to navigate towards the goal along the most Goal-directed/comfortable route. GPS trials involved no goal-directed navigation; instead, participants followed a dynamic arrow (Figure 1A). We only analyzed routes where participants successfully reached the goal ($M_{Goal\text{-}directed} = 3.37$, $M_{GPS} = 3.16$ routes). Comparing these conditions enabled us to contrast navigational signals associated with goal-directed navigation with matched motor control and optic flow, but no goal.

## fMRI acquisition and preprocessing

Participants were scanned with a 3T Siemens MRI scanner at the Rotman Research Institute at Baycrest. A high-resolution 3D MPRAGE T1-weighted pulse sequence image (160 axial slices, 1 mm thick, FOV = 256 mm) was first obtained to register functional maps against brain anatomy. Functional T2*-weighted images were acquired using echo-planar imaging (30 axial slices, 5 mm thick, TR = 2000 ms, TE = 30 ms, flip angle = 70 degrees, FOV = 200 mm). The native EPI resolution was 64 x 64 with a voxel size of 3.5mm x 3.5mm x 5.0mm. Images were first corrected for physiological motion using the Analysis of Functional NeuroImages (Cox, 1996).

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



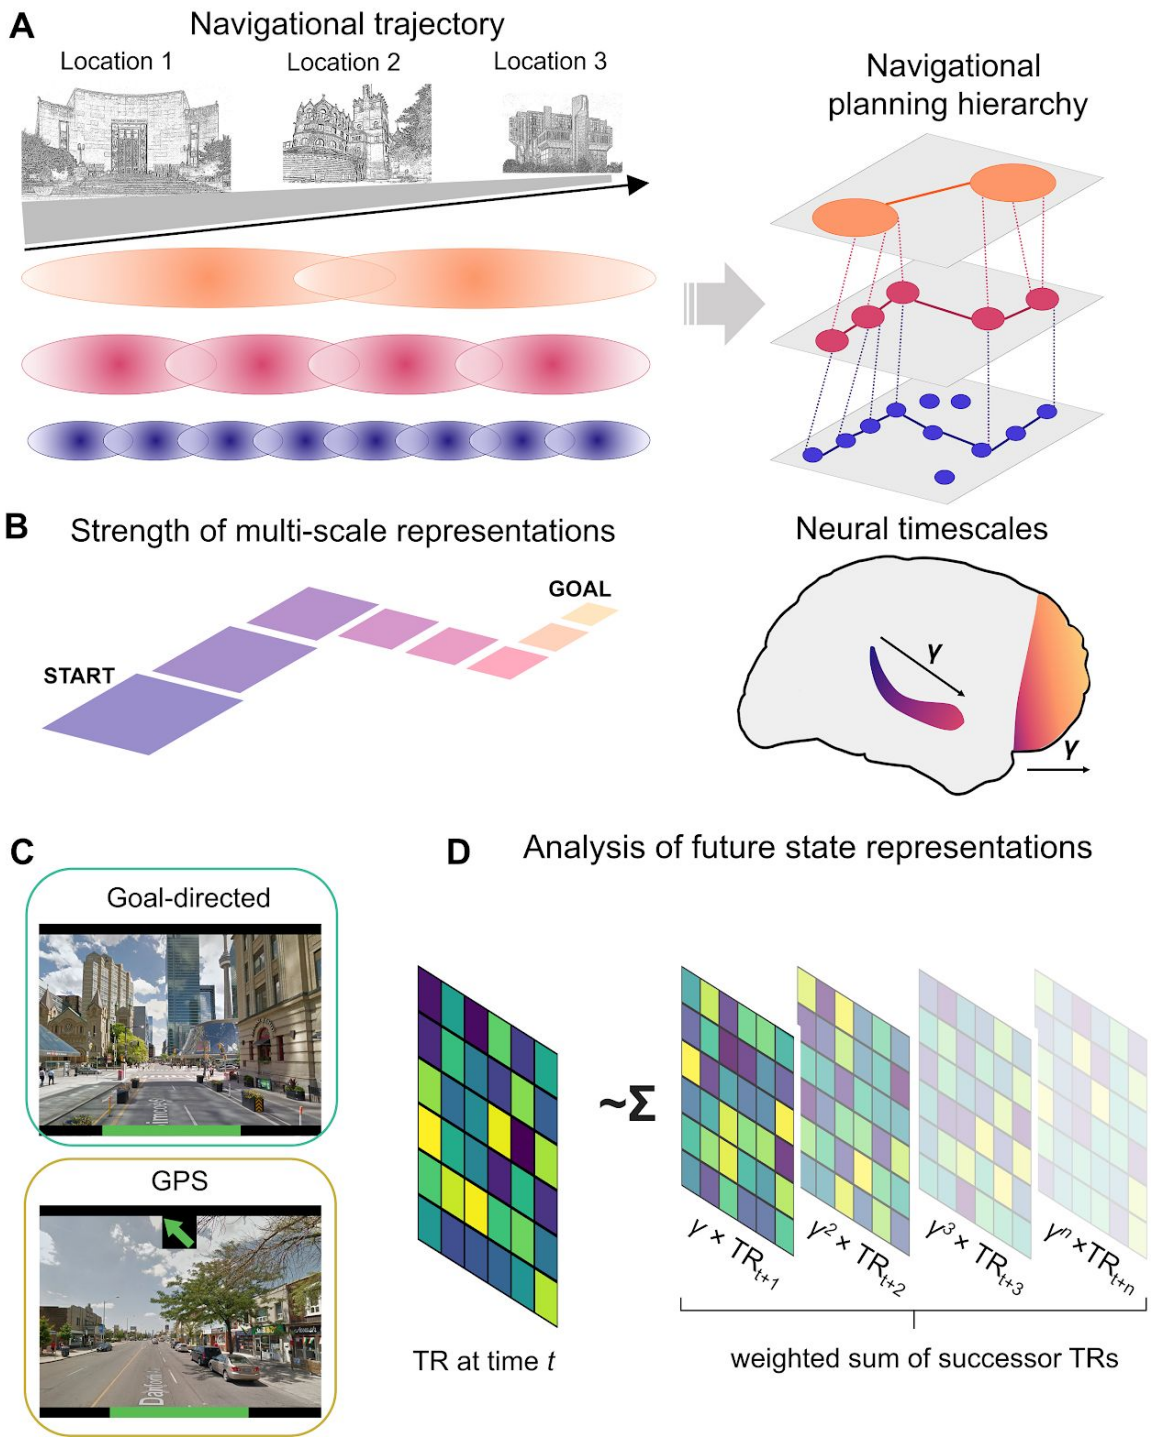


**Figure 1. Schematic of the hypothesis, task conditions, and analytic methods. A) Multiple scales of representation along a navigated route are activated simultaneously. Longer temporal horizons correspond to longer-range planning and greater scales of navigational representations. B) Predictive representations in the hippocampus and prefrontal cortex should proceed along a posterior-anterior gradient within the hippocampus and prefrontal cortex. C) Participants used the same keys to navigate Goal-directed routes and to follow the GPS dynamic arrow, but only**

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



**Goal-directed routes required goal-directed navigation. D) Analytic approach. The voxelwise pattern at each timepoint was correlated with the gamma-weighted sum of all future states (for gamma values of .1, .6, .8 and .9). With higher gamma values, the weighted future states remain above zero further into the future. Not displayed: We also computed similarity for each TR to goal, and similarity of each TR to mean of future TRs (equally weighted) within a given horizon (e.g. 10 TRs).**

All subsequent preprocessing steps were conducted using the statistical parametric mapping software SPM12 (Penny et al., 2011). Preprocessing involved slice timing correction, spatial realignment and co-registration with a resampled voxel size of 3 mm isotropic. No spatial smoothing was applied. The mean time-courses from participant-specific white matter and cerebrospinal fluid masks were regressed out of the functional images, alongside estimates of the 6 rigid body motion parameters from each EPI run. To further correct for the effects of motion which may persist despite standard processing (Power et al., 2012), an additional motion scrubbing procedure was added to the end of our preprocessing pipeline (Campbell et al., 2013). Using a conservative multivariate technique, time points that were outliers in both the six rigid-body motion parameter estimates and BOLD signal were removed, and outlying BOLD signal was replaced by interpolating across neighbouring data points. Motion scrubbing further minimizes any effects of motion-induced spikes on the BOLD signal, over and beyond standard motion regression, without leaving sharp discontinuities due to the removal of outlier volumes.

# Analysis

## Region of Interest analysis

We investigated the predictive similarity of each state to future representations in a set of regions of interest (ROIs). To do so, we first extracted voxelwise time courses across each navigated route and z-scored the values within each voxel. We then ran two predictive similarity analyses. First, we measured the correlation of each timepoint (TR) with the mean of successor TRs within a given horizon (e.g., correlation between TR at time t, and the mean of 10 following TRs). Second, we correlated the voxelwise pattern at each timepoint (TR) within each navigated route with a weighted sum of future TRs. The patterns at future TRs were weighted by different constant values ($\gamma$), corresponding to different predictive spatial scales. The specified $\gamma$ values were .1, .6, .8, and .9 (Fig. 1D). With increasing $\gamma$ values, timepoints further in the future remain weighted above zero.

As the average distance traversed within each TR was 25 meters, a $\gamma$ value of .1 meant that only each subsequent step (1 TR away) was weighted above zero, and steps farther in the distance contributed little-to-no weight to the sum of future representations. Note that we computed the predictive horizon using the unit of fMRI measurement, i.e., a TR of 2 seconds. Hence, depending on the speed of navigation, which was matched across conditions (Fig. 2C),

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



each step could cover a varying range of spatial distances (in meters) within and across subjects. Here we used the average distance traversed within a given horizon. For a ɣ value of .6, approximately 7 steps in the future were weighted above zero, corresponding to roughly 175 meters (Fig. 6C). For a value of .8, approximately 15 steps or 375 m were weighted above-zero, while this was the case for approximately 32 steps or 800 m for a ɣ value of .9 (Fig. 6C).

The TR-by-TR correlations within each route were averaged to derive the representation of future states on each trial. We first applied this analysis to *a priori* ROIs, including bilateral anterior and posterior hippocampi (aHPC, pHPC) and anterior and medial prefrontal cortical ROIs (antPFC, mPFC). As described in Brunec, Bellana et al. (2018), we divided the hippocampus into 6 anterior-posterior segments. We also examined the same measure in the mPFC and antPFC. The anterior PFC and medial PFC ROIs were defined as spheres surrounding peak voxels identified in preliminary findings from an fMRI adaptation of a known behavioral study of successor representations (Momennejad et al., 2017) reported in (Russek et al., 2018). The spheres were centered on an anterior prefrontal voxel (MNI coordinates x = 8, y = 68, z = 8) and a medial prefrontal voxel (MNI coordinates z = -22, y = 56, z = -10). These analyses were performed for each of the ROIs, as well as a searchlight within the prefrontal cortex.

## Prefrontal cortex searchlight analysis

In order to identify any gradients of predictive representation in the PFC, a custom searchlight analysis was performed within a prefrontal cortex mask (created in WFU PickAtlas). The analysis was restricted to grey matter voxels, and a spherical ROI with a 6mm radius was used to iteratively correlate each TR with the weighted sum of future states for voxels within each searchlight. The searchlight analysis was performed for four different values of ɣ: .1, .6., .8, and .9. The single-subject correlation maps were then compared against zero (AFNI *3dttest++*). The output z-score maps were thresholded at values corresponding to 5% false positive rates established by a cluster-size permutation simulation (AFNI *ClustSim*).

## Model-based analysis: The weighted sum of successor states

This section addresses the reasoning behind testing the successor representation hypothesis in terms of pattern similarity between a given state and the weighted sum of its successor states (Figure 1). Consider an environment that consists of *n* states, some of which lead to one another. Consider *T* to be the n x n matrix of transition probabilities for one-step transitions among these n states. In a deterministic environment, when there is a transition from a given state $S_i$ to state $S_j$, we assign 1 in the ith row and jth column of *T* (Supplementary Figure 1, left). The successor representation under a random policy can be then computed from *T* as follows (for comparison to policy-dependent SR see Momennejad, 2020):

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



$$(1)\ M = (I - \gamma T)^{-1}$$

Equation (2) expands equation 1 for computing the successor representation from state $s_1$ to the goal state $s_g$ from *T,* which is one cell in the SR matrix. Recall that *T* denotes the matrix of one-step transition probabilities among adjacent states, while SR contains multi-step dependencies among non-adjacent states. Here the parameter *t* refers to the number of steps (or the distance) between states. This parameter need not denote temporal steps, and can denote any type of sequential relationship among states.

$$(2)\ M(s_1,\ s_g)\ =\ \sum_{t=t1}^{tg} \gamma^t\ \ T(s_{t1},\ s_t)$$

Assume the starting state is $s_1$ and the goal state is $s_5$ (as in the Markov Decision Process (MDP) in Supplementary Figure 1). Expanding equation 2, the successor representation from state 1 to 5 is *the 5th element in the 1st row of the successor representation (Equation 3), and corresponds to the expected discounted number of times we expect to visit state 5, if we start from state 1:*

$$(3)\ M(s_1,\ s_5)\ =\ \gamma\ \ T(s_1,\ s_3)\ +\ \gamma^2\ T(s_3,\ s_5)$$

Note that equations 2 and 3 only capture 1 cell or element in the SR row associated with state $s_1$. In the successor representation framework, the *sth* row of the SR matrix (the M matrix in Dayan 1993's equations) is the representation we expect to observe when the agent is in state *s*. It denotes how often we expect to visit the current state's successors on average and given a discount. A given row of the successor representation includes the present state, and the weighted representation of successor states. Thus, at the moment when an agent is in state *s,* the row activation of successor states predicts the simultaneous activation of gamma-weighted representations. We take this simultaneous row activation as the *sum of all activated weighted states in the row* (Equation 4).

$$(4)\ Rep(s_i)\ = \sum_{i=t}^{n} \gamma^t\ Rep(s_t)$$

In short, the *1st row of the SR matrix* corresponds to the *representation that is simultaneously activated when the agent is in state 1, which is the sum of* $M(s_1, s_2)$, $M(s_1, s_3)$ , $M(s_1, s_4)$ , $M(s_1, s_5)$. Since we only have a goal-directed trajectory, this can be the weighted sum of representations of successor states (Equation 4). Each successor state is weighted by the discount factor (gamma, $\gamma$) to the power of its distance (here in the number of states) to the starting state. A simple prediction following this weighted sum view is that being in a given state along the trajectory activates the row associated with that state and hence the weighted sum of

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



successor states on that trajectory. This predicts neural similarity between the current state and the weighted sum of successor state representations.

Note that we did not have access to pretraining representations of the stimuli, e.g., the un-correlated representation of each location on the trajectory prior to being associated with specific paths (through lived experience in Toronto). Since we do not have these pretraining representations this method offers an approximation of the expected similarity structure. Therefore, as a general rule, we make the following prediction. In a goal-directed trajectory, and assuming the agent stays on path, we can assume that the transition probability between two adjacent states, e.g., $T(s_i, s_j)$, equals 1 (i.e., we have a deterministic MDP). We predict that Equation 3 approximates the pattern similarity of the TR in the *ith* state to the weighted sum of TRs that are its successor states. Note that the predictive horizon is the successor distance within which the discount parameter $\gamma$ is above zero (Figure 6). We hypothesize that different parts of the brain will show pattern similarity contingent with different values of the discount parameter $0 < \gamma < 1$, and thus different predictive horizons.

This is a first step towards testing the multi-scale predictive representation hypothesis in a realistic navigation setting. To improve prediction accuracy, future studies are needed that incorporate diverse paths through each state, to each goal, and to different goals. These studies should include a larger graph or MDP of the environment with different starting and goal locations. In order to study map-dependent and path-dependent changes in the representation of each location, a study design is needed where the participants learn a new environment. Such studies would enable us to compare pre-training and post-training neural correlations among the states or locations in the environment.

# Results

Participants navigated a set of distances they regularly traversed in everyday life ($M_{Goal\text{-}directed}$ = 3.5, $M_{GPS}$ = 2.5 km). After completing each route, participants rated how familiar each route felt, and how difficult they found it to navigate on a scale from 1-9 (where 1 would correspond to least familiar and least difficult, respectively). As expected, the average reported familiarity was higher in the Goal-directed condition (M = 7.0, SD = 1.44) than in the GPS condition (M = 3.0, SD = .51; t(18) = -10.53, p < .001; Figure 2A). The subjective difficulty was similar in the Goal-directed (M = 6.98, SD = 1.43) and GPS (M = 7.2, SD = 1.08) conditions, suggesting that all navigated routes were perceived to be similarly undemanding (t(18) = .827, p = .419; Figure 2B). There was also no difference in movement speed across the Goal-directed (M = 16.21, SD = 5.04) and GPS conditions (M = 17.02, SD = 1.48; t(18) = .719, p = .481; Figure 2C). GPS routes did, however, include more turns (M = 7.08, SD = 1.39) than Goal-directed routes (M = 5.86, SD = 1.78; t(18) = 3.04, p = .007; Figure 2D). This was the case despite the GPS routes being shorter than Goal-directed routes, on average (t(18) = -4.31, p < .001; Figure 2E).

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



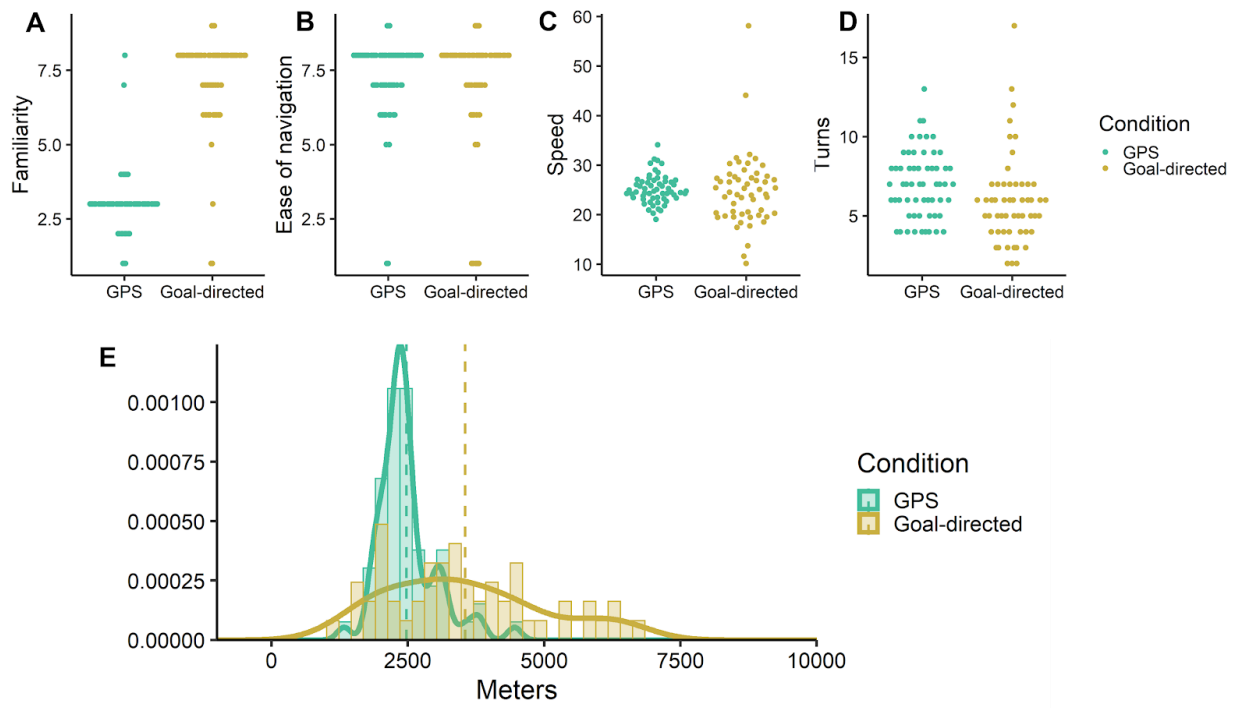


**Figure 2. Descriptive statistics for navigated distances in Goal-directed and GPS conditions. A) The Goal-directed routes were rated as more familiar by participants than GPS routes. Goal-directed and GPS routes were matched in B) ease of navigation and C) speed of travel. D) GPS routes included more turns, on average, than Goal-directed routes, but E) the Goal-directed routes tended to be longer than GPS routes.**

## Hippocampal and prefrontal gradients of near-future predictive representations

To investigate predictive representations along the hippocampal longitudinal axis and within the PFC, we first performed an analysis investigating the similarity between each timepoint (TR) and the average of future 1, 2, 3, 4, 5, or 10 TRs. As described in Brunec, Bellana et al. (2018), we divided the hippocampus into 6 anterior-posterior segments. We also examined the same measure in the mPFC and antPFC.

We first ran linear mixed effects models on these similarity measures in bilateral hippocampi for each of the routes travelled within each condition, including the average Fisher's z-transformed similarity on each route as the dependent variable, and axial segment (1-6), number of TRs (1-5), and hemisphere (L, R) as fixed effects. Participants were included as a random effect. The random intercept mixed effects models were implemented in R (R Core Team) using the packages lme4 (Bates, Maechler, Bolker, & Walker, 2015) and lmerTest (Kuznetsova et al., 2017) to assess significance. This produced a Type III ANOVA table with Satterthwaite's

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



method of approximating degrees of freedom. Where these included decimal numbers, they were rounded to the nearest integer. The similarity values for 10 TRs ahead were not entered in the present model due to the non-linear shift from 5 to 10 TR, but they are plotted in Figure 3.

We found a significant effect of axial segment ($F(5, 6796) = 45.38$, $p < .001$), driven by greater future representations in the anterior segments compared to posterior ones. There was also a main effect of condition ($F(1, 6796) = 1182.35$, $p < .001$), reflecting generally greater values in the Goal-directed (Figure 3A), compared to the GPS condition (Figure 3B), and a significant effect of the future horizon ($F(1, 6796) = 633.44$, $p < .001$), reflecting higher similarity values for states closer to the present. There was a main effect of hemisphere, reflecting higher values in the right compared to the left hemisphere ($F(1, 6796) = 6.97$, $p = .008$). There were significant interactions between axial segment and condition ($F(5, 6796) = 6.97$, $p < .001$), axial segment and future horizon ($F(20, 6796) = 2.13$, $p = .002$), and condition and future horizon ($F(4, 6796) = 13.16$, $p < .001$). The latter interaction is of particular interest as it suggests that the decline across different temporal horizons was greater in the GPS compared to the Goal-directed condition. There was no significant three-way interaction ($F < 1$).

We also ran the same models separately for the anterior and mPFC. In the antPFC, there was a significant main effect of condition ($F(1, 1222) = 363.76$, $p < .001$), as well as a main effect of future horizon ($F(4, 1222) = 48.36$, $p < .001$), but no condition by future horizon interaction ($F(4, 1222) = 1.18$, $p = .319$). In the mPFC, there was a significant effect of condition ($F(1, 1222) = 218.77$, $p < .001$) and a significant effect of future horizon ($F(4, 1222) = 114.82$, $p < .001$), but again no condition by future horizon interaction ($F(4, 1222) = 1.82$, $p = .122$).

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



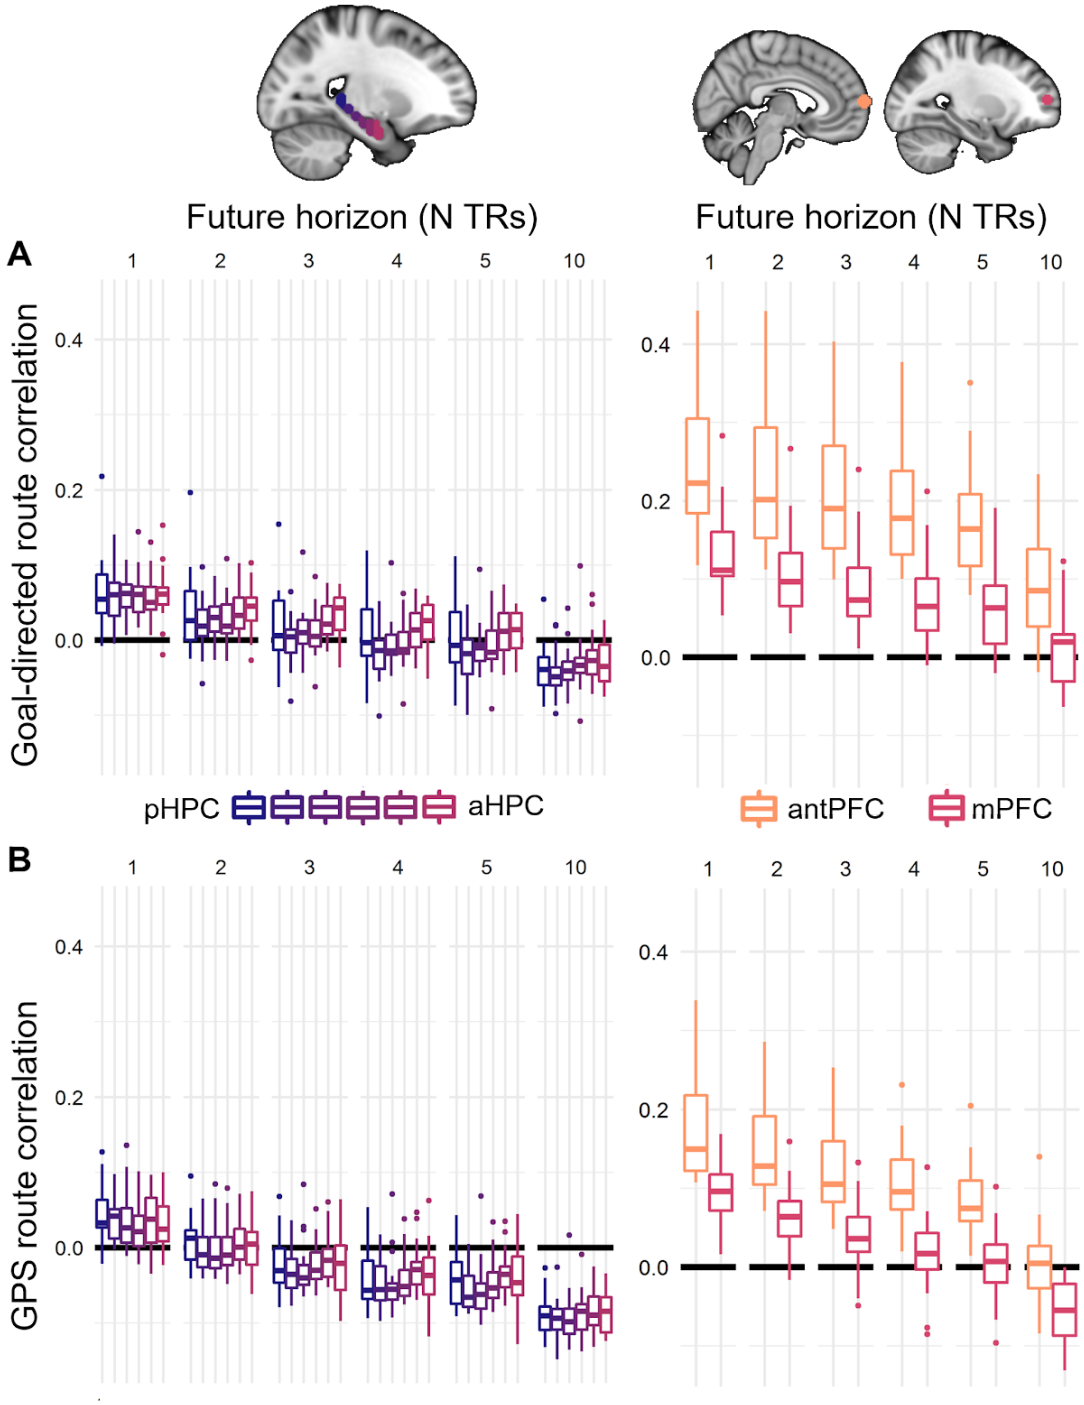


**Figure 3. Similarity of each TR to mean of future TRs (equally weighted).** **Average correlation between each timepoint and the average of future 1/2/3/4/5 or 10 timepoints in A) the Goal-directed condition and B) the GPS condition. As the average distance traversed within each TR was 25 m (meters).**

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



Comparing the representational similarity in the Goal-directed and GPS conditions against zero, we found that the anterior PFC displayed above-zero similarity for every predictive horizon including 10 steps ahead, in the Goal-directed condition (all p-values < .001), but only up to 5 steps in the GPS condition (all p-values for 1-5 steps < .001). In contrast, the medial PFC only displayed above-zero similarity up to 5 steps in the future on Goal-directed routes (p-values < .001) and three steps on GPS routes (p-values ≤ .002). The anterior-most hippocampal segment displayed above-zero similarity for up to 4 steps in the future (p-values ≤ .006) on Goal-directed routes and only one step on GPS routes (p < .001), while the posterior-most hippocampal segment displayed above-zero similarity for one step on Goal-directed routes (p < .001), and two steps on GPS routes (p-values ≤ .006).

## Model-based (weighted sum) predictive representations in ROIs

To investigate the similarity between each time point and ɣ-weighted representations of future states, we again ran a series of linear mixed effects models following the logic described above, including each route within each of the conditions. The models included Fisher's z-transformed representational similarity values as the dependent variable, with ɣ and condition as fixed effects and participant as a random effect. ɣ was modelled as an ordinal variable. For the hippocampus, the reported statistics and plotted values apply to the right hippocampus, but there was no significant difference between the left and right hippocampi (all ps > .34).

The first mixed effects model included all ROIs to determine whether the average correlation values differed across regions with different hypothesized future timescales. There was a significant main effect of ɣ, suggesting that all regions showed a decrease in correlation with increasing values of ɣ ($F(2, 1448) = 322.14$, $p < .001$), as well as a significant main effect of condition ($F(1, 1452) = 309.46$, $p < .001$), suggesting that correlations were generally greater in the Goal-directed, compared to the GPS condition (Figure 4A-B). There was a main effect of ROI ($F(3, 1448) = 547.38$, $p < .001$), confirming the prediction of strongest future representations in the antPFC, followed by mPFC, aHPC, and pHPC. There was also a significant interaction between ɣ and condition ($F(2, 1448) = 7.49$, $p < .001$), and a significant interaction between condition and ROI ($F(3, 1448) = 10.13$, $p < .001$).

Follow-up mixed effects models were run for values *within* each ROI. The significance levels were established against a Bonferroni-adjusted value of $\alpha = .0125$ (as 4 ROIs were investigated). In the antPFC, there was a significant main effect of ɣ, with significantly higher correlations for lower values of gamma ($F(2, 347) = 53.29$, $p < .001$). There was also a significant effect of condition, with significant higher correlations in the Goal-directed than the GPS condition ($F(1, 349) = 103.42$, $p < .001$). There was no significant ɣ x condition interaction ($F < 1$).

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



In mPFC, there was again a significant main effect of ɣ ($F(2, 350) = 106.39$), as well as a main effect of condition ($F(1, 352) = 83.19$, $p < .001$) in the same direction as the antPFC. There was no significant ɣ x condition interaction ($F(2, 350) = 3.44$, $p = .033$).

In the aHPC, there was a significant main effect of ɣ ($F(2, 348) = 151.90$, $p < .001$), a main effect of condition ($F(1, 350) = 128.05$, $p < .001$), as well as a ɣ x condition interaction ($F(2, 348) = 4.89$, $p = .008$). As in the mPFC, this interaction reflected a steeper slope across ɣ values in the GPS condition (-.16) than in the Goal-directed condition (-.12).
In the pHPC, there was a significant main effect of ɣ ($F(2, 349) = 218.38$, $p < .001$), a main effect of condition ($F(1, 351) = 87.99$, $p < .001$), and a significant ɣ x condition interaction ($F(2, 349) = 3.81$, $p = .023$), again reflecting a steeper slope in the GPS condition (-.17), compared to the Goal-directed condition (-.13).

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



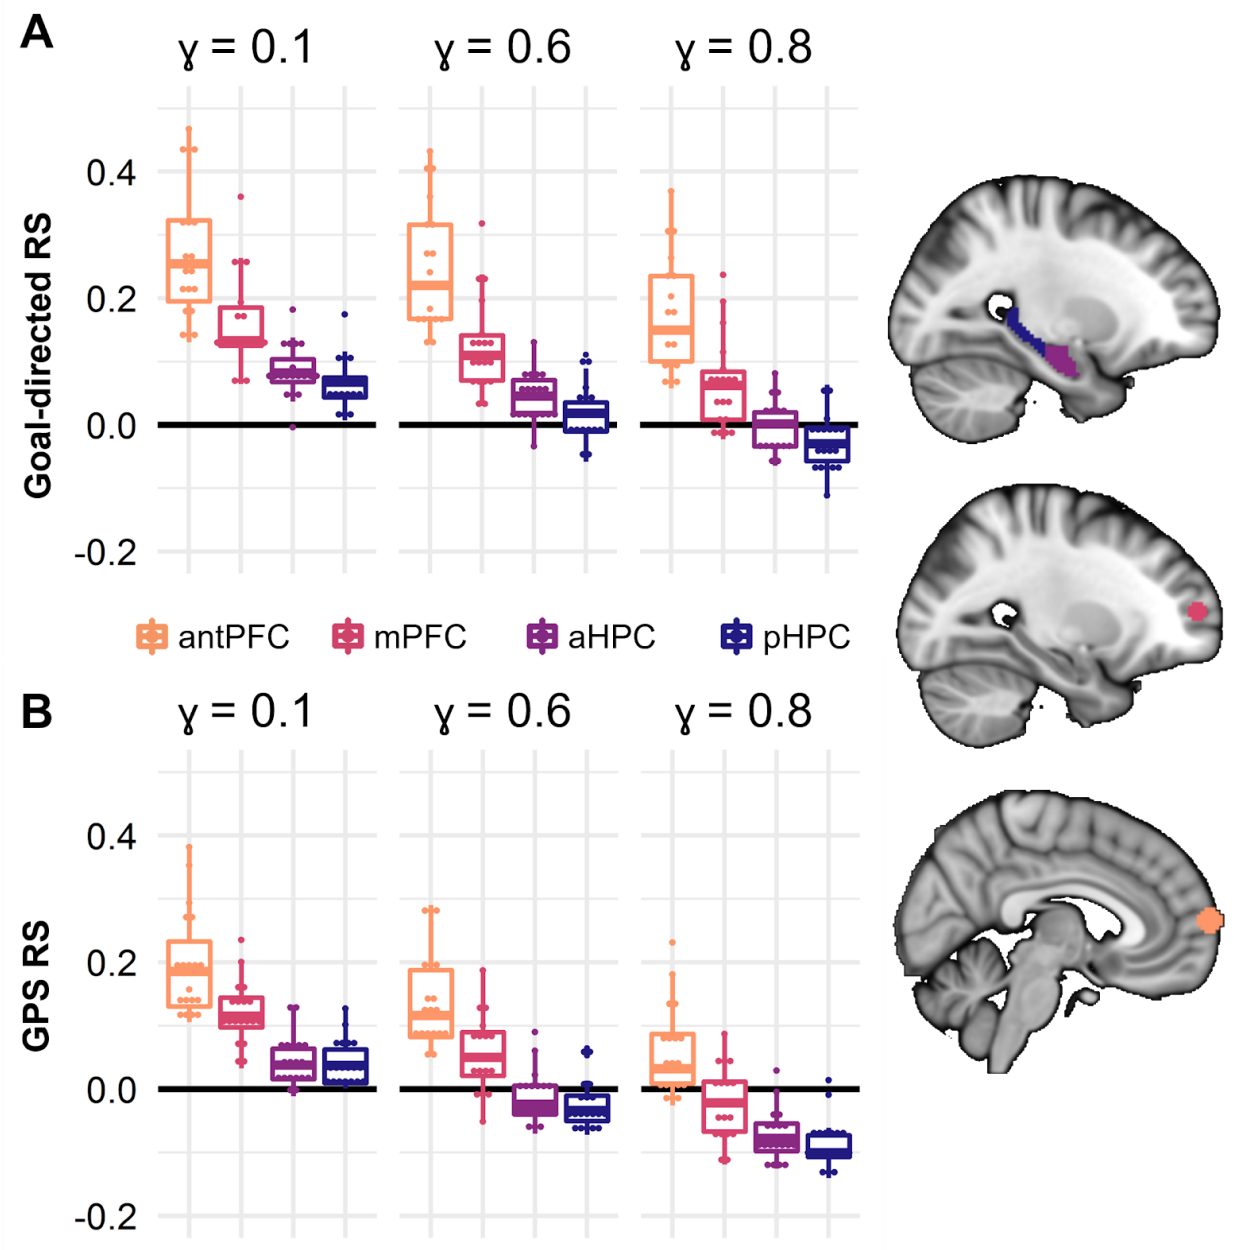


**Figure 4. Predictive similarity across temporal scales. Correlations between current timepoints and the ɣ-weighted sum of future states for different values of gamma, in the four specified ROIs in the A) Goal-directed and B) GPS conditions. ɣ=.1 only included 1-step (1 TR) away, ɣ=.6 reached approximately 6-7 steps in the future, corresponding to roughly 175 m, ɣ=.8, approximately 14 steps or 350 m ahead.**

To test for evidence of predictive representations, we tested these values against zero, with an adjusted value of α = .002 (24 comparisons in total). At ɣ = .1, the correlations in all ROIs were significantly above zero in both conditions. At ɣ = .6, all correlations in the Goal-directed condition were significantly above zero, with the exception of pHPC. In the GPS condition, however, correlations in neither the aHPC nor the pHPC were significantly above zero. At ɣ = .8,

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



values in both antPFC and mPFC remained significantly above zero in the Goal-directed condition, but only antPFC remained above zero in the GPS condition. For this value of ɣ, the values in aHPC and pHPC were not significantly above zero in either condition, and were actually significantly below zero in the pHPC. This significant negative correlation could reflect the differentiation of neural patterns across time, potentially as a manner of separating experience into fine-grained units.

## Representational similarity during goal-directed navigation is related to travelled path distance

If the hippocampus and PFC represent planning processes associated with the currently navigated route, these representations should be modulated by the route path distance. To test this, we included the path distance on each route as a factor in the model. Path distance was calculated as the summed change in longitude and latitude coordinates between each adjacent pair of TRs. To account for the contribution of time, we also regressed out the number of TRs on each route. The reported model fits thus account for the variability in the amount of time spent navigating. Here, we focused on the Goal-directed condition only, as we did not anticipate distance-related modulation in the GPS condition, where participants had no planned goal in mind. We excluded 9 Goal-directed routes from a total of 8 participants from this analysis as the recorded longitude and latitude information associated with their routes resulted in improbably long paths that diverged more than 1.5 km from the paths that the participants constructed ahead of the experiment. Including these paths, however, did not change the significance of any of the results. Prior to running these models, we mean-centered distances within each participant to account for different ranges travelled.

In the Goal-directed condition, there was a significant effect of ɣ ($F(2, 630) = 186.83$, $p < .001$), ROI ($F(3, 630) = 369.49$, $p < .001$). There was no significant main effect of path distance ($F(1, 645) = 1.06$, $p = .304$) but there were significant interactions between ROI and path distance ($F(3, 630) = 38.13$, $p < .001$) and ɣ and path distance ($F(2, 630) = 6.47$, $p = .002$; Figure 5), but no significant interaction between ɣ and ROI, nor a three-way interaction (both ps > .40). As predicted, we observed no main effect of path distance in the GPS condition ($F < 1$), nor any interactions with ROI ($F < 1$) or ɣ ($F < 1$). The main effects of ɣ ($F(2, 671) = 159.40$, $p < .001$) and ROI ($F(3, 671) = 186.99$, $p < .001$) remained significant, however.

To unpack the interaction between ROI and path distance, we ran a linear mixed effects model for each of the ROIs, predicting representational similarity from path distance and ɣ. In the antPFC, there were significant effects of ɣ ($F(2, 170) = 34.14$, $p < .001$) and path distance ($F(1, 174) = 112.14$, $p < .001$), but no interaction between the two ($F < 1$), suggesting that the effect of path distance did not change across different temporal horizons in this region. In the mPFC, the effects of ɣ ($F(2, 170) = 63.77$, $p < .001$) and path distance ($F(1, 173) = 87.27$, $p < .001$) were again significant, as well as the interaction between them ($F(2, 170) = 2.21$, $p = .113$). In

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



the aHPC, there was a significant effect of ɣ ($F(2, 170) = 84.58$, $p < .001$), a significant effect of path distance ($F(1, 177) = 60.51$, $p < .001$), and a weaker interaction between ɣ and path distance ($F(2, 170) = 3.55$, $p = .031$). Finally, in the pHPC, there were significant effects of ɣ ($F(2, 169) = 156.96$, $p < .001$), path distance ($F(1, 173) = 100.15$, $p < .001$), and a weaker interaction between the two ($F(2, 169) = 3.29$, $p = .040$).

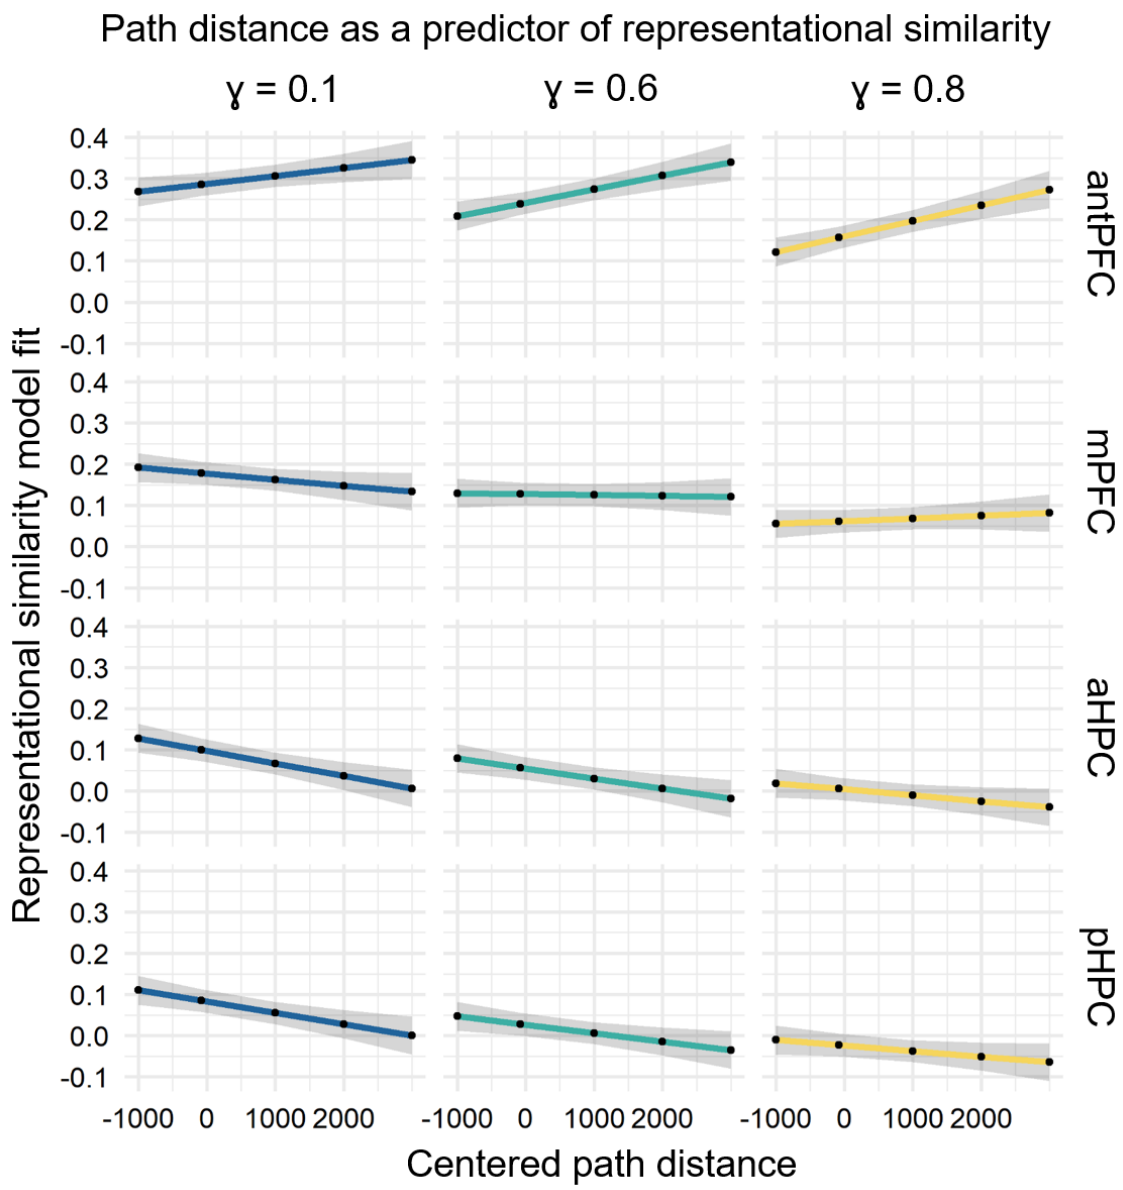


**Figure 5. Linear mixed effects model predicting representational similarity (y-axis) from path distance (x-axis), ɣ, and ROI. Voxelwise patterns in different ROIs interacted differently with path distance: in the antPFC, routes with longer path distances were associated with greater representational similarity, while the opposite trend was present in the hippocampus (both aHPC and pHPC). The plot depicts the model fit values and confidence intervals. These reflect the**

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



**relationships between the variables of interest after regressing out the effect of the number of TRs on each route and accounting for all other main effects and interactions. The plotted values were estimated using the *effects* package in R (Fox, 2003; Fox & Weisberg, 2011).**

To establish how specific these results were to the traversed paths, we re-ran the models but this time included the Euclidean distance from start to goal as a predictor instead. In the Goal-directed condition, the effects of ɣ and ROI remained significant (both ps < .001), but there was no main effect of Euclidean distance ($F < 1$), and no significant interaction between ɣ and Euclidean distance ($F(2, 631) = 1.93$, $p = .145$). There was an interaction between ROI and Euclidean distance ($F(3, 630) = 3.33$, $p = .019$), but no three-way interaction ($F < 1$). In the GPS condition, the effects of ɣ and ROI were again significant (ps < .001), and there was a weaker main effect of Euclidean distance ($F(1, 49) = 4.39$, $p = .041$), but no interactions between Euclidean distance and any other factor (all ps > .50).

## Predictive representations in prefrontal searchlights

Prefrontal cortex has a much larger volume than the hippocampus. In order to identify hierarchies of predictive representations comparable to hippocampal ROIs, we ran a searchlight analysis and computed similarity for voxels within every spherical searchlight (of 6mm radius). The searchlight analysis was performed for four values of ɣ (.1, .6, .8., .9) within each of the conditions. The thresholded z-score maps for different values of ɣ are displayed as overlays in Figure 6A, along with the average correlation maps within each condition (thresholded at .06; Figure 6B).

To capture the gradient of values from the anterior-most to the posterior-most segments of the PFC, we calculated the average value of representational similarity across voxels within each anterior-posterior slice (i.e., the y-direction). The slopes are plotted in Figure 7. These plots reveal a gradation of future state representations extending from posterior- to anterior-most slices of the PFC. This trend was reliable in both the Goal-directed and GPS conditions, but the representational similarity values were consistently greater in the Goal-directed condition.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



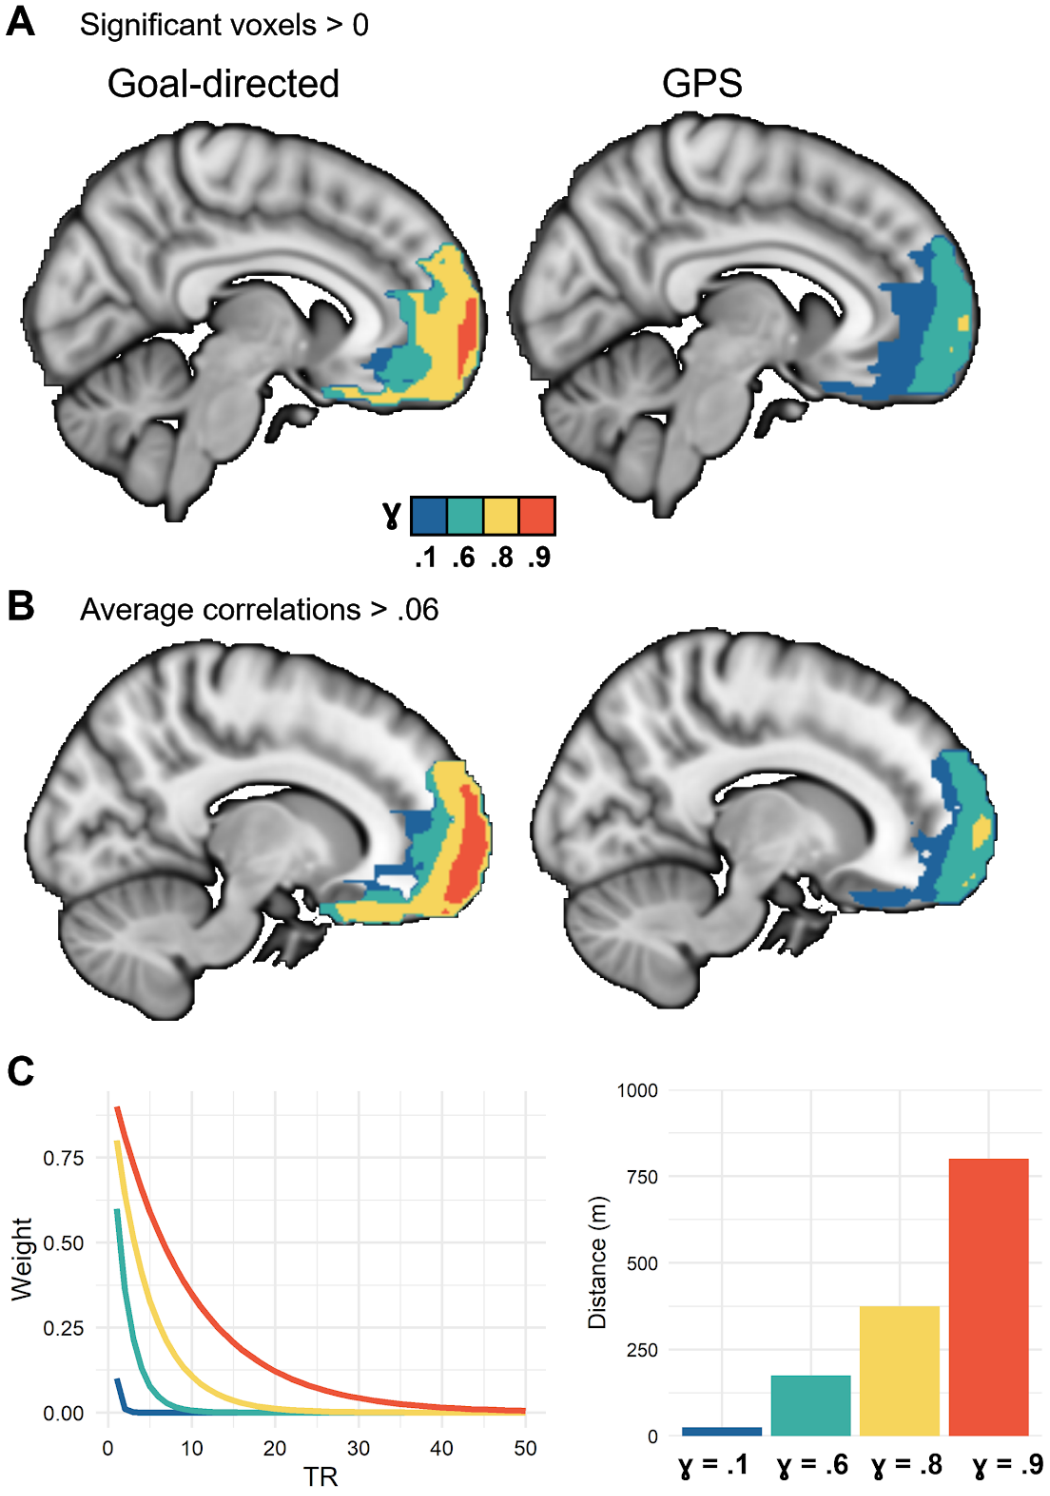


**Figure 6. One-sample T-tests for Goal-directed and GPS condition. A) Voxels with significant representations of future states in the Goal-directed and GPS conditions using a one-sample t-test against zero. B) Voxels with representational similarity (correlation) values above .06 for each value of γ. *See Supplementary Figure 2 for the same results controlled for***

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



***distance by adding it as a covariate within each condition*. C) Discounted weights corresponding to different gammas were applied to each successor TR. When the agent was in a given state (TR), the weighted successor states (TRs) were summed to estimate the predictive representation for that state. The average distance covered in each TR was approximately 25 m (24.8 m). Based on this, we computed approximate distances corresponding to predictive horizons for each discount parameter. Please note that the exact distances for each discount parameter differed across routes and participants depending on their speed. *γ=.1* only included 1-step (1 TR) away, *γ=.6* reached approximately 7 steps in the future, corresponding to roughly 175 m, *γ=.8*, approximately 15 steps or 375 m, *γ=.9* reached approximately 32 steps or 800 m ahead.**

## Representational similarity slope along PFC hierarchy

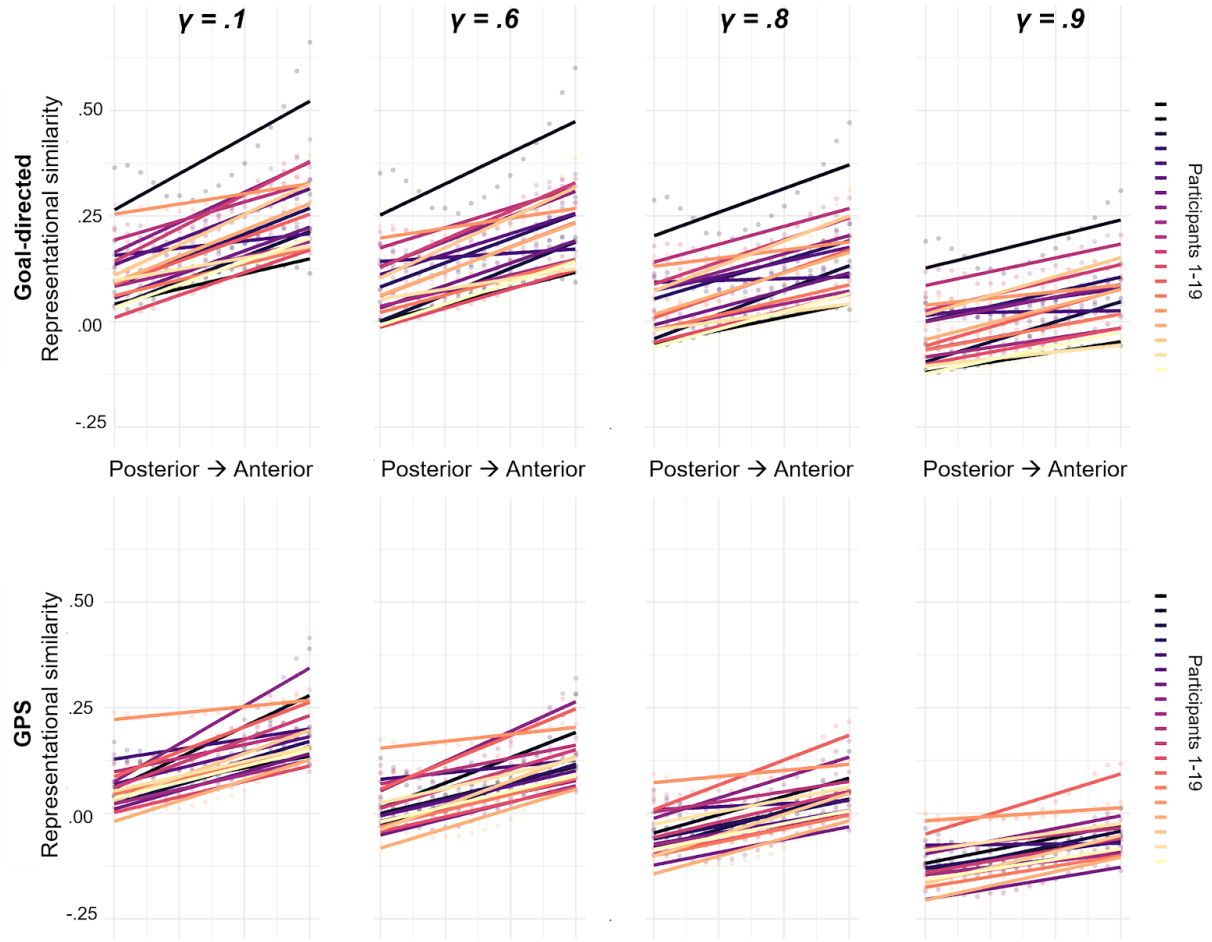


**Figure 7. Increasing predictive similarity along posterior-to-anterior PFC. In order to indicate which PFC regions displayed higher predictive similarity we computed the slope of correlations for posterior-to-anterior PFC slices for Goal-directed and GPS conditions. We computed these slopes for four values of γ, corresponding to gradients of low to high scales. Each line corresponds to predictive similarity results from one of 19 participants.**

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



To account for the proportion of different histologically-defined brain regions covered by each significant cluster, we calculated the % of overlap between each prefrontal Brodmann Area (BA) region and the significant voxels for each value of ɣ in each of the conditions. These percentages are reported in Table 1 and Figure 8. These percentages represent the proportion of each BA region covered by the significant thresholded clusters. We found the largest overlap between voxels in the anterior PFC (BA 10) and significant voxels in the searchlight analysis with various ɣ values. Following anterior and polar PFC was BA 11, corresponding to the orbitofrontal cortex, and then BA 25 and 32, corresponding to subgenual area or cingulate cortex and anterior cingulate cortex respectively. These regions were followed by smaller overlap in area 47, corresponding to the orbital part of the inferior frontal gyrus, areas 46and 9 corresponding to the dorsolateral PFC, and no overlap in area 45 corresponding to the inferior frontal gyrus.

| | **Goal-directed** | | | | **GPS** | | | |
|---|---|---|---|---|---|---|---|---|
| | **ɣ = .1** | **ɣ = .6** | **ɣ = .8** | **ɣ = .9** | **ɣ = .1** | **ɣ = .6** | **ɣ = .8** | **ɣ = .9** |
| **BA9** | 5.8% | 5.8% | 5.6% | 0% | 5.8% | 4.3% | 0% | 0% |
| **BA10** | 59.4% | 59.3% | 54.9% | 13.9% | 59.1% | 48.0% | 4.9% | 0% |
| **BA11** | 46.5% | 44.4% | 32.6% | 6.0% | 42.6% | 21.8% | 0% | 0% |
| **BA25** | 41.1% | 37.4% | 15.9% | 0% | 21.5% | 0% | 0% | 0% |
| **BA32** | 23.9% | 21.6% | 11.2% | 0% | 21.1% | 2.8% | 0% | 0% |
| **BA47** | 16.6% | 15.6% | 5.8% | 0% | 12.7% | 1.1% | 0% | 0% |
| **BA46** | 6.8% | 6.8% | 6.8% | 0% | 6.8% | 4.2% | 0% | 0% |
| **BA45** | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% |

**Table 1. Proportion of each prefrontal Brodmann Area accounted for by the significant prefrontal voxels.** Results were driven from the one-sample T-test results displayed in Figure 5A. (not matched for distance). Proportions are displayed for each value of ɣ within each condition.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



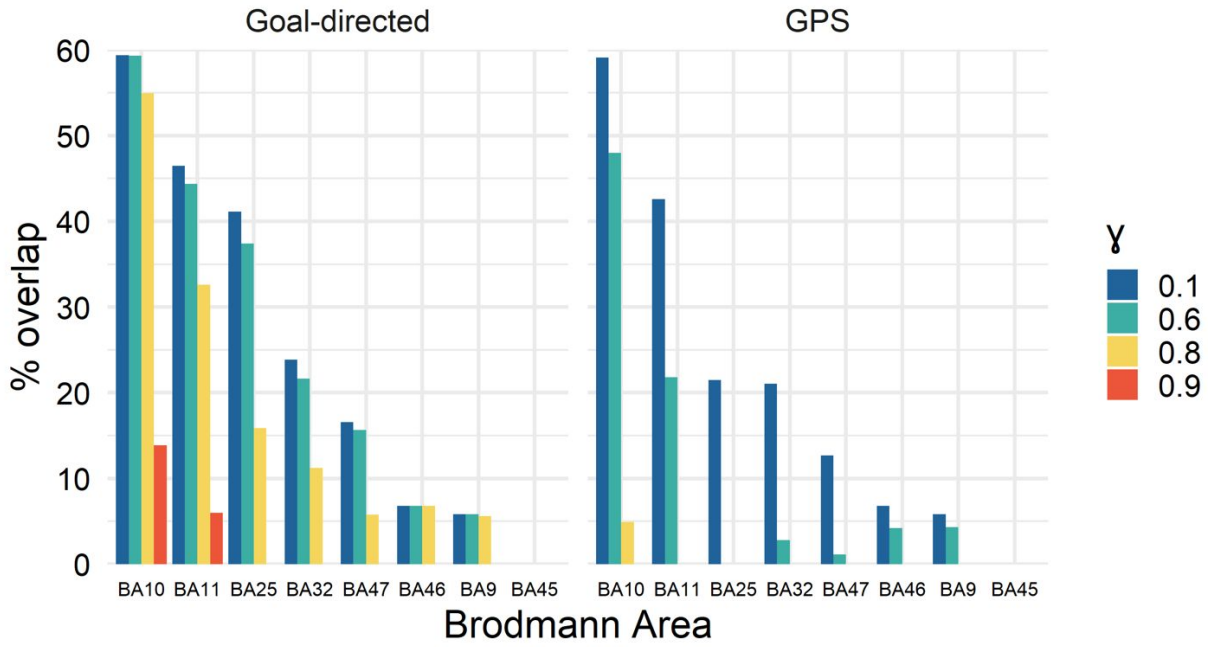


**Figure 8. Prefrontal cortex hierarchy in the Goal-directed and GPS conditions. Proportion of prefrontal Brodmann Areas accounted for by the significant PFC voxels in searchlight analysis are shown. Results were driven from the one-sample T-test results displayed in Figure 5A (not matched for distance). Colorbars reflect different discount values (γ) corresponding to different predictive horizons within each condition. (Blue: γ=.1, Green: γ=.6, Yellow: γ=.8, Red: γ=.9).**

## Controlling for Distance: Matched Distance Analysis

As discussed in the ROI sections, the distances were not matched between the two conditions (Figure 2E). To account for this difference, we conducted a matched analysis in which we manually selected pairs of routes with the minimum difference in distance for each participant, up to a kilometer (Figure 9A). We were unable to include 1 of the participants in this analysis as the distances in their Goal-directed and GPS routes were too different (with a difference in distance > 1.5 km). For the remaining 18 participants, there was no significant difference between the selected GPS and Goal-directed routes ($p = .215$). We ran a paired samples t-test comparing their prefrontal correlation maps for the two selected routes, with the participant-specific difference in distance for the two selected routes as a covariate. The brain maps of the average correlation values thresholded at .04 are presented in Figure 9B and the results of the 5% FPR corrected t-test in Figure 9C.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



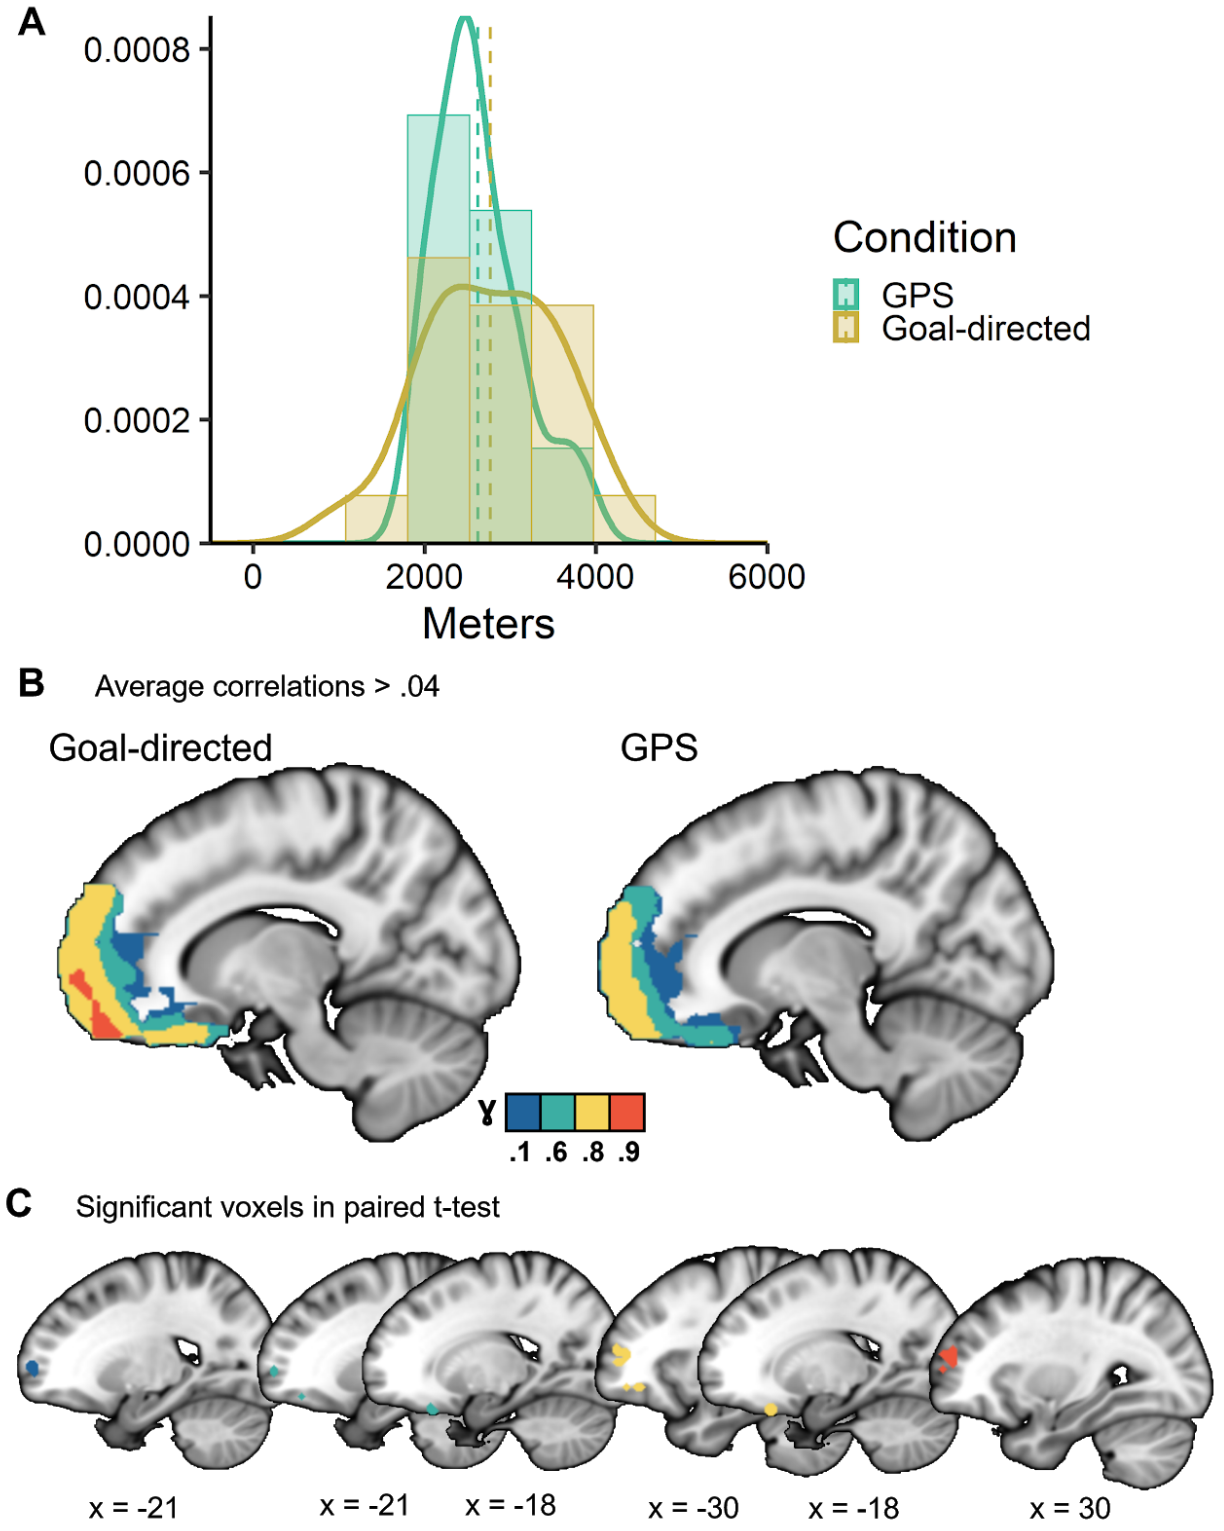


**Figure 9. Predictive representations for Goal-directed and GPS routes with matched distances. A) Distribution of distance-matched routes included in this analysis. B) Voxels with**

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



**average correlation values of > .04. C) Significant voxels in Goal-directed > GPS paired t-test, thresholded at t-value corresponding to 5% FPR. Colors reflect predictive horizons corresponding to different discount parameters (Blue: ɣ=.1, Green: ɣ=.6, Yellow: ɣ=.8, Red: ɣ=.9).**

We compared matched-distance searchlight results in the Goal-directed and GPS conditions. In this comparison, relatively few clusters significantly differed between the Goal-directed and GPS conditions. However, the comparison at each level of ɣ suggests that there is a set of clusters along the rostrocaudal extent of the PFC which differentiates between goal-directed and GPS-guided navigation (Table 1). Notably, while only orbitofrontal clusters were significantly different for smaller horizons, more dorsal and rostral/polar PFC clusters emerged in the comparison of larger horizons or scales–between the Goal-directed and GPS conditions. It is worth noting, however, to ensure matched distances between the Goal-directed and GPS condition we excluded individuals with a large difference between the distances in the two conditions. As a result, this analysis only included individual paths from 16 participants, likely resulting in increased noise and lower statistical power.

# Discussion

We investigated the hypothesis that relational knowledge – about navigational paths – is organized as multi-scale predictive representations in hippocampal and prefrontal hierarchies. We found evidence for such hierarchical representations in a task where participants navigated the city of Toronto virtually in both Goal-directed and GPS-guided conditions without knowing the goal. Both planned and guided paths covered realistically long distances (average 3 kilometers). We computed representational similarity between fMRI patterns corresponding to each location (TR) and all prospective locations (TRs) within given horizons along the path. Motivated by previous work on multi-scale predictive representations (Momennejad & Howard, 2018), we computed pattern similarity to discounted predictive horizons (25-875 meters)–i.e., sum of TRs weighed by a discount parameter into the future. These analyses revealed four main findings. First, fMRI similarity reflected longer predictive horizons for paths in the Goal-directed condition compared to GPS-guided paths. Second, similarity values in the anterior hippocampus and frontopolar cortex were significantly higher in the Goal-directed condition and for longer horizons. Third, predictive representations were organized along a posterior-anterior hierarchy of predictive horizons (25-175m) in the hippocampus with larger scales in gradually more anterior hippocampal regions. Fourth, similarity to future horizons was organized along a rostro-caudal hierarchy in the PFC with larger-scale horizons (25-875m) in gradually more polar regions (Figure 6). Overall, anterior PFC showed predictive similarity at the largest scales and posterior hippocampus the smallest, while the anterior hippocampus, pre-polar PFC, and orbitofrontal regions were in between.

These results support the hypothesis that prefrontal-hippocampal representations organize relational knowledge–in this case for spatial navigation–at different scales of generalization and

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



abstraction (Behrens et al., 2018; Momennejad & Howard, 2018). In the case of spatial navigation, this hierarchical representation in turn enables hierarchical planning and subgoal computation using graphs of the environment (Figure 1) abstracted at different scales (Ribas-Fernandes et al., 2018). Our proposal is that planning at larger scales may be enabled by larger and more abstract scales of predictive representations in anterior PFC (Figure 1, large scale graph). This higher level plan may be translated into more precise policies using representations in pre-polar PFC and anterior hippocampal regions (Figure 1, mid-scale graph), and finer scale trajectories are translated by hippocampal gradients down to the smallest predictive horizons of place fields (Figure 1, small scale graph). This proposal is also supported by previous findings.

Hierarchical structuring of fine-grained to coarse-grained representations would also enable abstraction and generalization. A similar approach can potentially generalize local information to derive schema (Figure 1A). We propose that local representations at small scales, supported by the hippocampus, are nested within large-scale, generalized representations in the PFC. Multiple scales of representation exist simultaneously, but the spatial precision of information differs across the scales: at larger scales the representations are more nebulous. The global structure of each route is thus represented in the PFC, but the precise information about individual locations is supported by the hippocampus.

Consistent with our proposal, recent work on cognitive maps in rodents, monkeys, and humans indicate PFC's involvement in *active* navigation and planning (Epstein et al., 2017), while earlier work on finer scale spatial representation had primarily focused on the hippocampus. The hippocampus is thought to support cognitive maps of space (O'Keefe & Nadel, 1978; Burgess, Maguire, & O'Keefe, 2002) as well as nonspatial relational structures (Bellmund et al., 2018; Garvert et al., 2017). Recent computational perspectives suggest that the hippocampus serves rapid statistical learning (Schapiro et al., 2016, 2017) to form and update a predictive map of the state space at multiple scales (Momennejad & Howard, 2018; Stachenfeld et al., 2017). As such, the hippocampus serves as a predictive map that organizes relational knowledge of spatial and non-spatial states (Garvert et al., 2017; McKenzie et al., 2014; Schuck et al., 2016). Notably, the long axis of the hippocampus is shown to support gradually larger spatiotemporal scales (Brunec et al., 2018; Nielson et al., 2015; Poppenk et al., 2013; Strange et al., 2014). Place cells along the dorsoventral extent of the rodent hippocampus display gradually larger place fields (Strange et al., 2014). Furthermore, a number of fMRI studies have focused on the role of posterior-anterior axis in spatio-temporal scales (Nielson et al., 2015; Peer et al., 2019), and inference on mnemonic relations (Collin et al., 2015; Schlichting & Preston, 2015).

An important aspect of our findings is that we observed predictive similarity reflecting gradually higher scales along the posterior-to-anterior gradient of the PFC hierarchy (Figures 6-9, Table 1). To compare predictive similarity in this gradient, we computed the slope of correlations for each predictive horizon (weighted sum of TRs) across posterior to anterior PFC slices. We found an overall effect of condition, where predictive similarity was generally higher in the Goal-directed vs. the GPS-guided condition especially for higher horizons. We also observed a

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



prefrontal gradient effect: more anterior PFC regions showed higher predictive similarity (correlation) values in general. Furthermore, we measured the proportional overlap between significant voxels in the searchlight analysis and voxels in different histologically-defined PFC regions. To do so, we calculated the % of overlap between each prefrontal Brodmann Area (BA) region and the significant voxels for each value of $\gamma$ in the Goal-directed and GPS-guided conditions. These percentages, reported in Table 1 and Figure 8, indicate the largest overlap in the anterior PFC (BA 10), orbitofrontal cortex (BA 11), and granular and anterior cingulate cortex (BA 25 and 32). These findings are consistent with the direction of the slope of predictive similarity in Figure 7.

The majority of PFC voxels that displayed significant predictive similarity were in polar or anterior PFC (BA 10), especially for larger predictive scales (Table 1, Figure 8). BA 10 is the largest cytoarchitectonic region of the human PFC, it has the largest volumetric and proportional difference between humans and other great apes, it is highly interconnected within the PFC, and its cells display longer decay times (Ramnani & Owen, 2004). Thus, the properties of BA 10 suggest a structurally well-connected region to support higher levels of abstraction. This includes supporting predictive representations with larger scales of integration, which can be thought of in terms of clustering of relational graphs with a higher radius. For temporal relations, this graph clustering or integration radius can be thought of in terms of longer decays or longer sustained memory leading to binding over longer time-scales. For spatial relations, this radius can be thought of in terms of associating locations that are farther apart. For relational structures, this radius can be thought of in terms of an increase in similarity among a cluster of associations within a given degree of separation.

It is noteworthy that while the present analyses were focused on representational similarity, we do not take them to mean the representations we're measuring are static, and simply there. It is possible that these representations are constructed from compressed representation, e.g., eigenvectors (Stachenfeld et al., 2017), inverse laplace transform (Ida Momennejad & Howard, 2018) (Momennejad, Howard, 2018), or generative models (Whittington et al., 2019). Future studies are required to shed light on prefrontal and medial temporal contributions to the process involved in integration, eigen-decomposition, generative models, and abstraction. In what follows we discuss the present representation-based findings.

A crucially non-spatial body of evidence from the study of goal-directed behavior and prospective memory is relevant here. These studies indicate a functional role for anterior or rostral prefrontal cortex in the encoding and retrieval of prospective task sets and goals (Gilbert, 2011; Haynes & Rees, 2006; Momennejad & Haynes, 2012, 2013). This frontopolar evidence fits well with the proposal that the PFC is organized in a rostrocaudal hierarchy (Badre & D'Esposito, 2007; Koechlin, 2011; Koechlin et al., 2003; Koechlin & Hyafil, 2007), with more anterior or rostral regions corresponding to higher levels of integration and relational abstraction (Bunge et al., 2003; Kalina Christoff et al., 2009; Momennejad & Haynes, 2013). Lesions to the frontopolar cortex do not impair usual navigation or performance on intelligence or working memory tests, but impair the patient's ability for multi-tasking and prospective memory (Burgess,

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



2000; Volle, Gonen-Yaacovi, Costello, Gilbert, & Burgess, 2011) such as completing a sequential plan for simple everyday tasks, e.g., plan a visit to multiple stores on a street to write a note and stamp and post it (Burgess, 2000).

Studies and models of the orbitofrontal cortex (OFC) have indicated the OFC as the brain's cognitive map of task-related state spaces that serves prediction, decision-making, and planning (Schuck et al., 2016; Wilson et al., 2014). However, some studies suggest that the involvement of the OFC is more tied with the anticipation of reward (Kahnt et al., 2010), reversal learning and reappraisal due to prediction errors (Boorman et al., 2009), and prediction of states-value associations (Wimmer & Büchel, 2019). One interpretation of these findings is that the ventral PFC and the OFC are well-suited to process state-value relations due to their connectivity to subcortical value systems. In contrast, the dorsal PFC is better suited to manage action policies due to its connectivity to the dorsal striatum and motor cortical regions.

Furthermore, OFC as well as anterior PFC have been suggested to support model-based reinforcement learning (Daw et al., 2011), where an animal unfolds a learned state-action-state associative model during goal-directed planning and decision-making. This finding has been replicated across different experiments (Daw & Dayan, 2014; McDannald et al., 2012, 2014; Pauli et al., 2019). It supports the idea that the OFC maintains task-relevant state-state relational maps that enable iterative value computation in planning and decision-making (Daw et al., 2005; Keiflin et al., 2013; Simon & Daw, 2011). Notably, recent work on the neural substrates of model-based behavior indicate a role for the hippocampus in model-based decision-making as well (Miller et al., 2017; Vikbladh et al., 2019). Consistently, the present predictive representations in the anterior hippocampus were the most similar to OFC representations. However, more anterior OFC regions yielded higher predictive similarity along larger predictive horizons, more similar to aPFC than the hippocampus (Figures 4, 8). These findings expand previous perspectives on OFC-hippocampal interactions in cognitive map-like representations (Keiflin et al., 2013; Schuck et al., 2016; Wikenheiser & Schoenbaum, 2016; Wood & Grafman, 2003).

While the present analyses were focused on spatial navigation, predictive representations are generalizable to non-spatial domains as well. Examples include relational knowledge and category generalization (Constantinescu et al., 2016; Garvert et al., 2017), abstraction and transfer (Cole et al., 2011), reward predictions (Takahashi et al., 2017), associative inference, and schema learning (Hebscher & Gilboa, 2016; McKenzie et al., 2014; Moscovitch & Melo, 1997; Spalding et al., 2018; van Kesteren et al., 2013; Yu, 2018; Zeithamova et al., 2012; Zeithamova & Preston, 2010). Previous work has proposed a hierarchy of time-scales in the brain (Chen et al., 2015) and indicated a role for hippocampal-prefrontal interactions in integrating episodes to build abstract schema (Schlichting & Preston, 2017). In previous modeling work, we have proposed a role for hierarchies of predictive representations along prefrontal and hippocampal gradients (Momennejad & Howard, 2018).

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



It is worth noting that the fMRI dataset used in the present analyses (Brunec, Bellana et al., 2018) was measured as participants moved through a virtual Google navigation of a city they lived in (Toronto). Importantly, here participants navigated realistically long spatial distances, between 1-5 kilometers, which allowed us to truly distinguish between different predictive scales. However, the data set has some caveats, some of which we addressed in our controlled analyses and some of which remain to be addressed by future studies.

The first caveat of the present dataset was that the navigated routes in the Goal-directed condition were, on average, significantly longer compared to the GPS-guided condition (Figure 2). To overcome this caveat, we first controlled for distance in one-sample t-tests to reveal regions with significant pattern similarity within a given horizon (Supplementary Figure 2). In a more conservative analysis, we reran the analyses excluding longer routes and including only Goal-directed routes that were within the range of distances in the GPS-guided condition. For longer predictive scales, more dorsal PFC regions displayed significantly different similarity between the two conditions (Figure 9). These controlled analyses suggest that our main findings are reliable (compare to Table 1 and Figures 7 and 8). However, future studies with a controlled design of traversed distances and more participants are needed to replicate these findings.

The second caveat is that, in the present design, the selection of routes did not include multiple past and future trajectories for each subgoal location, nor multiple past routes for each goal location. Future studies with such a design would allow for further testing of the graph structure of relational structures. Such a study would also advance previous work using routes with multiple paths (Balaguer et al., 2016; Chanales et al., 2017). Specifically, this design would allow us to dissociate pattern similarity due to the *memory of the past* from pattern similarity due to *predictive representation*. Measuring neural representations as participants navigate a full graph would also enable investigating compressed representations. For instance, we could study whether subgoal locations that appear on many paths have a pronounced predictive representation, or whether nearby locations are clustered as one subgoal by some brain regions. Such a design could be easily implemented in future fMRI studies, enabling a more thorough analysis of prefrontal-hippocampal interactions in abstraction and sub-goal processing during planning. Future studies can address the caveats and investigate the robustness of the present findings and theoretical proposal.

Theoretically, the interpretation we adopt here is that representational similarity between a given state and its frequently visited successors along the planned path is increased (Ezzyat & Davachi, 2014; Garvert et al., 2017; Momennejad et al., 2017; Stachenfeld et al., 2017). That said, it is possible that the similarity to successor states reported here is due to replay of previous trajectories or paths (Ambrose, Pfeiffer, & Foster, 2016; Momennejad et al., 2018; Wu & Foster, 2014). Our interpretation can also be discussed in terms of increased association, integration, abstraction, and clustering (Ritvo, Turk-Browne, & Norman, 2019). This similarity could also be due to the spread of activation across memory networks (Sievers & Momennejad, 2019). While there are clever analytic designs to hint one way or another, a clear-cut

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



dissociation of these hypotheses requires higher spatio-temporal resolutions such as electrophysiology, MEG, and other methods across species.

Future studies could test the temporal hierarchy of large-scale predictive representations in the PFC and the hippocampus for higher level plans (e.g., train from New York to Philadelphia) and smaller subgoal processing (e.g., walk to the bodega around the corner). They can also test the dynamics of goal and subgoal representation more closely. This analysis would complement previous electrophysiology and neuroimaging work showing goal representation in the hippocampus (Brown et al., 2016; Howard et al., 2014; Sarel et al., 2017; Tsitsiklis et al., 2019). Future studies with larger graphs–with multiple paths leading from one location to others–will be better suited to investigate the specifics of similarity to goal findings in hippocampal-prefrontal hierarchies.

*Summary* We present support for the hypothesis that predictive maps with different scales are structured in hippocampal-prefrontal hierarchies. We found that while posterior hippocampal regions supported smaller predictive scales–up to 100-200 meters–anterior prefrontal regions supported larger predictive horizons–which in this spatial navigation task extended to 875-900 meters. Our results support the idea that medial temporal-prefrontal representations underlie cognitive maps and hierarchical planning. The organizational principles of predictive hierarchies can be extended to non-spatial domains. Future studies can be specifically designed to further investigate planning, subgoal setting, and abstraction in spatial and non-spatial graphs.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



# References

Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron*, *91*(5), 1124–1136. https://doi.org/10.1016/j.neuron.2016.07.047

Badre, D., & D'Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *Journal of Cognitive Neuroscience*, *19*(12), 2082–2099. https://doi.org/10.1162/jocn.2007.19.12.2082

Balaguer, J., Spiers, H., Hassabis, D., & Summerfield, C. (2016). Neural Mechanisms of Hierarchical Planning in a Virtual Subway Network. *Neuron*, *90*(4), 893–903. https://doi.org/10.1016/j.neuron.2016.03.037

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1). https://doi.org/10.18637/jss.v067.i01

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, *100*(2), 490–509. https://doi.org/10.1016/j.neuron.2018.10.002

Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, *362*(6415), eaat6766. https://doi.org/10.1126/science.aat6766

Boorman, E. D., Behrens, T. E. J., Woolrich, M. W., & Rushworth, M. F. S. (2009). How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron*, *62*(5), 733–743. https://doi.org/10.1016/j.neuron.2009.05.014

Brown, T. I., Carr, V. A., LaRocque, K. F., Favila, S. E., Gordon, A. M., Bowles, B., Bailenson, J. N., & Wagner, A. D. (2016). Prospective representation of navigational goals in the human hippocampus. *Science (New York, N.Y.)*, *352*(6291), 1323–1326. https://doi.org/10.1126/science.aaf0784

Brun, V. H., Solstad, T., Kjelstrup, K. B., Fyhn, M., Witter, M. P., Moser, E. I., & Moser, M.-B. (2008). Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus*, *18*(12), 1200–1212. https://doi.org/10.1002/hipo.20504

Brunec, I. K., Bellana, B., Ozubko, J. D., Man, V., Robin, J., Liu, Z.-X., Grady, C., Rosenbaum, R. S., Winocur, G., Barense, M. D., & Moscovitch, M. (2018). Multiple Scales of Representation along the Hippocampal Anteroposterior Axis in Humans. *Current Biology*, *28*(13), 2129-2135.e6. https://doi.org/10.1016/j.cub.2018.05.016

Bunge, S. A., Kahn, I., Wallis, J. D., Miller, E. K., & Wagner, A. D. (2003). Neural circuits subserving the retrieval and maintenance of abstract rules. *Journal of Neurophysiology*, *90*(5), 3419–3428. https://doi.org/10.1152/jn.00910.2002

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, *35*(4), 625–641. https://doi.org/10.1016/s0896-6273(02)00830-9

Burgess, P. W. (2000). Strategy application disorder: The role of the frontal lobes in human multitasking. *Psychological Research*, *63*(3–4), 279–288.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



Campbell, K., Grigg, O., Saverino, C., Churchill, N., & Grady, C. (2013). Age differences in the intrinsic functional connectivity of default network subsystems. *Frontiers in Aging Neuroscience*, *5*, 73.

Chanales, A. J. H., Oza, A., Favila, S. E., & Kuhl, B. A. (2017). Overlap among Spatial Memories Triggers Repulsion of Hippocampal Representations. *Current Biology: CB*, *27*(15), 2307-2317.e5. https://doi.org/10.1016/j.cub.2017.06.057

Chen, J., Hasson, U., & Honey, C. J. (2015). Processing Timescales as an Organizing Principle for Primate Cortex. *Neuron*, *88*(2), 244–246. https://doi.org/10.1016/j.neuron.2015.10.010

Christoff, K., & Gabrieli, J. D. E. (2000). The frontopolar cortex and human cognition: Evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology*, *28*(2), 168–186.

Christoff, K, Prabhakaran, V., Dorfman, J., Zhao, Z., Kroger, J. K., Holyoak, K. J., & Gabrieli, J. D. (2001). Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *NeuroImage*, *14*(5), 1136–1149. https://doi.org/10.1006/nimg.2001.0922

Christoff, Kalina, Keramatian, K., Gordon, A. M., Smith, R., & Mädler, B. (2009). Prefrontal organization of cognitive control according to levels of abstraction. *Brain Research*, *1286*, 94–105. https://doi.org/10.1016/j.brainres.2009.05.096

Cole, M. W., Etzel, J. A., Zacks, J. M., Schneider, W., & Braver, T. S. (2011). Rapid transfer of abstract rules to novel contexts in human lateral prefrontal cortex. *Frontiers in Human Neuroscience*, *5*, 142. https://doi.org/10.3389/fnhum.2011.00142

Collin, S. H. P., Milivojevic, B., & Doeller, C. F. (2015). Memory hierarchies map onto the hippocampal long axis in humans. *Nature Neuroscience*, *18*(11), 1562–1564. https://doi.org/10.1038/nn.4138

Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science (New York, N.Y.)*, *352*(6292), 1464–1468. https://doi.org/10.1126/science.aaf0941

Contreras, M., Pelc, T., Llofriu, M., Weitzenfeld, A., & Fellous, J. (2018). The ventral hippocampus is involved in multi-goal obstacle-rich spatial navigation. *Hippocampus*, *28*(12), 853–866. https://doi.org/10.1002/hipo.22993

Cox, R. W. (1996). RW Cox. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*, 162–173.

Daw, N. D., & Dayan, P. (2014). The algorithmic anatomy of model-based evaluation. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *369*(1655). https://doi.org/10.1098/rstb.2013.0478

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*(6), 1204–1215. https://doi.org/10.1016/j.neuron.2011.02.027

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711. https://doi.org/10.1038/nn1560

Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, *5*(4), 613–624.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



Duncan, K. D., & Schlichting, M. L. (2018). Hippocampal representations as a function of time, subregion, and brain state. *Neurobiology of Learning and Memory*, *153*, 40–56.

Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: Spatial navigation and beyond. *Nature Neuroscience*, *20*(11), 1504–1513. https://doi.org/10.1038/nn.4656

Ezzyat, Y., & Davachi, L. (2014). Similarity breeds proximity: Pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron*, *81*(5), 1179–1189. https://doi.org/10.1016/j.neuron.2014.01.042

Fox, J. (2003). Effect displays in R for generalised linear models. *Journal of Statistical Software*, *8*(15), 1–27.

Fox, J., & Weisberg, S. (2011). *An {R} Companion to Applied Regression* (2nd ed.). Sage. http://socserv.socsci.mcmaster.ca/jfox/Books/Companion

Garvert, M. M., Dolan, R. J., & Behrens, T. E. (2017). A map of abstract relational knowledge in the human hippocampal-entorhinal cortex. *ELife*, *6*. https://doi.org/10.7554/eLife.17086

Gilbert, S. (2011). *Decoding the content of delayed intentions. 31*, 2888–2894.

Haynes, J.-D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews. Neuroscience*, *7*(7), 523–534. https://doi.org/10.1038/nrn1931

Hebscher, M., & Gilboa, A. (2016). A boost of confidence: The role of the ventromedial prefrontal cortex in memory, decision-making, and schemas. *Neuropsychologia*, *90*, 46–58. https://doi.org/10.1016/j.neuropsychologia.2016.05.003

Howard, L. R., Javadi, A. H., Yu, Y., Mill, R. D., Morrison, L. C., Knight, R., Loftus, M. M., Staskute, L., & Spiers, H. J. (2014). The hippocampus and entorhinal cortex encode the path and Euclidean distances to goals during navigation. *Current Biology: CB*, *24*(12), 1331–1340. https://doi.org/10.1016/j.cub.2014.05.001

Javadi, A.-H., Emo, B., Howard, L., Zisch, F., Yu, Y., Knight, R., Pinelo Silva, J., & Spiers, H. (2017). Hippocampal and prefrontal processing of network topology to simulate the future. *Nature Communications*, *8*, 14652. https://doi.org/10.1038/ncomms14652

Jung, M., Wiener, S., & McNaughton, B. (1994). Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *The Journal of Neuroscience*, *14*(12), 7347–7356. https://doi.org/10.1523/JNEUROSCI.14-12-07347.1994

Kahnt, T., Heinzle, J., Park, S. Q., & Haynes, J.-D. (2010). The neural code of reward anticipation in human orbitofrontal cortex. *Proceedings of the National Academy of Sciences*, *107*(13), 6010–6015. https://doi.org/10.1073/pnas.0912838107

Keiflin, R., Reese, R. M., Woods, C. A., & Janak, P. H. (2013). The Orbitofrontal Cortex as Part of a Hierarchical Neural System Mediating Choice between Two Good Options. *Journal of Neuroscience*, *33*(40), 15989–15998. https://doi.org/10.1523/JNEUROSCI.0026-13.2013

Kjelstrup, K. B., Solstad, T., Brun, V. H., Hafting, T., Leutgeb, S., Witter, M. P., Moser, E. I., & Moser, M.-B. (2008). Finite Scale of Spatial Representation in the Hippocampus. *Science*, *321*(5885), 140–143. https://doi.org/10.1126/science.1157086

Koechlin, E. (2011). Frontal pole function: What is specifically human? *Trends in Cognitive Sciences*, *15*(6), 241; author reply 243. https://doi.org/10.1016/j.tics.2011.04.005

Koechlin, E., & Hyafil, A. (2007). Anterior prefrontal function and the limits of human

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



decision-making. *Science (New York, N.Y.)*, *318*(5850), 594–598. https://doi.org/10.1126/science.1142995

Koechlin, E., Ody, C., & Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science (New York, N.Y.)*, *302*(5648), 1181–1185. https://doi.org/10.1126/science.1088545

Kringelbach, M. L., & Rolls, E. T. (2004). The functional neuroanatomy of the human orbitofrontal cortex: Evidence from neuroimaging and neuropsychology. *Progress in Neurobiology*, *72*, 341–372. https://doi.org/10.1016/j.pneurobio.2004.03.006

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, *82*(13), 1–26.

Leutgeb, J. K., Leutgeb, S., Moser, M.-B., & Moser, E. I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science (New York, N.Y.)*, *315*(5814), 961–966. https://doi.org/10.1126/science.1135801

Lohnas, L. J., Duncan, K., Doyle, W. K., Thesen, T., Devinsky, O., & Davachi, L. (2018). Time-resolved neural reinstatement and pattern separation during memory decisions in human hippocampus. *Proceedings of the National Academy of Sciences*, *115*(31), E7418–E7427. https://doi.org/10.1073/pnas.1717088115

McDannald, M. A., Jones, J. L., Takahashi, Y. K., & Schoenbaum, G. (2014). Learning theory: A driving force in understanding orbitofrontal function. *Neurobiology of Learning and Memory*, *108*, 22–27. https://doi.org/10.1016/j.nlm.2013.06.003

McDannald, M. A., Takahashi, Y. K., Lopatina, N., Pietras, B. W., Jones, J. L., & Schoenbaum, G. (2012). Model-based learning and the contribution of the orbitofrontal cortex to the model-free world. *The European Journal of Neuroscience*, *35*(7), 991–996. https://doi.org/10.1111/j.1460-9568.2011.07982.x

McKenzie, S., Frank, A. J., Kinsky, N. R., Porter, B., Rivière, P. D., & Eichenbaum, H. (2014). Hippocampal Representation of Related and Opposing Memories Develop within Distinct, Hierarchically Organized Neural Schemas. *Neuron*, *83*(1), 202–215. https://doi.org/10.1016/j.neuron.2014.05.019

Miller, K. J., Botvinick, M. M., & Brody, C. D. (2017). Dorsal hippocampus contributes to model-based planning. *Nature Neuroscience*, *20*(9), 1269–1276. https://doi.org/10.1038/nn.4613

Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement learning. *Nature Human Behaviour*, *1*(9), 680–692. https://doi.org/10.1038/s41562-017-0180-8

Momennejad, Ida. (2020). Learning Structures: Predictive Representations, Replay, and Generalization. *Current Opinion in Behavioral Sciences*, *32*, 155–166. https://doi.org/10.1016/j.cobeha.2020.02.017

Momennejad, Ida, & Haynes, J.-D. (2012). Human anterior prefrontal cortex encodes the “what” and “when” of future intentions. *NeuroImage*, *61*(1), 139–148. https://doi.org/10.1016/j.neuroimage.2012.02.079

Momennejad, Ida, & Haynes, J.-D. (2013). Encoding of prospective tasks in the human prefrontal cortex under varying task loads. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *33*(44), 17342–17349.

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



https://doi.org/10.1523/JNEUROSCI.0492-13.2013

Momennejad, Ida, & Howard, M. W. (2018). Predicting the Future with Multi-scale Successor Representations. *BioRxiv*, 449470. https://doi.org/10.1101/449470

Momennejad, Ida, Otto, A. R., Daw, N. D., & Norman, K. A. (2018). Offline replay supports planning in human reinforcement learning. *ELife*, *7*, e32548. https://doi.org/10.7554/eLife.32548

Moscovitch, M., & Melo, B. (1997). Strategic retrieval and the frontal lobes: Evidence from confabulation and amnesia. *Neuropsychologia*, *35*(7), 1017–1034. https://doi.org/10.1016/s0028-3932(97)00028-6

Nielson, D. M., Smith, T. A., Sreekumar, V., Dennis, S., & Sederberg, P. B. (2015). Human hippocampus represents space and time during retrieval of real-world memories. *Proceedings of the National Academy of Sciences*, *112*(35), 11078–11083. https://doi.org/10.1073/pnas.1507104112

O'Keefe, J., & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press. https://repository.arizona.edu/handle/10150/620894

Pauli, W. M., Gentile, G., Collette, S., Tyszka, J. M., & O'Doherty, J. P. (2019). Evidence for model-based encoding of Pavlovian contingencies in the human brain. *Nature Communications*, *10*(1), 1–11. https://doi.org/10.1038/s41467-019-08922-7

Peer, M., Ron, Y., Monsa, R., & Arzy, S. (2019). Processing of different spatial scales in the human brain. *ELife*, *8*.

Penny, W. D., Friston, K. J., Ashburner, J. T., Kiebel, S. J., & Nichols, T. E. (2011). *Statistical parametric mapping: The analysis of functional brain images*. Elsevier.

Poppenk, J., Evensmoen, H. R., Moscovitch, M., & Nadel, L. (2013). Long-axis specialization of the human hippocampus. *Trends in Cognitive Sciences*, *17*(5), 230–240. https://doi.org/10.1016/j.tics.2013.03.005

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage*, *59*(3), 2142–2154.

R Core Team. (n.d.). *R: A language and environment for statistical computing* (3.5.1) [Computer software]. R Foundation for Statistical Computing. https://www.R-project.org/

Ramnani, N., & Owen, A. M. (2004). Anterior prefrontal cortex: Insights into function from anatomy and neuroimaging. *Nature Reviews. Neuroscience*, *5*(3), 184–194. https://doi.org/10.1038/nrn1343

Ribas-Fernandes, J. J. F., Shahnazian, D., Holroyd, C. B., & Botvinick, M. M. (2018). Subgoal- and Goal-related Reward Prediction Errors in Medial Prefrontal Cortex. *Journal of Cognitive Neuroscience*, 1–16. https://doi.org/10.1162/jocn_a_01341

Ruediger, S., Spirig, D., Donato, F., & Caroni, P. (2012). Goal-oriented searching mediated by ventral hippocampus early in trial-and-error learning. *Nature Neuroscience*, *15*(11), 1563–1571. https://doi.org/10.1038/nn.3224

Russek, E., Momennejad, I., Botvinick, M. M., Gershman, S. J., & Daw, N. D. (2018). FMRI evidence for the successor representation in human value computation. *Society for Neuroscience*, 360.02.

Sarel, A., Finkelstein, A., Las, L., & Ulanovsky, N. (2017). Vectorial representation of spatial

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



goals in the hippocampus of bats. *Science*, *355*(6321), 176–180. https://doi.org/10.1126/science.aak9589

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, *16*(4), 486–492. https://doi.org/10.1038/nn.3331

Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: A neural network modelling approach to reconciling episodic memory with statistical learning. *Phil. Trans. R. Soc. B*, *372*(1711), 20160049. https://doi.org/10.1098/rstb.2016.0049

Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., & Botvinick, M. M. (2016). Statistical learning of temporal community structure in the hippocampus. *Hippocampus*, *26*(1), 3–8. https://doi.org/10.1002/hipo.22523

Schlichting, M. L., Mumford, J. A., & Preston, A. R. (2015). Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nature Communications*, *6*, 8151. https://doi.org/10.1038/ncomms9151

Schlichting, M. L., & Preston, A. R. (2015). Memory integration: Neural mechanisms and implications for behavior. *Current Opinion in Behavioral Sciences*, *1*, 1–8. https://doi.org/10.1016/j.cobeha.2014.07.005

Schlichting, M. L., & Preston, A. R. (2017). The Hippocampus and Memory Integration: Building Knowledge to Navigate Future Decisions. In D. E. Hannula & M. C. Duff (Eds.), *The Hippocampus from Cells to Systems: Structure, Connectivity, and Functional Contributions to Memory and Flexible Cognition* (pp. 405–437). Springer International Publishing. https://doi.org/10.1007/978-3-319-50406-3_13

Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron*, *91*(6), 1402–1412. https://doi.org/10.1016/j.neuron.2016.08.019

Sievers, B., & Momennejad, I. (2019). SAMPL: The Spreading Activation and Memory PLasticity Model. *BioRxiv*, 778563. https://doi.org/10.1101/778563

Simon, D. A., & Daw, N. D. (2011). Neural Correlates of Forward Planning in a Spatial Decision Task in Humans. *The Journal of Neuroscience*, *31*(14), 5526–5539. https://doi.org/10.1523/JNEUROSCI.4647-10.2011

Spalding, K. N., Schlichting, M. L., Zeithamova, D., Preston, A. R., Tranel, D., Duff, M. C., & Warren, D. E. (2018). Ventromedial Prefrontal Cortex Is Necessary for Normal Associative Inference and Memory Integration. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *38*(15), 3767–3775. https://doi.org/10.1523/JNEUROSCI.2501-17.2018

Spiers, H. J., & Gilbert, S. J. (2015). Solving the detour problem in navigation: A model of prefrontal and hippocampal interactions. *Frontiers in Human Neuroscience*, *9*. https://doi.org/10.3389/fnhum.2015.00125

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653. https://doi.org/10.1038/nn.4650

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



Strange, B. A., Witter, M. P., Lein, E. S., & Moser, E. I. (2014). Functional organization of the hippocampal longitudinal axis. *Nature Reviews. Neuroscience*, *15*(10), 655–669. https://doi.org/10.1038/nrn3785

Takahashi, Y. K., Stalnaker, T. A., Roesch, M. R., & Schoenbaum, G. (2017). Effects of inference on dopaminergic prediction errors depend on orbitofrontal processing. *Behavioral Neuroscience*, *131*(2), 127–134. https://doi.org/10.1037/bne0000192

Tsitsiklis, M., Miller, J., Qasim, S. E., Inman, C. S., Gross, R. E., Willie, J. T., Smith, E. H., Sheth, S. A., Schevon, C. A., Sperling, M. R., Sharan, A., Stein, J. M., & Jacobs, J. (2019). Single-neuron representations of spatial memory targets in humans. *BioRxiv*, 523753. https://doi.org/10.1101/523753

van Kesteren, M. T. R., Beul, S. F., Takashima, A., Henson, R. N., Ruiter, D. J., & Fernández, G. (2013). Differential roles for medial prefrontal and medial temporal cortices in schema-dependent encoding: From congruent to incongruent. *Neuropsychologia*, *51*(12), 2352–2359. https://doi.org/10.1016/j.neuropsychologia.2013.05.027

Vikbladh, O. M., Meager, M. R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., Burgess, N., & Daw, N. D. (2019). Hippocampal Contributions to Model-Based Planning and Spatial Memory. *Neuron*, *102*(3), 683-693.e4. https://doi.org/10.1016/j.neuron.2019.02.014

Volle, E., Gonen-Yaacovi, G., Costello, A. de L., Gilbert, S. J., & Burgess, P. W. (2011). The role of rostral prefrontal cortex in prospective memory: A voxel-based lesion study. *Neuropsychologia*, *49*(8), 2185–2198. https://doi.org/10.1016/j.neuropsychologia.2011.02.045

Whittington, J. C., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. (2019). The Tolman-Eichenbaum Machine: Unifying space and relational memory through generalisation in the hippocampal formation. *BioRxiv*, 770495.

Wikenheiser, A. M., & Schoenbaum, G. (2016). Over the river, through the woods: Cognitive maps in the hippocampus and orbitofrontal cortex. *Nature Reviews Neuroscience*, *advance online publication*. https://doi.org/10.1038/nrn.2016.56

Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, *81*(2), 267–279. https://doi.org/10.1016/j.neuron.2013.11.005

Wimmer, G. E., & Büchel, C. (2019). Learning of distant state predictions by the orbitofrontal cortex in humans. *Nature Communications*, *10*(1), 1–11. https://doi.org/10.1038/s41467-019-10597-z

Wood, J. N., & Grafman, J. (2003). Human prefrontal cortex: Processing and representational perspectives. *Nature Reviews. Neuroscience*, *4*(2), 139–147. https://doi.org/10.1038/nrn1033

Wu, X., & Foster, D. J. (2014). Hippocampal Replay Captures the Unique Topological Structure of a Novel Environment. *The Journal of Neuroscience*, *34*(19), 6459–6469. https://doi.org/10.1523/JNEUROSCI.3414-13.2014

Yu, L. Q. (2018). Neural Mechanisms Underlying Schemas and Inferences. *Journal of Neuroscience*, *38*(37), 7930–7931. https://doi.org/10.1523/JNEUROSCI.1345-18.2018

Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, *75*(1), 168–179. https://doi.org/10.1016/j.neuron.2012.05.010

Zeithamova, D., & Preston, A. R. (2010). Flexible memories: Differential roles for medial temporal lobe and prefrontal cortex in cross-episode binding. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(44), 14676–14684. https://doi.org/10.1523/jneurosci.3250-10.2010

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.



# Supplementary Material

## Supplementary Methods

**Supplementary Figure 1.** Assume an environment with the structure in the graph or Markov Decision Process (MDP) below. The starting state is $s_1$ and the goal state is $s_5$ (as in the MDP below), where we observe the highest reward. On the left we have the 6x6 transition probability matrix T, with probability 1 for every 1-step connection between any 2 adjacent states. On the right we have the successor representation matrix, which can be computed As follows:

$$M = (I - \gamma T)^{-1}$$

Note that the *1st row of the SR matrix* corresponds to the *representation when the agent is in state 1*. Consider an environment with 6 states and relationships depicted in the Markov Decision Process (MDP) in Supplementary Figure 1. In reinforcement learning, computing the value ($V$) of a state, under policy π, can be arrived at by multiplying the associated SR row by the vector of rewards (a 6 item vector storing the observed reward for each state).

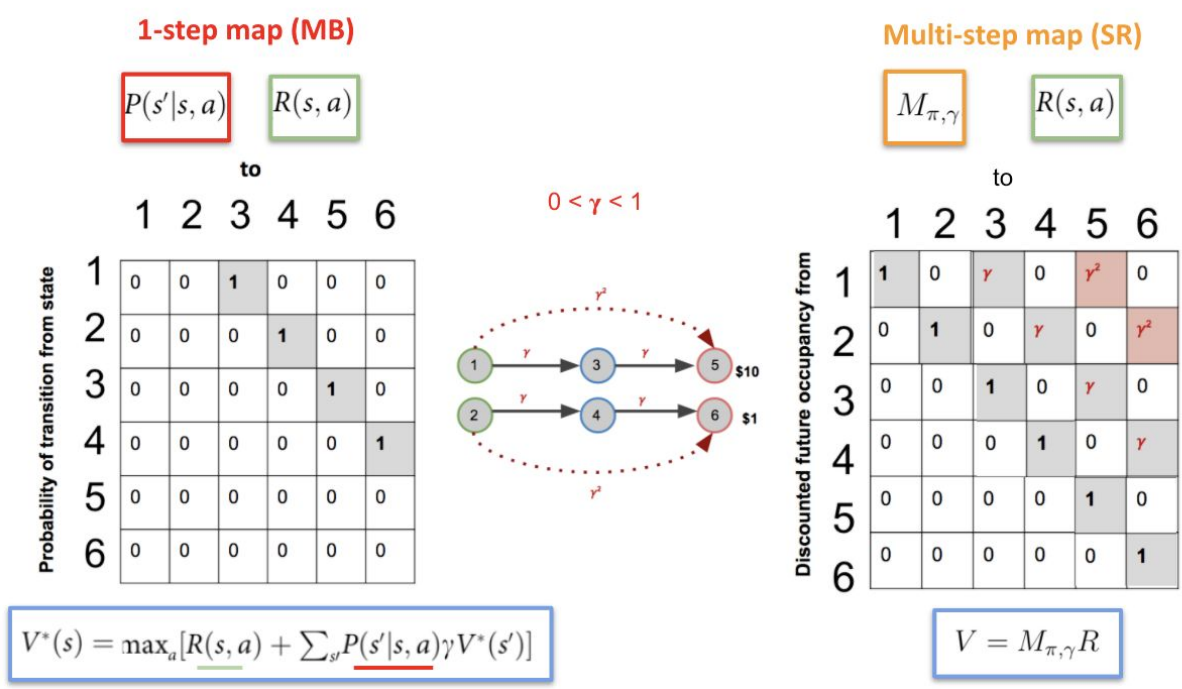

bioRxiv preprint doi: https://doi.org/10.1101/786434; this version posted July 14, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-ND 4.0 International license.


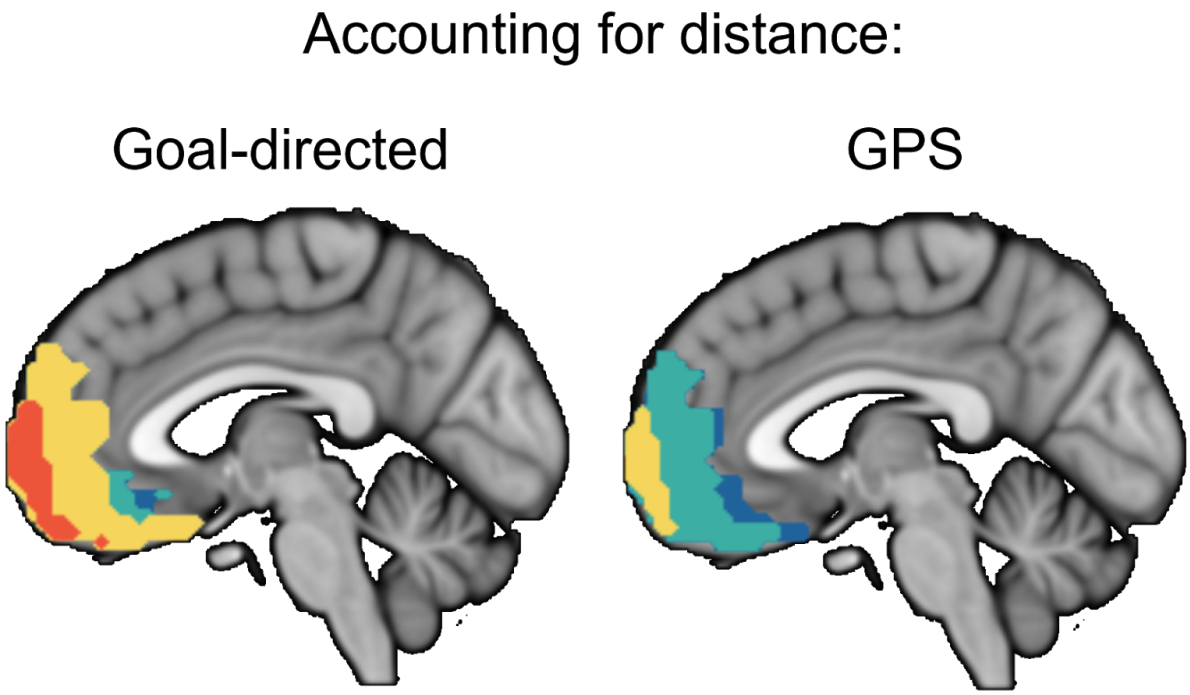
**Supplementary Figure 2. One-sample T-tests with distance as covariate.** The results look very similar to running to running a T-test on Goal-directed routes vs. zero, and GPS routes vs. zero. The mean distance, per participant, per condition, was included as a covariate.