Idiosyncratic codes in distributed minds

Ivette Colón (ycolon@wisc.edu)[1] **& Timothy Rogers**[2]
[1]Department of Psychology, University of Wisconsin - Madison
[2]Wisconsin Institute for Discovery

Abstract

Debates about the neural organization of semantic knowledge have long contrasted modular accounts, which posit anatomically specialized regions, with distributed accounts, which emphasize patterns of activity spanning many neural populations. Although multivariate neuroimaging has provided strong evidence for distributed representations, both perspectives often presume that the representational solutions supporting cognition are largely shared across individuals. Here, we challenge that premise by asking whether semantic representations are not only distributed, but also stable within individuals and meaningfully idiosyncratic across people. Using individualized multivariate fMRI decoding across two scanning sessions, we show that representations distinguishing faces, places, and objects are widely distributed across cortex, stable within individuals over time, and heterogeneous across participants. We then use these individualized representational profiles to guide transcranial magnetic stimulation (TMS), showing that stimulating non-canonical, decoding-defined sites selectively disrupts face similarity judgments in a manner comparable to stimulation of a canonical face-selective region. Together, these findings provide causal evidence that semantic representations arise from distributed neural systems whose organization is both stable and person-specific, calling for theories of cognition that make room not only for shared structure, but for the idiosyncratic ways meaning is represented across brains.

Keywords: individual differences; distributed representation; neural decoding; face processing; fMRI; TMS

Introduction

Debates about the organization of cognitive representations in the brain have long been framed around a tension between modular and distributed accounts (Haxby, 2012). Modular theories propose that distinct cognitive functions are supported by specialized neural subsystems, often localized to circumscribed cortical regions. In the visual domain, faces, places, and objects have served as canonical examples, with regions such as the fusiform face area, parahippocampal place area, and lateral occipital complex frequently cited as evidence for functional specialization (Epstein and Kanwisher, 1998; Grill-Spector et al., 2001; Kanwisher et al., 1997).

By contrast, connectionist and distributed accounts challenge a strict modular interpretation. Rather than assigning representations to discrete locations in the brain, these frameworks emphasize patterns of activity distributed across many units, with information emerging from their collective organization (Rogers and McClelland, 2004). On this view, robustness and generalization arise not from specialization alone, but from redundancy and overlap. Empirical neuroimaging work using multivariate analyses has provided strong support for this perspective, showing that category information can be decoded from distributed patterns spanning ventral temporal cortex and beyond, even when canonical category-selective regions are excluded (C. R. Cox and Rogers, 2021; Haxby, 2012; Haxby et al., 2001, 2014).

Notably, however, both perspectives often share a deeper assumption: that the representational solutions supporting cognition are largely the same across individuals. Modular accounts tend to assume homologous functional components, while distributed accounts often emphasize representational geometries expressed across many brains (Haxby et al., 2020). In practice, this assumption of universality underwrites much of cognitive neuroscience, motivating group-level averaging, anatomical alignment to common templates, and the interpretation of population-level effects as revealing the neural basis of cognition (Frost and Goebel, 2012; Laumann et al., 2015).

Increasingly, this assumption appears difficult to sustain. Across domains, neural and behavioral data indicate that individuals differ systematically in how information is represented, even when outward performance appears similar (Colón and Rogers, 2025; Dubois and Adolphs, 2016; Finn et al., 2015). Importantly, these differences are often stable within individuals over time, suggesting that they reflect meaningful representational structure rather than noise (Gordon et al., 2017). From this perspective, the critical theoretical question is not only whether representations are modular or distributed, but whether they are shared across people in a functionally interchangeable way.

With these issues in mind, we report two studies addressing the following questions:

- Are neural representations that support face, place, and object categorization confined to canonical category-selective regions, or are they distributed across cortex? (Study 1)
- Are these distributed representations stable within individuals across time, or do they primarily reflect transient or noisy variation? (Study 1)
- To what extent are representational solutions shared across individuals versus individualized in their neural organization? (Study 1)
- Do non-canonical, individualized representational sites play a causal role in guiding semantic behavior, comparable to canonical category-selective regions? (Study 2)

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Figure 1: Example of controlled face, place, and object stimuli used in this work; Study 2 uses only faces.

Study 1 applies whole-brain multivariate decoding to assess whether neural representations differentiating faces, places, and objects are (i) distributed across cortex, (ii) stable over time within an individual, and (iii) reliably different across participants. Building on these results, Study 2 uses transcranial magnetic stimulation (TMS) to assess whether brain areas that appear to encode stimulus information in a manner idiosyncratic to an individual are causally involved in processing face-specific content.

Study 1: fMRI Decoding

We apply individualized whole-brain multivariate decoding to each participant to identify voxel patterns that reliably separate semantic domains, examine whether these patterns extend beyond canonical category-selective regions, and test whether representational solutions generalize more strongly within individuals over time than across individuals.

Study 1: Methods

General protocol Participants completed four sessions. In Session 1, they performed a computer-based triadic comparison task to familiarize them with the stimulus set and for additional analyses not reported in this work. Sessions 2 and 3 consisted of fMRI scanning while participants completed an image-encoding task using a slow event-related design and an active non-semantic baseline task to keep them engaged (number judgment). In Session 4, participants completed a modified triadic comparison task while undergoing transcranial magnetic stimulation (TMS; see Study 2). All sessions used the same primary stimulus set, with additional foil stimuli introduced only in the TMS session.

Participants Seventeen participants completed both MRI sessions; however one fell asleep in the scanner and was excluded from further analysis. Participants were aged 18-35 (M = 26.3 years). All procedures were approved by the University of Wisconsin–Madison IRB. Informed consent was obtained.

Stimuli We generated stimuli across three domains (faces, places, and objects) using an image-generating AI (DALLE-3) and systematically varying five binary attributes in the prompt (e.g., "Generate a realistic image of a young, Black woman at midday in a rural location"). All images varied in age of the depicted item (e.g., older vs younger/newer), setting (urban/rural), and time of day (midday/evening). Within each domain, stimuli additionally varied on two domain-specific features: building type (house/church) and size (big/small) for places; category (vehicle/furniture) and size (big/small) for objects; and apparent gender (men/women) and race (Black/White) for faces. We note that these binarized categories do not capture the full diversity of gender and race, but were included to reflect socially salient distinctions that shape how faces are organized and perceived. For each attribute combination within a domain, two exemplars were created, yielding 32 images per domain (2^5; see Fig. 1). An additional set of 32 with the same feature combinations were used as foils in the TMS study.

Triadic comparison task Participants completed a 1-hour computer-based triadic comparison task with a brief mid-session break. On each trial, a reference item appeared above two choice items, and participants selected the choice item most similar to the reference using any criteria. The task comprised 612 trials: 200 face-only, 200 place-only, and 200 object-only triads, plus 12 attention-check trials (four per domain) in which one choice item was identical to the reference. Triads were generated in advance; all participants viewed the same set of triads, presented in a randomized order. Although not the subject of this work, the responses in this task will be used in future representational similarity learning analyses (C. R. Cox et al., 2024).

MRI task and acquisition Participants completed two MRI scanning sessions (~2–2.5 hours each) after standard MRI safety screening. During scanning, they performed a combined image-encoding and n-back number judgment task. On each trial, a face, place, or object image was presented for 6000 ms, followed by a 2000 ms fixation, an n-back number comparison task, and a final 2000 ms fixation. In this slow event-related design, participants were instructed to study each image for a subsequent memory test. Images were pseudorandomly selected and balanced across domains and binarized attributes. During the n-back task, participants indicated whether each number was greater or less than the preceding number via mouse response; this task maintained engagement and provided a measure of attentional performance. At the end of each block, participants completed a "memory test" by selecting the two images most similar to those previously viewed from an array of eight foil images within the same domain. Each session comprised 12 runs (6 min 24 s each), with 24 trials per run, amounting to 76.8 minutes of functional MRI data and 288 image encoding trials per session.

MRI data were collected on a 3T GE Discovery MR750 scanner. Functional images were acquired with a TR of 1600 ms and analyzed at the whole-cortex level. We preprocessed fMRI using standard AFNI procedures (R. W. Cox, 1996), and restricted analyses to cortical gray matter voxels.

Decoding analyses Our slow event-related fMRI design across two sessions allowed us to examine how category information evolves over time in distributed neural patterns. For

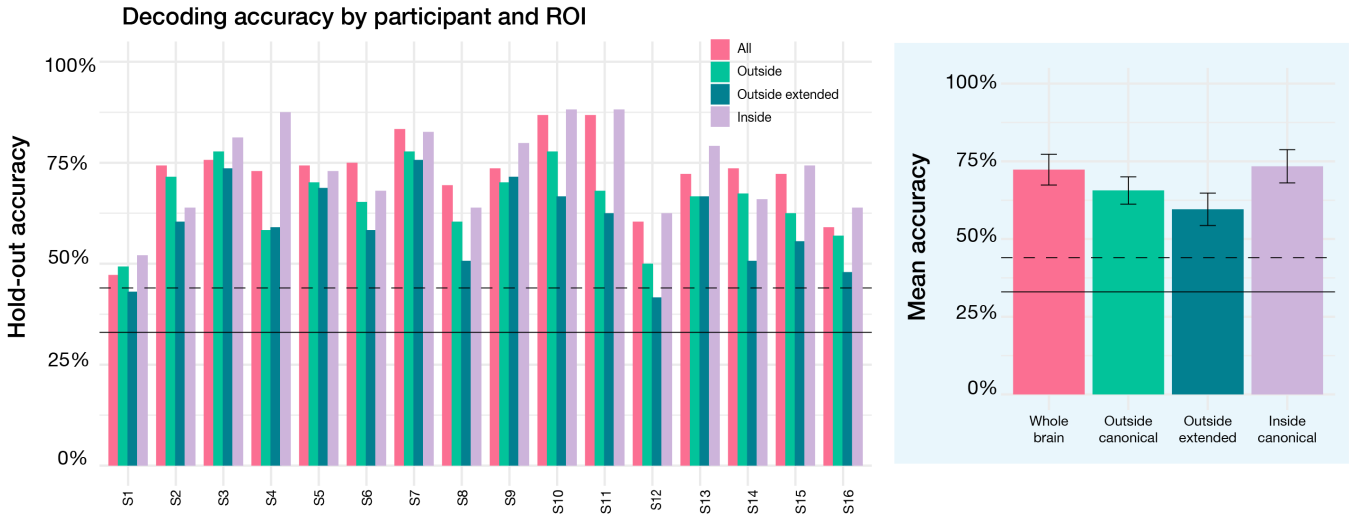

Figure 2: Scan 1 decoding accuracy across participants and ROIs. Decoding from all ROIs is above chance (44%) for most participants. There is no significant difference between decoding from whole brain vs. within canonical regions ($p = .072$).

each trial, BOLD responses were extracted at TRs 1–10 following stimulus onset, averaged across repeated presentations of each stimulus, and organized into stimulus-by-voxel matrices. For each participant, scan, and TR, voxel values were z-scored across stimuli prior to decoding.

We used IteratedLASSO (Z. Li et al., 2025) followed by ridge regression to identify distributed voxel patterns distinguishing faces, places, and objects. Performance was assessed using repeated 90/10 random-split cross-validation (20 iterations), with balanced category representation in the test set. Decoding was performed independently at each TR and over sliding two-TR windows to characterize the temporal profile of category information. Decoding accuracy peaked at TRs 6–8; therefore, responses averaged across this window were used for primary analyses.

Study 1: Results

Decoding within and beyond canonical regions To test whether decoding was driven solely by canonical category-selective regions, we compared decoding accuracy across complementary canonical and non-canonical voxel sets using the same decoding procedure applied to the first scan of each participant. Decoding was performed using four voxel masks: (i) all cortical gray matter (whole brain), (ii) canonical category-selective regions (combined FFA, PPA, and LOC; Julian et al., 2012), (iii) cortical gray matter outside canonical regions, and (iv) cortical gray matter outside canonical regions plus a 5 mm exclusion radius.

For most participants, we are able to decode stimulus category significantly above chance for each scan separately (here, decoding significance is above 44%, given the binomial probability for 96 items with $p < .04$). Across participants, mean decoding accuracy was 72.31% (SD = 10.09%) for Scan 1 and 66.93% (SD = 11.96%) for Scan 2.

Decoding accuracy was above chance for most participants in all regions. Across participants, mean accuracy for Scan 1 (Scan 2) was 73.34% (70.06%) within canonical regions, 65.62% (60.88%) outside canonical regions, and 59.54% (55.36%) outside the extended mask. Paired t-tests on Scan 1 revealed significantly higher accuracy within canonical regions than outside canonical regions ($t(15) = 3.59$, $p = .0027$) and outside the extended mask ($t(15) = 7.34$, $p < .0001$), but no difference between canonical and whole-brain decoding ($t(15) = 0.68$, $p = .5099$). Crucially, robust above-chance decoding from non-canonical regions indicates that category information is not confined to traditional category-selective areas, suggesting that meaningful representational signal extends beyond canonical visual processing regions.

Stability and individuality of representations By comparing neural representations across two scanning sessions, we can ask not only how stable an individual's representations are over time, but also how distinctive they are relative to those of other people. Our central question with this analysis is whether representations derived from a person's own data provide a better account of their later neural responses than representations derived from others. Addressing this question is complicated by the fact that even after registration to a common template, brains differ in size, shape, voxel sampling, and scan-specific coverage, such that corresponding functional signals may not occupy identical voxel locations across scans or individuals (Schleim and Roiser, 2009).

To overcome this challenge, we used bipartite graph matching to establish voxelwise correspondences between scans within and between participants (Y. Li et al., 2020). This procedure allows us to evaluate cross-session and cross-participant generalization on a common footing. The resulting pattern shows that although category information generalizes across brains to some extent, decoding is reliably strongest when models are trained and tested within the same individual

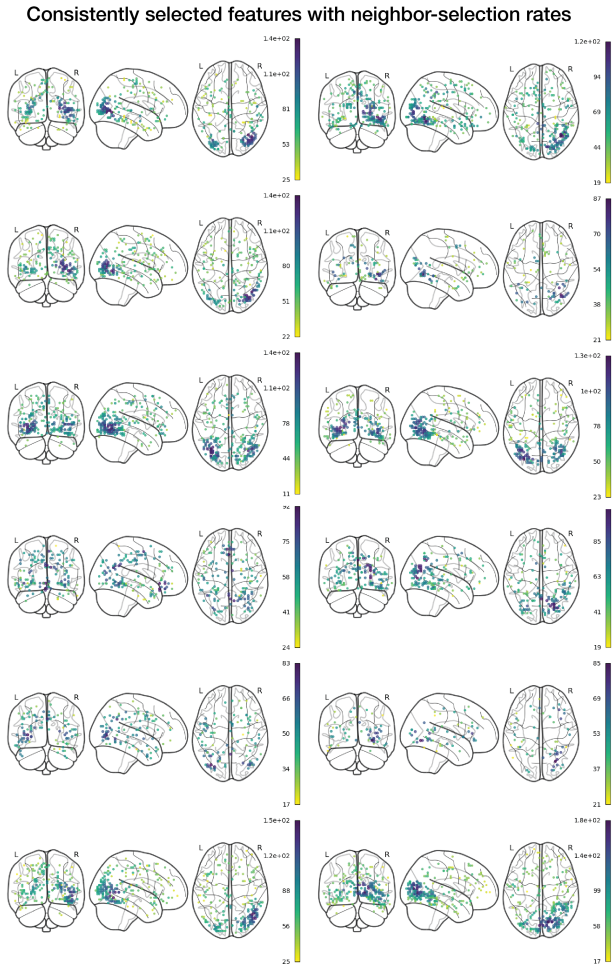

Figure 3: Consistently selected voxels across participants with above-chance decoding in both scans. Points colored by the number of neighboring voxels selected within a 20mm radius.

(see diagonal in Fig. 4). Across participants, within-subject cross-session decoding exceeded between-subject decoding by 22.5% on average (95% CI [16.21, 28.71], $p < .00001$), indicating that neural representations are not only stable over time, but also meaningfully individualized.

Study 1: Discussion

Using individualized multivariate decoding across two fMRI sessions, we show that neural representations differentiating categories are distributed across cortex, stable over time, and heterogeneous across participants. Category information was recoverable well beyond canonical category-selective regions, indicating that these regions do not exhaust the neural substrates of semantic representation. Representational solutions generalized more strongly within individuals than across participants— even after accounting for anatomical correspondence— indicating that individual variability reflects stable representational structure.

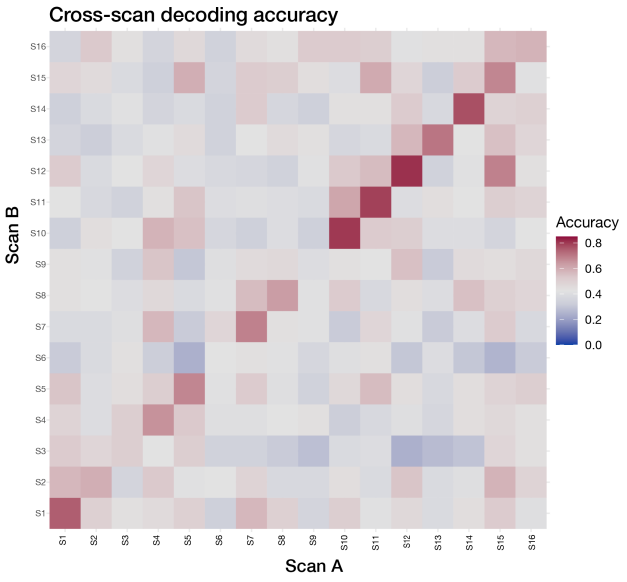

Figure 4: Cross-session and cross-participant decoding accuracy. Diagonal indicates within-participant generalization between scans, and reflects greater accuracy within a participant vs. across participants ($p < .00001$).

Study 2: TMS

In Study 1, we identified reliable, distributed, and individualized neural representations that jointly support discrimination among faces, places, and objects. However, the ability to decode category information alone does not establish whether these patterns play a causal role in guiding behavior. In Study 2, we therefore asked whether the individually defined, non-canonical cortical sites identified through decoding are causally involved in semantic similarity judgments. To address this question, we used transcranial magnetic stimulation (TMS) to perturb these sites and compared their behavioral effects to stimulation of a canonical face-selective region.

TMS provides a means of testing causal contributions of brain regions by transiently disrupting neural activity at targeted sites, effectively creating a reversible "virtual lesion" (Pascual-Leone, 1999). Neuropsychological and neurostimulation research on face processing suggests that faces are supported by distributed networks rather than isolated modules: damage or stimulation to regions such as the occipital face area (OFA) can impair face recognition even when other canonical areas remain intact (Pitcher et al., 2009; Solomon-Harris et al., 2013). However, most prior work has focused on stimulating canonical category-selective regions, leaving open the question of whether non-canonical, individualized representational sites also play a functional role.

In this study, we focus on faces because face processing provides a well-characterized test case where both canonical and distributed accounts make strong, competing predictions, and because socially relevant facial distinctions place especially high demands on representational precision.

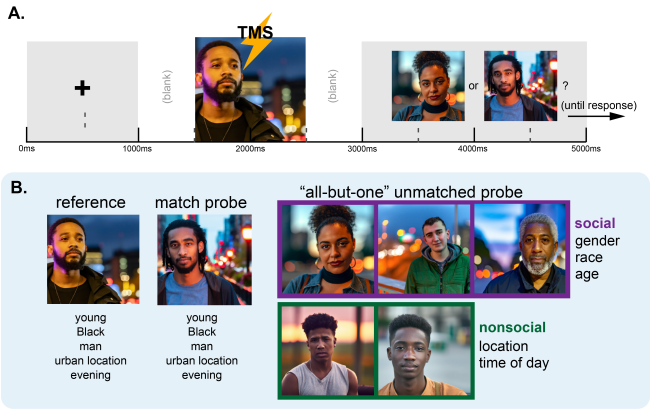

Figure 5: A) An example trial of the TMS task, with 1000 ms stimulation during reference item presentation. B) Example of stimuli in social and non-social trials.

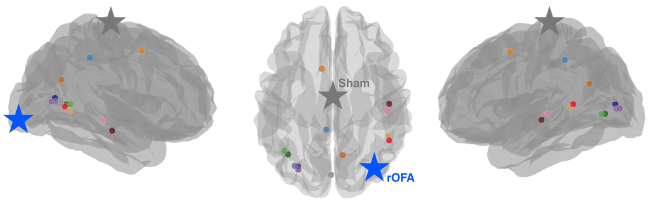

Figure 6: All TMS stimulation sites. Blue stars indicate location of right OFA stimulation site. Gray stars indicate Sham stimulation site. All other dots are each participant's IND stimulation site.

Study 2: Methods

TMS task and stimulation sites Participants completed a modified triadic comparison task while receiving TMS. On each trial, a reference face image was presented during stimulation, followed by two choice face images. One choice matched the reference on all attributes, while the other differed on a single attribute. Trials were categorized as *social* when the differing attribute was person-related (age, race, or gender) and *nonsocial* when it reflected contextual information (location or time of day). Reaction time and accuracy were recorded for each trial.

Each participant completed separate blocks of stimulation at three sites: (i) a canonical right OFA location (Pitcher et al., 2011), (ii) a vertex control site ("sham"), and (iii) an individualized, non-canonical cortical site identified via decoding. Individualized stimulation points ("IND") were derived from each participant's fMRI data from Study 2 by identifying voxels that reliably contributed to category decoding across independent scans. Specifically, we focused on voxels that were consistently selected as informative for distinguishing faces from places and objects in both scans, and that formed spatially coherent clusters rather than isolated points. Candidate sites were chosen from these clusters to maximize functional reliability while remaining anatomically accessible for stimulation and non-overlapping with canonical face-selective regions used in prior work.

To ensure reliable targeting, TMS analyses were restricted to participants who showed above-chance decoding accuracy in both fMRI sessions, yielding a final sample of 12 participants. TMS was delivered at 10 Hz for 1000 ms at 40% stimulator output using neuronavigation to maintain accurate coil placement based on each participant's anatomical MRI. Coil position was continuously monitored to ensure consistent stimulation across blocks.

Behavioral effects of stimulation To assess whether stimulation of individualized neural sites produces behavioral effects comparable to stimulation of a canonical face area, we analyzed reaction time and accuracy during the TMS triadic comparison task. Stimulation location was the primary manipulation, with three levels: sham, stimulation of the occipital face area (OFA), and stimulation of the individualized location (IND). Trials were further categorized based on whether the critical distractor differed from the reference in a social attribute (race, age, or gender) or a nonsocial attribute (location or time of day). Reaction times were analyzed using a linear mixed-effects model with fixed effects of stimulation condition, trial type (social vs. nonsocial), and block order, including their interactions. Random intercepts were included for participants and trial triads to account for baseline differences in response speed and stimulus difficulty. Accuracy was analyzed using a generalized linear mixed-effects model with a comparable fixed-effects structure and random intercepts for participants and items.

Study 2: Results

In the sham condition, participants were faster and more accurate on social than nonsocial trials, consistent with the greater salience of person-related information in face processing. Notably, stimulation effects were selective to social information: relative to sham, stimulation of both OFA and IND reliably slowed responses on social trials, with minimal effects on nonsocial trials (OFA b = 0.19; IND b = 0.08; sham b = -0.27; see Fig. 7). Accuracy showed the same pattern, where both OFA and IND stimulation sites selectively reduced performance for social judgments while largely sparing nonsocial judgments (OFA b = -0.95; IND b= -1.09; sham b = 2.73).

Stimulation of individualized, decoding-based sites produced behavioral effects that closely mirrored those observed when stimulating a canonical face-selective region. Although IND sites were defined based on multivariate patterns that jointly distinguished faces, places, and objects, their disruption affected behavior in a content-specific manner, selectively impairing judgments that depended on socially relevant facial information.

Together, these findings support a distributed account of face processing in which socially relevant information is supported by networks that extend beyond canonical face-selective regions and are partly individualized. Disrupting

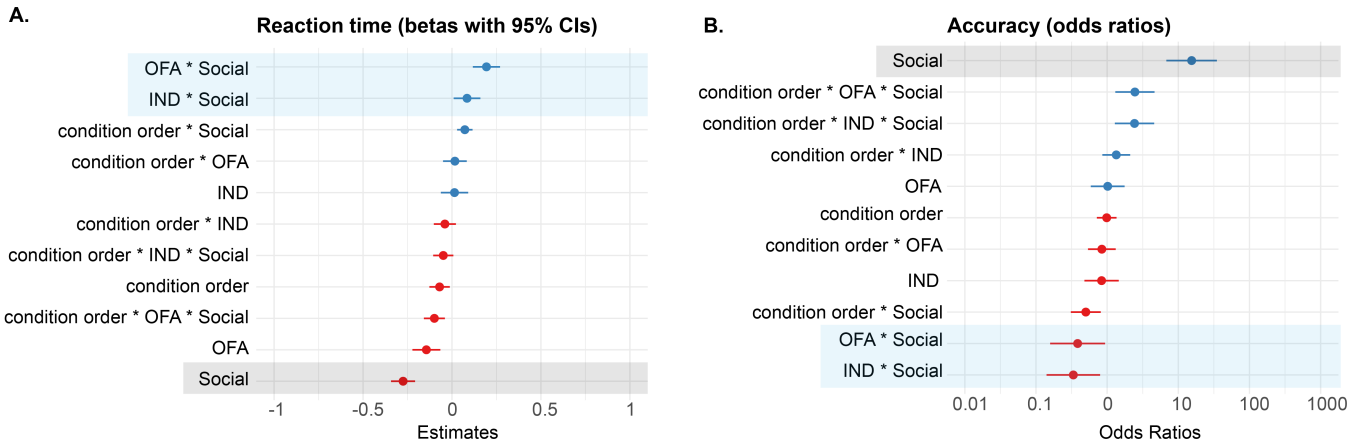

Figure 7: A) Beta estimates from the log RT model. Social trials under sham stimulation are processed quickly (highlighted in gray); stimulation to both OFA and IND regions significantly slows performance for processing social trials (blue). B) Accuracy for choosing the matched item reflected as odds ratios. Social trials under sham stimulation are typically very accurate (gray); stimulation to both OFA and IND regions significantly reduces accuracy for social trials (blue).

either canonical or non-canonical nodes within this network selectively impairs behavior when task demands rely on fine-grained social distinctions, providing causal evidence that distributed representational systems contribute meaningfully to semantic behavior. Overall, Study 2 provides causal evidence that individualized, distributed neural representations contribute meaningfully to behavior.

Discussion

Across two fMRI sessions in Study 1, we found that representations differentiating faces, places, and objects are broadly distributed across cortex, extending well beyond canonical category-selective regions. Although decoding accuracy was highest within traditional regions such as FFA, PPA, and LOC, substantial category information was recoverable outside these areas, consistent with distributed and network-based accounts of visual representation. These distributed representations were not only robust, but stable within individuals over time and markedly individualized across participants. Even after establishing voxel-wise correspondences across brains, decoding models generalized far better within an individual than between individuals, indicating that representational heterogeneity reflects stable structure.

Distributed representations are often assumed to be shared in form, differing only in strength or location across individuals (Haxby, 2012). Our results support a different picture: representations may be distributed and individualized, with each person expressing a reliable but idiosyncratic configuration of representational weights across cortex. Study 2 provides causal leverage on this interpretation. By targeting TMS to individualized, non-canonical sites derived from each participant's decoding solution, we tested whether these idiosyncratic representational nodes play a functional role in behavior. Stimulation of these sites selectively disrupted face similarity judgments, producing effects that closely mirrored those observed when stimulating a canonical face-selective region (OFA). Importantly, these effects were content-specific: disruption was observed primarily when decisions depended on *socially* relevant facial attributes, rather than nonsocial contextual features.

These findings should not be interpreted as evidence that face processing is localized to either canonical or individualized sites. Rather, TMS likely perturbs a broader representational state supported by coordinated activity across a distributed network. Disrupting any sufficiently connected node within this network— canonical or not— can reduce the fidelity of the representational pattern when task demands place heavy weight on that information. Faces, as socially salient stimuli, may place particularly strong demands on such distributed integration, rendering their representations more sensitive to partial network disruption. This interpretation of our results aligns naturally with connectionist perspectives, in which representations derive their explanatory power from graded, overlapping structure rather than discrete modules (Rogers and McClelland, 2004).

Rather than asking where faces, places, or objects "live" in the brain, the present results suggest that representational content emerges from weighted patterns distributed across many regions, patterns that are stable within individuals yet idiosyncratic in their configuration. Canonical regions may reflect places where these patterns are most consistently aligned across people, but they are not the only place in the brain supporting meaningful representations. By showing that an individual's own representational nodes have selective behavioral consequences, this work highlights the importance of treating representational structure as both distributed and personal, and suggests that understanding cognition will ultimately require theories that accommodate not only shared structure, but the richly idiosyncratic ways in which minds and brains make meaning.

References

Colón, Y. I., & Rogers, T. T. (2025). Finding motifs in mental representations of faces, places, and objects. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*.

Cox, C. R., & Rogers, T. T. (2021). Finding distributed needles in neural haystacks. *Journal of Neuroscience*, *41*(5), 1019–1032.

Cox, C. R., Rogers, T. T., Shimotake, A., Kikuchi, T., Kunieda, T., Miyamoto, S., Takahashi, R., Matsumoto, R., Ikeda, A., & Lambon Ralph, M. A. (2024). Representational similarity learning reveals a graded multidimensional semantic space in the human anterior temporal cortex. *Imaging Neuroscience*, *2*, 1–22.

Cox, R. W. (1996). Afni: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research*, *29*(3), 162–173.

Dubois, J., & Adolphs, R. (2016). Building a science of individual differences from fmri. *Trends in cognitive sciences*, *20*(6), 425–443.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598–601.

Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., Papademetris, X., & Constable, R. T. (2015). Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature neuroscience*, *18*(11), 1664–1671.

Frost, M. A., & Goebel, R. (2012). Measuring structural–functional correspondence: Spatial variability of specialised brain regions after macro-anatomical alignment. *Neuroimage*, *59*(2), 1369–1381.

Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., Ortega, M., Hoyt-Drazen, C., Gratton, C., Sun, H., et al. (2017). Precision functional mapping of individual human brains. *Neuron*, *95*(4), 791–807.

Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision research*, *41*(10-11), 1409–1422.

Haxby, J. V. (2012). Multivariate pattern analysis of fmri: The early beginnings. *Neuroimage*, *62*(2), 852–855.

Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding neural representational spaces using multivariate pattern analysis. *Annual review of neuroscience*, *37*(1), 435–456.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*(5539), 2425–2430.

Haxby, J. V., Guntupalli, J. S., Nastase, S. A., & Feilong, M. (2020). Hyperalignment: Modeling shared information encoded in idiosyncratic cortical topographies. *elife*, *9*, e56601.

Julian, J. B., Fedorenko, E., Webster, J., & Kanwisher, N. (2012). An algorithmic method for functionally defining regions of interest in the ventral visual pathway. *Neuroimage*, *60*(4), 2357–2364.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of neuroscience*, *17*(11), 4302–4311.

Laumann, T. O., Gordon, E. M., Adeyemo, B., Snyder, A. Z., Joo, S. J., Chen, M.-Y., Gilmore, A. W., McDermott, K. B., Nelson, S. M., Dosenbach, N. U., et al. (2015). Functional system and areal organization of a highly sampled individual human brain. *Neuron*, *87*(3), 657–670.

Li, Y., Mark, B., Raskutti, G., Willett, R., Song, H., & Neiman, D. (2020). Graph-based regularization for regression problems with alignment and highly correlated designs. *SIAM journal on mathematics of data science*, *2*(2), 480–504.

Li, Z., Colón, Y. I., Mukherjee, K., & Rogers, T. T. (2025). Iterated lasso reveals highly distributed and variable representations of faces, places, and objects. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*.

Pascual-Leone, A. (1999). Transcranial magnetic stimulation: Studying the brain–behaviour relationship by induction of 'virtual lesions'. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *354*(1387), 1229–1238.

Pitcher, D., Charles, L., Devlin, J. T., Walsh, V., & Duchaine, B. (2009). Triple dissociation of faces, bodies, and objects in extrastriate cortex. *Current Biology*, *19*(4), 319–324.

Pitcher, D., Walsh, V., & Duchaine, B. (2011). The role of the occipital face area in the cortical face perception network. *Experimental brain research*, *209*(4), 481–493.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. MIT press.

Schleim, S., & Roiser, J. P. (2009). Fmri in translation: The challenges facing real-world applications. *Frontiers in Human Neuroscience*, *3*, 845.

Solomon-Harris, L. M., Mullin, C. R., & Steeves, J. K. (2013). Tms to the "occipital face area" affects recognition but not categorization of faces. *Brain and cognition*, *83*(3), 245–251.