Parallel systems for social and spatial reasoning within the cortical apex

Ben Deen* and Winrich A. Freiwald
The Rockefeller University, New York, NY, USA

Abstract. What is the cognitive and neural architecture of systems for high-level reasoning and memory in humans? We ask this question using deep neuroimaging of individual human brains on various tasks involving reasoning and memory about familiar people, places, and objects. We find that thinking about people and places elicits responses in distinct but parallel networks within high-level association cortex, spanning the frontal, parietal, and temporal lobes. Person- and place-preferring brain regions were systematically yoked across multiple cortical zones. These regions were strongly category-selective, across visual, semantic, and episodic tasks, and were specifically functionally connected to other parts of association cortex with similar category preferences. These results demonstrate that selectivity for content domain is a widespread feature of high-level association cortex in humans. They support a theoretical model in which reasoning about people and places are supported by parallel cognitive and neural mechanisms.

How are human reasoning systems structured in the mind and brain? At the cognitive level, some have argued for domain-specific systems specialized for learning and reasoning about specific classes of input, like people, objects, and places[1-3], while others have emphasized the role of domain-general processes[4,5]. Parallel debates exist in the neural domain, regarding the function of corticohippocampal learning systems[6] and the default mode or apex network – a set of regions distributed across cortex, and thought to be positioned at the top of a rough hierarchy of anatomical connectivity in mammalian cortex[7]. Some have argued for domain-specific components of the apex network[8-10], while others have treated the network as a domain-general reasoning, decision-making, or memory system[11-13].

We introduce a new theoretical framework that aims to reconcile these views. The ability to understand people and places have both been argued to employ abstract, internal models – theories of mind[14,15] and cognitive maps[16,17] – and to rely on hippocampal-dependent information about specific familiar people and places[18,19]. We hypothesize that separate but parallel cognitive mechanisms are employed for learning and reasoning about people and places, supported by parallel subnetworks within the apex network[20], interacting closely with the medial temporal lobe. These parallel systems support internal cognitive models of specific people and places, which are used for a range of functions: storing semantic information, structuring episodic memories, reasoning about events in the present, predicting and simulating future events, and making choices in navigation and social decision making.

Here, we test foundational predictions of this theory in a human fMRI experiment, using deep imaging of individual human brains, across a range of tasks engaging reasoning and memory about people, places, and objects. To maximize the strength of memory-related responses, we use tasks involving closely familiar people and places for each participant. Tasks involved visual perception (viewing images of familiar and unfamiliar faces and scenes, and generic objects); semantic judgment (answering questions about personality traits of familiar people, spatial properties of familiar places, and physical properties of generic objects); and episodic simulation (imagining familiar people talking about common conversation topics, navigating through familiar places, and physical interactions of generic objects); along with

* Correspondence should be addressed to benjamin.deen@gmail.com

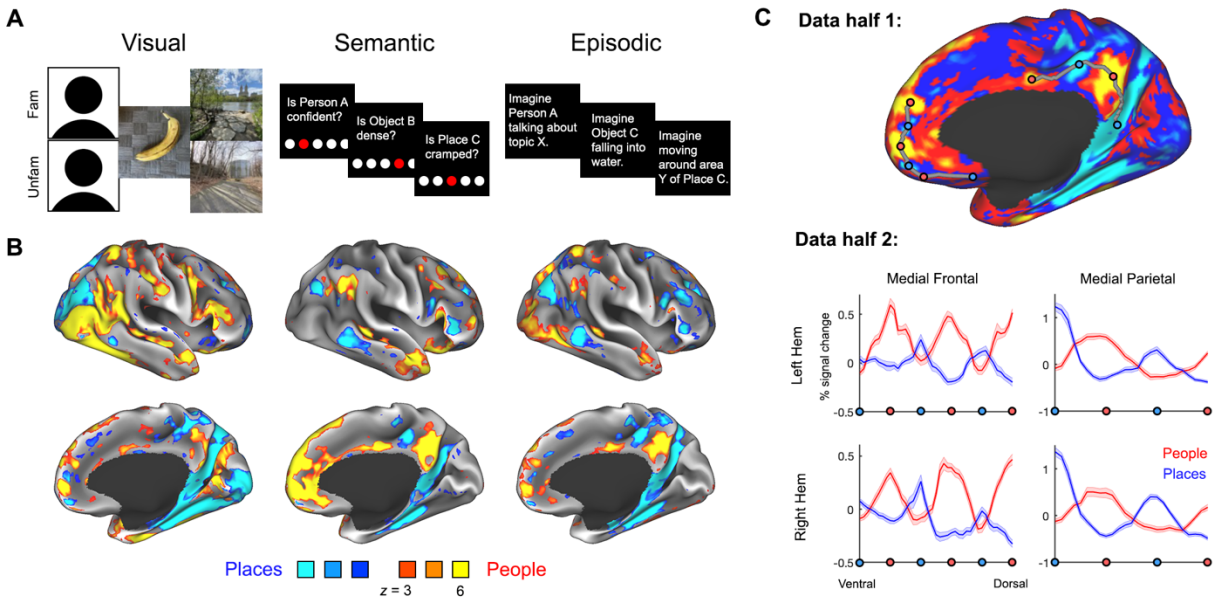

Figure 1: Parallel responses to people and places across high-level association cortex. A) Schematic of conditions from visual, semantic, and episodic tasks (face images removed from preprint). **B)** Whole-brain general linear model-based responses to people versus places, from one representative participant (thresholded at a False Discovery Rate of $q < .01$ to correct for multiple comparisons across coordinates). **C)** Split-half analysis of alternating responses to people and places in medial prefrontal and parietal cortex. Upper: anchor coordinates identified in data half 1. Lower: responses extracted along cortical geodesics between anchor coordinates in data half 2. Shaded regions show standard error across runs.

localizer scans for mental state reasoning or theory of mind (ToM), language comprehension, and dynamic visual perception (Fig 1A; S1). Optimized data acquisition and preprocessing methods yielded data with high temporal signal to noise ratio (tSNR, mean 158), including in regions of signal loss, such as the anterior medial temporal lobes (Fig S2).

We tested three predictions: 1) Separate but anatomically parallel networks across high-level association cortex are engaged during social and spatial reasoning tasks. 2) These networks include components in parts of cortex with anatomical connectivity to the medial temporal lobe, including the temporal pole (TP, social network) and parahippocampal and retrosplenial cortex (PHC and RSC, spatial network). 3) Person- and place-preferring areas of cortex are functionally connected within network, but not between.

Comparing responses to people and places, we observed a set of regions across association cortex with parallel anatomical organization (Figs 1B, S6-8), including in medial prefrontal cortex (MPFC), medial parietal cortex (MPC), temporo-parietal junction (TPJ), and superior frontal gyrus (SFG), bilaterally (whole-brain general linear model-based analysis, corrected for temporal autocorrelation using prewhitening with an ARMA(1,1) model, and corrected for multiple comparisons across coordinates using a false discovery rate of $q < .01$). Responses to people were additionally observed in the superior temporal sulcus (STS), orbitofrontal cortex, and temporal pole, and responses to places in parahippocampal cortex. This pattern of activity was observed consistently across visual, semantic, and episodic tasks (Fig 1B) and across individuals (Figs S6-8). While task-generalizing person responses were largely confined to areas associated with the apex network, place responses extended into regions associated with the frontoparietal network[21], including the lateral intraparietal sulcus, middle temporal gyrus, and parts of middle and inferior frontal gyri. Place responses in TPJ, MPC, and PHC also extended posteriorly into visual cortex.

Within zones associated with the cortical apex, regions responsive to people and places were systematically yoked, with adjacent or alternating parts of cortex exhibiting opposing stimulus preferences. In the TPJ, a pair of regions was typically observed, with a person-preferring area just anterior to a place-preferring area. In MPFC and MPC, a series of interdigitated responses to people and places were observed along a ventral to dorsal axis, with 3-5 (MPFC) or 2-3 (MPC) pairs of regions observed in individual participants. To formalize this pattern of alternation, we performed a split-half analysis (Fig 1C). In data from even runs, we used whole-brain statistical maps, combined across tasks, to identify peak coordinates (anchors) of the alternating response to people and places – two pairs of areas in parietal cortex, and three pairs in prefrontal cortex, in each participant and hemisphere. We then extracted responses (% signal change) to people and places in left-out data from odd runs, along a path defined by geodesics on the cortical surface between anchors. This analysis verified an alternating pattern of response to people and places along a ventral to dorsal axis in medial prefrontal and parietal cortex (Fig 1C).

In addition to yoked responses across much of association cortex, anatomically separate responses to people and places were observed in regions with known anatomical connections to the hippocampal formation. Place-preferring areas were observed in parahippocampal and retrosplenial cortices, consistent with prior studies on the perception and memory of places[22,23]. A person-preferring area was reliably observed within the temporal pole, a primate-specific brain region with bidirectional connections with entorhinal cortex[24,25]. An anterior temporal person representation has long been hypothesized, but hasn't been clearly demonstrated in individual human brains, possibly due to fMRI signal dropout in the anterior medial temporal lobes[26]. These areas provide candidate anatomical relays between category-sensitive regions of association cortex and the hippocampus.

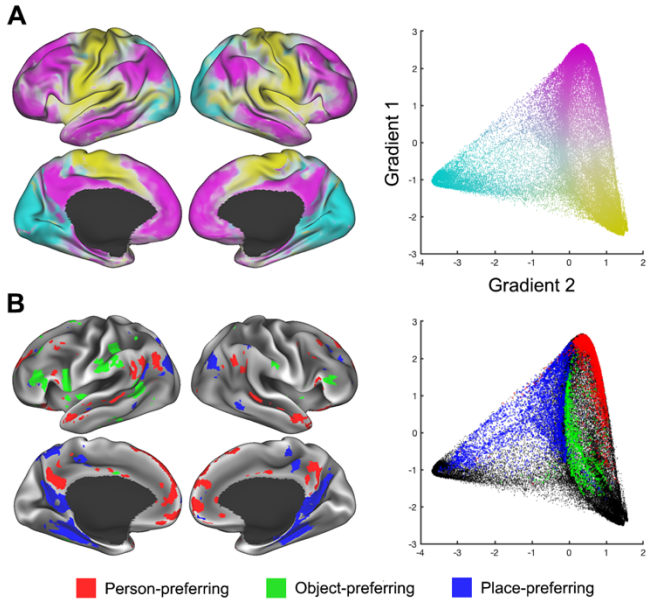

Figure 2: Category-sensitive regions are clustered in functional connectivity distance space. A) Diffusion embedding of resting-state functional connectivity data reveals large-scale cortical gradients, separate the cortical apex (pink) from visual (cyan) and somatomotor (yellow) cortices. **B)** Person-, place-, and object-preferring regions (coordinate-wise $Z > 2.3$ across visual, semantic, and episodic tasks), shown on the cortical surface and in the diffusion embedding. All results from one representative participant.

Could preferential responses to tasks involving people and places be driven by stimulus confounds, rather than abstract stimulus category? Visual stimuli were not controlled for various mid-level features, such as spatial frequency content and rectilinearity. Episodic conditions differed on self-reported level of detail, visual imagery, and emotional experience (Fig S5). Semantic conditions, while less straightforwardly confounded than visual or episodic, may differ in semantic features of the adjectives used. However, we focus here on responses that generalize across the three tasks. Such effects cannot be explained by a confound specific to any one task, and are more parsimoniously explained as an effect of stimulus domain.

How are person-, place-, and object-preferring areas situated along a connectivity-based cortical hierarchy? As for people and places, we identified a set

of object-preferring regions across frontal, parietal, and temporal cortex, responding to objects over places across multiple tasks (Figs S9-11). To position these areas along an estimated cortical hierarchy, we computed low-dimensional embeddings of graphs defined by resting-state correlation distance in individual participants, using the diffusion maps algorithm to effectively capture both local and global distance structure[7,27]. A principal gradient of functional connectivity distance was identified, separating unimodal sensorimotor cortices from a transmodal cortical apex, along with a secondary gradient separating visual from somatomotor and auditory cortices (Fig 2A).

We next identified person-, place-, and object-preferring coordinates in this space, only including coordinates that responded significantly across visual, semantic, and episodic tasks (Fig 2B, S20; $Z > 2.3$ for each task). The three sets of regions, while distributed and interdigitated on the cortical surface, were each clustered in connectivity distance space. Person-preferring regions were positioned at or near the cortical apex. Place-preferring areas spanned nearly the whole length of the principal gradient, stretching from visual cortex up to the cortical apex, consistent with prior findings associating these areas with multiple functional networks[23,28]. On the cortical surface, place responses often straddled boundary regions between visual and apex cortex. By contrast, object-preferring areas were typically clustered toward the lower end of the principal gradient, adjacent to somatomotor cortex. This mapping of brain areas involved in different reasoning domains onto separate regions of connectivity distance space may reflect constraints imposed by the specific reasoning problems faced in each domain. For instance, reasoning about people and places involves a stronger demand for abstract representations that integrate information across long timescales, relative to reasoning about objects and physical interactions.

To what extent are person- and place-preferring areas specifically involved in processing information within their domain? We addressed this question with a region-of-interest (ROI)

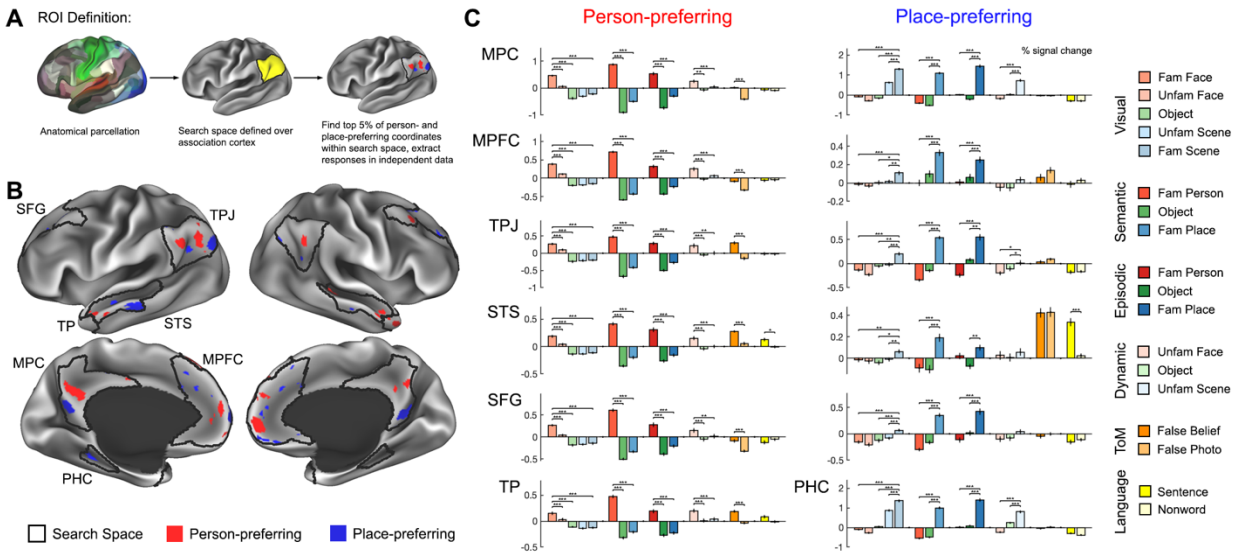

Figure 3: Region-of-interest (ROI)-based analysis reveals category-selective responses. A) ROIs were defined as the top 5% of person- or place-preferring coordinates within anatomical search spaces. **B)** Search spaces and example functional ROIs from one representative participant. **C)** Responses (% signal change) extracted from functionally defined ROIs, across all task conditions. Error bars show standard error across runs. * $P < .05/7 = .0071$, ** $P < 10^{-3}$, *** $P < 10^{-4}$ (linear mixed model across runs, with participant included as random effect). Abbreviations: MPC, medial parietal cortex; MPFC, medial prefrontal cortex; TPJ, temporo-parietal junction; STS, superior temporal sulcus; SFG, superior frontal gyrus; TP, temporal pole; PHC, parahippocampal cortex.

analysis, extracting responses from functionally defined ROIs across all six tasks and 17 conditions. ROIs were defined as the top 5% of maximally person- or place-preferring coordinates (semantic task) within anatomical search spaces capturing core areas of the cortical apex: MPC, MPFC, TPJ, STS, SFG, TP, and PHC (Fig 3A-B). For the semantic task, a leave-one-run-out cross-validation procedure was used, to ensure that independent data were used to define regions and extract responses.

The ROI analysis found that responses in category-sensitive subregions of association cortex were strongly selective for their preferred stimulus category. Person-preferring areas in MPC, MPFC, TPJ, STS, SFG, and TP responded significantly more strongly to images of familiar people versus unfamiliar people, objects, and familiar scenes; judgment about familiar people versus objects and familiar places; and imagining events involving familiar people versus objects and familiar places (all 42 P's < .0071 = .05 / 7, applying Bonferroni correction across the full set of 7 ROIs; individual values in Table S3; linear mixed effects model across runs, with participant included as a random effect). Similarly, place-preferring areas in MPC, MPFC, TPJ, STS, SFG, and PHC responded significantly more strongly to images of familiar scenes versus unfamiliar scenes, objects, and familiar people; judgment about familiar places versus objects and familiar people; and imagining events involved familiar places versus objects and familiar people (all 42 P's < .0071; individual values in Table S4). Of these areas, the only region with a response profile that countered our theoretical expectation was place-responsive STS, which had a stronger response to conditions involving language than to place conditions. Across ROIs, responses to non-preferred stimulus categories were typically at or below baseline. Strong domain selectivity was not only observed for the focal, maximally responsive areas used in this analysis, but across a range of ROI sizes, from the top 5-40% of category-preferring coordinates (Fig S13). Object-preferring regions also showed category-selective responses (Fig S14). These results demonstrate that selectivity for abstract stimulus category is a widespread feature of high-level association cortex.

Outside of the STS, person- and place-preferring areas were not modulated by the language task, despite the presence of nearby regions of association cortex

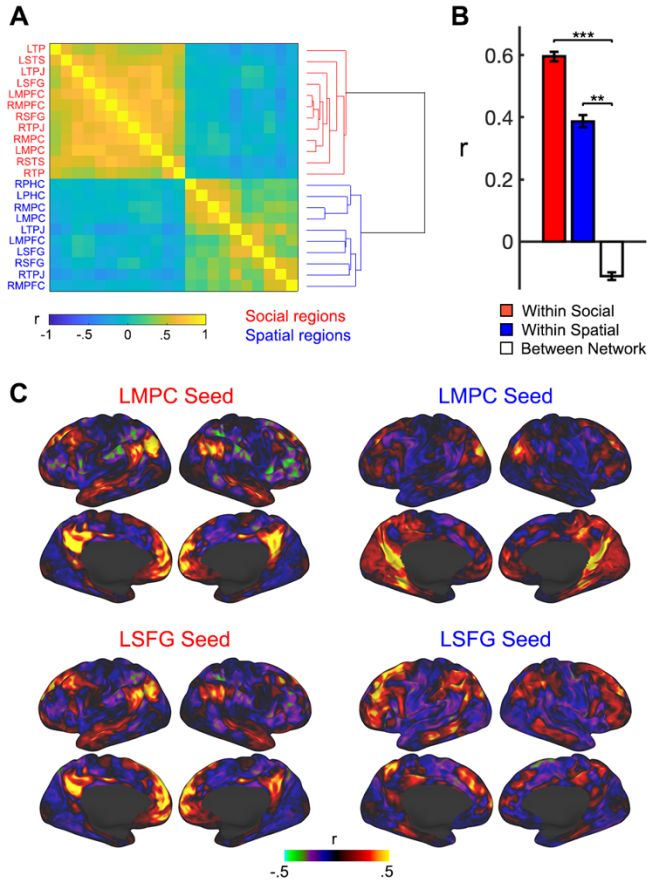

Figure 4: Person- and place-preferring areas constitute functionally coupled networks across association cortex. A) Matrix of resting-state correlations between functionally-defined ROIs, and hierarchical clustering results. **B)** Mean correlations within social and spatial networks, and between the two. Error bars show standard error across correlation values. ** $P < 10^{-3}$, *** $P < 10^{-4}$ (permutation test). **C)** Whole-brain correlation maps from person- and place-preferring LMPC and LSFG seeds, in one representative participant. Abbreviations: MPC, medial parietal cortex; MPFC, medial prefrontal cortex; TPJ, temporo-parietal junction; STS, superior temporal sulcus; SFG, superior frontal gyrus; TP, temporal pole; PHC, parahippocampal cortex.

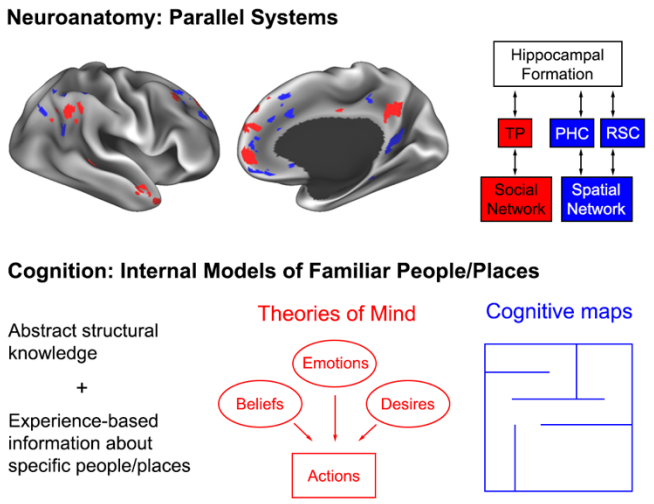

Figure 5: Theoretical model, parallel systems for social and spatial reasoning. Upper: Systems for reasoning about people and places have a parallel anatomical organization. Both are positioned with the cortical apex and interact closely with the hippocampal formation. Lower: on the cognitive level, these systems are argued to implement internal models of familiar people and places, which integrate abstract knowledge with experience-dependent information.

showing strong and selective responses during language comprehension (Fig S15). All person-preferring regions, but no place-preferring regions, were additionally sensitive to the ToM task, responding more when participants answered questions about false mental versus physical representations. Conversely, regions defined as maximally responsive in the ToM task also responded to familiar people over places (Fig S16). This suggests a common neural substrate for long term memory-based understanding of familiar people, and for abstract mental state judgments about unfamiliar people[8,29].

Having defined category-sensitive ROIs in association cortex, we next directly probed their functional connectivity, to determine whether these sets of regions constitute functionally coupled networks. Resting-state correlations were substantially stronger within person- or place-preferring areas than between (Fig 4A-B; permutation test; $P < 10^{-4}$, person vs between; $P < 10^{-3}$, place vs between). Hierarchical clustering of regions based on correlation distance revealed a dominant two-cluster solution, separating person- and place-preferring regions. These results justify the description of these sets of areas as social and spatial "networks" – functionally coupled regions distributed across cortex, with common stimulus selectivity.

How do social and spatial networks relate to previously described large-scale functional networks[21]? We next assessed whole-brain functional connectivity patterns of person- and place-preferring MPC and SFG subregions (Figs 4C, S16-19; other regions showed similar patterns). Person-preferring areas reliably showed functional connectivity with an apex subnetwork. Place-preferring areas showed a mixed pattern of functional connectivity, positively correlated with regions associated with apex, frontoparietal, and visual networks. Specific place-preferring areas differed in the strength of their network associations, with regions of MPC and PHC showing stronger correlations with the apex network, and regions of MPFC, SFG, and TPJ showing stronger correlations with the frontoparietal network. This difference was reflected in the hierarchical clustering solution, in which place-preferring regions were split into subclusters, with lower but still positive correlations between subclusters. The mixed pattern of connectivity shown by place-preferring areas may reflect the tendency of these areas, as noted above, to cut across multiple previously described large-scale networks.

Taken as a whole, these results demonstrate that social and spatial reasoning elicit domain-specific responses in parallel subnetworks within high-level association cortex. In addition to previously characterized place-preferring areas, we identify unknown responses in apex network areas SFG and MPFC. Involvement of MPFC in place understanding has been hypothesized[30] and supported by lesion evidence[31]. In addition to previously described socially responsive regions of association cortex[8,29], we identify a person-preferring region of temporal

pole, which could plausibly provide a direct interface between the social network and the hippocampal formation[24,25].

These results are consistent with our high-level theoretical perspective, proposing parallel cognitive and neural mechanisms for social and spatial reasoning, implemented by parallel networks in the cortical apex and medial temporal lobe. This parallel in neural systems may reflect similarities in the reasoning problems posed by understanding people and places. In navigating the physical world, we decide how to move around in an uncertain, constantly changing environment, in order to execute plans and optimize rewards. In navigating the social world, we decide what to say and do around people with uncertain, constantly changing internal states, in order to achieve personal and collective goals. In both cases, the mind arguably solves the problem by integrating abstract structural models (theories of mind, cognitive maps) with hippocampal-dependent information about specific people and places[15,17-19,32].

Our theoretical perspective suggests that areas in the cortical apex play a role in a range of psychological functions that are sometimes treated separately – from reasoning, to semantic and episodic memory, to decision-making. While our results are consistent with this perspective and provide evidence for domain-specificity and parallel anatomical organization, they don't provide direct evidence for a role in multiple functions. We refer to tasks as visual, semantic, and episodic based on the behavior involved, but it's possible that semantic or episodic memories were elicited during any of the tasks. Post-scan self-report, while variable across individual participants, found that memories were recalled at similar rates during familiar person and place conditions across tasks (Fig S4). Nevertheless, we do make the positive claim that the social episodic task – imagining familiar people discussing a common topic – involves episodic processing by design. The finding of no response to this task in place-preferring areas like PHC and MPC suggests against the hypothesis that these areas are broadly involved in episodic processing[10], instead favoring a domain-specific view. Human episodic recall and future-thinking typically involve both a spatial context, and social content – information about people, their relationships, internal states, and actions. On this view, episodic processing is expected to engage multiple apex subnetworks, with separate subnetworks processing information from different domains.

This work identifies a novel organizing principle in high-level association cortex: parallel, yoked responses to people and places. This finding inspires myriad questions for future study, including: how did these parallel subnetworks evolve, and how do they develop? The finding that these areas respond across a range of tasks, including visual tasks more easily amenable to studying infants and nonhuman primates, suggests an approach for addressing these questions.

Methods

Participants. Ten human participants (5 male, 5 female; age 28-40) were scanned using fMRI. Participants were healthy with normal or corrected vision, right-handed, and native English speakers. The experimental protocol was approved by the Rockefeller University Institutional Review Board, and participants provided written, informed consent.

Tasks. We used a range of perceptual and cognitive tasks involving familiar people and places. Participants were asked to choose six of their top ten most familiar people and places, and tasks involved processing these six people and places. Main tasks included: visual perception of familiar and unfamiliar faces and scenes, and generic objects; semantic judgment about familiar people and places, and generic objects; and episodic simulation of events involving familiar people and places, and generic objects (Fig 1A, S1). Additionally, localizer tasks from the existing literature were run, including tasks eliciting theory of mind (ToM; reading and answering questions about stories involving false beliefs or false physical representations)[33]; language comprehension (reading sentences or nonword lists)[34]; and dynamic visual perception (watching videos of moving faces, objects, and scenes)[35]. In each of three scans, one main task and one localizer were run (order: visual and ToM, semantic and dynamic perception, episodic and language). Stimulus presentation scripts can be found at https://osf.io/5yjgh/.

In the visual perception task, participants viewed serially presented naturalistic images of faces, objects, and scenes. For each participant, we obtained 20 images each of six familiar faces and scenes. Familiar face images were obtained directly from friends or family members, without the participant seeing them. Control images were defined as six yoked unfamiliar faces and scenes. Familiar and unfamiliar faces were matched on age group (young adult, middle-aged, old), race, and gender; familiar and unfamiliar scenes were matched on rough semantic category (e.g. outdoor street view; building interior). Object images were of six generic objects with varying physical properties - a banana, a baseball, a feather, a rock, a sponge, and a wrench. Face and object images contained no clear spatial structure (e.g. corners), and had minimal contextual cues beyond the background. Scene images contained no people. All five image categories were matched, for each participant, on the mean and variance across images of luminance, root-mean-square contrast, and saturation (in CIE-Lch space, with a D65 illuminant): all P's > .05, one-way ANOVA and Bartlett test. Images were presented at 768 x 768 resolution, 12.8 x 12.8 degrees of visual angle, for 1.85s each with a 150ms interstimulus interval, in 18s blocks of images of one identity. Participants performed a one-back task, pressing a button when an image was repeated. A post-scan questionnaire verified that participants could recognize the person or place in the familiar images (mean: 100% for faces, 84% for scenes), but not the unfamiliar images (mean: 0% for faces, 2% for scenes, Fig S4).

In the semantic task, participants rated traits of familiar people and places, and generic objects, on a 0 to 4 scale. This included personality traits of people (e.g. confident, angry, intelligent), spatial or navigational properties of places (e.g. cramped, large, has walls), and physical properties of objects (e.g. soft, heavy, rough; Table S1). Participants rated by moving an icon left or right, over 18s blocks of four questions for a given identity, for a total of 20 questions per condition and identity.

Prior to the episodic task, participants listed five common conversation topics for each familiar person, and five familiar subregions of each familiar place. In the scan, they were asked to imagine familiar people talking about common topics, navigating specific subregions of

familiar places, and objects engaged in physical interactions (e.g. rolling or sliding down a hill, falling into water). Participants were specifically asked to generate a novel event, rather than remember a past event. Imagination blocks lasted 18s, with a 3s verbal prompt, and a 1s hand icon at the end of the block, which participants responded to with a button press. After the scan, participants were asked (for 2/5 of trials) questions about difficulty, detail, visual imagery, emotional response, and memory recall (Fig S5). Additionally, a free response description was collected, to ensure that participants could describe what they imagined. 9/10 participants were able to describe all events; the remaining one recalled 75% of events.

Across the three main tasks, blocks were separated by 4s of resting fixation, and presented in five 8-13 minute runs per task, with palindromic block orders, counterbalanced across runs and participants. Fixation blocks were included in the beginning, middle, and end of the experiment to estimate a resting baseline. Localizer tasks were split into 4-5 minute runs, with four runs for theory of mind and language tasks, and six for dynamic perception. Scanner task performance was high, indicating sustained attention during scan sessions (Fig S3). After each scan, participants were asked to rate on a 0-10 scale to what extent they recalled memories of events during the person and place conditions (Fig S4).

Behavioral data analysis. To evaluate participants' self-reported experience in the episodic task, we compared behavioral responses across each pair of conditions (person, object, place), for each of the five questions. Statistics were computed using a linear mixed effects model (MATLAB's fitlme, unpaired, two-tailed comparisons) across items or event ratings, with random intercepts for participant.

MRI acquisition. Participants were scanned on a Siemens 3T Prisma across three 2.5-hour sessions, which included task acquisitions as well as 40 minutes of high-resolution anatomical images, 60 minutes of resting-state acquisitions, and spin echo acquisitions for distortion correction. Three each of T1- and T2-weighted anatomical images were acquired at 8mm resolution. Task and resting-state data were acquired using a multiband, multi-echo EPI pulse sequence, optimized to boost temporal signal-to-noise ratio (tSNR) throughout the brain and specifically in the anterior medial temporal lobes (TR = 2s, TE = 14.4, 33.9, 53.4, 72.9, and 92.4ms, 2.4x2.4x2.5mm resolution, 48 oblique axial slices with near whole brain coverage, multiband acceleration 3x, GRAPPA acceleration 2x, interleaved slice acquisition). 3-4 parameter-matched spin echo acquisitions were acquired per scan, between every four runs of task. Raw MRI data can be found at https://openneuro.org/datasets/ds003814.

MRI preprocessing. Data were preprocessed and analyzed using a custom pipeline, integrating software elements from multiple software packages: FSL (6.0.3), Freesurfer (7.1.1), AFNI, Connectome Workbench 1.5, tedana 0.0.10, and Multimodal Surface Matching (MSM). The code is available at https://github.com/bmdeen/fmriPermPipe/releases/tag/v2.0.1, with dataset-specific wrapper scripts at https://github.com/bmdeen/identAnalysis.

Anatomical images were preprocessed using an approach based on the Human Connectome Project (HCP) pipeline[36]. The three images for each modality were linearly registered using FLIRT[37] and averaged; registered from T2- to T1- weighted; aligned to ACPC orientation using a rigid-body registration to MNI152NLin6Asym space; and bias-corrected using the sqrt(T1*T2) image[38]. Cortical surface reconstructions and subcortical parcellations were generated using Freesurfer's recon-all[39]. Surface-based registration (MSMSulc) was used to

register individual surfaces to fsLR average space[40]. This registration was used to project the HCP multimodal cortical parcellation[41] onto individual surfaces.

Functional data were preprocessed using a pipeline tailored to multi-echo data, aiming to optimize tSNR while maintaining high spatial resolution. Motion parameters were first estimated using MCFLIRT[42]. Intensity outliers were removed using AFNI's 3dDespike, and slice timing correction was performed using FSL's slicetimer. Motion correction was then applied, in combination with topup-based distortion correction[43], and rigid registration to a functional template image, in a single-shot transformation with one linear interpolation, to minimize spatial blurring. Multi-echo ICA was performed using tedana, with manual adjustments, to remove non-blood-oxygen-dependent (non-BOLD) noise components[44]. Data were intensity normalized and resampled to an individual-specific CIFTI space aligned with the anatomical template image, with 32k density fsLR surface coordinates, and 2mm volumetric subcortical coordinates. Registration between the functional and anatomical templates was computed using boundary-based registration (bbregister)[45]. Surface data and subcortical volumetric data were both smoothed with a 2mm-FWHM Gaussian kernel. For resting-state data, the global mean signal was removed via linear regression, to diminish global respiratory artifacts not removed by multi-echo ICA[46].

Whole-brain analysis. Whole-brain statistical analyses were performed in individuals using AFNI's 3dREMLfit, modeling temporal autocorrelation with a coordinate-wise ARMA(1,1) model[47]. Results were thresholded using a false discovery rate of $q < .01$ (two-tailed), to correct for multiple comparisons across ~90K CIFTI coordinates[48]. We compared responses to people versus places, and objects versus places, across tasks (boxcar regressors convolved with canonical double gamma hemodynamic response function).

Interdigitation analysis. To test for spatially interdigitated responses, we used part of the dataset (even runs), combined across visual, semantic, and episodic tasks, to identify person- and place-sensitive parts of medial prefrontal and parietal cortex. We identified at least 3 pairs of regions in medial prefrontal cortex, and 2 pairs in medial parietal cortex, in each participant and hemisphere, and hand-drew anchor coordinates over these responses. We then defined paths corresponding to geodesics on the cortical surface between anchors, and extracted person and place responses (% signal change) in left-out data from odd runs. To combine these responses across participants, we interpolated responses along each path to standardized coordinate frame with 7 equidistant coordinates between each anchor.

Connectivity gradient analysis. To position functional responses along large-scale connectional gradients in cortex, we computed a low-dimensional embedding of resting-state functional connectivity similarity, using the diffusion maps algorithm[27]. Following Margulies et al.[7], we computed a pairwise correlation matrix between resting-state time series from ~59K cortical surface coordinates, kept only the top 10% of values per row, zeroed negative values, and computed cosine similarity between rows to generate a positive, symmetric similarity matrix. We then generated a diffusion map embedding, using $\alpha = .5$ (Fokker-Planck normalization), and automated estimation of diffusion time via spectral regularization, by multiplying Laplacian eigenvalues λ_i by $\frac{1}{1-\lambda_i}$. Across participants, the space spanned by the first two principal directions included previously described anatomical gradients[7]: one between sensory cortex and transmodal association cortex, and one between sensorimotor and visual cortex (Fig 4). Because

the orientation of these two gradients in 2D space varied across participants, a Procrustes transformation was used to map each participant's diffusion embedding to the first participant's, using the top three dimensions. To overlay functional responses on the diffusion embedding, we identified coordinates that were responsive to either 1) people over places, 2) places over people, or 3) objects over places, across visual, semantic, and episodic tasks (coordinate-wise $Z > 2.3$ for each task). These coordinates were displayed in color in diffusion space. Only a small proportion (<.2% across participants) of coordinates were responsive to the object contrast and person or place contrasts; these were colored as object-preferring.

Region-of-interest (ROI) analysis. ROI-based analyses were conducted to assess how person- and place-sensitive regions of the apex network respond across a range of task conditions. ROIs were defined as the top 5% of person- or place-preferring coordinates (semantic task) within anatomical search spaces defining parts of the apex network: medial frontal cortex, medial parietal cortex, temporo-parietal junction, superior frontal gyrus, superior temporal sulcus, parahippocampal cortex, and temporal pole. Search spaces were composed of regions from the HCP parcellation (Table S2). To extract responses to the task used to define ROIs, a leave-one-run-out analysis was performed, in which ROIs were defined in all but one run of data, and responses extracted in the left out run. Statistics were performed on percent signal change values across runs, using a linear mixed effects model (MATLAB's fitlme), with participant included as a random effect. Similar analyses were performed for object-sensitive regions, and for the language and theory of mind localizers. Analyses were also performed for a range of ROI sizes, from 5 to 40% (Fig S13).

For ROIs defined as person-preferring or by the language or ToM tasks, responses to people were compared with places and objects, across visual, semantic, episodic, and dynamic tasks. Additionally, responses to familiar/unfamiliar faces were compared. For place-preferring ROIs, responses to places were compared with people and objects across tasks, and responses to familiar/unfamiliar faces were compared. For object-preferring ROIs, responses to objects were compared with people and places, across tasks. For all ROIs, we compared responses to false belief versus false photo and sentence versus nonword list conditions. All comparisons were paired and one-tailed.

ROI-based functional connectivity analysis. Resting-state correlations were computed between person- and place-preferring areas defined similarly to the ROI analysis, but using separate search spaces for each hemisphere. Place-preferring STS subregions were not included, because such responses were not reliably observed across participants. The regions were hierarchically clustered based on correlation distance, and regions were ordered in a way that minimized distance between adjacent pairs, without separating clusters (Fig 4A; MATLAB's optimalleaforder). Within- versus between-network correlations were compared using a permutation test, permuting regions (10,000 iterations). Whole-brain correlation maps, with functional ROIs as seeds, were computed for person- and place-preferring regions of left MPC and SFG, within each individual participant. These two regions were chosen for presentation because the functional connectivity of other ROIs generally resembled one of these two.

References

1 Spelke, E. S. & Kinzler, K. D. Core knowledge. *Developmental science* **10**, 89-96 (2007).

2 Scholl, B. J. & Leslie, A. M. Modularity, development and 'theory of mind'. *Mind Language* **14**, 131-153 (1999).

3 Duchaine, B., Cosmides, L. & Tooby, J. Evolutionary psychology and the brain. *Current opinion in neurobiology* **11**, 225-230 (2001).

4 McClelland, J. L. & Thompson, R. M. Using domain-general principles to explain children's causal reasoning abilities. *Developmental Science* **10**, 333-356 (2007).

5 Elman, J. L. *et al. Rethinking innateness: A connectionist perspective on development*. (MIT press, 1996).

6 Eichenbaum, H. Is the rodent hippocampus just for 'place'? *Current opinion in neurobiology* **6**, 187-195 (1996).

7 Margulies, D. S. *et al.* Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences* **113**, 12574-12579 (2016).

8 Saxe, R. & Kanwisher, N. People thinking about thinking people: The role of the temporo-parietal junction in "theory of mind". *Neuroimage* **19**, 1835-1842 (2003).

9 Silson, E. H., Steel, A., Kidder, A., Gilmore, A. W. & Baker, C. I. Distinct subdivisions of human medial parietal cortex support recollection of people and places. *Elife* **8**, e47391 (2019).

10 DiNicola, L. M., Braga, R. M. & Buckner, R. L. Parallel distributed networks dissociate episodic and social functions within the individual. *Journal of neurophysiology* **123**, 1144-1179 (2020).

11 Barrett, L. F. & Satpute, A. B. Large-scale brain networks in affective and social neuroscience: towards an integrative functional architecture of the brain. *Current opinion in neurobiology* **23**, 361-372 (2013).

12 Binder, J. R. & Desai, R. H. The neurobiology of semantic memory. *Trends in cognitive sciences* **15**, 527-536 (2011).

13 Spreng, R. N., Mar, R. A. & Kim, A. S. The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis. *Journal of cognitive neuroscience* **21**, 489-510 (2009).

14 Premack, D. & Woodruff, G. Does the chimpanzee have a theory of mind? *Behav. Brain Sci.* **1**, 515-526 (1978).

15 Baker, C. L., Jara-Ettinger, J., Saxe, R. & Tenenbaum, J. B. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour* **1**, 0064 (2017).

16 Tolman, E. C. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189 (1948).

17 Peer, M., Brunec, I. K., Newcombe, N. S. & Epstein, R. A. Structuring knowledge with cognitive maps and cognitive graphs. *Trends in cognitive sciences* (2020).

18 Tse, D. *et al.* Schemas and memory consolidation. *Science* **316**, 76-82 (2007).

19 Rubin, R. D., Watson, P. D., Duff, M. C. & Cohen, N. J. The role of the hippocampus in flexible cognition and social behavior. *Frontiers in human neuroscience* **8**, 742 (2014).

20 Braga, R. M. & Buckner, R. L. Parallel interdigitated distributed networks within the individual estimated by intrinsic functional connectivity. *Neuron* **95**, 457-471. e455 (2017).

21 Thomas Yeo, B. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology* **106**, 1125-1165 (2011).
22 Epstein, R. A. & Baker, C. I. Scene perception in the human brain. *Annual review of vision science* **5**, 373-397 (2019).
23 Steel, A., Billings, M. M., Silson, E. H. & Robertson, C. E. A network linking scene perception and spatial memory systems in posterior cerebral cortex. *Nature communications* **12**, 1-13 (2021).
24 Insausti, R., Amaral, D. G. & Cowan, W. The entorhinal cortex of the monkey: II. Cortical afferents. *Journal of Comparative Neurology* **264**, 356-395 (1987).
25 Moran, M., Mufson, E. & Mesulam, M. M. Neural inputs into the temporopolar cortex of the rhesus monkey. *Journal of Comparative Neurology* **256**, 88-103 (1987).
26 Olson, I. R., Plotzker, A. & Ezzyat, Y. The enigmatic temporal pole: a review of findings on social and emotional processing. *Brain* **130**, 1718-1731 (2007).
27 Coifman, R. R. & Lafon, S. Diffusion maps. *Applied and computational harmonic analysis* **21**, 5-30 (2006).
28 Baldassano, C., Beck, D. M. & Fei-Fei, L. Differential connectivity within the parahippocampal place area. *Neuroimage* **75**, 228-237 (2013).
29 Thornton, M. A. & Mitchell, J. P. Consistent neural activity patterns represent personally familiar people. *Journal of cognitive neuroscience* **29**, 1583-1594 (2017).
30 McCormick, C., Ciaramelli, E., De Luca, F. & Maguire, E. A. Comparing and contrasting the cognitive effects of hippocampal and ventromedial prefrontal cortex damage: a review of human lesion studies. *Neuroscience* **374**, 295-318 (2018).
31 Bertossi, E., Aleo, F., Braghittoni, D. & Ciaramelli, E. Stuck in the here and now: Construction of fictitious and future experiences following ventromedial prefrontal damage. *Neuropsychologia* **81**, 107-116 (2016).
32 Whittington, J. C. *et al.* The Tolman-Eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell* **183**, 1249-1263. e1223 (2020).
33 Dodell-Feder, D., Koster-Hale, J., Bedny, M. & Saxe, R. fMRI item analysis in a theory of mind task. *Neuroimage* **55**, 705-712 (2011).
34 Fedorenko, E., Behr, M. K. & Kanwisher, N. Functional specificity for high-level linguistic processing in the human brain. *Proceedings of the National Academy of Sciences* **108**, 16428-16433 (2011).
35 Pitcher, D., Dilks, D. D., Saxe, R. R., Triantafyllou, C. & Kanwisher, N. Differential selectivity for dynamic versus static information in face-selective cortical regions. *Neuroimage* **56**, 2356-2363 (2011).
36 Glasser, M. F. *et al.* The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* **80**, 105-124 (2013).
37 Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**, 825-841 (2002).
38 Glasser, M. F. & Van Essen, D. C. Mapping human cortical areas in vivo based on myelin content as revealed by T1-and T2-weighted MRI. *Journal of Neuroscience* **31**, 11597-11616 (2011).
39 Dale, A. M., Fischl, B. & Sereno, M. I. Cortical surface-based analysis: I. Segmentation and surface reconstruction. *Neuroimage* **9**, 179-194 (1999).

40 Robinson, E. C. *et al.* MSM: a new flexible framework for Multimodal Surface Matching. *Neuroimage* **100**, 414-426 (2014).
41 Glasser, M. F. *et al.* A multi-modal parcellation of human cerebral cortex. *Nature* **536**, 171-178 (2016).
42 Power, J. D., Plitt, M., Kundu, P., Bandettini, P. A. & Martin, A. Temporal interpolation alters motion in fMRI scans: Magnitudes and consequences for artifact detection. *PloS one* **12**, e0182939 (2017).
43 Andersson, J. L., Skare, S. & Ashburner, J. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage* **20**, 870-888 (2003).
44 Kundu, P., Inati, S. J., Evans, J. W., Luh, W.-M. & Bandettini, P. A. Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI. *Neuroimage* **60**, 1759-1770 (2012).
45 Greve, D. N. & Fischl, B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* **48**, 63 (2009).
46 Power, J. D. *et al.* Ridding fMRI data of motion-related influences: Removal of signals with distinct spatial and physical bases in multiecho data. *Proceedings of the National Academy of Sciences* **115**, E2105-E2114 (2018).
47 Olszowy, W., Aston, J., Rua, C. & Williams, G. B. Accurate autocorrelation modeling substantially improves fMRI reliability. *Nature communications* **10**, 1-11 (2019).
48 Benjamini, Y. & Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* **57**, 289-300 (1995).

Data and code availability. Raw data are available at https://openneuro.org/datasets/ds003814. Stimulus materials are available at https://osf.io/5yjgh/. Analysis code is available at https://github.com/bmdeen/fmriPermPipe/releases/tag/v2.0.1 (generic analysis tools) and https://github.com/bmdeen/identAnalysis (dataset-specific wrapper scripts).

Acknowledgements. We thank the staff at the Cornell Citigroup Biomedical Imaging Center for assistance with data acquisition, and Charles Lynch for helpful discussion on pulse sequence design. This work was supported by fellowships from the Helen Hay Whitney and Leon Levy foundations (B.D.), and the Center for Brains, Minds & Machines, funded by National Science Foundation STC award CCF-1231216 (W.A.F.).

Author Contributions. B.D. and W.A.F. conceived and designed the experiment. B.D. collected and analyzed the data. B.D. and W.A.F. wrote the paper.