



Dynamic neural architecture for social knowledge retrieval

Yin Wang^{a,1}, Jessica A. Collins^b, Jessica Koski^c, Tehila Nugiel^c, Athanasia Metoki^a, and Ingrid R. Olson^{a,1}

^aDepartment of Psychology, Temple University, Philadelphia, PA 19122; ^bFrontotemporal Dementia Unit, Department of Neurology, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02114; and ^cDepartment of Psychology, University of Texas at Austin, Austin, TX 78712

Edited by Rebecca Saxe, Massachusetts Institute of Technology, Cambridge, MA, and accepted by Editorial Board Member Michael S. Gazzaniga February 16, 2017 (received for review December 27, 2016)

Social behavior is often shaped by the rich storehouse of biographical information that we hold for other people. In our daily life, we rapidly and flexibly retrieve a host of biographical details about individuals in our social network, which often guide our decisions as we navigate complex social interactions. Even abstract traits associated with an individual, such as their political affiliation, can cue a rich cascade of person-specific knowledge. Here, we asked whether the anterior temporal lobe (ATL) serves as a hub for a distributed neural circuit that represents person knowledge. Fifty participants across two studies learned biographical information about fictitious people in a 2-d training paradigm. On day 3, they retrieved this biographical information while undergoing an fMRI scan. A series of multivariate and connectivity analyses suggest that the ATL stores abstract person identity representations. Moreover, this region coordinates interactions with a distributed network to support the flexible retrieval of person attributes. Together, our results suggest that the ATL is a central hub for representing and retrieving person knowledge.

person knowledge | anterior temporal lobe | person identity node | semantic memory | social neuroscience

As social creatures, it is essential that we develop a rich storehouse of knowledge about other members of our social network, such as who they are, how they look and sound, where they live, and what they do for a living. However, little is known about how and where such “person knowledge” is represented, stored, and retrieved in the brain. This inquiry is challenging because person knowledge is highly multimodal and multifaceted, being linked to both abstract features such as personality and social status as well as more concrete features such as eye color; in addition, familiar individuals are associated with detailed episodic and semantic memories (e.g., memories of shared experiences and biographic information) (1, 2). The neural circuit for person knowledge must therefore have the ability to combine multiple sources of information into an abstract representation accessible from multiplicative cues.

An influential theory by Burton and Bruce (3) proposes that person recognition is achieved through a hierarchical process that begins with the activation of modality-specific recognition units that selectively respond to the presence of a known face, name, or voice. This information is then sent to an amodal person identity node (PIN) that integrates information from the modality-specific recognition units into a multimodal representation for that individual. Excitation of the PIN ultimately allows the retrieval of person-specific semantic information independently of stimulus modality (4, 5). A similar design is embedded in the “hub-and-spoke” theory of semantic knowledge, which proposes that different features of a concept (such as its color or taste) are distributed throughout the brain (the “spokes”) and that a centralized “hub” integrates these features into a coherent, modality-invariant concept (6–8).

Here, we tested the hypothesis that a region in the anterior temporal lobe (ATL) has a function akin to a person identity node (1, 9, 10), subserving access to abstract person identity representations that can be retrieved from multiple cues, such as

when you see a photo of Graceland and it cues your knowledge of Elvis (study 1). Next, we asked whether the ATL acts as a neural switchboard, performing in concert with other brain regions to enable the retrieval of different facets of person knowledge in a flexible and context-appropriate manner (study 2). We focus on the ATL because multiple lines of evidence from neuropsychology, electrophysiology, and neuroimaging have documented the critical role of the ATL in person identification (4, 5, 11–16), person-related learning (10, 17–21), semantic memory (6–8), and abstract social knowledge (1, 22–33). Individuals with ATL damage due to resection or stroke have multimodal person recognition deficits (34), lose access to stored knowledge about familiar people (35, 36), and have difficulties learning information about new people (4, 22, 37, 38). A subregion of the ATL contains a face-sensitive patch, first identified in monkeys, and more recently in humans (9).

Across two fMRI studies, 50 participants learned biographical information about a group of fictitious people for 2 d and then completed a person memory test in the MRI scanner on day 3. In study 1, we used stimuli from different categories commonly associated with familiar people (e.g., faces, names, homes, and objects) to cue memories for specific individuals (Fig. 1A). We compared the similarity of response patterns elicited by stimuli from different categories but associated with the same individual to test the abstract person representation properties of the ATL. In study 2, we asked participants to recollect specific content of fictitious people’s biographies (Fig. 1B) and examined whether the ATL coordinates the retrieval of different aspects of person knowledge by recruiting the activation of different brain regions depending on task requirements.

Significance

Knowledge about other people is critical for group survival and may have unique cognitive processing demands. Here, we investigate how person knowledge is represented, organized, and retrieved in the brain. We show that the anterior temporal lobe (ATL) stores abstract person identity representation that is commonly embedded in multiple sources (e.g. face, name, scene, and personal object). We also found the ATL serves as a “neural switchboard,” coordinating with a network of other brain regions in a rapid and need-specific way to retrieve different aspects of biographical information (e.g., occupation and personality traits). Our findings endorse the ATL as a central hub for representing and retrieving person knowledge.

Author contributions: Y.W., J.A.C., J.K., and I.R.O. designed research; Y.W., T.N., and A.M. performed research; I.R.O. was principal investigator of the laboratory; Y.W. analyzed data; and Y.W., J.A.C., and I.R.O. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. R.S. is a Guest Editor invited by the Editorial Board.

See Commentary on page 4042.

¹To whom correspondence may be addressed. Email: mirroneuronwang@gmail.com or iolson@temple.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1621234114/-DCSupplemental.

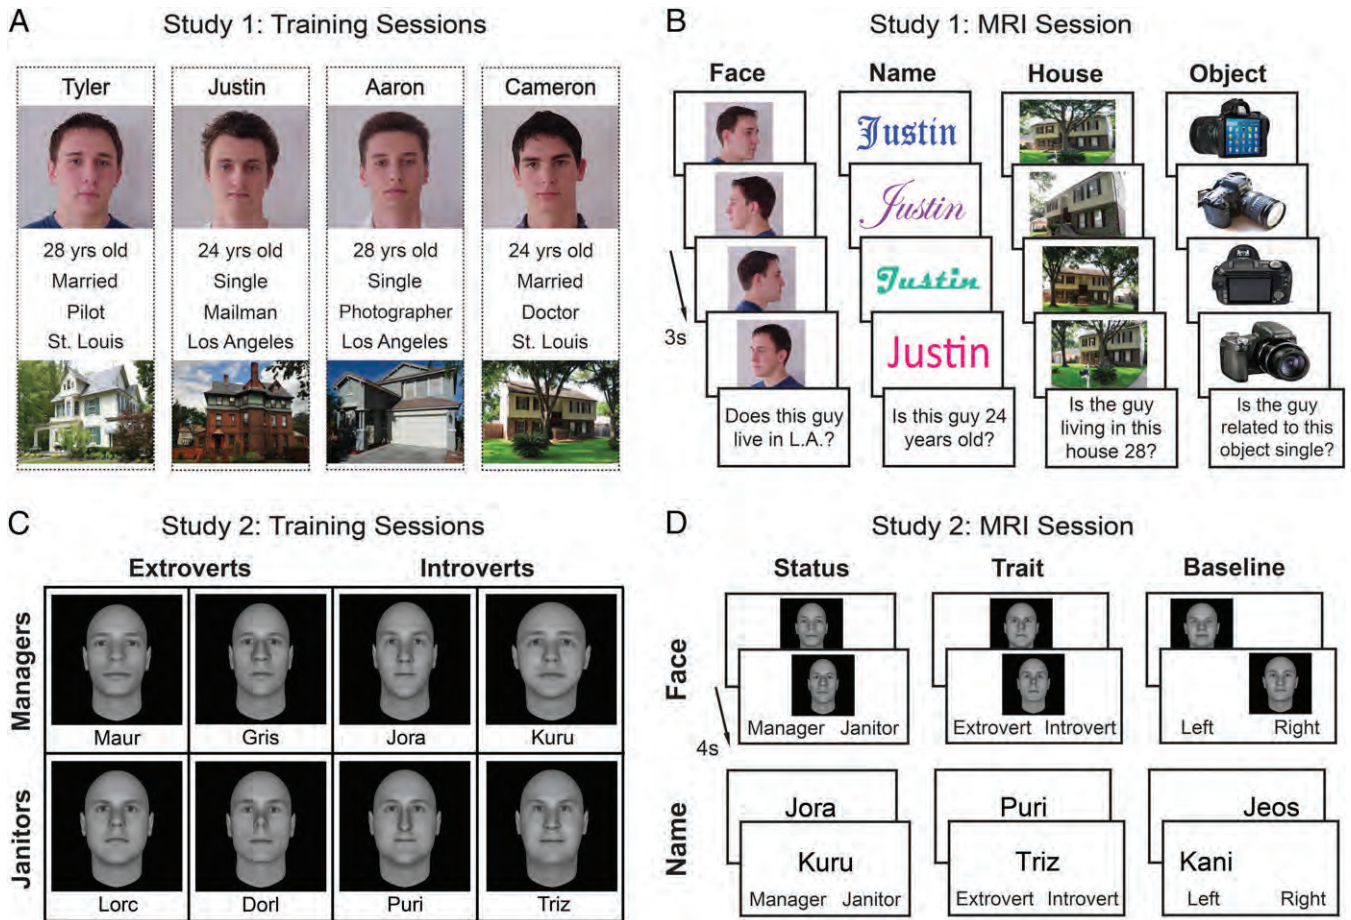


Fig. 1. Training material and fMRI task depiction. (A) In study 1, participants learned biographical details (i.e., face, name, age, marital status, occupation, city of residence, and current house) of four fictitious males. (B) In the scanner, participants completed a person memory task where they first viewed a series of stimuli that cued a particular fictitious male and then answered a question about that person. Each run presented stimuli only in one specific category (i.e., face run, name run, house run, or object run). (C) In study 2, participants learned different biographic details (i.e., face, name, occupational status, and personality trait) associated with a different set of fictitious males. Status and trait information were manipulated in a 2 × 2 factorial design. (D) In the scanner, participants completed a person memory task in which they had to indicate the status or personality trait associated with the learned males cued by either a particular face or name. Participants also performed a control condition task (i.e., nonmemory baseline) where they saw unfamiliar faces and names and indicated the spatial location of the stimuli on the screen.

In study 1, findings from a series of multivariate pattern analyses (MVPAs) consistently suggested that portions of the ATL represent person knowledge in an abstract form that is divorced from the ground state of a person’s face or name. Specifically, the ATL contains multivoxel patterns for facial identity that are similarly sensitive to identity accessed through a person’s name, an image of their home, or even an object associated with that person’s occupation. In study 2, we found that the ATL person identity node is embedded in a neural circuit that is consistently engaged during person memory tasks. Multivariate analyses suggested that different content areas of person knowledge—social status, personality traits, and identity—were represented in discrete nodes within this distributed person identification circuit. Connectivity analyses further revealed that the ATL may serve as a “neural switchboard” and is capable of coordinating the flow of person-specific information between sensory brain regions that encode incoming cues and other nodes of this circuit that are engaged when retrieving specific person knowledge content.

Results

Study 1: The ATL Is a Person Identity Node. In study 1, we used two-way MVPA cross-category (CC) classification (39–42) to identify

regions in which abstract and category-invariant person information is represented. The classifier was first trained to discriminate individual identities using one category of cues (e.g., that person’s face) and subsequently tested on a different category of cues (e.g., that person’s name), and vice versa (*Materials and Methods*). In this way, only the abstract conceptual information that was general to both stimulus categories was informative to the classifier. Three different decoding approaches were performed: regions of interest (ROI-based), combinatorial (43, 44), and whole-brain searchlight (11).

In the ROI-based analyses, we tested whether identity-specific information in face-selective regions can generalize across different categories of cues that have been linked to a specific person (e.g., Bill’s face ↔ Bill’s name, Bill’s face ↔ Bill’s house, Bill’s face ↔ Bill’s object). A separate functional localizer was used to identify face-selective regions in both hemispheres including the ATL, occipital face area (OFA), fusiform face area (FFA), amygdala (AMY), and orbitofrontal cortex face patch (OFC), as well as a place-selective region [i.e., parahippocampal place area (PPA)] and a control region in early visual cortex (V1) (Fig. S1A). In the current paradigm, the face-selective ATL region was the only ROI that showed above-chance CC classification accuracy for “face ↔ name” [left ATL: $t_{(23)} = 3.114, P = 0.030$;

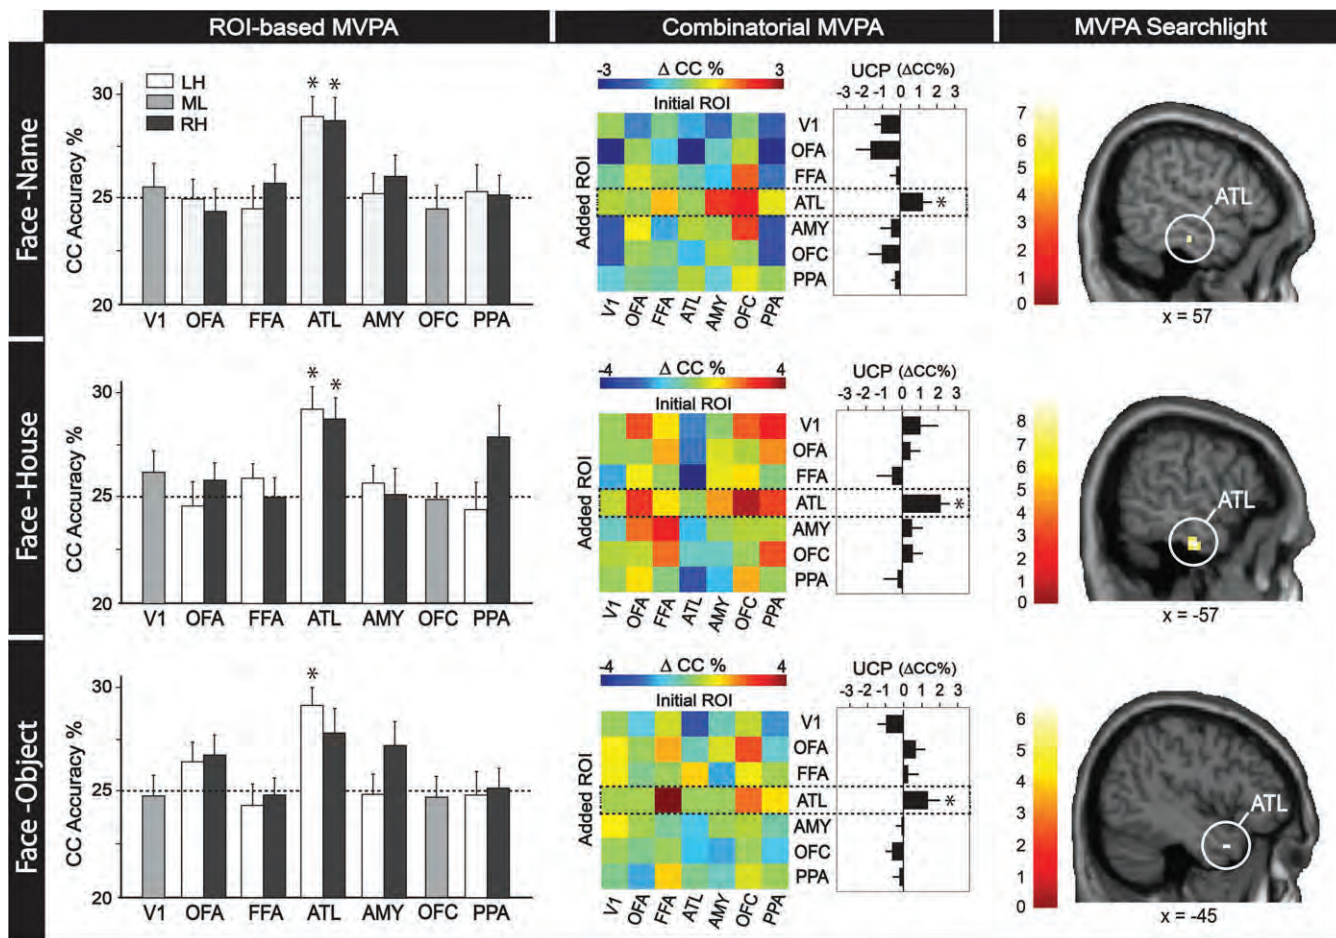


Fig. 2. MVPA analyses from study 1. Each column contains data from each of three decoding approaches (i.e., ROI-based, combinatorial, and searchlight); each row contains data from each of three cross-category (CC) classifications (i.e., face-name, face-house, and face-object). Three different decoding strategies were used to assess whether results were robust and generalized across different analytic methods. Asterisks indicate statistical significance, and error bars denote SE. Abbreviations: LH, left hemisphere; ML, midline; RH, right hemisphere; UCP, unique combinatorial performance.

right ATL: $t_{(23)} = 3.941, P = 0.006$], “face ↔ house” [left ATL: $t_{(23)} = 3.952, P = 0.006$; right ATL: $t_{(23)} = 4.346, P = 0.001$], and “face ↔ object” [left ATL: $t_{(23)} = 3.555, P = 0.012$] identity representation (Fig. 2). None of the face-selective ROIs were able to cross-classify individual identities based on “nonface” stimulus pairs (i.e., “name ↔ house,” “name ↔ object,” “house ↔ object”), likely due to the face-selective nature of these ROIs (Fig. S2).

Next, we used a combinatorial MVPA analysis to validate the unique role of the face-selective ATL region (among all face-selective ROIs) for the category-invariant representation of person identity. Combinatorial analyses allow the estimation of the unique information carried within an ROI (43, 44). For every ROI in the present analysis, we calculated the average increase (or decrease) in CC classification accuracy when paired with other ROIs, referred to as the unique combinatorial performance (UCP). A positive UCP means that an ROI carries information that contributes to classification beyond what is obtained from other ROIs. A negative UCP means that information within an ROI is largely redundant with that of other ROIs. As expected, the results suggested that only the ATL displayed significantly positive UCP for face ↔ name [UCP = 1.027, $t_{(5)} = 2.879, P = 0.018$], face ↔ house [UCP = 2.083, $t_{(5)} = 4.471, P = 0.004$], and face ↔ object [UCP = 1.273, $t_{(5)} = 2.063, P = 0.047$] identity representation (Fig. 2).

Finally, we conducted searchlight analyses to explore whether any regions in the brain beyond our face-selective ROIs can

represent person identity across stimulus categories. Consistent with the previous two approaches, significant clusters in the ATL for face ↔ name, face ↔ house, and face ↔ object CC classification (Fig. 2) were revealed in this analysis. Although we did not find any regions that represented identity across name ↔ house, name ↔ object, and house ↔ object at the standard threshold, further exploratory analyses suggest that the ATL was the only subthreshold region showing CC properties for all nonface CC pairs when using a lenient threshold (Fig. S2 and Table S1).

Taken together, the CC MVPA analyses in study 1 suggest that the ATL may be the only brain region that holds an abstract representation of person identity that can be accessed via multiple stimulus categories. Previous studies using face or voice stimuli have found amodal representations of person identity in unimodal and multimodal face-processing regions, including the FFA, superior temporal sulcus, and the ATL (5, 16, 42, 45). In the present paradigm, the only region consistently demonstrating classification of identity across all stimulus types was the ATL. In addition, we found that the bilateral ATLS engage in cross-modal representation, whereas hemispheric asymmetry has been previously reported in the ATLS for unimodal stimuli (1, 31, 46, 47).

Study 2: The ATL Is the Hub of Flexible Person Memory Retrieval.

Univariate analyses of the neural circuit for person memory. In study 2, we first used univariate analyses to identify brain regions that

were strongly engaged when retrieving information associated with an individual (Fig. S3). The contrast of “all memory conditions > baseline” revealed a set of brain regions that included the ATL, inferior parietal lobe (IPL), posterior cingulate cortex (PCC), and hippocampus (HIPP). A similar set of brain regions was found for the contrasts of “status memory > baseline” and “trait memory > baseline.” When comparing memory conditions cued by two different categories (face vs. name), we found that the right fusiform gyrus (i.e., FFA) was activated more during face-cued memory, whereas left fusiform gyrus [also known as the visual word form area (VWFA)] was activated more during name-cued memory.

MVPA analyses for person knowledge representations. The univariate analyses delineated a network of brain regions (ROIs) important for person knowledge retrieval (i.e., left ATL, right FFA, left VWFA, left IPL, PCC, and left HIPP) (Fig. S1B). Next, we used MVPA analyses to examine whether distinct regions within this network represent different aspects of knowledge associated with a person (e.g., that person’s status, personality traits, or identity). As knowledge content associated with an individual per se should be invariant to changes in the stimulus category used to cue that person, we collapsed our MVPA analysis of memory

condition across the two cue categories (i.e., face and name). Here, we used the three decoding approaches used in study 1 and a “leave-one-run-out” cross-validation (CV) classification scheme to identify brain regions sensitive to different aspects of person knowledge.

For status and personality-trait knowledge, ROI-based analyses suggested that the IPL was the only region within the person-knowledge network that accurately represented learned people’s occupational status [i.e., manager vs. janitor; $t_{(25)}=4.500$, $P < 0.001$], whereas the PCC was the only region that accurately represented trait information [i.e., whether a person was extroverted vs. introverted; $t_{(25)} = 3.192$, $P = 0.012$] (Fig. 3). Combinatorial analyses confirmed these findings by showing the only positive UCP in IPL for status information [UCP = 2.372, $t_{(4)} = 3.767$, $P = 0.01$] and the highest positive UCP in PCC for personality trait knowledge [UCP = 2.628, $t_{(4)} = 3.982$, $P = 0.008$]. Searchlight analyses further identified significant clusters in IPL for status classification and PCC for trait classification (Table S2). These consistent results are in line with previous neuroimaging studies suggesting a critical role of the IPL in representing social status (48, 49) and of the PCC in representing personality traits (50–52).

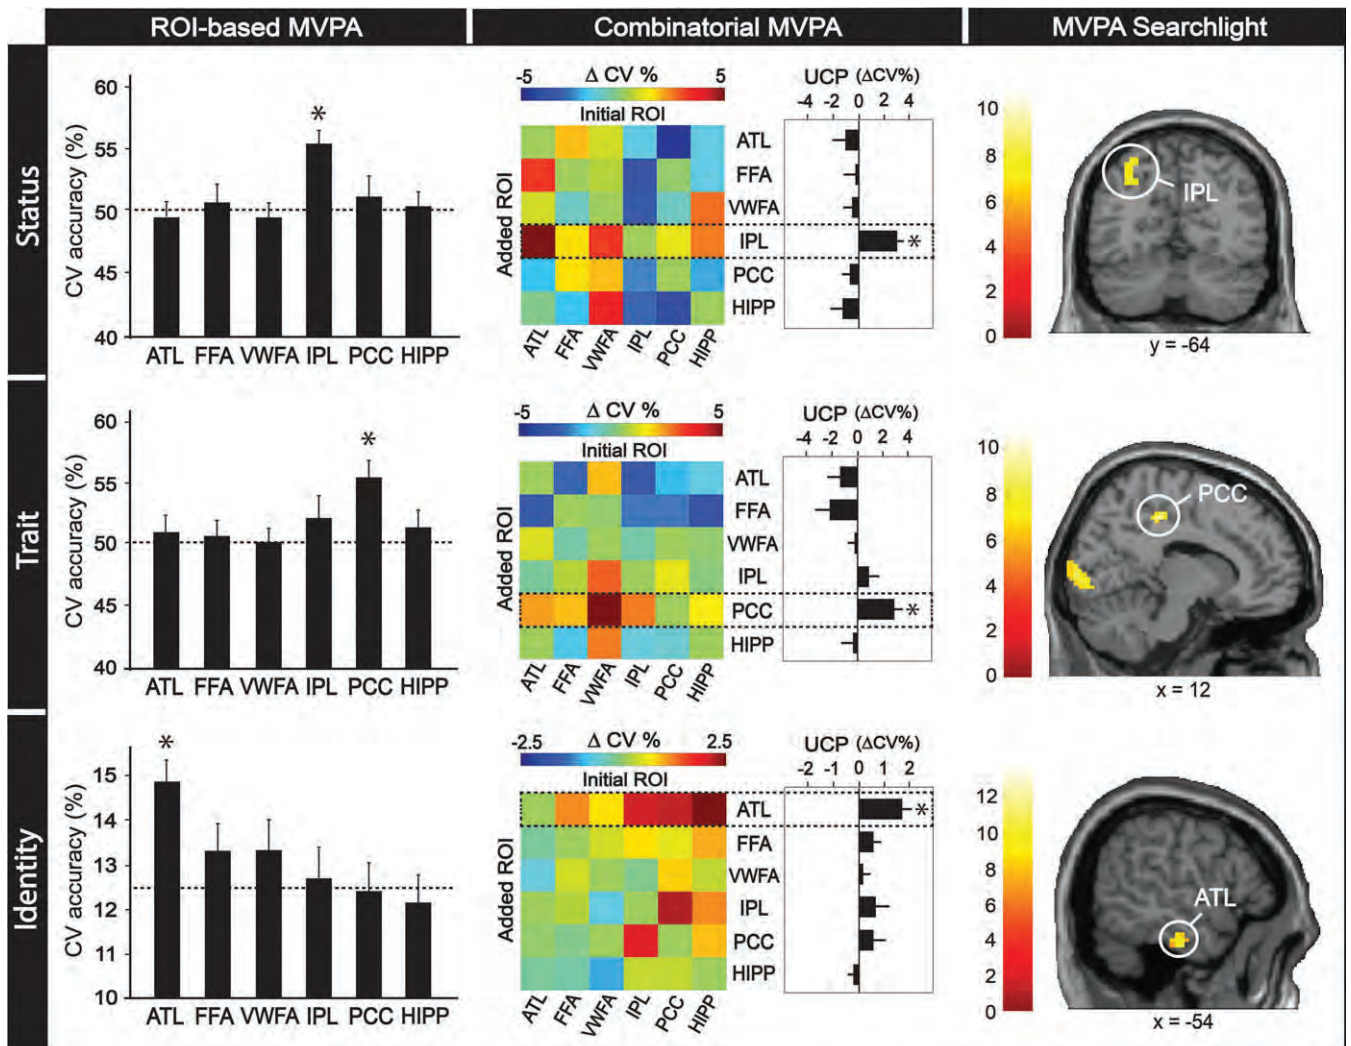


Fig. 3. MVPA analyses from study 2. Each column contains data from each of three decoding approaches (i.e., ROI-based, combinatorial, and searchlight); each row contains data from cross-validation (CV) classification for three types of knowledge representations (i.e., status, personality trait, and identity). Asterisks indicate statistical significance, and error bars denote the SE. Abbreviations: UCP, unique combinatorial performance.

For identity knowledge, we observed that the ATL was the only region that could accurately distinguish the eight learned people in ROI-based analyses [$t_{(25)} = 3.166$, $P = 0.012$], combinatorial analyses [UCP = 1.811, $t_{(4)} = 5.232$, $P = 0.003$], and searchlight analyses (Fig. 3 and Table S2). Along with the results of study 1, our MVPA analyses provide convergent evidence suggesting that the ATL is a critical region for person identity representation, regardless of whether the decoding was performed across categories (e.g., face/name/scene/object in study 1) or with categories collapsed together (e.g., face/name in study 2).

Overall, the MVPA analyses in study 2 suggest that a distributed network represents specific content of person knowledge: status information was represented in the IPL, trait information in the PCC, and, consistent with the results of study 1, person identity was represented in the ATL.

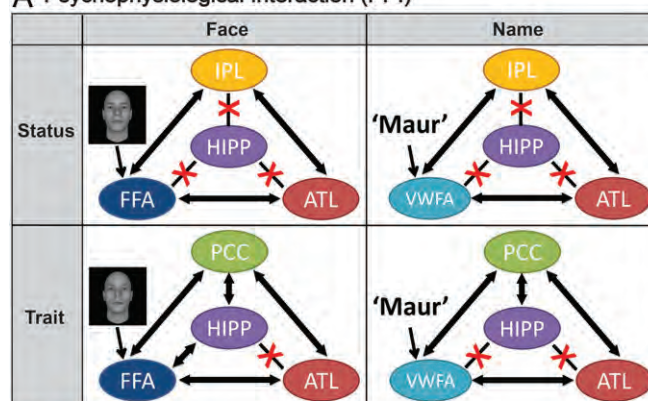
Functional connectivity analyses for neural dynamics of person memory. To elucidate the underlying neural dynamics supporting person knowledge retrieval, we performed a psychophysiological interaction (PPI) analysis to examine functional connectivity between nodes of the person-memory network during the retrieval of person-specific knowledge. PPI measures task-dependent interactions between different brain regions and by examining the similarity of activity patterns (“connectivity”) between a seed and other brain areas as a function of specific task demands (53). We took a comprehensive approach to defining the seed regions for our PPI analysis using regions identified by our univariate and MVPA analyses as being important for the distributed representation of person-related knowledge (Fig. S1B). We defined three types of seeds for the PPI analyses: sensory cue regions (FFA, VWFA), knowledge representation regions (IPL, PCC), and the ATL (Fig. S4).

Because information generally flows from posterior to more anterior regions during person perception (54), we first used sensory regions (FFA or VWFA) as seeds to explore where in the brain the sensory input is forwarded to, and then used regions in association cortex that may be involved in knowledge representation (IPL or PCC) as seeds to examine which brain area(s) are engaged in retrieving knowledge content. The results of our PPI analysis suggest that sensory regions had stronger coupling with the IPL and ATL during the retrieval of a person’s status, but with the PCC and ATL during the retrieval of a person’s personality traits (Fig. S4A). Knowledge representation regions showed enhanced coupling with the FFA and ATL when cued by faces, but with the VWFA and ATL when cued by names (Fig. S4B). The hippocampus also showed increased coupling with these two groups of seed regions, but not in all conditions (Fig. 4A and Tables S3 and S4).

We also used the ATL as a seed region to more specifically test whether its functional connectivity with other nodes of the person-memory network changed as a function of task demands. The ATL showed increased connectivity with the IPL during the retrieval of a person’s status, but with the PCC during the retrieval of a person’s personality traits. The functional connectivity of the ATL with the FFA increased during memory retrieval when cued by faces, but with the VWFA when cued by names (Fig. S4C and Table S5). The functional connectivity between the ATL and the hippocampus did not change in any conditions.

In sum, the results of our PPI analyses suggest that functional activity in secondary sensory regions (i.e., FFA, VWFA), knowledge representation regions in association cortex (i.e., IPL, PCC), and the ATL is tightly coupled during the retrieval of person-specific information, and that different regions within this circuit are differentially coupled depending on the content of the knowledge being retrieved. In contrast, the hippocampus, a region that has an undisputed role in the formation of episodic memories (55), does not appear to interact with other regions during the retrieval of person knowledge (Fig. 4A).

A Psychophysiological Interaction (PPI)



B Dynamic Causal Modelling (DCM)

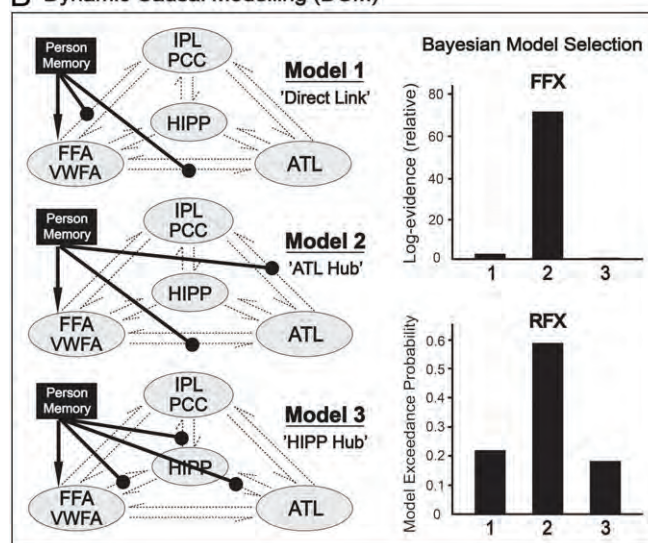


Fig. 4. Psychophysiological interaction (PPI) and dynamic causal modeling (DCM) analyses. (A) Overall summary of functional coupling between network regions in four different memory conditions. Bold double arrows indicate increased connectivity between two regions, whereas red crosses indicate nonsignificant connectivity change. For more detailed results, see Fig. S4, and Tables S3–S5. (B) Three distinct models were compared in DCM, and the optimal one was determined by Bayesian model selection. Both FFX and RFX analysis suggested model 2 (ATL hub) as the optimal model. For more details of model specification, see Fig. S5.

Dynamic causal modeling to test the “ATL-hub” theory. PPI analysis does not provide information about the direction of causal influences between source and target regions, nor whether the connectivity is mediated by other regions. To gain traction on these issues, we performed dynamic causal modeling (DCM) to explore detailed information processing dynamics within the aforementioned network. We tested three models that might support the implementation of person knowledge retrieval. In model 1, retrieval is implemented by a direct link between sensory regions and knowledge representation regions, which could be formed by associative learning during the 2-d training session that preceded scanning. Alternatively, in model 2, the ATL acts as a “switchboard-like” hub to flexibly coordinate the flow of information when retrieving different content of person knowledge. In model 3, the hippocampus, rather than the ATL, acts as a domain-general memory hub (see *Materials and Methods* and Fig. S5 for details regarding the model specifications).

Bayesian model selection clearly selected model 2 as the optimal model (Fig. 4B): it has the highest relative group log evidence

in fixed-effects analysis (71.29) as well as the highest exceedance probability in random-effects analysis (59.23%). An analysis of absolute model fits confirmed that model 2 provides the highest explained variance among the three models, accounting for $24 \pm 6\%$ (mean \pm SD) of the observed variance. Therefore, our DCM support a neural mechanism of person memory in which the ATL serves as the central hub coordinating the person-information retrieval from discrete sources.

Discussion

The results of the two studies reported here suggest that the ATL plays a critical role in representing and retrieving person knowledge. Using MVPA CC classification, study 1 provided strong evidence that the ATL face patch has a function akin to a person identity node, representing abstract conceptually invariant person identity information. Research in nonhuman primates has shown that cells in anterior-ventral temporal cortex are highly sensitive to particular facial identities as well as to facial familiarity (56, 57). Previous studies in humans using intracranial recording (16) or fMRI analyses (15, 42, 45) have suggested that the ATL can distinguish between different people using their faces, voices, or names. Our study extends these findings by using sophisticated multivariate analyses and a wider range of stimulus categories. Based on our results, we suggest that the ATL may be the only region that can merge such a wide variety of multicategorical person information.

The results of study 2 shed light onto the neural architecture of person knowledge. Study 2 provides strong support for the idea that person knowledge is organized in a hub-and-spoke manner such that specific features of person knowledge are neurally distributed and portions of the ATL serve as a convergence zone or hub (6–8). MVPA analyses suggested that status and personality trait information are stored separately in portions of the IPL and PCC, respectively, and PPI analyses showed that the ATL acts in concert with these regions to flexibly retrieve different aspects of person knowledge. Moreover, when a particular person memory was cued by a face, there was enhanced coupling between the FFA and ATL, whereas if the same memory was cued by a name, there was enhanced coupling between the visual word form area and the ATL. Critically, we performed DCM to directly test the hub-and-spoke accounts (i.e., model 2 and 3) against an alternative “distributed-only” account (5) (i.e., model 1), and the ATL-hub model was dominantly favored in the comparison.

Functional coupling as measured by PPI and DCM does not necessitate a direct anatomical connection between regions (53). However, the ATL is viewed as a convergence zone precisely because of its unusual pattern of white matter connectivity (58). It receives direct input from the ventral visual stream, which includes the FFA and the VWFA, via the inferior longitudinal fasciculus. Portions of the ATL are also directly connected to the PCC via a limbic pathway, the cingulum bundle, whereas superior aspects of the ATL are structurally interconnected with the IPL by a pathway most closely associated with language, the middle longitudinal fasciculus (59). Thus, the results of the PPI and DCM analyses rest on a verified neurostructural ground truth.

The hub-and-spoke model is a popular account of general semantic knowledge (6–8), and our DCM findings generally support that model. However, one aspect of our findings is challenging for this model. The hub-and-spoke model predicts that, within the ATL, there should be a single person identity node that links all features of all stimuli associated with a particular person. Across ROI-based and searchlight analyses in study 1, we did find that the ATL contains convergence zones (60) that bind different categories of cues related to an identity; however, we were unable to identify a single site—a “master node”—within the ATL that linked all four categories (faces, houses, names, and objects) and survived all six CC pairings.

These findings are consistent with neuropsychological evidence that patients with ATL lesions often have person identification deficits from one or two modalities (46, 61), but no patient has been shown to have impairments in all (three or more) modalities at the same time (4). Moreover, our ATL ROI seemed to be biased toward faces: we found CC classifications only worked well for face-related pairings (i.e., face \leftrightarrow name, face \leftrightarrow house, and face \leftrightarrow object; Fig. 2), but not for nonface pairing (i.e., name \leftrightarrow house, name \leftrightarrow object, and house \leftrightarrow object; Fig. S2). One possibility is that known faces serve as singular categories (62), grouping and linking the bits of biographical knowledge and descriptors that define a person, but in a manner that is inherently asymmetrical. Another possibility is that our training regime created a person representation biased toward faces.

It is important to consider that biographical knowledge bears all of the signatures of semantic memory with little to none of the signatures of episodic memory, such as spatiotemporal context. It has been proposed that the learning of semantic associations can bypass the hippocampus entirely and, instead, rely on the ATL (63). Discrete hippocampal damage slows the learning rate of semantic associations but does not impair it completely (64). Indeed, individuals with hippocampal damage can learn new semantic associations through a type of incidental learning called “fast mapping,” whereas patients with ATL lesions cannot (65). Consistent with this small literature, our results suggest that the hippocampus plays a nonsignificant role in the retrieval of biographical information. The results of our univariate analysis showed that the hippocampus was active across all memory conditions so it was included in our network analyses. However, the functional connectivity analysis suggested that it had only minimal interactions with other brain regions during person knowledge retrieval (Fig. 4A), and DCM ruled out any critical role of the hippocampus in person knowledge retrieval (Fig. 4B). We speculate that the hippocampus may be important during the initial acquisition of semantic associations, but within a short amount of time, it is no longer needed and, instead, the ATL becomes the nexus for retrieval (21). Thus, the ATL, which is geographically proximal to the hippocampus, and structurally interconnected with the hippocampus, may play a vital role in orchestrating the cascade of neural events that ultimately result in the retrieval of person knowledge.

There are several limitations of the current investigation that should be noted. First, across the two studies reported here, we tested different groups of subjects and defined “ATL” ROIs in separate ways (i.e., “faces > houses” or “memory > baseline”). We are, therefore, unable to examine whether these ROIs in the separate groups correspond to the same anatomical region. Visual inspection of the each individual’s ROIs in both groups suggest that these regions fall in roughly the same portion of the ATLs (Fig. S1); however, whether the same population of neurons in the ATL of a given individual has both amodal and hub-like properties remains an open question. Second, we only tested the ATL-hub model for two person attributes (i.e., status and trait). It will be interesting for future researchers to assess whether the ATL similarly coordinates the retrieval of other attributes, such as physical (i.e., eye color, body shape) or evaluative characteristics (i.e., likeability, reputation). If our model is generalizable, one would expect to see increased connectivity between the ATL and area V4 during the retrieval of eye color, extrastriate body area during the retrieval of body characteristics, and orbitofrontal cortex during the retrieval of social-evaluative information. Third, the present study only focused on a subset of social knowledge—the biographical information about others. Although biographical information is important, several other types of social knowledge contribute to social behavior, including knowledge about other people’s mental states and goals, semantic knowledge about people that is not specific to particular individuals (e.g., people need to eat, people tend to like chocolate), and

knowledge about social etiquette and norms (66). Future research is needed to elucidate the neural architectures supporting each knowledge type and how they relate to one another.

Materials and Methods

Subjects. A total of 24 subjects (9 males; $M_{\text{age}} = 23.21$) participated in study 1, and a new cohort of 26 subjects (21 males; $M_{\text{age}} = 20.38$) participated in study 2. All subjects were native English speakers, right-handed, and had no history of psychological or neurological disorders. They were financially compensated and gave informed consent in a manner approved by the Institutional Review Board of Temple University. The sample size ($n \sim 25$) was chosen based on previous studies using similar training paradigms and analytic approaches (10, 41, 67, 68).

Materials.

Study 1. Biographies of four fictitious males were provided during training sessions (Fig. 1A). Each male's face was professionally photographed (i.e., full-frontal, neutral expressions, uniform lighting/background) and was linked to six pieces of biographical information: a name, age, marital status, occupation, city of residence, and a house. The first five pieces of biographical information were presented as words and house was presented as a color image, procured from publically available sources on the Internet. Proper names were taken from the social security database of popular names (<https://www.ssa.gov/cgi-bin/popularnames.cgi>) and each had two syllables. A critical manipulation was made purposely on biographies: each male's face, name, house, and occupation were distinctive ("unique bios"), whereas their age, marital status, and city of residence could be shared ("common bios") with one other person. Biographical information was randomly assigned to each face such that each participant learned different associations for any given face.

In the fMRI session, a new set of stimuli was constructed for each fictitious male (Fig. 1B). Specifically, the learned faces/houses in training sessions were rephotographed from five different camera angles (0° , $\pm 45^\circ$, and $\pm 90^\circ$), and the learned names were presented in five different fonts/colors. In object runs, participants viewed daily-used objects that were representative of each male's occupation from five different viewpoints (e.g., if the fictitious male was a doctor, a stethoscope was shown).

Study 2. Biographies of eight fictitious males were provided in training sessions (Fig. 1C). Participants learned four pieces of information about each male: their face, name, occupation, and one personality trait. Eight distinct faces were generated by FaceGen Modeler 3.5, which allowed us to control low-level visual features such as color, brightness, and illumination. Eight artificial names were adapted from a previous study (69). The attractiveness and likeability of these artificial faces and names were carefully controlled by asking a different group of 20 participants to assign ratings. Social status and personality traits were manipulated in a 2×2 factorial design: each fictitious male was associated with an occupation title and salary that suggested high or low status (e.g., "Maur is a manager earning more than \$9,000 per month" or "Lorc is a janitor earning less than \$1,000 per month") and with personality trait descriptions that denoted high or low sociability (e.g., "Jora is an introverted, quiet *bookworm*" or "Gris is an extroverted, chatty *party animal*"). In real life, there are infinite attributes associated with an individual (e.g., attractiveness, reputation, temper, athleticism, etc.); here, we only manipulated "social status" and "sociability traits" because they are the two principal dimensions of person concepts and impression formation (70). Again, biographical information was randomly combined such that each participant learned different occupation/personality characteristics of any given face/name.

In the fMRI session, the eight learned faces or names were presented during the "memory" conditions (Fig. 1D). In addition, we generated eight novel faces/names for the "nonmemory" baseline conditions. These faces/names matched the learned faces/names in regard of low-level visual features, phonemic properties (e.g., start with the same consonant), and subjective ratings of attractiveness/likeability. Participants could easily identify them as unfamiliar during the practice session right before the fMRI session.

Procedure.

Behavioral training sessions. Participants were told to learn biographical information about four males in study 1 (Fig. 1A) and eight males in study 2 (Fig. 1C). The training protocol was adapted from the literature (10, 71). Training was conducted over 2 d, with the first day session lasting ~ 30 min, and the second day session lasting 20 min. During each training session, participants first completed "show" trials in which they viewed slides containing a face image, a name, a house (only in study 1), and a short para-

graph containing other biographical information. Each slide was presented three times. There was no time limit on slide presentation; participants pressed a button when they wished to move to the next screen. Next, participants completed "naming" trials in which they viewed a previously learned face and were asked to type that person's name. After responding, participants were told whether they were correct or incorrect, and the correct biographical information for that individual was presented again on the screen. For naming trials, each male was presented six times in a random order. Finally, participants completed "matching" trials in which they were presented either with learned faces or names and were asked to select the corresponding biographical information (e.g., occupation) from all options presented below (e.g., pilot/mailman/photographer/doctor in study 1; manager/janitor in study 2). The matching phase consisted of blocks of 40 trials in study 1 (or 32 trials in study 2) and participants received correctness feedback after each trial as well as an accumulative accuracy at the end of each block. Participants were considered fully trained if they performed a matching block with accuracy over 95% in day 1 and 100% in day 2. If participants did not reach this level of accuracy, they had to do extra matching blocks until they reached the appropriate level (e.g., overlearning ensured). On day 3, all participants were retested using a paper and pencil test immediately before the scan, to ensure that everyone could accurately recall the learned biographical information.

fMRI session in study 1. The fMRI session was scheduled on day 3 and consisted of two parts: two runs of a functional localizer and eight runs of the "person memory" task. Both tasks used a block design. In the localizer task, participants were instructed to pay attention to the images and respond whenever the same image was presented twice in a row (one-back task). Each run consisted of 30 blocks, evenly divided between three alternating stimulus types: "faces," "places," or "fixation cross." Each stimulus was presented for 800 ms followed by a 200-ms interstimulus interval. All faces and places stimuli were adapted from a previous study in our laboratory (10).

In the person memory task, participants were asked to retrieve knowledge about the four males they had learned about during training (Fig. 1B). Participants first observed a series of stimuli related to a unique aspect of the "target" male's identity (i.e., their face, name, house, or occupation) and then responded to a question regarding a common aspect of that male's biography (i.e., their age, city of residence, or marital status). For instance, after viewing five images of Tyler's house (presented at different vantage points), participants had to respond to a question regarding Tyler's age: "True or False: the person associated with this house is 28?"

The person memory task consisted of eight runs (i.e., two runs for each category). Only one category (face/name/object) was consistently used to cue target males across the whole run. Each run consisted of 16 blocks of stimulus presentation (4 males \times 4 repetitions) and each block contained five cue stimuli related to the same individual (3 s each) and one true/false question about that person (9 s). Note that face/object stimuli were only displayed in full-frontal view during training, whereas during the fMRI session the face/object cue stimuli were presented from five different vantage points. Similarly, the name cue stimuli here were presented in five different fonts and colors. These manipulations aimed to vary the low-level visual properties of each male's cue stimuli while holding the high-level person identity constant, and also avoid visual habituation (56). The object cue stimuli were related to each male's occupation. Five pictures depicting a daily-used object associated with each male's occupation were presented. Note that these objects were not shown during training, because occupation information was presented in written form during training. They were used here to ensure participants thoroughly learned each male and their associated biographical information. The cue stimuli order (within a block), the target male order (within a run), and the category order (across runs) alternated and were randomized across participants.

fMRI session in study 2. Participants completed six runs of a person memory task in the scanner on day 3 (Fig. 1D). In each run, we adopted a 3×2 factorial block-design where the factors were "task type" (status memory, trait memory, or nonmemory baseline) and "cue category" (name or face of the target male). In status/trait memory blocks, participants viewed faces or names of the eight learned males and were asked to indicate their social status (manager/janitor) or personality trait (extroverted/introverted). In the "nonmemory baseline" blocks, participants viewed eight novel faces/names they never learned before and were asked to indicate the spatial location of the cue stimuli appearing on the screen (left/right). At the beginning of each block, participants were informed of the upcoming task type (presented for 3 s). Each run consisted of 18 blocks (i.e., 3 task types \times 2 cue types \times 3 repetitions), with eight trials in each block to cover the whole set of faces/names. To make the memory and nonmemory conditions similar in terms of cognitive demand, participants had to respond in 4 s for status/trait trials but

in 2 s for baseline trials. The target male order (within a block) and the task order (within a run) alternated and were randomized across participants.

Data Acquisition and Preprocessing. The fMRI session was conducted at the Temple University Hospital on a 3-T Siemens Verio scanner, equipped with a 12-channel head coil. In study 1, functional images were acquired using a gradient-echo echo-planar imaging (EPI) sequence [repetition time (TR) = 3,000 ms; echo time (TE) = 20 ms; field of view = 240 × 240; matrix size = 80 × 80; flip angle = 90°]. Sixty-one interleaved slices (3 × 3 × 2.5-mm voxels) were acquired aligned to 30° tilted from the anterior commissure–posterior commissure line, with full brain coverage. These imaging parameters (i.e., short TE, tilted slices) were optimized for mitigating susceptibility artifacts around ATL and OFC (72), and were validated by pilot scans as well as previous studies in the laboratory (10, 28). In study 2, we adopted the same pulse sequence except the TR (2,000 ms) and 40 interleaved slices (3 × 3 × 3.5-mm voxels). In study 1 the visual stimulus was delivered by E-Prime software; in study 2, stimuli were delivered by Cogent toolbox running under Matlab R2014b. Responses were recorded using a four-button fiber optic response pad system.

To remove sources of noise and artifact, functional data were corrected for slice timing, realigned, unwarped, normalized to the EPI template [Montreal Neurological Institute (MNI) space, resampled to 3 × 3 × 3 mm], spatially smoothed (8 mm), high-pass filtered at 128 s, and prewhitened by means of an autoregressive model AR (1) using SPM8 software.

Univariate General Linear Model Analyses. Subject-specific parameter estimates (β weights) for each condition were derived through a general linear model (GLM). For each subject, the data were best-fitted at every voxel using a combination of effects of interest. These were delta functions representing the onset of each of the experiment conditions, convolved with the SPM8 hemodynamic response function. The six motion regressors and memory error trials were also included as nuisance regressors. Next, subject-specific β weights were entered into a group-level random-effect GLM to allow statistical inference. Statistics maps were generated using a voxel-level familywise error (FWE)-corrected threshold of $P < 0.05$. Stereotaxic coordinates are reported in MNI space and regional labels were derived using the automated anatomical labeling (AAL) atlas in xjView.

In study 1, we set up one GLM for each of the four categories (face, name, house, and object) across two runs using unsmoothed data. This was prepared for subsequent MVPA CC classification. β maps for each target male (i.e., Tyler, Justin, Aaron, and Cameron) were extracted from each block of a run. Each male had four β maps per run and eight β maps in total per category.

In study 2, four different GLMs were set up across six runs, based on different analytic purposes. One GLM was set up for a univariate analysis of the main effect of task types. β maps for each task type regressor (i.e., face–status, face–trait, face–baseline, name–status, name–trait, and name–baseline) were extracted and contrasts of interest were conducted upon them (i.e., all memory > baseline, status memory > baseline, trait memory > baseline, face memory > name memory, and name memory > face memory) (Fig. S3 and Table S6). This GLM was also used for DCM analysis. For MVPA decoding of status, personality traits, and identity representations, three GLMs were set up separately using unsmoothed data. β maps for each status regressor (i.e., manager or janitor), each trait regressor (i.e., extrovert or introvert), and each learned male regressor (i.e., eight males) were extracted from status memory blocks, trait memory blocks, and all memory blocks, respectively.

ROIs Localization. Study 1 used an established functional localizer (10) to localize subject-specific ROIs. Face-selective regions (i.e., OFA, FFA, ATL, AMY, and OFC) were defined by individuating the peaks showing greater activity for faces than for places (“faces > places”; $P < 0.05$, uncorrected). A place-selective region (i.e., PPA) was defined by the opposite contrast (“places > faces”). To rule out effects driven by the low-level perceptual features of our stimuli, we additionally defined a control V1 ROI in early visual cortex around the voxel showing the greatest activation for all visual stimuli (“faces + places” > “fixation”) (Fig. S1A).

In study 2, subject-specific ROIs were functionally localized using a series of contrasts between task types in univariate analysis: ATL (“all memory > baseline”), FFA (“face memory > name memory”), VWFA (“name memory > face memory”), IPL (“status memory > baseline”), PCC (“trait memory > baseline”), and the hippocampus (“all memory > baseline”) (Fig. S1B). We stress that, even though ROI definition and subsequent MVPA analysis were based on the same data, our analyses did not suffer from circular logic (73), because the voxel selection criteria were based on “task types” in univariate analyses where all eight learned males were modeled by the same regressor whereas

MVPA analysis aimed to classify each learned male. In other words, the voxel selection procedure used no information about the specific biographies of each learned male (74, 75). Similar rationale also applies to our PPI (32, 76) and DCM analysis (i.e., DCM inferences were made via model comparison, which were conditioned on prespecified regions but not biased by the selection of the regions per se) (77, 78).

A 6-mm spherical mask was generated for all ROIs (68, 79) in study 1 and study 2, centered on MNI coordinates with the highest activation within each peak (Table S7), and confirmed by anatomical AAL atlas. In study 1, we localized bilateral areas for each ROI (except midline areas such as OFC and V1) for ROI-based MVPA analyses; then we merged both hemispheres into one unified mask to represent each ROI for combinatorial MVPA analyses. In study 2, we generated a unilateral mask for each ROI because only one hemisphere survived in univariate analysis. In study 2, the same ROIs were used in subsequent MVPA, PPI, and DCM analysis.

MVPA Analyses. In both studies, we adopted three MVPA decoding approaches: ROI-based, combinatorial, and searchlight analysis. All were implemented by the Decoding Toolbox (80) using a support vector machine as a classifier. Combinatorial analysis was executed in the same way as the ROI-based approach, except that conjunction masks were created between all possible combinations of two ROIs. Searchlight analysis was executed with a radius of 4 voxels (i.e., 12 mm) across the whole brain.

In study 1, a two-way CC classification scheme was used for each subject. The classifier was first trained on one category (e.g., face runs) and subsequently tested on the other category (e.g., name runs); the reverse decoding direction was also performed (e.g., first trained on name runs and then tested on face runs). The average two-way CC classification accuracy was then calculated for each ROI for each CC pair and compared with chance performance (identity of four males = 0.25). In total, six CC classification analyses were performed in a pairwise fashion between any two of the four categories (i.e., face–name, face–house, face–object, name–house, name–object, house–object).

Study 2 used a leave-one-run-out CV scheme in which the classifier was trained on five runs of data and tested on the remaining untrained run. This procedure was repeated six times, each time using a different test run, and the average CV accuracy was calculated for each ROI and compared with chance performance (status = 0.5, personality trait = 0.5, identity of eight fictitious males = 0.125).

For group-level inference, ROI-based analysis used one-tailed t tests with a Bonferroni correction of $P < 0.05$. For combinatorial analysis, we used one-tailed t tests of $P < 0.05$ to test whether the average change of CC/CV accuracy (by combining a particular ROI with each of the rest ROIs) was significantly above 0% (i.e., increase the prediction) (43). For searchlight analysis, a voxel-level FWE-corrected threshold of $P < 0.05$ was used. If null results were found, a lenient voxel-level threshold ($P < 0.0001$, uncorrected) was then used to reveal subthreshold results for exploratory purposes (Fig. S2 and Table S1).

PPI Analyses. We carried out PPI analysis in SPM8 (53) to examine functional coupling between the six ROIs identified using univariate analyses during the person memory task. Seed regions included sensory input ROIs (VWFA/FFA) or knowledge representations ROIs (IPL/PPC) or the ATL (identity representation); and four experimental conditions were examined (i.e., face/name × status/personality trait) vs. baseline. For instance, in one PPI analysis, we used the FFA as the seed region and explored its “coupling” brain regions during the retrieval of social status; in another PPI analysis, we set the IPL as a seed region and explored its coupling regions during face–status conditions. In total, 12 PPI analyses were performed for each subject (Fig. S4 and Tables S3–S5).

For each PPI analysis, we created a new GLM with (i) a “physiological” regressor in which the seed region’s time course (i.e., first eigenvariate) was deconvolved to estimate the underlying neural activity; (ii) a “psychological” regressor in which task type trials (vs. baseline trials) were convolved with the canonical hemodynamic response function; and (iii) a “PPI interaction” regressor in which the psychological regressor was multiplied by the physiological regressor. We used this interaction regressor to identify voxels in which functional activity covaried in a task-dependent manner with the seed region. Subject-level PPI analyses were run to generate SPM contrast images similar to a subject-level GLM model, and these contrast images were entered into a group-level random-effects GLM and thresholded at $P < 0.05$ (voxel-level FWE corrected).

DCM. We performed DCM analysis using DCM10 in SPM8 (81). Three distinct, but not mutually exclusive, models were tested (Fig. 4B and Fig. S5):

(i) Model 1 hypothesizes that person knowledge is retrieved via a direct link from sensory input areas (FFA/VWFA) to knowledge representation areas (IPL/PCC), with the ATL being mainly implemented in identity recognition and therefore not involved in specific knowledge content retrieval. In model 1, one can expect two separate information flows during face–status memory: “FFA → IPL” (for status knowledge retrieval) and “FFA → ATL” (for implicit person identification). This model is conceptually similar to the distributed-only account in semantic memory research. (ii) Model 2 assumes that the ATL serves as a hub that coordinates activity within the network during person memory retrieval: a person’s identity is first recognized in the ATL, which then directs information flow to knowledge representation areas (IPL/PCC). In model 2, the information flows for face–status memory and name–trait memory can be expected as “FFA → ATL → IPL” and “VWFA → ATL → PCC,” respectively. Note that we did not set ATL’s modulation to hippocampus in this model because MVPA analyses suggest that the hippocampus does not contain any person-knowledge representations and PPI analyses found no functional connectivity between the ATL and hippocampus during person memory retrieval. (iii) In model 3, the hippocampus is the central hub: all person knowledge (including identity) is coordinated by this domain-general hub. The information flows during face–status memory would be “FFA → hippocampus → IPL (for status knowledge)” and “FFA → hippocampus → ATL (for implicit person identification).” For each model, the six ROIs were all set to be bidirectionally connected (Fig. S5A), and their time courses (i.e., first eigenvariate) were extracted for each subject individually.

To prevent poor model fit, we calculated the mean explained variance of each model across subjects by using the SPM function

- Olson IR, McCoy D, Klobusicky E, Ross LA (2013) Social cognition and the anterior temporal lobes: A review and theoretical framework. *Soc Cogn Affect Neurosci* 8(2): 123–133.
- Barton JJS, Corrow SL (2016) Recognizing and identifying people: A neuropsychological review. *Cortex* 75:132–150.
- Burton AM, Bruce V (1992) I recognize your face but I can’t remember your name: A simple explanation? *Br J Psychol* 83(Pt 1):45–60.
- Blank H, Wieland N, von Kriegstein K (2014) Person recognition and the brain: Merging evidence from patients and healthy individuals. *Neurosci Biobehav Rev* 47: 717–734.
- Perrodin C, Kayser C, Abel TJ, Logothetis NK, Petkov CI (2015) Who is that? Brain networks and mechanisms for identifying individuals. *Trends Cogn Sci* 19(12):783–796.
- Patterson K, Nestor PJ, Rogers TT (2007) Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat Rev Neurosci* 8(12): 976–987.
- Martin A (2007) The representation of object concepts in the brain. *Annu Rev Psychol* 58(1):25–45.
- Lambon Ralph MA (2013) Neurocognitive insights on conceptual knowledge and its breakdown. *Philos Trans R Soc Lond B Biol Sci* 369(1634):20120392.
- Collins JA, Olson IR (2014) Beyond the FFA: The role of the ventral anterior temporal lobes in face processing. *Neuropsychologia* 61(1):65–79.
- Collins JA, Koski JE, Olson IR (2016) More than meets the eye: The merging of perceptual and conceptual knowledge in the anterior temporal face area. *Front Hum Neurosci* 10(May):189.
- Kriegeskorte N, Formisano E, Sorger B, Goebel R (2007) Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proc Natl Acad Sci USA* 104(51): 20600–20605.
- Nestor A, Plaut DC, Behrmann M (2011) Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proc Natl Acad Sci USA* 108(24):9998–10003.
- Von Der Heide RJ, Skipper LM, Olson IR (2013) Anterior temporal face patches: A meta-analysis and empirical study. *Front Hum Neurosci* 7(February):17.
- Anzellotti S, Caramazza A (2014) The neural mechanisms for the recognition of face identity in humans. *Front Psychol* 5(Jun):672.
- Blank H, Kiebel SJ, von Kriegstein K (2015) How the human brain exchanges information across sensory modalities to recognize other people. *Hum Brain Mapp* 36(1):324–339.
- Abel TJ, et al. (2015) Direct physiologic evidence of a heteromodal convergence region for proper naming in human left anterior temporal lobe. *J Neurosci* 35(4): 1513–1520.
- von Kriegstein K, Giraud AL (2006) Implicit multisensory associations influence voice recognition. *PLoS Biol* 4(10):e326.
- Tsukiura T, Suzuki C, Shigemune Y, Mochizuki-Kawai H (2008) Differential contributions of the anterior temporal and medial temporal lobe to the retrieval of memory for person identity information. *Hum Brain Mapp* 29(12):1343–1354.
- Tsukiura T, et al. (2010) Dissociable roles of the anterior temporal regions in successful encoding of memory for person identity information. *J Cogn Neurosci* 22(10): 2226–2237.
- Skipper LM, Ross LA, Olson IR (2011) Sensory and semantic category subdivisions within the anterior temporal lobes. *Neuropsychologia* 49(12):3419–3429.
- Nieuwenhuis IL, et al. (2012) The neocortical network representing associative memory reorganizes with time in a process engaging the anterior temporal lobe. *Cereb Cortex* 22(11):2622–2633.

“spm_dcm_fmri_check.” To determine the optimal model, fixed-effects (FFX) and random-effects (RFX) group analyses were implemented by Bayesian models selection (82). In the FFX case, one assumes that the optimal model is identical across the population. It uses group log evidence to quantify the relative goodness of models, which is the exponentiated sum of the log model evidence of each subject-specific model. Usually, a difference in a group log evidence of 3 is taken as statistically strong evidence. Thus, if the group log evidence of one model is bigger than the other models’ by 3 or more, that model would be considered by FFX analysis to be the optimal model (78).

Because the FFX analysis is vulnerable to outlier subjects, we also implemented an RFX analysis, which accounts for the heterogeneity of the model structure across subjects. It uses hierarchical Bayesian modeling that estimates the parameters of a Dirichlet distribution over the probabilities of all models considered. These probabilities define a multinomial distribution over model space, enabling the computation of the posterior probability of each model given the data of all subjects and the models considered. The results of RFX analysis are reported in terms of the exceedance probability that one model is more likely than any other model. The optimal model in RFX analysis would be considered to be the one with the largest exceedance probability.

ACKNOWLEDGMENTS. We thank Ashley Unger for assistance in data collection, David V. Smith for helpful comments, and two anonymous reviewers for their time and effort in reviewing this manuscript. This work was supported by NIH Grant R01 MH091113 (to I.R.O.).

- Olson IR, Plotzker A, Ezzyat Y (2007) The enigmatic temporal pole: A review of findings on social and emotional processing. *Brain* 130(Pt 7):1718–1731.
- Schurz M, Radua J, Aichhorn M, Richlan F, Perner J (2014) Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neurosci Biobehav Rev* 42: 9–34.
- Mitchell JP, Heatherton TF, Macrae CN (2002) Distinct neural systems subserve person and object knowledge. *Proc Natl Acad Sci USA* 99(23):15238–15243.
- Frith U, Frith CD (2003) Development and neurophysiology of mentalizing. *Philos Trans R Soc Lond B Biol Sci* 358(1431):459–473.
- Zahn R, et al. (2007) Social concepts are represented in the superior anterior temporal cortex. *Proc Natl Acad Sci USA* 104(15):6430–6435.
- Zahn R, et al. (2009) The neural basis of human social values: Evidence from functional MRI. *Cereb Cortex* 19(2):276–283.
- Ross LA, Olson IR (2010) Social cognition and the anterior temporal lobes. *Neuroimage* 49(4):3452–3462.
- Simmons WK, Reddish M, Bellgowan PSF, Martin A (2010) The selectivity and functional connectivity of the anterior temporal lobes. *Cereb Cortex* 20(4):813–825.
- Gilbert SJ, Swencionis JK, Amodio DM (2012) Evaluative vs. trait representation in intergroup social judgments: Distinct roles of anterior temporal lobe and prefrontal cortex. *Neuropsychologia* 50(14):3600–3611.
- Wong C, Gallate J (2012) The function of the anterior temporal lobe: A review of the empirical evidence. *Brain Res* 1449:94–116.
- Greven IM, Downing PE, Ramsey R (2016) Linking person perception and person knowledge in the human brain. *Soc Cogn Affect Neurosci* 11(4):641–651.
- Pobric G, Lambon Ralph MA, Zahn R (2016) Hemispheric specialization within the superior anterior temporal cortex for social and nonsocial concepts. *J Cogn Neurosci* 28(3):351–360.
- Gainotti G (2007) Different patterns of famous people recognition disorders in patients with right and left anterior temporal lesions: A systematic review. *Neuropsychologia* 45(8):1591–1607.
- Ellis AW, Young AW, Critchley EMR (1989) Loss of memory for people following temporal lobe damage. *Brain* 112(Pt 6):1469–1483.
- Gainotti G, Marra C (2011) Differential contribution of right and left temporo-occipital and anterior temporal lesions to face recognition disorders. *Front Hum Neurosci* 5(June):55.
- Tsukiura T, et al. (2002) Neural basis of the retrieval of people’s names: Evidence from brain-damaged patients and fMRI. *J Cogn Neurosci* 14(6):922–937.
- Todorov A, Olson IR (2008) Robust learning of affective trait associations with faces when the hippocampus is damaged, but not when the amygdala and temporal pole are damaged. *Soc Cogn Affect Neurosci* 3(3):195–203.
- Peelen MV, Atkinson AP, Vuilleumier P (2010) Supramodal representations of perceived emotions in the human brain. *J Neurosci* 30(30):10127–10134.
- Oosterhof NN, Wiggett AJ, Diedrichsen J, Tipper SP, Downing PE (2010) Surface-based information mapping reveals crossmodal vision-action representations in human parietal and occipitotemporal cortex. *J Neurophysiol* 104(2):1077–1089.
- Fairhall SL, Caramazza A (2013) Brain regions that represent amodal conceptual knowledge. *J Neurosci* 33(25):10552–10558.
- Awwad Shiekh Hasan B, Valdes-Sosa M, Gross J, Belin P (2016) “Hearing faces and seeing voices”: Amodal coding of person identity in the human brain. *Sci Rep* 6:37494.
- Clithero JA, Carter RM, Huettel SA (2009) Local pattern classification differentiates processes of economic valuation. *Neuroimage* 45(4):1329–1338.
- Smith DV, Clithero JA, Rorden C, Karnath HO (2013) Decoding the anatomical network of spatial attention. *Proc Natl Acad Sci USA* 110(4):1518–1523.