Neuron

Article

Neural geometry in the human hippocampus enables generalization across spatial position and gaze

Highlights

- Hippocampal neurons track the positions of virtual agents and gaze in a simple video game
- Place codes for self, other, and gaze are not identical but are not orthogonal either
- Partially overlapping maps enable both generalization and differentiation

Authors

Assia Chericoni, Chad Diao, Xinyuan Yan, ..., Joshua Jacobs, Sameer A. Sheth, Benjamin Y. Hayden

Correspondence

benhayden@gmail.com

In brief

Chericoni et al. show that neurons in the human hippocampus track the positions of multiple agents in a simple pursuit-based video game. They identified place codes for each agent and for gaze. The codes were systematically related in a way that allowed generalization while preserving the ability to differentiate.

Chericoni et al., 2027, Neuron *115*, 1–15
January 20, 2027 © 2026 Elsevier Inc. All rights reserved, including those for text and data mining, AI training, and similar technologies.
https://doi.org/10.1016/j.neuron.2026.08.004

CellPress

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

Neuron

CellPress

Article

Neural geometry in the human hippocampus enables generalization across spatial position and gaze

Assia Chericoni,[1] Chad Diao,[1] Xinyuan Yan,[1] Taha S. Ismail,[1] Elizabeth A. Mickiewicz,[1] Melissa Franch,[1] Ana G. Chavez,[1] Danika L. Paulo,[1] Vaishnav Krishnan,[2,3,4,5] Mohamed Hegazy,[4] Alica M. Goldman,[4] Lu Lin,[4] Gabriela Tantillo Sepulveda,[4] Garrett P. Banks,[1,6] Nisha Giridharan,[1] Mohammed Hasen,[1] Eleonora Bartoli,[1,3] Nicole R. Provenza,[1,2,3,7,8] Seng Bum Michael Yoo,[9,10,11] Jay Hennig,[2,3,8] Joshua Jacobs,[12] Sameer A. Sheth,[1,2,3,5,8,13,14] and Benjamin Y. Hayden[1,2,3,8,14,15,*]

[1]Department of Neurosurgery, Baylor College of Medicine, Houston, TX 77030, USA
[2]Department of Neuroscience, Baylor College of Medicine, Houston, TX 77030, USA
[3]Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA
[4]Department of Neurology, Baylor College of Medicine, Houston, TX 77030, USA
[5]Department of Psychiatry and Behavioral Sciences, Baylor College of Medicine, Houston, TX 77030, USA
[6]Department of Neurosurgery, Michael E. DeBakey VA Medical Center, Houston, TX 77030, USA
[7]Department of Bioengineering, Rice University, Houston, TX 77005, USA
[8]Neuroengineering Initiative, Rice University, Houston, TX 77005, USA
[9]Center for Neuroscience Imaging Research, Sungkyunkwan University, Suwon 16419, South Korea
[10]Department of Biomedical Engineering, Sungkyunkwan University, Suwon 16419, South Korea
[11]Department of Intelligence Health Precision Convergence, Sungkyunkwan University, Suwon 16419, South Korea
[12]Department of Neurology, University of Chicago, Chicago, IL 60637, USA
[13]Gordon and Mary Cain Pediatric Neurology Research Foundation Laboratories, Jan and Dan Duncan Neurological Research Institute, Texas Children's Hospital, Houston, TX 77030, USA
[14]These authors contributed equally
[15]Lead contact
*Correspondence: benhayden@gmail.com
https://doi.org/10.1016/j.neuron.2026.08.004

SUMMARY

Hippocampal neurons track the positions of self, others, and gaze direction. However, it is unclear how their respective neural codes differ enough to avoid confusion while allowing for abstraction across spatial frames. We recorded from populations of hippocampal neurons while human participants performed a joystick-controlled virtual prey-pursuit task. We found that neurons have mixed selective responses that map the positions of self, prey, and predator, as well as gaze direction. Their codes occupied mostly orthogonal subspaces, but the geometric structures of these subspaces allowed them to be aligned by simple linear transformations. Moreover, their geometries supported some degree of generalization, such that a linear rule learned on one agent could be transferred to another. This scheme enables reliable individuation and abstraction across both agent identity and viewpoint. Together, these findings suggest that hippocampal spatial knowledge is structured as a family of geometrically related manifolds that can be related through simple linear transformations.

INTRODUCTION

Our ability to track ourselves and others is implemented with the help of neural maps.[1–3] For example, hippocampal place cells exhibit increased firing rates when individuals enter specific locations, and social place cells do the same for conspecifics.[4,5] How do we simultaneously track self and other or two different others? At first blush, the problem appears trivial: we can solve it with separate neurons (e.g., labeled line codes) or orthogonal coding axes. However, this solution does not account for our ability to generalize across agents and contexts.[6,7] Imagine you are a child playing tag. While running, you see one of your friends slip on a muddy spot on the grass. Even without stepping there yourself, you infer that the floor is slippery and adjust your trajectory to avoid it. This capacity for cross-individual generalization sits in tension with the need to distinguish representations of self and others.[8]

Along with spatial mapping, the hippocampus tracks gaze direction.[9–15] Some accounts suggest that hippocampal responses reflect the observer's position in the world; others suggest that they encode gaze location or attentional sampling of space.[16–22] Resolving this issue is challenging because gaze and position are typically tightly coupled.[2] If the hippocampus represents both self location and gaze location, it faces a

All rights are reserved, including those for text and data mining, AI training, and similar technologies.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

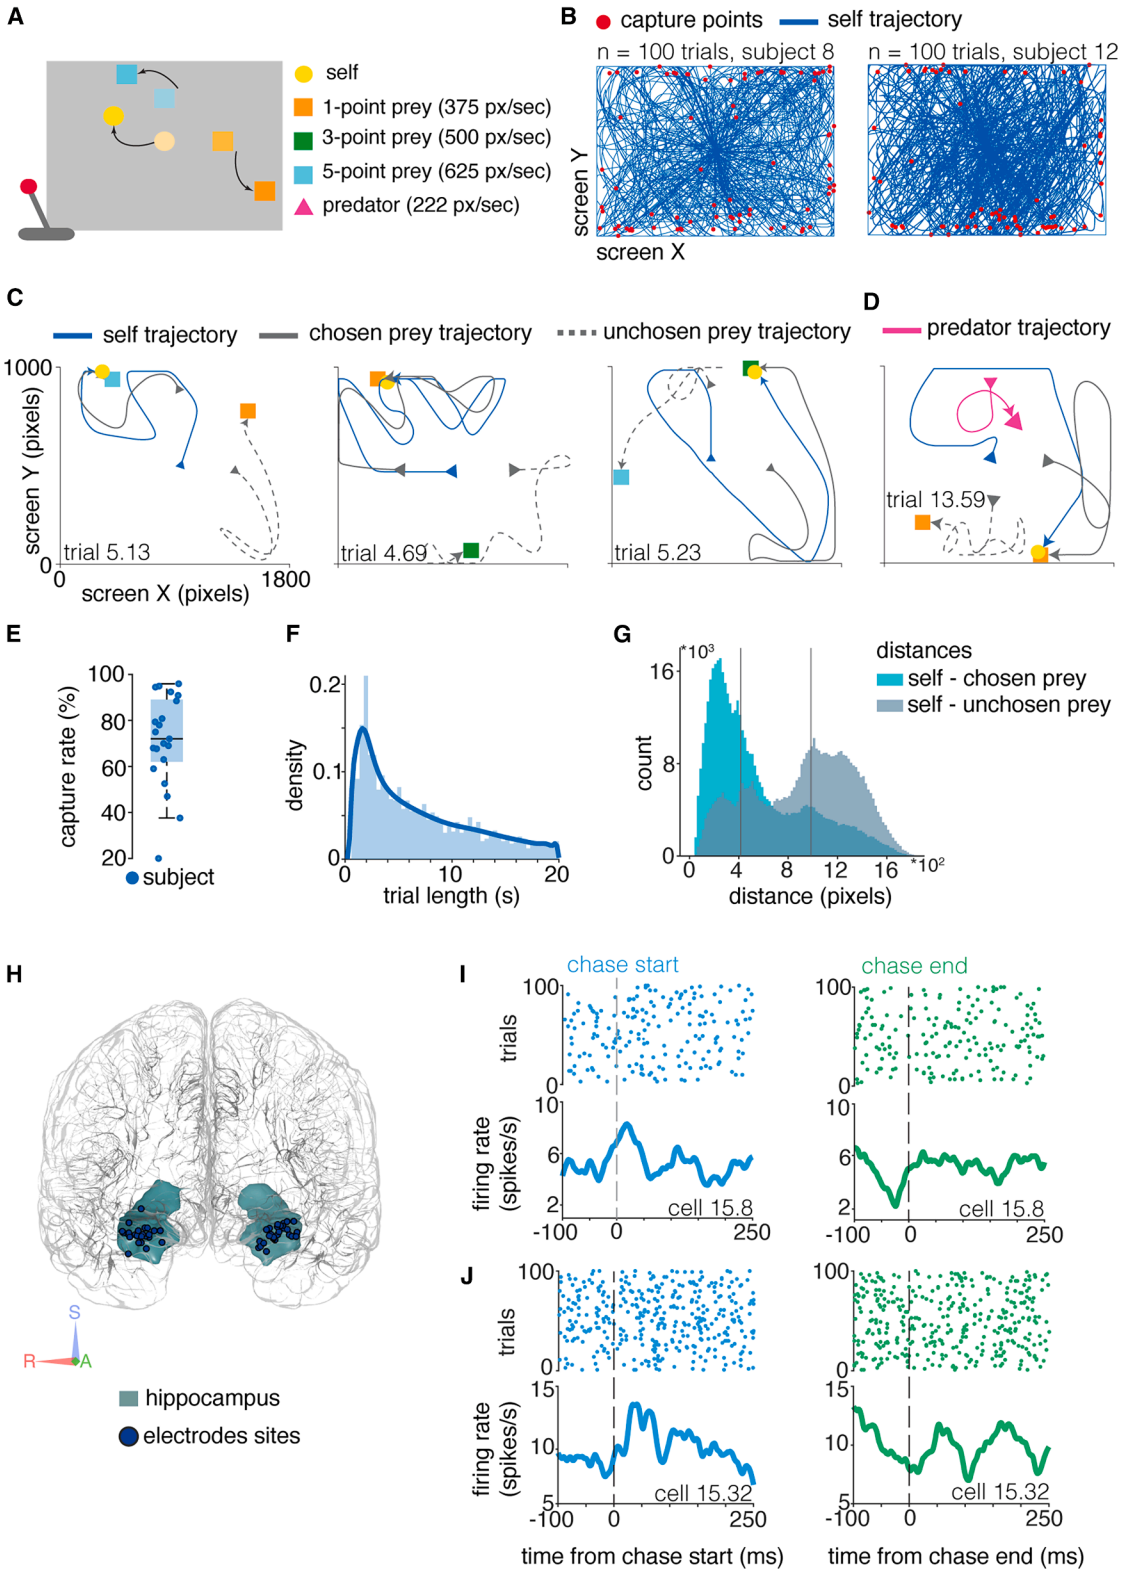


Figure 1. Task design and neural recordings

(A) Schematic of pursuit task design. Prey were colored squares. Self, yellow circle; prey, colored squares. For some participants, we also included a predator (STAR Methods).

(legend continued on next page)

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

problem analogous to that posed by multi-agent tracking, namely, how to maintain separable representations while preserving shared structure across reference frames.

Recent advances in theories of manifold coding suggest a solution.[23–26] Rather than assigning each variable to independent neurons or orthogonal coding axes, population activity can be placed on a manifold that can be projected in different ways depending on whether generalization or differentiation is needed.[6,7,27] This principle helps solve generalization-differentiation trade-offs in decision-making, motor control, working memory, and context-dependent behavior.[28–37] Here, we propose that this logic extends to navigation. Some evidence indicates that population geometry principles apply to hippocampal spatial codes.[38–43] Building on this work, we proposed that different maps can be encoded as partially overlapping maps, implemented in orthogonalized or semi-orthogonalized linearly transformable subspaces, allowing the brain to reuse spatial structure across mapped variables.[7] We therefore hypothesized that the hippocampus represents self location, the location of others, and gaze using semi-overlapping population maps arranged as semi-orthogonal linearly transformable subspaces.

We examined single-neuron activity in humans performing a continuous, naturalistic navigation task in a virtual two-dimensional open field.[44–46] Human participants used a joystick to pursue moving prey while sometimes simultaneously avoiding a predator, creating a setting in which the positions of self, multiple others, and gaze were all behaviorally relevant and changing. We found that individual hippocampal neurons carried information about the positions of the self-avatar, prey, predator, and gaze and used partially distinct response fields for each. These positional signals were organized into orthogonalized population subspaces with a consistent linear relationship that preserved aspects of the internal geometry across space representations, supporting abstraction. Decoders trained to read out spatial variables for one agent or viewpoint successfully transferred to others across multiple spatial partitions (i.e., cross-condition generalization performance, CCGP). Together, these results indicate that the hippocampus implements a flexible spatial code characterized by semi-overlapping spatial manifolds that can be re-centered on different agents and viewpoints.

RESULTS

Prey pursuit in humans

Human participants (n = 21) performed the prey-pursuit task[46] (Figure 1A; STAR Methods). On each trial, the participant used a joystick to continuously move the position of an avatar (yellow circle) in a rectangular field displayed on a computer screen (Video S1; Figures 1A and 1B). The participant had up to 20 s to capture fleeing prey (colored squares) to obtain points. Prey avoided the avatar with a deterministic strategy that combined repulsion from the avatar's current position with repulsion from the walls of the field (STAR Methods). The prey items were drawn randomly on each trial from a set of three that differed in maximum velocity and reward size. Over the course of each session, participants' avatars typically spanned the entire field (Figure 1B).

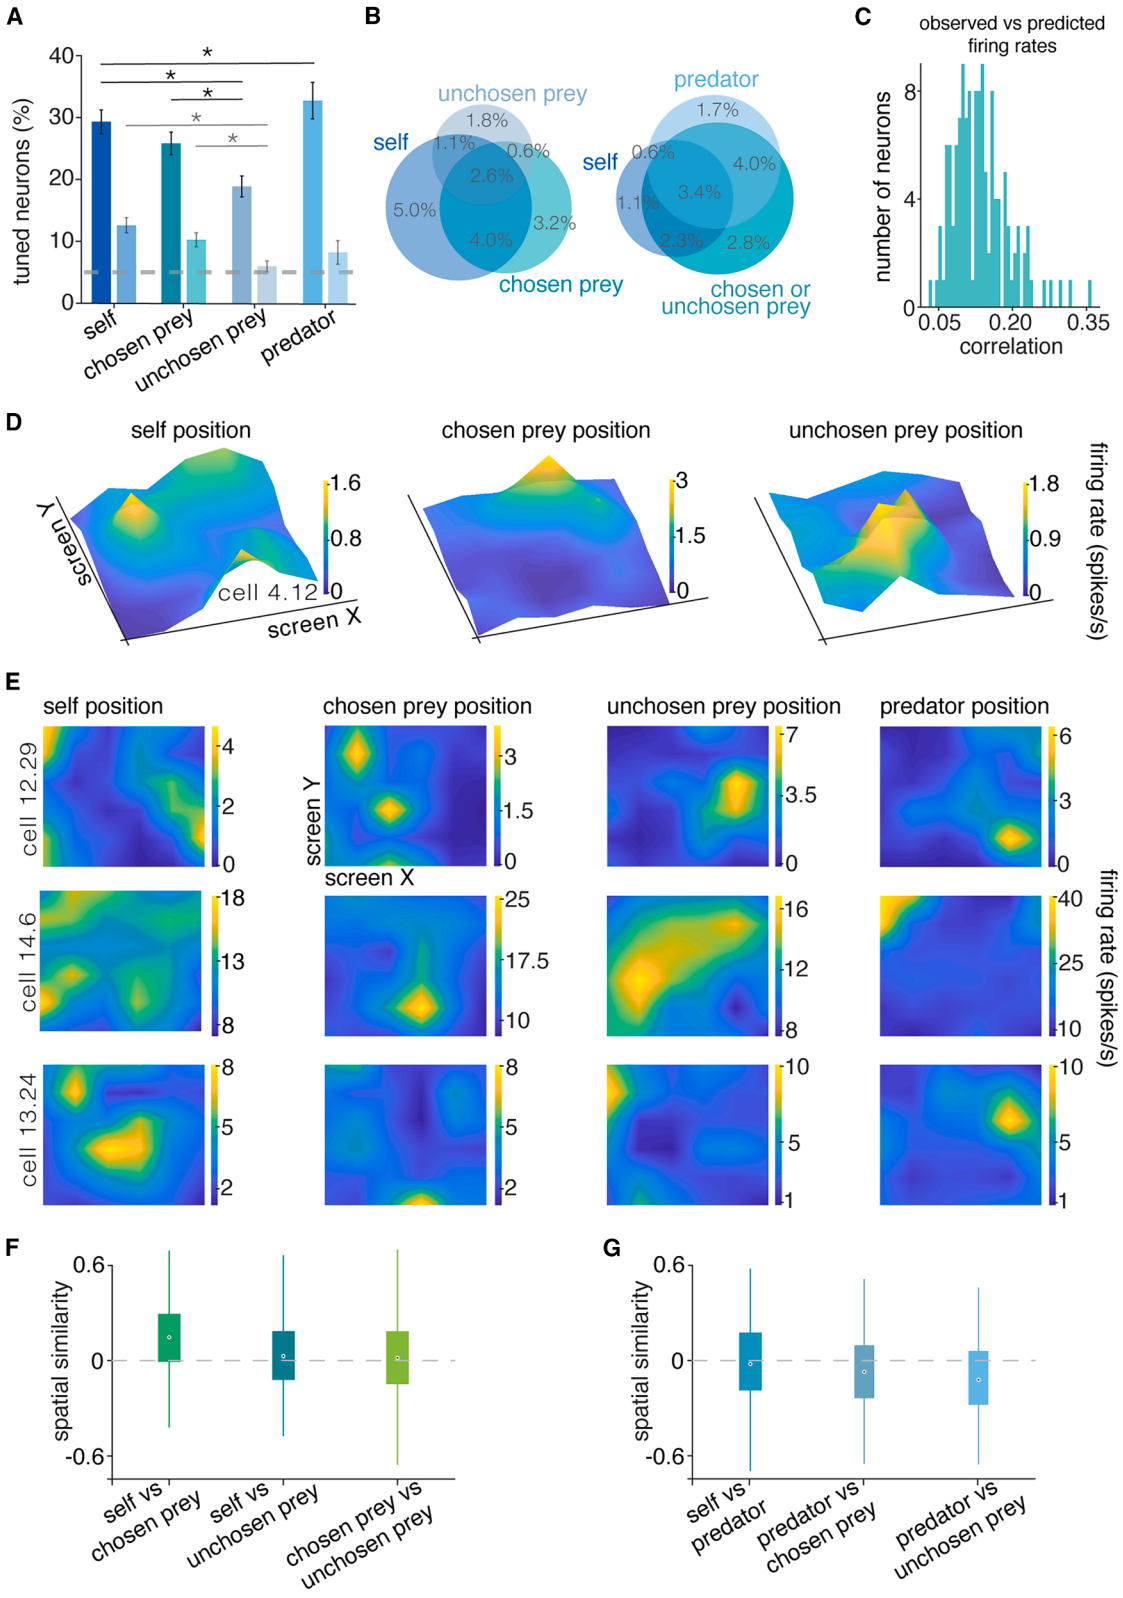
Each trial began with one or two prey appearing at one or two of the cardinal points (Figure 1A). In two-prey trials (from 50% to 81% of trials, depending on the participant), the participant was free to decide which prey to pursue at any moment (Figures 1C and 1D). In trials with two prey, participants successfully captured one of the prey in 71.25% of trials (Figure 1E) and, on successful trials, did so in an average of 7.51 s (SD: 1.89 s; Figure 1F) with an average reaction time of 0.87 s (SD: 0.07 s), defined as the interval between the agents' appearance on the screen and the participant's first move. Here and throughout, we use agent simply to refer to the self-avatar and the prey, as a neutral descriptive label. Participants' performance did not depend significantly on prey type (trial length $\times$ reward level, p = 0.07; reaction time $\times$ reward level, p = 0.24). As expected from the nature of the prey-pursuit task, self and chosen prey trajectories were positively correlated (x: r = 0.42; y: r = 0.31; both $p < 0.001$; r denotes Pearson's correlation coefficient), whereas those of self and unchosen prey were anti-correlated (x: r = −0.31; y: r = −0.25; both $p < 0.001$). Consistent with this, the self-avatar remained closer to the chosen prey than to the unchosen prey (mean Euclidean distance ± SD: 549.5 ± 391.6 vs. 920.3 ± 409.7 pixels; Wilcoxon signed-rank test, $p < 0.001$; Figure 1G).

In five participants, we employed a variant of the task in which, in addition to the prey, there was also a predator pursuing the participant's avatar (Figure 1D). In this case, the predator used a simple distance-minimizing pursuit strategy (STAR Methods). Our patients successfully evaded capture by the predator in 95.2% of trials.

We recorded responses from 726 neurons in the hippocampus while participants performed this task (average n = 34.6 neurons per participant, Figure 1H). Of these neurons, 176 were recorded in the variant of the task with a predator. Figures 1I and 1J show responses of two neurons with clear event-aligned responses; overall, however, neurons had complex selectivities that were not easily described as being about trial start and end. We therefore developed a more sophisticated approach to understanding how neural responses related to the positions of avatars in the task.

(B) Example sessions from two participants; red dots indicate points of prey capture.
(C and D) Typical example trials; arrowheads indicate direction of travel. Blue line, self-avatar trajectory; gray line, chosen prey trajectory; dashed gray line, unchosen prey trajectory; magenta line, predator trajectory.
(E) Distribution of performance (proportion of trials in which prey was captured) across our 21 participants.
(F) Histogram of trial durations; each trial ended either with the successful capture of the prey or after 20 s.
(G) Histograms depicting the distribution of distances between self and both prey across patients and trials; gray vertical lines represent the medians of each distribution.
(H) Glass brain showing recording locations of neurons for all patients.
(I and J) Example responses from two neurons aligned to trial start (blue) and trial end (green).
See also Video S1.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004



(legend on next page)

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

CellPress

Hippocampal maps for positions of self, prey, and predators

Because we cannot assume *a priori* the shape of spatial tuning functions, we adopted the Poisson generalized linear model (GLM) (linear-nonlinear [LN]/GLM) developed by Hardcastle et al.[47] (STAR Methods). We found that 29.5% (n = 214/726) of neurons tracked the position of the self-avatar on the monitor (Figure 2A). Because the participants freely chose one prey to pursue, the two prey could be differentiated based on their status as either chosen or unchosen (defined as whether they were eventually captured, Figures 1C and 1D). We found that 26.3% of neurons (n = 191/726) tracked the position of the chosen prey and 19.0% (n = 138/726) tracked the position of the unchosen prey (Figure 2A; both proportions are greater than chance, one-sided binomial test, $p < 0.05$). Although the proportions of neurons tuned to the self and chosen prey position were comparable (chi-squared test for proportions: $p = 0.18$), both were greater than the proportion tuned to the unchosen prey (self-unchosen prey, $p < 0.001$; chosen prey-unchosen prey, $p < 0.001$; Figure 2A). In the neurons that we recorded during predator trials, 33.0% (n = 58/176) tracked the position of the predator (Figure 2A). This proportion is higher than that of the unchosen prey ($p < 0.001$) but comparable to those for self and chosen prey ($p = 0.37$ and $p = 0.08$, respectively; Figure 2A).

To account for temporal autocorrelation between neuronal firing and behavior, we performed permutation testing by circularly shuffling the spike trains within each session (Zhang et al.[48]; STAR Methods). This criterion reduced the proportion of tuned neurons to 12.7% (92/726) for self position, 10.3% (75/726) for chosen prey, 6.1% (44/726) for unchosen prey, and 9.7% (17/176) for predator position (shaded bars in Figures 2A and S1A). Tuning to self, chosen prey, and predator remained above chance (one-sided binomial tests: $p < 0.001$, $p < 0.001$, and $p = 0.007$, respectively), whereas unchosen prey tuning no longer exceeded chance ($p = 0.11$). The relative pattern across variables was largely preserved: comparable proportions of neurons encoded self and chosen prey ($p = 0.16$), with both exceeding unchosen prey tuning ($p = 0.003$ and $p < 0.001$), and predator tuning was comparable to self tuning ($p = 0.44$). At the subject level, spatial tuning for self, prey, and predator positions remained evident in multiple participants, although fewer subjects contributed to each effect following permutation testing (Figure S1B).

We found no evidence for discrete populations of neurons corresponding to the different agents. Instead, populations overlapped, so that if a neuron participated in mapping one agent, it likely participated in mapping others (Figures 2B and S1C). Following permutation testing, 5.7% (n = 41/726) of the neurons encoded the position of any two agents, whereas 2.6% (n = 19/726) encoded all three. During predator trials, 6.9% of the neurons (n = 12/176) tracked the position of two or three agents, with 3.4% (6/176) tracking all four agents.

We evaluated the quality of fit of each neuron's firing structure through a Pearson correlation between the observed firing rates and those predicted by the LN/GLM (Figures 2C and S1D). Correlations were positive (mean = 0.13), confirming that the estimated tuning profiles captured meaningful spatial structure. Note that although the magnitude of this correlation was low, it was around half the approximate maximum observable correlation under a position-only Poisson model, which we directly simulated (Figure S2).

Example tuning surfaces illustrate the range of response fields we observed (Figures 2D and 2E). Tuning was non-localized; we observed multiple peaks. Together, these results indicate that, although a minority of neurons showed activity that resembles classic place-cell activity in its shape, these neurons are part of a larger and more heterogeneous distribution of tuning functions.[47,49,50]

Next, we asked whether hippocampal neurons had different maps for the different agents (Figures 2F and 2G). To quantify the relationship between maps, we used a "spatial similarity index" (spatial efficiency metric [SPAEF][51]; STAR Methods), which measures the correlation between spatial representations. Although SPAEF is not a well-known metric, it offers advantages for comparing spatial maps because it jointly captures spatial correlation, variability, and distributional overlap, thereby providing a more holistic measure of map similarity than correlation alone. A SPAEF value of zero indicates full orthogonality between maps, whereas values closer to +1 or −1, respectively, indicate correlation or anti-correlation between maps. The mean SPAEF between self and chosen prey maps was 0.13. This value was greater than zero ($p < 0.001$, Figure 2F; STAR Methods) and below the noise ceiling (0.22, quantified by half-splits bootstrapping). Thus, maps for self and others are positively correlated. Indeed, the mean SPAEFs between self and unchosen prey maps and between chosen and unchosen prey maps were also greater than zero (self-unchosen prey SPAEF = 0.04, $p < 0.001$; chosen-unchosen prey SPAEF = 0.006; $p < 0.001$; Figure 2F). Not surprisingly, all SPAEF values were below the noise ceiling (self-unchosen prey ceiling = 0.21; chosen-unchosen prey ceiling, SPAEF = 0.19; $p < 0.001$). In the subset of neurons in which we had predator

Figure 2. Hippocampal mapping functions for self and prey

(A) Proportion of neurons with significant tuning for agent positions according to the LN/GLM. Proportions do not sum to 100: one neuron may be tuned to more than one agent's position. Dashed gray line: chance. error bars: binomial standard errors. Asterisks and horizontal bars denote significant pairwise differences. Shaded bars indicate the proportion of neurons remaining significant after circular time-shift validation.

(B) Venn diagram showing overlap between neurons' selectivities for neurons that remained significant after circular time-shift validation.

(C) Distribution of correlations between observed and LN/GLM-predicted firing rates for neurons that remained significant after circular time-shift validation. Positive correlations indicate that the LN/GLM reliably captured each neuron's spatial firing pattern.

(D) Heatmaps showing tuning of three neurons to the positions of the self, chosen prey, and unchosen prey. True firing rates are shown.

(E) More examples of tuned neurons of each type. Each row corresponds to a single neuron, while each column represents firing relative to different agents (self, chosen prey, unchosen prey, and predator when present).

(F and G) Boxplots representing the median spatial similarity between maps across different agent representations. Error bars: standard error. Spatial similarity values near zero indicate orthogonality (low similarity) between maps. All collected neurons are included in this analysis.

See also Figures S1 and S2.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

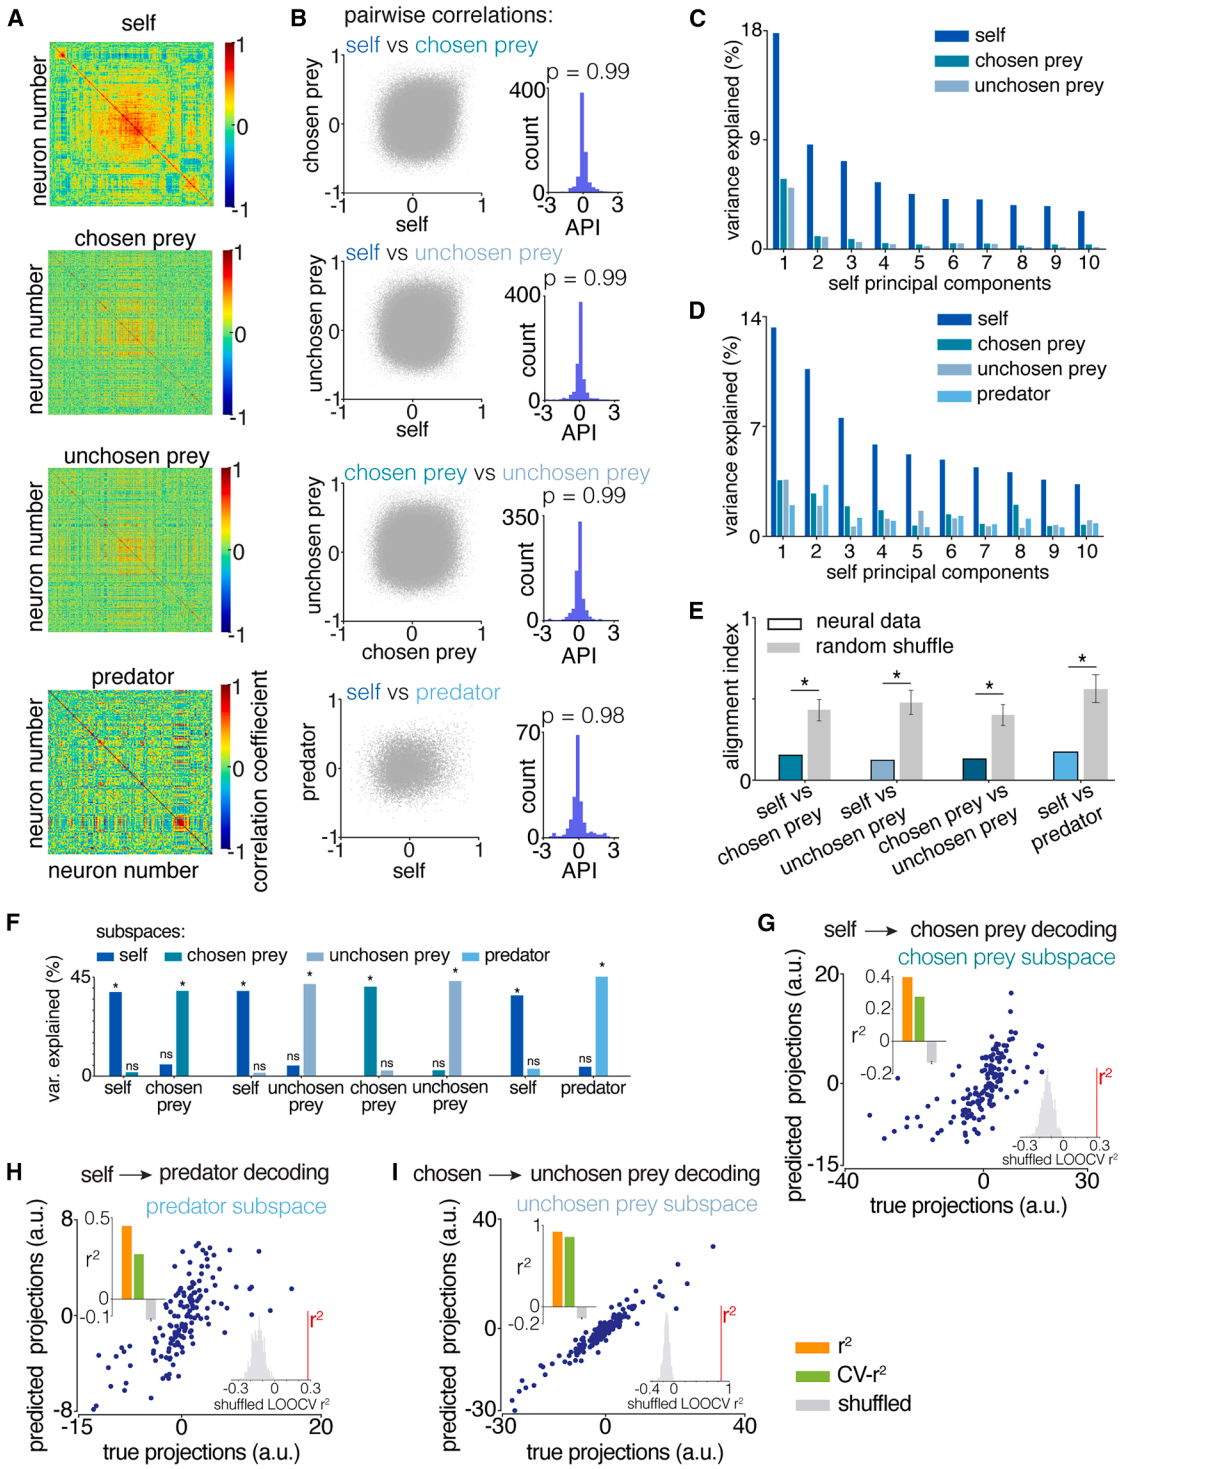


Figure 3. Population-level organization separates and links spatial maps for self and others

(A) Correlation matrices for self, prey, and predator tuning functions across neurons. Each *x* and *y* axis represents individual neurons, and the color (*z* axis) indicates the Pearson correlation coefficient between neuron pairs. The order of the neurons is identical across matrices, determined by hierarchical clustering of the self matrix using centroid linkage. The correlation structure differs across matrices.

(B) Relationship between pairwise correlation values across each pair of matrices. Each point represents a neuron pair. The weak correlation indicates that the similarity structure among neurons for self does not predict the similarity structure among neurons for prey and predators. Histogram: distribution of agent preference index, quantifying the strength of neural activity for self vs. others. Distributions are unimodal.

(legend continued on next page)

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

data, we found that the SPAEF between self and predator was slightly negative (SPAEF = −0.01, $p < 0.001$; Figure 2G).

Hippocampal spatial maps occupy separable yet linearly related subspaces

To assess the degree of orthogonality among encoding subspaces, we used a method developed to study motor regions[29] that projects activity onto low-dimensional subspaces that capture most of the variance between agents' representations (STAR Methods). We first computed the correlation matrices for self, prey, and predator tuning functions (Figure 3A). These matrices are, essentially, plots of all possible correlation coefficients between the tuning functions for all possible pairs of neurons. In the first row of Figure 3A, the entries of the matrix are sorted by a dendrogram with centroid linkage, thereby creating a hotspot in the center. The critical finding is that, when sorting the other matrices using the same indices, this hotspot nearly disappears, indicating that correlation *across* matrices was close to zero. In other words, maps for self and others tended to be distinct. Indeed, when we compared the entries of the self matrix against the entries of the prey and predator matrices, as well as the entries of the chosen prey matrix against the ones of the unchosen prey matrix, we found little correlation (self-chosen prey $r^2 = 0.02$; self-unchosen prey $r^2 = 0.008$; self-predator $r^2 = 0.01$; chosen-unchosen prey $r^2 = 0.007$; all $p < 0.001$; Figure 3B).

Why are spatial tuning functions mostly uncorrelated? One possibility is that neurons form distinct pools, each selectively representing a particular agent (this would be called a labeled line code). However, given the overlapping maps observed in our tuning analysis (especially Figure 2B), we expected this not to be the case. To confirm this prediction, for each neuron, we calculated an agent-preference index (API) that quantifies the modulation of neural activity for one agent compared with the modulation for another agent (this is analogous to what Elsayed et al.[29] call the "epoch preference index"). A neuron that is concerned with the representation of one or another agent would have an index of +1 or −1, whereas a neuron selective for multiple agents would have an intermediate index of 0. Therefore, the presence of self-, prey-, and predator-only neurons would result in a bimodal distribution of APIs across the population of neurons. In reality, all the API distributions clustered tightly around zero. Indeed, we can statistically test for bimodality (Hartigan dip test for a bimodal distribution, all $p = 0.99$; histograms in Figure 3B). As a result, any differences between these representations are likely to come from how their activity patterns are arranged relative to each other in population subspaces, not from functional categoricality.[52–54]

To identify the relationship between these activity patterns, we used a principal component (PC) analysis.[27,29] If the coding subspaces for self and others are orthogonal, self-PCs should capture little to no prey and predator variance. Indeed, we found that the top ten self-PCs explained 70.7% of the variance of self maps, but they captured very little prey and predator map variance (12.4% of chosen prey variance; 9.6% of unchosen prey variance; 12.6% of predator variance; Figures 3C and 3D). We validated these results with half-split cross-validation (STAR Methods). Across splits, self-prey correlations did not exceed the within-self noise floor (self-chosen prey $r^2 = 0.008$, $p = 0.48$; self-unchosen prey $r^2 = 0.005$, $p = 1$). The chosen-unchosen prey and self-predator correlations were higher than baseline ($r^2 = 0.005$, $p < 0.001$ and $r^2 = 0.02$, $p = 0.02$, respectively).

We next computed the "alignment index," which measures the degree of rotation between two subspaces. The alignment index is defined as the variance of prey and predator tuning functions captured by the top ten self-PCs, normalized by the variance captured by the top ten PCs for each corresponding agent (STAR Methods; Elsayed et al.[29]; Yoo et al.[45]). Values near zero indicate orthogonal subspaces, whereas values near one indicate strong alignment. We found that the alignment between self and the other agents, as well as between chosen and unchosen prey, was greater than zero (self-chosen prey = 0.15; self-unchosen prey = 0.12; self-predator = 0.17; chosen-unchosen = 0.13; all $p < 0.001$; Figure 3E). Thus, self and other representations lie in semi-orthogonal (meaning, between orthogonal and collinear) neural subspaces. Among other things, these results tell us that, although predator-related analyses were limited to a subset of patients and neurons, the evidence we do have suggests that they were tracked in the same manner as other agents.

The observed semi-orthogonality may be partially driven by separate neuronal pools, with mixed-selective neurons contributing to the collinearity and single-agent neurons contributing to the orthogonality. To test this possibility, we computed alignment indices separately for self-only neurons and neurons tuned to all three agents (8.7% and 10.5% of the population, respectively; Figure S1C). In both cases, the alignment indices remained intermediate (self-only neurons: self-chosen prey = 0.33, self-unchosen prey = 0.28; all-three neurons: self-chosen prey = 0.27, self-unchosen prey = 0.24; all $p < 0.001$), indicating that semi-orthogonal organization was not driven by separate neuronal subpopulations (see Figure S3 for additional analyses). At the subject level, the overall pattern of weak positive correlations between self, prey, and predator representations was

(C) Percentage of self, chosen, and unchosen prey variance explained by the top ten self-PCs.
(D) Percentage of self, prey, and predator variance explained by the top ten self-PCs.
(E) Alignment index for neural and random data. Self and others representations are more orthogonal than expected by chance.
(F) Percentage variance explained by the self and prey subspaces. Cross-map projections significantly reduce the variance explained.
(G) Chosen prey subspace projections predicted from self subspace coordinates using a linear transformation, plotted against observed chosen prey projections. Each point represents a spatial condition (36 spatial bins across the four-dimensional subspace). Histogram: shuffled cross-validated R^2 distribution. Decoding performance is significantly greater than chance. Bar plot: linear decoder performance; training R^2, LOOCV R^2, and shuffled control for self-chosen prey linear mapping. Error bars: 95% confidence interval of the shuffled distribution.
(H) Same as in (G), but for predator projections predicted from self subspace coordinates.
(I) Same as in (G), but for unchosen prey projections predicted from chosen prey subspace coordinates.
See also Figures S6 and S7.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

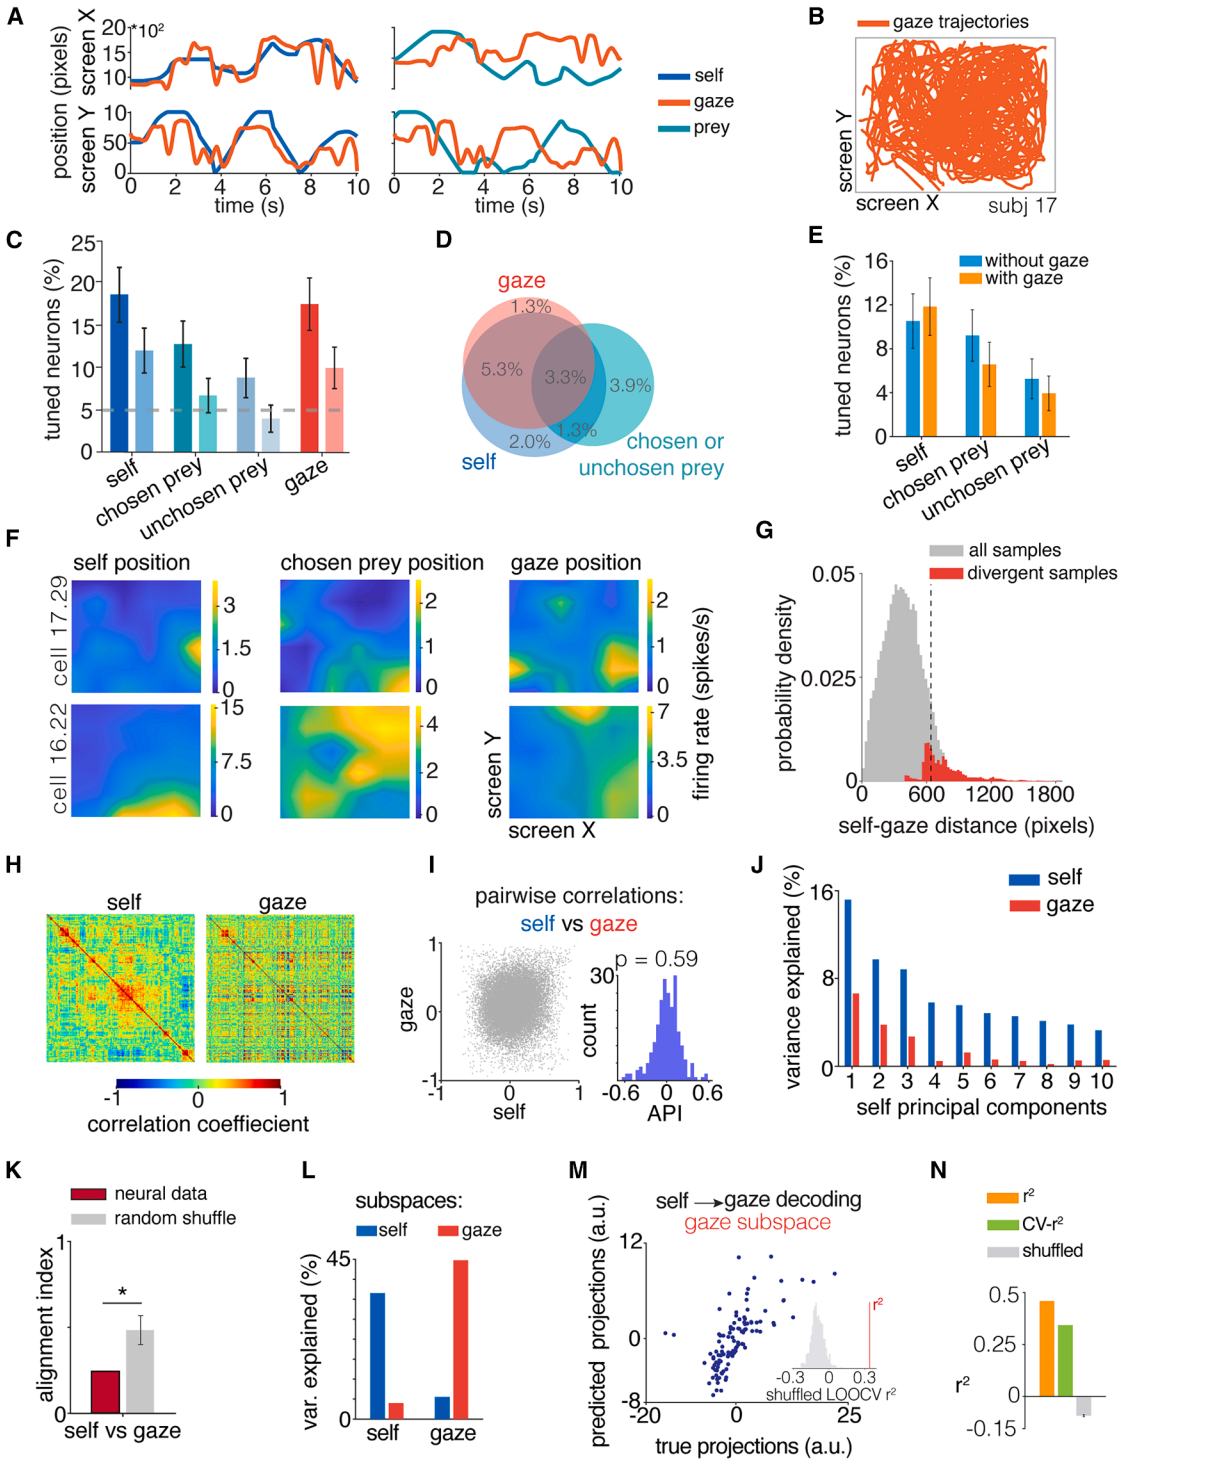


Figure 4. Neural effects of gaze on hippocampal neurons

(A) Two examples of 10-s gaze trajectories (red line) aligned to the position of the self (blue line) and chosen prey (cyan line). The *x* and *y* axes are separated.

(B) Gaze trajectory across a session for a typical patient.

(C) The LN/GLM reveals separable effects of gaze on neuronal firing alongside main effects of self and prey position. Shaded bars indicate the proportion of neurons remaining significant after circular time-shift validation.

(D) Gaze did not recruit a distinct neuronal population. Only neurons that remained significant after circular time-shift validation are included.

(legend continued on next page)

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

preserved across participants, consistent with a semi-orthogonal organization of spatial representations (Figures S4A and S4B). Similarly, subject-level alignment indices showed some variability across participants but remained significantly lower than shuffled controls in a substantial fraction of subjects (Figures S5 and S6A).

Subspaces are distinct but linearly transformable

This semi-orthogonal structure suggests that the spatial representations may be organized in a way that allows self and other maps to be separated. Such separation would not be achievable if the spatial maps were fully collinear. Therefore, we solved a generalized eigenvalue problem that identifies the dimensions capturing self variance and then constructs an orthogonal set of dimensions for prey and predator[29,44,55,56] (STAR Methods). We found that each agent's spatial representation contains self-, prey-, and predator-specific dimensions that can be arranged into mutually orthogonal subspaces (Figure 3F). Specifically, we identified four self-dimensions that captured 38.1% of the variance in self maps; four dimensions for the chosen and unchosen prey that captured 38.7% and 41.6%, respectively, of the variance in chosen and unchosen prey maps; and four predator dimensions that explained 44.9% of the variance in predator maps (Figure 3F). (We used four dimensions per agent because this was the minimum dimensionality required to explain at least 35% of the variance under the orthogonality constraint; STAR Methods.) Each self, prey, and predator subspace was specific to its corresponding agent: projecting self maps onto the prey and predator subspaces (and vice versa) explained only a minimal portion of the variance (bootstrap, all $p > 0.05$; Figure 3F). Thus, maps that are mixed at the neuron level are separable at the population level.

Even orthogonal subspaces, if structured appropriately, can have a relationship defined by a linear transformation (typically a rotation). This transformability allows the subspaces to be easily transformed when generalization is favored but kept distinct when it is not.[29,36,57] To test for linear transformability, we trained a linear decoder to predict one agent's spatial map from another's within the corresponding orthogonal subspaces (Figures 3G–3I). Specifically, we used self maps projected onto the four-dimensional self subspace to predict the activity patterns in the chosen and unchosen prey subspaces (STAR Methods). Self and prey maps were indeed related through a linear mapping (self-chosen prey $R^2 = 0.40$, Figure 3G; self-unchosen prey $R^2 = 0.38$). Leave-one-out cross-validation (LOOCV) confirmed that predictions on held-out conditions remained significantly above the shuffled baseline (self-chosen prey LOOCV $R^2 = 0.29$, $p < 0.001$, Figure 3G, see insets; self-unchosen prey LOOCV $R^2 = 0.24$, $p < 0.001$, data not shown). Predator maps were also linearly predictable from self maps ($R^2 = 0.46$; LOOCV $R^2 = 0.28$, $p < 0.001$; Figure 3H). Finally, chosen and unchosen prey subspaces were linearly related ($R^2 = 0.92$; LOOCV $R^2 = 0.85$, $p < 0.001$; Figure 3I). Thus, although the hippocampus uses separate codes for different agents, these codes are linearly transformable.

Distinct yet complementary coding of gaze and self position

We recorded eye movements from six patients while they were performing the prey-pursuit task (Figures 4A and 4B; STAR Methods). Gaze direction was correlated with the positions of the prey, so we leveraged the ability of the LN/GLM to deal with the resulting multicollinearity.[47] We focused on the subset of neurons collected from patients with sufficiently large numbers of trials for follow-up analyses ($n = 4$, see STAR Methods). We found that a comparable proportion of neurons encoded gaze position (17.1%, $n = 26/152$; Figure 4C) and self position (18.4%, $n = 28/152$; Figure 4C); after circular time-shift validation, these proportions were reduced to 9.9% ($n = 15/152$, $p = 0.009$) and 11.8% ($n = 18/152$, $p < 0.001$), respectively (Figures 4C, S7A, and S7B). In these participants, a significant fraction of neurons was also tuned to the position of the chosen prey (12.5%, $n = 19/152$; Figure 4C) and unchosen prey (8.6%, $n = 13/152$; Figure 4C); the corresponding validated percentages were 6.6% ($n = 10/152$, $p = 0.23$) and 3.9% ($n = 6/152$, $p = 0.77$), respectively. As above, these populations were not segregated: 9.9% of neurons ($n = 15/152$; Figures 4D and S7C) encoded the combination of two or three predictors.

We refit the GLM with and without gaze as a predictor. Including gaze did not substantially improve model performance for self or prey tuning, indicating that gaze position was not a trivial reflection of body position but instead contributed distinct

(E) Including gaze as a regressor modestly decreased model performance for self and prey, indicating that gaze-avatar correlations are not a major confound. As in (D), only neurons that remained significant after circular time-shift validation are included.
(F) Example heatmaps show responses of two neurons to gaze position; self and chosen prey maps are also shown.
(G) Histogram of the Euclidean distances between gaze and self trajectories. Gray: full-distance distribution across all samples. Orange: samples retained after selecting distances above the 90th percentile separately for each patient. Dashed line: average 90th-percentile threshold across patients.
(H) Correlation matrices for self and gaze tuning functions across neurons. The x and y axes represent individual neurons, whereas the z axis represents the strength of correlation between pairs of neurons. Neuron order is identical for self and gaze (centroid-linkage dendrogram from the self matrix).
(I) Relationship between pairwise correlation values across self and gaze matrices (on the left). Each point represents a neuron pair. The weak correlation indicates that the similarity structure among neurons for self does not predict the similarity structure among neurons for gaze. Histogram: distribution of the epoch preference index. The distribution is not significantly bimodal.
(J) Percentage of self and gaze variance explained by the top ten self-PCs.
(K) Alignment index for neural and random data. Self and gaze representations are more orthogonal than expected by chance.
(L) Percentage of variance explained by the self and gaze subspaces. Cross-map projections significantly reduce the variance explained.
(M) Gaze subspace projections predicted from self subspace coordinates using a linear transformation, plotted against the observed gaze projections. Each point represents a spatial condition. Inset: shuffled cross-validated R^2 distribution; significance demonstrates that self and gaze subspaces are related by a linear transformation.
(N) Linear decoder performance: training R^2, LOOCV R^2, and shuffled control for self-gaze linear mapping. Error bars: 95% confidence interval.
See also Figures S5–S7.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

information to hippocampal firing (Figures 4E and S7D). Model fit quality, measured as the Pearson correlation between observed and GLM-predicted firing rates, was comparable across models (with gaze: r = 0.12; without gaze: r = 0.12). For single neurons, gaze tuning was spatially structured and comparable to self, prey, and predator tuning (Figure 4F; unchosen prey and predator data not shown).

Nevertheless, gaze and self trajectories *were* naturally coupled during the prey-pursuit task. We further tested whether self tuning could instead reflect a gaze confound, focusing on moments in which gaze and self diverged. We visualized the full distribution of distances between self and gaze trajectories (mean Euclidean distance ± SD: self = 384.88 ± 97.96 pixels; gray distribution in Figure 4G) and adopted a threshold of 10%, wherein we restricted our analysis to the most divergent 10% of samples (i.e., 10% of samples per patient; orange distribution in Figure 4G). Refitting our LN/GLM on this subset of samples, we found that self position tuning survived (8.6%, n = 13/152 neurons; binomial test, p = 0.04), despite the substantially reduced number of samples included in the analysis. Note that the identical proportion of self-tuned neurons in this subsampled analysis (8.6%) and unchosen-prey-tuned neurons in the full dataset (8.6%) is a coincidence.

How are self and gaze maps organized relative to each other? Using the same dendrogram-sorting approach described above, we found differences (Figure 4H) and, consistent with visual intuition, low correlations ($r^2 = 0.02$, $p < 0.001$; Figure 4I). This pattern was consistent at the subject level (all $r^2 < 0.16$; Figure S4C) and unimodal (Hartigan's dip test, $p > 0.05$; histogram in Figure 4I), arguing against distinct neuronal subpopulations.

Gaze and self occupy distinct neural axes: the top ten self-PCs explained 65.7% of self variance but only 17.2% of gaze variance (Figure 4J). A cross-validation procedure tells us that this is not due to noise and that there is a small but significant collinearity ($r^2 = 0.02$, $p < 0.001$). The alignment index between self and gaze was greater than zero but below the shuffled control (0.24, $p < 0.001$; Figure 4K), indicating semi-orthogonality. We found similar results at the subject level (Figure S6B). Statistics for gaze/prey were similar (chosen prey-gaze = 0.19; unchosen prey-gaze = 0.18).

Subspaces for gaze and position are also distinct but linearly transformable

We next asked whether self and gaze subspaces could be separated into mutually orthogonal dimensions. Solving the generalized eigenvalue problem identified distinct self- and gaze-specific dimensions (Figure 4L). Projecting maps across subspaces explained only a minimal portion of the variance, confirming that each subspace captured variance specific to its corresponding variable (bootstrap, all $p > 0.05$; Figure 4L). Finally, we tested whether these orthogonal subspaces were nonetheless linearly related. Using projections onto the self subspace to predict gaze subspace activity via a linear transformation revealed significant decoding performance (Figures 4L and 4M). R^2 values exceeded the shuffled controls (self-gaze R^2 = 0.55; Figure 4L), indicating that although self and gaze occupy separable neural subspaces, their geometry permits a systematic linear mapping between them. Cross-validation confirmed that predictions on held-out data remained significantly above the shuffled baseline (self-gaze LOOCV R^2 = 0.40; $p < 0.001$, Figures 4M and 4N). Thus, hippocampal representations of gaze and self position are organized in separable yet geometrically coordinated subspaces.[29,36]

This geometric organization corresponds to what Elsayed and colleagues[29] define as an "orthogonal but linked" structure, in which separable subspaces share a consistent relational geometry. Such configurations are considered a hallmark of abstract (or disentangled) representations: neural populations preserve a shared structure of space across coordinate frames while isolating agent-specific components.[27] This geometry enables, though does not by itself guarantee, generalization across agents and viewpoints.

Neural population geometry enables generalization across agents and gaze positions

The above results show that the population code *can* support generalization across agents but do not demonstrate that such generalization actually occurs. We tested this using CCGP[7] (Figure 5A). We trained a linear classifier (support vector machine, SVM; STAR Methods) on population activity for self position to discriminate left from right screen locations and then applied the same classifier to chosen prey and unchosen prey population activity to test for generalization across agents. When trained on self activity, the classifier generalized well to chosen prey (CCGP = 0.73; Figure 5B) and unchosen prey activity (CCGP = 0.32; Figure 5B), with both values significant (two-tailed permutation test, both $p < 0.001$; Figure 5B).

The value lower than 0.5 for unchosen prey was particularly interesting. It reflected an inverted decoding axis. In other words, although they are generalizable, the maps used for self position and for unchosen prey are reversed, like photographic negatives. We observed a similar axis inversion when training the SVM on chosen-prey-related firing and testing on unchosen prey (CCGP = 0.26; $p < 0.001$; Figure 5B). (We did not have enough data to test the predator condition.) Despite inter-subject variability, the main CCGP patterns were reproduced in a substantial fraction of subjects (Figure S8).

Because the positions of self and chosen prey were naturally coupled, this generalization may be explained by spatial proximity between agents. We ruled out this confound by performing the CCGP analysis only on the subset of samples in which the distributions of distances between self and each prey were matched (see distance distribution overlap in Figure 1G). We found that the generalization of self to both chosen and unchosen prey representations remained significantly different from chance (CCGP = 0.66 and 0.40, respectively; both $p < 0.001$), although values were reduced compared with the original analysis (both $p < 0.001$). Therefore, proximity may contribute to the generalization of self to prey representations, but it does not fully explain it.

We then trained a classifier on population activity associated with gaze location and evaluated its performance on self- and prey-related activity. The gaze-activity-trained classifier generalized above chance to self (CCGP = 0.79; $p < 0.001$; Figure 5C) and chosen prey (CCGP = 0.62; $p < 0.001$; Figure 5C). Generalization to unchosen prey representations was not significant

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

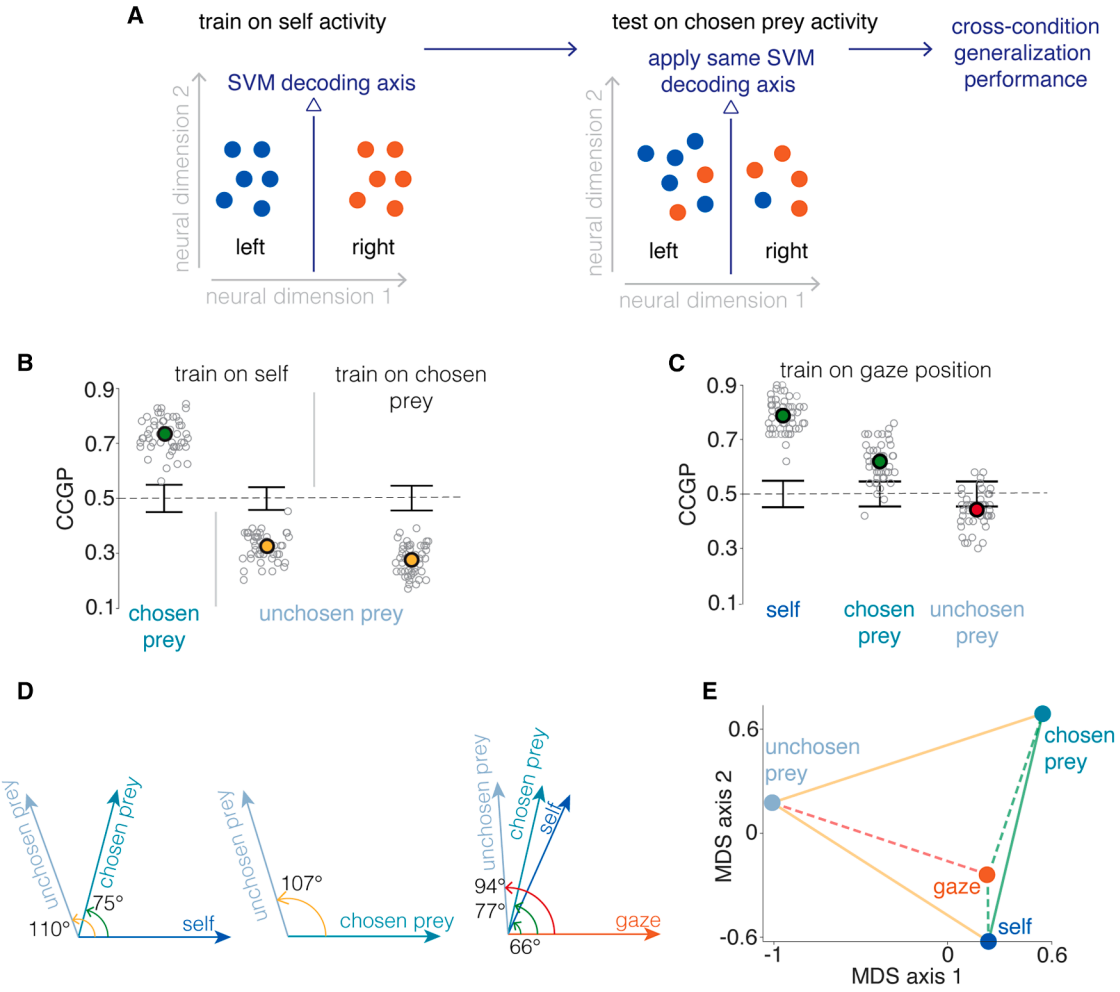


Figure 5. CCGP across agents and gaze

(A) Schematic of the CCGP analysis. A linear SVM is trained to classify left vs. right position on the screen using population activity for one agent (here, self). The same decoding axis is then applied to activity from a different agent (here, chosen prey), and CCGP quantifies how well decoding performance transfers across conditions (STAR Methods).

(B) A linear SVM trained on self position activity generalized to chosen prey position and with an inverted axis to unchosen prey position (left). A linear SVM trained on chosen prey position activity generalized with an inverted axis to unchosen prey position (right). Error bars: ±2 standard deviations around chance. Gray circles: CCGP values from individual resampling iterations; colored circles: mean. Green: significantly above chance; yellow: significantly below chance (inverted axis); red: no significant CCGP.

(C) Same as in (B), but here the classifier was trained on gaze positions and tested on the self and prey population vectors.

(D) Angles between decoding axes of linear classifiers, each trained and tested separately on self, both prey and gaze, quantifying similarity of population coding directions. Decoding axes are determined as the SVMs' weight vectors.

(E) Representation of the neural geometry obtained by applying MDS to the set of pairwise cosine distances between decoding axes. The distance between points reflects the angular separation between decoding directions in the neural space. Links between points reflect the CCGP results. Continuous lines: CCGP across agents; dashed lines: CCGP across agents and gaze.

See also Figures S8 and S9.

(CCGP = 0.44; p = 0.98; Figure 5C). At the subject level, generalization of gaze to self and chosen prey representations was preserved in most subjects (Figure S9). However, CCGP between gaze and unchosen prey was variable across subjects, with only two subjects showing no significant generalization (Figure S9). Therefore, caution is warranted in interpreting this specific effect.

These results were not tied to the specific screen partition (left vs. right), as repeating the same analysis using alternative partitions (top vs. bottom and diagonal vs. anti-diagonal splits) yielded comparable results. Indeed, a decoder trained on self spatial representations consistently generalized to those of the chosen prey above chance (top-bottom: 0.70; diagonal: 0.72; anti-diagonal: 0.67; all $p < 0.001$), whereas generalization to the unchosen prey remained significant, with inverted axes (top-bottom: 0.31; diagonal: 0.33; anti-diagonal: 0.32; all $p < 0.001$). Similarly, generalization from chosen to unchosen prey also remained significant, with inverted axes across all

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

partition schemes (top-bottom: 0.29; diagonal: 0.32; anti-diagonal: 0.34; all $p < 0.001$).

We quantified the angles between the decoding axes for each pair of agents and gaze representations by computing the angle between corresponding SVM weight vectors (Figure 5D; STAR Methods). To visualize these relationships, we computed the pairwise cosine distances between weight vectors and projected them into a two-dimensional space using multidimensional scaling (MDS; Figure 5E). In this projection, the distances between points are proportional to the angular distances between decoding axes in neural space. Consistent with the MDS visualization, self and gaze showed the greatest alignment (cosine similarity = 0.398), whereas gaze and unchosen prey formed a near-orthogonal pair (cosine similarity = −0.066).

DISCUSSION

In a virtual prey-pursuit task, neurons in the human hippocampus track positions of multiple moving items (self, two prey, and a predator in a subset of patients) and gaze simultaneously. How do single neurons represent multiple pieces of information without confusion? We find that the hippocampus utilizes semi-orthogonal, transformable subspaces rather than separate neuronal populations. Specifically, representations for self, others, and gaze occupy neural dimensions that are partly, but not fully, orthogonal. This interpretation is supported by two complementary analyses. First, the decoder-transfer results provide the strongest evidence for generalization because they directly test linear transformability across representations. Second, alignment indices and covariance correlations are more modest in magnitude and, rather than directly quantifying representational overlap, provide supporting evidence for a shared representational geometry.

This semi-orthogonal organization allows for simultaneous generalization and differentiation, depending on readout.[6,29,37,58] Together, these findings suggest that principles of manifold statistics apply to navigational codes in the hippocampus just as they do to prefrontal representations.[39–41,43,48,59] Note that our population-level analyses rely on pseudo-populations constructed by pooling neurons across subjects. As a result, the observed generalization does not prove within-subject generalization.

Our observation that gaze position was encoded at the single-neuron level with a prevalence comparable to spatial coding may provide additional perspective on ongoing debates about spatial and gaze representations in the primate hippocampus.[12,16–22] Prior work in nonhuman primates and humans has demonstrated hippocampal sensitivity to gaze and view, often interpreted as reflecting "spatial view" or scene-based coding that complements traditional place fields.[11,18,60] Conversely, others have argued that the human (and, more generally, primate) hippocampus may track position, much like that of rodents.[9,61,62] Our results raise the possibility that hippocampal neurons represent both gaze and navigational position. Nonetheless, this interpretation remains speculative because our task uses an overhead virtual view and differs from traditional place-cell paradigms, including the third-person viewpoint, the absence of vestibular signals, and reduced opportunities for path integration and large-scale route planning.

The results presented here are somewhat surprising in light of recent findings reporting little modulation of hippocampal firing by the spatial location of task-relevant stimuli presented on a screen.[63] However, these two sets of results are less contradictory than they may appear. Our prey-pursuit task, which involves continuous navigation, places greater demands on spatial representation than discrete paradigms, which may drive more consistent and reliable spatial modulation that is not observed in discrete tasks.[44] Moreover, the LN/GLM approach[47] provides a flexible framework for capturing complex tuning properties that simpler regression methods may fail to detect.

There are major ongoing debates about the nature and existence of place cells in primates.[16,17,64,65] Our results have only a tangential relation to these debates. The neurons in our sample have a single feature in common with place cells, namely, selectivity for position. However, they lack many of the defining features of place cells, including small, localized place fields and specificity to real-world navigation. Indeed, place cells have primarily been characterized during real-world movement, whereas our task involved virtual movement viewed from an overhead perspective. Nonetheless, our non-place-cell neurons carry a rich spatial code. Thus, we speculate that navigation may involve complementary place-cell and non-place-cell spatial coding systems.[66] Unlike place cells, this second group of neurons has the ability to combine representations for multiple agents and gaze within a shared (semi-orthogonal) population code that supports generalization. Whether these represent distinct cell classes remains unresolved.[67]

A growing body of work argues that neural computations are best understood as projections of a high-dimensional manifold.[25,31,68] In motor and cognitive systems, such population-level structure has proven powerful for explaining flexibility, generalization, and learning. Some recent manifold-oriented analyses suggest that the hippocampus may obey similar principles.[39,40,43,59] Similar population-level approaches have already yielded important insights into non-navigational hippocampal representations.[7,38,42,69–72] As large-scale recordings become routine, population-geometric frameworks may provide a principled way to unify classic spatial coding phenomena with modern theories of abstraction, generalization, and flexible cognitive maps.

RESOURCE AVAILABILITY

Lead contact

Requests for further information and resources should be directed to, and will be fulfilled by, the lead contact, Benjamin Y. Hayden (benhayden@gmail.com).

Materials availability

This study did not generate new unique reagents.

Data and code availability

De-identified human data have been deposited at Figshare at https://doi.org/10.6084/m9.figshare.33075104 and are publicly available as of the date of publication.

All original code has been deposited at GitHub at https://github.com/assiachericoni/Neural-representations-for-self-and-others-in-the-human-hippocampus and is publicly available as of the date of publication.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

ACKNOWLEDGMENTS

This research was supported by the McNair Foundation. We thank Joshua Adkinson, Justin Fine, Victoria Gates, George Kokalas, Raissa Mathura, and Andrew Watrous for invaluable assistance.

AUTHOR CONTRIBUTIONS

A.C. and B.Y.H. conceptualized the study and designed the experimental procedure and downstream analysis. B.Y.H. and S.A.S. supervised the work. A.C. prepared the figures. A.C., E.A.M., T.S.I., M.F., and A.G.C. collected the data. A.C. and C.D. curated the data. S.A.S., D.L.P., G.P.B., N.G., and M. Hasen performed the intracranial stereotactic surgery. V.K., M. Hegazy, A.M.G., L.L., and G.T.S. provided neurological care for the study participants. A.C. and B.Y.H. wrote the original draft. X.Y., J.H., and J.J. contributed to the interpretation of the findings and their framing within the existing literature. E.B., N.R.P., S.B.M.Y., and S.A.S. contributed to data interpretation and manuscript drafting.

DECLARATION OF INTERESTS

S.A.S. has consulting agreements with Boston Scientific, Zimmer Biomet, Koh Young, Abbott, and Neuropace. S.A.S. is a co-founder of Motif Neurotech.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS
  - Human intracranial neurophysiology
  - Ethics approval
- METHOD DETAILS
  - Electrode visualization
  - The prey-pursuit task
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Linear-nonlinear model (LN/GLM)
  - Spatial similarity index
  - Neural subspaces analysis
  - Cross-correlation generalization performance

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.neuron.2026.08.004.

Received: March 16, 2026
Revised: May 29, 2026
Accepted: August 10, 2026

REFERENCES

1. Ekstrom, A.D., Spiers, H.J., Bohbot, V.D., and Rosenbaum, R.S. (2018). Human Spatial Navigation (Princeton University Press). https://doi.org/10.2307/j.ctvc773wg.
2. Epstein, R.A., Patai, E.Z., Julian, J.B., and Spiers, H.J. (2017). The cognitive map in humans: spatial navigation and beyond. Nat. Neurosci. *20*, 1504–1513. https://doi.org/10.1038/nn.4656.
3. Moser, E.I., Moser, M.-B., and McNaughton, B.L. (2017). Spatial representation in the hippocampal formation: a history. Nat. Neurosci. *20*, 1448–1464. https://doi.org/10.1038/nn.4653.
4. Danjo, T., Toyoizumi, T., and Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. Science *359*, 213–218. https://doi.org/10.1126/science.aao3898.
5. Omer, D.B., Maimon, S.R., Las, L., and Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. Science *359*, 218–224. https://doi.org/10.1126/science.aao3474.
6. Barak, O., Rigotti, M., and Fusi, S. (2013). The Sparseness of Mixed Selectivity Neurons Controls the Generalization–Discrimination Trade-Off. J. Neurosci. *33*, 3844–3856. https://doi.org/10.1523/JNEUROSCI.2753-12.2013.
7. Bernardi, S., Benna, M.K., Rigotti, M., Munuera, J., Fusi, S., and Salzman, C.D. (2020). The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. Cell *183*, 954–967.e21. https://doi.org/10.1016/j.cell.2020.09.031.
8. Behrens, T.E.J., Muller, T.H., Whittington, J.C.R., Mark, S., Baram, A.B., Stachenfeld, K.L., and Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. Neuron *100*, 490–509. https://doi.org/10.1016/j.neuron.2018.10.002.
9. Ekstrom, A.D., Kahana, M.J., Caplan, J.B., Fields, T.A., Isham, E.A., Newman, E.L., and Fried, I. (2003). Cellular networks underlying human spatial navigation. Nature *425*, 184–188. https://doi.org/10.1038/nature01964.
10. Jutras, M.J., Fries, P., and Buffalo, E.A. (2013). Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. Proc. Natl. Acad. Sci. USA *110*, 13144–13149. https://doi.org/10.1073/pnas.1302351110.
11. Killian, N.J., Jutras, M.J., and Buffalo, E.A. (2012). A map of visual space in the primate entorhinal cortex. Nature *491*, 761–764. https://doi.org/10.1038/nature11587.
12. Hoffman, K.L., Dragan, M.C., Leonard, T.K., Micheli, C., Montefusco-Siegmund, R., and Valiante, T.A. (2013). Saccades during visual exploration align hippocampal 3–8 Hz rhythms in human and non-human primates. Front. Syst. Neurosci. *7*, 43. https://doi.org/10.3389/fnsys.2013.00043.
13. Leonard, T.K., Mikkila, J.M., Eskandar, E.N., Gerrard, J.L., Kaping, D., Patel, S.R., Womelsdorf, T., and Hoffman, K.L. (2015). Sharp wave ripples during visual exploration in the primate hippocampus. J. Neurosci. *35*, 14771–14782. https://doi.org/10.1523/JNEUROSCI.0864-15.2015.
14. Leonard, T.K., and Hoffman, K.L. (2017). Sharp-wave ripples in primates are enhanced near remembered visual objects. Curr. Biol. *27*, 257–262. https://doi.org/10.1016/j.cub.2016.11.027.
15. Meister, M.L.R., and Buffalo, E.A. (2018). Neurons in primate entorhinal cortex represent gaze position in multiple spatial reference frames. J. Neurosci. *38*, 2430–2441. https://doi.org/10.1523/JNEUROSCI.2432-17.2018.
16. Mao, D., Avila, E., Caziot, B., Laurens, J., Dickman, J.D., and Angelaki, D.E. (2021). Spatial modulation of hippocampal activity in freely moving macaques. Neuron *109*, 3521–3534.e6. https://doi.org/10.1016/j.neuron.2021.09.032.
17. Martinez-Trujillo, J. (2025). Why do primates have view cells instead of place cells? Trends Cogn. Sci. *29*, 226–229. https://doi.org/10.1016/j.tics.2024.12.007.
18. Nau, M., Navarro Schröder, T., Bellmund, J.L.S., and Doeller, C.F. (2018). Hexadirectional coding of visual space in human entorhinal cortex. Nat. Neurosci. *21*, 188–190. https://doi.org/10.1038/s41593-017-0050-8.
19. Payne, H.L., and Aronov, D. (2025). Remote activation of place codes by gaze in a highly visual animal. Nature *643*, 1037–1043. https://doi.org/10.1038/s41586-025-09101-z.
20. Piza, D.B., Corrigan, B.W., Gulli, R.A., Do Carmo, S., Cuello, A.C., Muller, L., and Martinez-Trujillo, J. (2024). Primacy of vision shapes behavioral strategies and neural substrates of spatial navigation in marmoset hippocampus. Nat. Commun. *15*, 4053. https://doi.org/10.1038/s41467-024-48374-2.
21. Watrous, A.J., and Ekstrom, A.D. (2014). The Spectro-Contextual Encoding and Retrieval Theory of Episodic Memory. Front. Hum. Neurosci. *8*, 75. https://doi.org/10.3389/fnhum.2014.00075.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

22. Wirth, S., Baraduc, P., Planté, A., Pinède, S., and Duhamel, J.-R. (2017). Gaze-informed, task-situated representation of space in primate hippocampus during virtual navigation. PLoS Biol. *15*, e2001045. https://doi.org/10.1371/journal.pbio.2001045.

23. Ebitz, R.B., and Hayden, B.Y. (2021). The population doctrine in cognitive neuroscience. Neuron *109*, 3055–3068. https://doi.org/10.1016/j.neuron.2021.07.011.

24. Langdon, C., Genkin, M., and Engel, T.A. (2023). A unifying perspective on neural manifolds and circuits for cognition. Nat. Rev. Neurosci. *24*, 363–377. https://doi.org/10.1038/s41583-023-00693-x.

25. Perich, M.G., Narain, D., and Gallego, J.A. (2025). A neural manifold view of the brain. Nat. Neurosci. *28*, 1582–1597. https://doi.org/10.1038/s41593-025-02031-z.

26. Vyas, S., Golub, M.D., Sussillo, D., and Shenoy, K.V. (2020). Computation Through Neural Population Dynamics. Annu. Rev. Neurosci. *43*, 249–275. https://doi.org/10.1146/annurev-neuro-092619-094115.

27. Johnston, W.J., Fine, J.M., Yoo, S.B.M., Ebitz, R.B., and Hayden, B.Y. (2024). Semi-orthogonal subspaces for value mediate a binding and generalization trade-off. Nat. Neurosci. *27*, 2218–2230. https://doi.org/10.1038/s41593-024-01758-5.

28. Alleman, M. (2024). Parsing the Geometry of Distributed Representations (Columbia University).

29. Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. Nat. Commun. *7*, 13239. https://doi.org/10.1038/ncomms13239.

30. Fine, J.M., Maisson, D.J.-N., Yoo, S.B.M., Cash-Padgett, T.V., Wang, M.Z., Zimmermann, J., and Hayden, B.Y. (2023). Abstract value encoding in neural populations but not single neurons. J. Neurosci. *43*, 4650–4663. https://doi.org/10.1523/JNEUROSCI.1954-22.2023.

31. Gallego, J.A., Perich, M.G., Miller, L.E., and Solla, S.A. (2017). Neural Manifolds for the Control of Movement. Neuron *94*, 978–984. https://doi.org/10.1016/j.neuron.2017.05.025.

32. Libby, A., and Buschman, T.J. (2021). Rotational dynamics reduce interference between sensory and memory representations. Nat. Neurosci. *24*, 715–726. https://doi.org/10.1038/s41593-021-00821-9.

33. Parthasarathy, A., Herikstad, R., Bong, J.H., Medina, F.S., Libedinsky, C., and Yen, S.-C. (2017). Mixed selectivity morphs population codes in prefrontal cortex. Nat. Neurosci. *20*, 1770–1779. https://doi.org/10.1038/s41593-017-0003-2.

34. Tang, C., Herikstad, R., Parthasarathy, A., Libedinsky, C., and Yen, S.-C. (2020). Minimally dependent activity subspaces for working memory and motor preparation in the lateral prefrontal cortex. eLife *9*, e58154. https://doi.org/10.7554/eLife.58154.

35. Xie, Y., Hu, P., Li, J., Chen, J., Song, W., Wang, X.-J., Yang, T., Dehaene, S., Tang, S., Min, B., et al. (2022). Geometry of sequence working memory in macaque prefrontal cortex. Science *375*, 632–639. https://doi.org/10.1126/science.abm0204.

36. Yoo, S.B.M., and Hayden, B.Y. (2020). The Transition from Evaluation to Selection Involves Neural Subspace Reorganization in Core Reward Regions. Neuron *105*, 712–724.e4. https://doi.org/10.1016/j.neuron.2019.11.013.

37. Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. Nat. Neurosci. *17*, 1500–1509. https://doi.org/10.1038/nn.3776.

38. Esparza, J., Sebastián, E.R., and de la Prida, L.M. (2023). From cell types to population dynamics: Making hippocampal manifolds physiologically interpretable. Curr. Opin. Neurobiol. *83*, 102800. https://doi.org/10.1016/j.conb.2023.102800.

39. Fenton, A.A., Lytton, W.W., Barry, J.M., Lenck-Santini, P.-P., Zinyuk, L.E., Kubík, S., Bureš, J., Poucet, B., Muller, R.U., and Olypher, A.V. (2010). Attention-like modulation of hippocampus place cell discharge. J. Neurosci. *30*, 4613–4625. https://doi.org/10.1523/JNEUROSCI.5576-09.2010.

40. Fenton, A.A. (2024). Remapping revisited: how the hippocampus represents different spaces. Nat. Rev. Neurosci. *25*, 428–448. https://doi.org/10.1038/s41583-024-00817-x.

41. Levy, E.R.J., Carrillo-Segura, S., Park, E.H., Redman, W.T., Hurtado, J.R., Chung, S., and Fenton, A.A. (2023). A manifold neural population code for space in hippocampal coactivity dynamics independent of place fields. Cell Rep. *42*, 113142. https://doi.org/10.1016/j.celrep.2023.113142.

42. Nieh, E.H., Schottdorf, M., Freeman, N.W., Low, R.J., Lewallen, S., Koay, S.A., Pinto, L., Gauthier, J.L., Brody, C.D., and Tank, D.W. (2021). Geometry of abstract learned knowledge in the hippocampus. Nature *595*, 80–84. https://doi.org/10.1038/s41586-021-03652-7.

43. Tang, W., Shin, J.D., and Jadhav, S.P. (2023). Geometric transformation of cognitive maps for generalization across hippocampal-prefrontal circuits. Cell Rep. *42*, 112246. https://doi.org/10.1016/j.celrep.2023.112246.

44. Yoo, S.B.M., Tu, J.C., and Hayden, B.Y. (2021). Multicentric tracking of multiple agents by anterior cingulate cortex during pursuit and evasion. Nat. Commun. *12*, 1985. https://doi.org/10.1038/s41467-021-22195-z.

45. Yoo, S.B.M., Tu, J.C., Piantadosi, S.T., and Hayden, B.Y. (2020). The neural basis of predictive pursuit. Nat. Neurosci. *23*, 252–259. https://doi.org/10.1038/s41593-019-0561-6.

46. Chericoni, A., Fine, J.M., Ismail, T.S., Delgado Salazar, G., Franch, M.C., Mickiewicz, E., Chavez, A.G., Bartoli, E., Paulo, D.L., Krishnan, V., et al. (2026). Neural basis of compositional control. Nature. Published online August 12, 2026. https://doi.org/10.1038/s41586-026-10896-8.

47. Hardcastle, K., Maheswaranathan, N., Ganguli, S., and Giocomo, L.M. (2017). A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. Neuron *94*, 375–387.e7. https://doi.org/10.1016/j.neuron.2017.03.025.

48. Zhang, H., Rich, P.D., Lee, A.K., and Sharpee, T.O. (2023). Hippocampal spatial representations exhibit a hyperbolic geometry that expands with experience. Nat. Neurosci. *26*, 131–139. https://doi.org/10.1038/s41593-022-01212-4.

49. Harland, B., Contreras, M., Souder, M., and Fellous, J.-M. (2021). Dorsal CA1 hippocampal place cells form a multi-scale representation of megaspace. Curr. Biol. *31*, 2178–2190.e6. https://doi.org/10.1016/j.cub.2021.03.003.

50. Varga, V., Petersen, P., Zutshi, I., Huszar, R., Zhang, Y., and Buzsáki, G. (2024). Working memory features are embedded in hippocampal place fields. Cell Rep. *43*, 113807. https://doi.org/10.1016/j.celrep.2024.113807.

51. Koch, J., Demirel, M.C., and Stisen, S. (2018). The SPAtial EFficiency metric (SPAEF): Multiple-component evaluation of spatial patterns for optimization of hydrological models. Geosci. Model Dev. *11*, 1873–1886. https://doi.org/10.5194/gmd-11-1873-2018.

52. Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat. Neurosci. *17*, 1784–1792. https://doi.org/10.1038/nn.3865.

53. Blanchard, T.C., Piantadosi, S.T., and Hayden, B.Y. (2018). Robust mixture modeling reveals category-free selectivity in reward region neuronal ensembles. J. Neurophysiol. *119*, 1305–1318. https://doi.org/10.1152/jn.00808.2017.

54. Posani, L., Wang, S., Muscinelli, S.P., Paninski, L., and Fusi, S. (2026). Rarely categorical, highly separable representations along the cortical hierarchy. Nature. https://doi.org/10.1038/s41586-026-10668-4.

55. Boumal, N., Mishra, B., Absil, P.-A., and Sepulchre, R. (2014). Manopt, a Matlab toolbox for optimization on manifolds. J. Mach. Learn. Res. *15*, 1455–1459.

56. Cunningham, J.P., and Ghahramani, Z. (2015). Linear dimensionality reduction: Survey, insights, and generalizations. J. Mach. Learn. Res. *16*, 2859–2900.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

57. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat. Neurosci. *17*, 440–448. https://doi.org/10.1038/nn.3643.

58. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78–84. https://doi.org/10.1038/nature12742.

59. Nakai, S., Kitanishi, T., and Mizuseki, K. (2024). Distinct manifold encoding of navigational information in the subiculum and hippocampus. Sci. Adv. *10*, eadi4471. https://doi.org/10.1126/sciadv.adi4471.

60. Rolls, E.T. (1999). Spatial view cells and the representation of place in the primate hippocampus. Hippocampus *9*, 467–480. https://doi.org/10.1002/(SICI)1098-1063(1999)9:4<467::AID-HIPO13>3.0.CO;2-F.

61. Kunz, L., Brandt, A., Reinacher, P.C., Staresina, B.P., Reifenstein, E.T., Weidemann, C.T., Herweg, N.A., Patel, A., Tsitsiklis, M., Kempter, R., et al. (2021). A neural code for egocentric spatial maps in the human medial temporal lobe. Neuron *109*, 2781–2796.e10. https://doi.org/10.1016/j.neuron.2021.06.019.

62. Stangl, M., Topalovic, U., Inman, C.S., Hiller, S., Villaroman, D., Aghajan, Z.M., Christov-Moore, L., Hasulak, N.R., Rao, V.R., Halpern, C.H., et al. (2021). Boundary-anchored neural mechanisms of location-encoding for self and others. Nature *589*, 420–425. https://doi.org/10.1038/s41586-020-03073-y.

63. Mackay, S., Reber, T.P., Bausch, M., Boström, J., Elger, C.E., and Mormann, F. (2024). Concept and location neurons in the human brain provide the 'what' and 'where' in memory formation. Nat. Commun. *15*, 7926. https://doi.org/10.1038/s41467-024-52295-5.

64. Robertson, R.G., Rolls, E.T., and Georges-Fran ois, P. (1998). Spatial View Cells in the Primate Hippocampus: Effects of Removal of View Details. J. Neurophysiol. *79*, 1145–1156. https://doi.org/10.1152/jn.1998.79.3.1145.

65. Rueckemann, J.W., and Buffalo, E.A. (2017). Spatial responses, immediate experience, and memory in the monkey hippocampus. Curr. Opin. Behav. Sci. *17*, 155–160. https://doi.org/10.1016/j.cobeha.2017.08.008.

66. Stefanini, F., Kushnir, L., Jimenez, J.C., Jennings, J.H., Woods, N.I., Stuber, G.D., Kheirbek, M.A., Hen, R., and Fusi, S. (2020). A Distributed Neural Code in the Dentate Gyrus and in CA1. Neuron *107*, 703–716.e4. https://doi.org/10.1016/j.neuron.2020.05.022.

67. Luo, X., Mok, R.M., and Love, B.C. (2024). The inevitability and superfluousness of cell types in spatial cognition. eLife *13*, RP99047.

68. Chung, S., and Abbott, L.F. (2021). Neural population geometry: An approach for understanding biological and artificial neural networks. Curr. Opin. Neurobiol. *70*, 137–144. https://doi.org/10.1016/j.conb.2021.10.010.

69. Boyle, L.M., Posani, L., Irfan, S., Siegelbaum, S.A., and Fusi, S. (2024). Tuned geometries of hippocampal representations meet the computational demands of social memory. Neuron *112*, 1358–1371.e9. https://doi.org/10.1016/j.neuron.2024.01.021.

70. Courellis, H.S., Minxha, J., Cardenas, A.R., Kimmel, D.L., Reed, C.M., Valiante, T.A., Salzman, C.D., Mamelak, A.N., Fusi, S., and Rutishauser, U. (2024). Abstract representations emerge in human hippocampal neurons during inference. Nature *632*, 841–849. https://doi.org/10.1038/s41586-024-07799-x.

71. Feather, J., and Chung, S. (2023). Unveiling the benefits of multitasking in disentangled representation formation. Trends Cogn. Sci. *27*, 699–701. https://doi.org/10.1016/j.tics.2023.05.010.

72. Whittington, J.C.R., Muller, T.H., Mark, S., Chen, G., Barry, C., Burgess, N., and Behrens, T.E.J. (2020). The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. Cell *183*, 1249–1263.e23. https://doi.org/10.1016/j.cell.2020.10.024.

73. Chaure, F.J., Rey, H.G., and Quian Quiroga, R. (2018). A novel and fully automatic spike-sorting implementation with variable number of features. J. Neurophysiol. *120*, 1859–1871. https://doi.org/10.1152/jn.00339.2018.

74. Dale, A.M., Fischl, B., and Sereno, M.I. (1999). Cortical Surface-Based Analysis. I. Segmentation and Surface Reconstruction. Neuroimage *9*, 179–194. https://doi.org/10.1006/nimg.1998.0395.

75. Joshi, A., Scheinost, D., Okuda, H., Belhachemi, D., Murphy, I., Staib, L.H., and Papademetris, X. (2011). Unified framework for development, deployment and robust testing of neuroimaging algorithms. Neuroinformatics *9*, 69–84. https://doi.org/10.1007/s12021-010-9092-8.

76. Groppe, D.M., Bickel, S., Dykstra, A.R., Wang, X., Mégevand, P., Mercier, M.R., Lado, F.A., Mehta, A.D., and Honey, C.J. (2017). iELVis: An open source MATLAB toolbox for localizing and visualizing human intracranial electrode data. J. Neurosci. Methods *281*, 40–48. https://doi.org/10.1016/j.jneumeth.2017.01.022.

77. Magnotti, J.F., Wang, Z., and Beauchamp, M.S. (2020). RAVE: Comprehensive open-source software for reproducible analysis and visualization of intracranial EEG data. Neuroimage *223*, 117341. https://doi.org/10.1016/j.neuroimage.2020.117341.

78. Li, X., Morgan, P.S., Ashburner, J., Smith, J., and Rorden, C. (2016). The first step for neuroimaging data analysis: DICOM to NIfTI conversion. J. Neurosci. Methods *264*, 47–56. https://doi.org/10.1016/j.jneumeth.2016.03.001.

79. Jenkinson, M., and Smith, S. (2001). A global optimisation method for robust affine registration of brain images. Med. Image Anal. *5*, 143–156. https://doi.org/10.1016/S1361-8415(01)00036-6.

80. Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage *17*, 825–841. https://doi.org/10.1016/s1053-8119(02)91132-8.

81. Yang, A.I., Wang, X., Doyle, W.K., Halgren, E., Carlson, C., Belcher, T.L., Cash, S.S., Devinsky, O., and Thesen, T. (2012). Localization of dense intracranial electrode arrays using magnetic resonance imaging. Neuroimage *63*, 157–165. https://doi.org/10.1016/j.neuroimage.2012.06.039.

82. Hart, P.E., Nilsson, N.J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. IEEE Trans. Syst. Sci. Cybern. *4*, 100–107. https://doi.org/10.1109/TSSC.1968.300136.

83. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nat. Neurosci. *21*, 1281–1289. https://doi.org/10.1038/s41593-018-0209-y.

84. Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995–999. https://doi.org/10.1038/nature07140.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Analysis code and data repository	This paper	Github: https://github.com/assiachericoni/Neural-representations-for-self-and-others-in-the-human-hippocampus; Figshare: https://doi.org/10.6084/m9.figshare.33075104
Software and algorithms		
MATLAB	MathWorks	https://www.mathworks.com; RRID:SCR_001622
Wave_clus spike sorting	Chaure et al.[73]	https://github.com/csn-le/wave_clus
FreeSurfer version 7.4.1	Dale et al.[74]	https://surfer.nmr.mgh.harvard.edu; RRID:SCR_001847
BioImage Suite version 3.5β1	Joshi et al.[75]	http://bioimagesuiteweb.github.io/webapp
iELVis (electrode visualization toolbox)	Groppe et al.[76]	https://github.com/iELVis/iELVis
RAVE (R Analysis and Visualization of iEEG)	Magnotti et al.[77]	https://rave.wiki; RRID:SCR_019040
Manopt toolbox for optimization on manifolds	Boumal et al.[55]	https://manopt.org; Version 8.0
Other		
Ad-Tech Behnke-Fried sEEG depth electrodes (8 microwires per probe)	Ad-Tech Medical	https://www.adtechmedical.com
Blackrock Microsystems Neuroport system (512-channel, 30 kHz)	Blackrock Microsystems	https://blackrockneurotech.com/products/neuroport/
AlphaOmega microdrive	AlphaOmega	https://www.alphaomega-eng.com
ROSA ONE Brain robotic arm	Zimmer Biomet	https://www.zimmerbiomet.com
Kinevo 900 surgical microscope	Carl Zeiss Meditec	N/A
LabJack U6	LabJack	https://labjack.com
Ethylene oxide sterilization system	Bioseal	N/A

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Human intracranial neurophysiology

Experimental data were recorded from 21 adult participants (7 males and 14 females) undergoing intracranial monitoring for epilepsy. Age ranged from 20 to 68 years old (mean ± SD = 39.9 ± 12.5). Although biological sex was obtained from the clinical record, the study was not designed or powered to evaluate sex-based differences, and the number of participants was insufficient for meaningful statistical comparisons between sexes. Information on gender identity, race, ethnicity, ancestry, and socioeconomic status was not collected for the purposes of this study. Consequently, the potential influence of these variables on the reported findings could not be evaluated. The hippocampus was not a seizure focus area of any patients included in the study. Single neuron data were recorded from stereotactic (sEEG) probes, specifically AdTech Medical probes in a Behnke-Fried configuration. Each patient had an average of 3 probes terminating in left and right hippocampus. Electrode locations are verified by co-registered pre-operative MRI and post-operative CT scans. Each probe includes 8 microwires, each with 8 contacts, specifically designed for recording single-neuron activity. Single neuron data were recorded using a 512-channel Blackrock Microsystems Neuroport system sampled at 30 kHz. To identify single neuron action potentials, the raw traces were spike-sorted using the WaveClus sorting algorithm[73] and then manually evaluated. Noise was removed and each signal was classified as multi or single unit using several criteria: consistent spike waveforms, waveform shape (slope, amplitude, trough-to-peak), and exponentially decaying ISI histogram with no ISI shorter than the refractory period (1 ms). The analyses here used only single unit activity.

Ethics approval

The experiment's procedures were conducted by the standards set by the Declaration of Helsinki and approved by the Institutional Review Board for Baylor College of Medicine, Houston, TX (H-18112). Participants provided written informed

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

consent after the experimental procedure had been fully explained and were reminded of their right to withdraw at any time during the study.

METHOD DETAILS

Electrode visualization

Electrodes were localized using the software pipeline intracranial Electrode Visualization (iELVis[76]) and plotted across patients on an average brain using Reproducible Analysis & Visualization of iEEG (RAVE[77]). For each patient, DICOM images of the preoperative T1 anatomical MRI and the postoperative Stealth CT scans were acquired and converted to NIfTI format.[78] The CT was aligned to MRI space using FSL.[79,80] The resulting coregistered CT was loaded into BioImage Suite[75] (version 3.5β1) and the electrode contacts were manually localized. Electrodes coordinates were converted to patient native space using iELVis MATLAB functions[75,81] and plotted on the Freesurfer[74] (version 7.4.1) reconstructed brain surface. Microelectrode coordinates are taken from the first (deepest) macro contact on the Ad-Tech Behnke Fried depth electrodes. RAVE was used to transform each patient's brain and electrode coordinates into MNI152 average space.[77] The average coordinates were plotted together on a glass brain with the hippocampus segmentation and colored by patient.

The prey-pursuit task

Task description

The task used here is similar to one we have previously used in macaques[44,45] and the same as the one we have used in humans in Chericoni et al.[46] At trial start, two or three shapes appeared on a gray computer monitor placed directly in front of the subject. The yellow circle (15-pixels in diameter) was an avatar that represented the subject and began at the center of the screen. Subject position was determined by the joystick and was limited by the screen boundaries. A square (30 pixels in length) represented the prey(s). Prey movement was determined by a simple algorithm (see below). Each trial ended with either the successful capture of the prey or after 20 s, whichever came first. Successful capture was defined as any spatial tangency between the avatar circle and the prey square. Capture resulted in an immediate numeric reward. Prey movement was generated interactively using an A-star algorithm.[82] Specifically, for every frame (16.67 ms), we computed the cost of 15 possible future positions the prey could move to in the next time-step. These 15 positions were equally spaced on the circumference of a circle centered on the current position of the prey, with a radius equal to the maximum distance the prey could travel within one time-step. The cost was based on two factors: the position in the field and the position of the subject's avatar. The field that the prey moved in had a built-in bias for cost, which made the prey more likely to move toward the center. The cost due to distance from the avatar of the subject was transformed using a sigmoidal function: the cost became zero beyond a certain distance so that the prey did not move, and it became greater as distance from the avatar of the subject decreased. From these 15 positional costs, the position with the lowest cost was selected for the next movement. If the next movement was beyond the screen range (1,800 × 1,000 resolution), then the position with the next lowest cost was selected until it was within the screen range. The maximum subject speed was 23 pixels per frame (each frame = 16.67 ms).

To ensure sufficient time of pursuit, the minimum distance between the initial position of each subject avatar and prey was 400 pixels. The maximum and minimum speeds of the prey remained the same across subjects. In trials with predators, a predator (triangle shape) appeared on 50% of trials. Capture by the predator led to points loss. Predators came in five different types (indicated by color) indicating different levels of points loss, ranging from 1 to 5 points. The algorithm of the predator is to minimize the distance between itself and player. Unlike the prey, the predator algorithm is governed by this single rule. The design of the task reflects primarily the desire to have a rich and variegated virtual world with opportunities for choices at multiple levels that is neither trivially simple nor overly complex.

Experimental apparatus

Patients played at least 100 trials (average 119 trials) of the prey-pursuit task using a joystick (Figures 1A and 1B). The joystick was a commercially available joystick with a built-in potentiometer (Logitech Extreme Pro 3D). The joystick position was read out by a custom-coded program in Matlab running on the stimulus-control computer. The joystick was controlled by an algorithm that detected the positional change of the joystick and limited the maximum pixel movement to within 23 pixels in 16.67 ms. Task events were synchronized to the neural recording system via comments, sent through analog port, from the computer playing the task to the Neural Signal Processor (NSP) at 30 kHz.

Eye tracking

A subset of six patients performed the task while wearing a binocular, head-mounted eye-tracking system (ISCAN, Inc.) that recorded the patient's field of view (scene camera) at 60 Hz and eye movements via two eye cameras. Before the experiment, we performed a point-of-regard calibration procedure to map the patients' eye positions to the matrix of the head-mounted scene camera (640 × 480 pixel space). The calibration was performed on the same monitor that the patients viewed during the experiment. Patients were instructed to view five calibration points on the computer monitor, consisting of a center point and four outer points. The eye-calibration software synchronized the monocular eye-movement data with the scene camera, superimposing a crosshair indicating the center of gaze onto the scene video frames in real-time. If the participant looked outside the field of view of the scene camera frames, gaze position could not be estimated, and no crosshair was displayed on the corresponding frame.

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

We recorded scene camera videos using a CORENTSC DVR (I.O. Industries). This DVR recorded all videos at 60 frames per second and sent pulses to the Blackrock Cerebus neural recording system (neural signal processor; NSP) for every captured frame, as well as the start and end of video recording, thus synchronizing video data to neural activity by providing a timestamp (in neural time) for every frame. Eye position data was synchronized to neural activity via serial cable going directly from the eye tracking computer to the NSP.

To reconstruct gaze trajectories within the virtual environment, we utilized DeepLabCut[83] to train a neural network to automatically label relevant objects in the frames. These included the crosshair (patient's point of gaze) as well as eight landmarks defining the monitor geometry: the four screen corners and the midpoints of the top, bottom, left, and right edges. The DeepLabCut output included the location coordinates for these designated landmarks and crosshair found in each frame. Using these landmarks, a projective geometric transformation (homography) was computed for each video frame via MATLAB's *fitgeotform2d* function. The resulting transformation matrix was applied to the crosshair coordinates to map the gaze position from the head-mounted camera's frame of reference to the task's monitor coordinate system, allowing us to determine gaze position throughout the experiment and identify moments when the gaze was following either prey (Figures 4A and 4B).

QUANTIFICATION AND STATISTICAL ANALYSIS

Linear-nonlinear model (LN/GLM)

To test the selectivity of neurons for spatial experimental variables, we constructed GLMs with navigational variables.[47,84] The GLMs estimated the spike rate of one neuron during time bin t as an exponential function of the weighted sum of the relevant value of each variable at time t, for which the weights are determined by a set of coefficients w_i. The estimated firing rates from the GLMs can be expressed as follows:

$$r = e^{\left(\sum X_i^T w_i\right)/dt}$$

Here, r denotes the predicted firing rate across T time points for a single neuron, and the variables of interest were the 2D positions of the self, both prey, the predator, and gaze, each expressed as an XY coordinate on the screen. This formulation allowed us to model neural activity continuously without time-locking to specific events. Each position variable was discretized into a 6×6 grid (36 bins), and each grid was represented using one-hot encoded columns in the design matrix X. Thus, each neuron was assigned a design matrix whose rows corresponded to the number of samples (concatenated across successful trials) and whose columns corresponded to the total number of spatial bins across variables (e.g., predictors being self + prey position, then X would have 36×2 columns). Our spatial binning (6x6 grid) may appear somewhat coarse, but it represents the finest practical scale usable given the two-dimensional structure of the space and the amount of data we have. Indeed, spatial bins span ∼5 cm of screen space. Ultimately, this choice reflects a tradeoff between spatial resolution and occupancy sampling. Using a finer spatial grid would substantially reduce occupancy per bin and increase noise in firing-rate estimates, thereby reducing the reliability of the inferred spatial maps. We therefore selected a binning resolution that balanced spatial specificity with robust estimation of tuning structure.

Unlike conventional tuning-curve analysis, GLMs do not assume a parametric shape for tuning. Instead, the weights defining each 2D tuning surface were optimized by maximizing the Poisson log-likelihood of the observed spike train. Because all predictors were 2D grids, we applied separate smoothness regularization to the parameters of each variable, enforcing smoothness across both spatial dimensions and reducing overfitting.

Regularization strengths were selected by 10-fold cross-validation, and optimization used Matlab's fminunc with analytic gradients and Hessians. For each test fold, model performance was quantified by the spike-normalized increase in log-likelihood relative to a null model predicting only the mean firing rate of the training data. To characterize predictive accuracy, we also computed the correlation between predicted and empirical firing rates after Gaussian smoothing.

Forward model selection

Model selection followed a forward stepwise procedure based on cross-validated log-likelihood. We first fit all single-variable models, one for each spatial predictor. The best single-variable model was identified as the one producing the largest spike-normalized increase in log-likelihood relative to a null model that predicted only the mean firing rate. Next, we evaluated models containing two variables, but only those that included the best single variable. A more complex model was retained only if it produced a significant improvement in cross-validated log-likelihood relative to the simpler model (Wilcoxon signed-rank test across folds, one-tailed, $\alpha = 0.05$). This procedure continued in the same manner for models with additional spatial variables: a model was accepted only if it significantly outperformed the best model from the previous step. If adding further variables failed to provide a significant improvement at any stage, the forward search terminated early. After the final model was identified, we also compared its log-likelihood increase against the null model baseline; neurons whose best model failed to exceed the null model significantly were classified as not tuned to any of the spatial variables considered.

Time-shift permutations

To assess whether the observed model fits depended on the temporal relationship between neuronal activity and behavior, we generated surrogate datasets by circularly shifting each neuron's spike train relative to the behavioral variables while preserving the temporal autocorrelation structure of the spikes. For each surrogate dataset, we repeated the model-fitting procedure described

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

above. Model performance was quantified as the increase in log-likelihood relative to a null model predicting a constant firing rate (mean firing rate; ΔLLH = LLH_model − LLH_null). Repeating this procedure 500 times generated a time-shift null distribution of ΔLLH values for each neuron. A neuron was classified as significantly tuned only if the empirical ΔLLH obtained from the unshifted data exceeded the 95th percentile of its neuron-specific null distribution.

Spatial similarity index

To compare the similarity between two spatial representations, we used the spatial efficiency measure (SPAEF) that prior literature suggests being more robust than the 2D spatial correlation.[51] It quantifies the similarity between two maps as follows:

$$SPAEF = 1 - \sqrt{(A-1)^2 + (B-1)^2 + (C-1)^2}$$

where A is the Pearson correlation between two maps, B is the ratio between the coefficients of variation for each map and C is the activity similarity measured by histogram profiles. Values near −1 indicate anticorrelated maps (one tends to be high when the other is low), 0 indicates uncorrelated maps and 1 indicates perfect matching between the two. Although SPAEF is not mathematically constrained to [–1, 1], values outside this range are uncommon and did not occur in our dataset.

To quantify the similarity between self, prey and predator representation for each neuron map, we computed SPAEF separately for each neuron and each pair of agents (see "data preprocessing" for spatial map construction). This produced one SPAEF value per neuron per agent pair. To determine if the SPAEF values were significantly different from zero, we built a shuffled distribution over 1000 permutations. We then used a two-sided test to extract p values, considering significance at $p < 0.05$. To determine the noise ceiling, we performed 1000 within-agent half-splits. That is, for each iteration and each neuron, we randomly split the session concatenated trials into two independent halves, computed the spatial tuning maps separately for each half and each spatial variable (self, prey and predator positions) and then measured the spatial similarity between the two resulting maps. This process estimates the upper bound of explainable variance by accounting for the inherent noise in the data, providing a reference against which observed spatial similarities can be compared.

Neural subspaces analysis

Data preprocessing

We computed a 2-dimensional spatial map for each neuron and each spatial variable (x and y coordinates of self, chosen prey, unchosen prey, predator and gaze position on the screen used to perform the task) and used these maps to construct population activity matrices. Specifically, the screen field was divided into a 6x6 grid (36 bins).

To estimate the spatial map of each neuron, we first concatenated all trials (ended with successful capture of the prey) into a continuous time series of the agent's x and y coordinates and the corresponding spike counts. For each bin, we extracted all time points whose (x, y) position fell within that bin and computed the mean firing rate across those samples. This produced a 6×6 spatial map for each neuron and each spatial variable. Maps were smoothed with a Gaussian filter ($\sigma = 0.5$) and vectorized into 36-element column vectors ("neuron spatial maps"). Stacking these vectors across neurons and participants, for each agent (that is, each spatial variable) we obtained a pseudo-population activity matrix $M_{agent} \in \mathbb{R}^{N\times 36}$ where N is the number of recorded neurons. Because neurons differ widely in mean firing rate and response scale, each row (corresponding to one neuron) was Z-scored across spatial bins, so that each row had zero mean and unit variance. We used these matrices as input for all the subsequent subspace analysis.

Correlation matrices visualization

To visualize the similarity structure across neurons, we computed neuron-neuron correlation matrices for each spatial variable (Figures 3A and 3G). Starting from the population activity matrix (M_{agent}), we extracted the Pearson correlation coefficient between all pairs of neurons, yielding an $N \times N$correlation matrix (C_{agent}) in which each entry reflects the similarity of two neurons' spatial tuning profiles. To compare the correlation structure across agents (self, prey, predator) and gaze, we ordered neurons using a hierarchical clustering algorithm (dendrogram with centroid linkage applied to one reference matrix (for instance, self) and used the resulting neuron order to display all other agents' correlation matrices. This enabled visual comparison of the population correlation structure across agents and gaze. To quantify differences in correlation structure, we vectorized the unique pairwise neuron-neuron correlations (upper triangle without the diagonal) for each agent's matrix and computed the Pearson correlation between these vectors (Figures 3B and 4H).

Agent preference index (API)

To quantify whether single neurons activity was more strongly driven by one agent (or gaze) compared to another, we adapted the epoch-preference index of Elsayed and colleagues[29] to compute an agent preference index. For each neuron i, we first measured its tuning strength to each agent by taking the 95th-5th percentile range of its unnormalized spatial map for that agent ($S_{self}(i)$, $S_{chosen\ prey}(i)$, $S_{unchosen\ prey}(i)$, $S_{predator}(i)$, $S_{gaze}(i)$). Using unnormalized maps preserves absolute response magnitude, which is required for comparing tuning strengths. Because different agents (or gaze) can have different overall tuning magnitudes across the population, we normalized each neuron's tuning strength by the mean tuning strength across all neurons for that agent or gaze ($\overline{S}_{self}, \overline{S}_{chosen\ prey}, \overline{S}_{unchosen\ prey}, \overline{S}_{predator}, \overline{S}_{gaze}$). The API for neuron i was then defined as the difference between these normalized tuning strengths. For instance, the API between self and prey for neuron i would be:

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

$$API_{self-prey} = \frac{S_{self}(i)}{\overline{S}_{self}} - \frac{S_{prey}(i)}{\overline{S}_{prey}}$$

Positive values indicate stronger tuning for self relative to prey, and negative values indicate the opposite. Notably, a bimodal distribution of this index across neurons would indicate the presence of two neural subpopulations. For instance, one preferentially representing self position and another the prey position. We assessed bimodality using the Hartigan dip test (Figures 3B and 4H).

Principal component analysis

To assess whether two or more agents and gaze representations shared a common low-dimensional structure, we performed principal component analysis (PCA). For ease of description, we describe this analysis using the self vs chosen prey comparison as an example, although the same procedure was applied to all other combinations of agents or gaze (Figures 3C, 3D, and 4I). We applied PCA to the *Z*-scored and transposed self population matrix (36 spatial bins x *N* neurons) and extracted the top ten self-PCs. We then projected the chosen prey population matrix onto these self-PCs and quantified the proportion of prey variance each self-PC captured. This revealed how much covariance is shared by self and prey representation.

Cross-condition covariance correlation

To determine whether the covariance structure of two representations reflected genuine shared structure or noise sampling, we compared cross-agent covariance against within-agent covariance, which provides an estimate of the noise ceiling. For ease of description, we refer to the self vs. chosen prey comparison, but the same procedure was applied to all other agent and gaze pairs. The within-agent baseline was generated as follows: for each participant, concatenated trials were randomly divided into two equal splits (500 iterations), and spatial maps for self were computed for each subset. These maps were concatenated across participants for each subset, yielding two independent estimates of the self representation. After z-scoring each subset, we computed the corresponding covariance matrices, and the correlation between these matrices provided the expected covariance similarity due purely to noise.

We then estimated cross-agent (self-prey) correlation using the same trial splits. For each iteration, covariance matrices for self and chosen prey were computed using the identical subset of trials, and their correlation was taken as that iteration's cross-agent similarity. The mean cross-agent correlation over 500 iterations was then compared to the within-self baseline. One-sided *p* values were obtained as the fraction of iterations in which the within-self correlation exceeded the self-prey correlation, allowing us to assess whether the cross-agent similarity was greater than expected from sampling noise.

Subspace overlap analysis

To quantify how much of the neural population structure was shared between two agents, we computed the alignment index.[29] For clarity, we describe the self vs. chosen prey comparison, but the same procedure was applied to all other agent and gaze pairs (Figures 3E and 4J). We took the ten top self-PCs and asked how much of the chosen prey variance fell within this self-defined subspace. The index was computed as follow:

$$AI = \frac{tr\left(D_{self}^{T} C_{prey} D_{self}\right)}{\sum_{i}^{10} \sigma_{prey}(i)}$$

Where C_{prey} is the covariance of the prey population matrix (M_{prey}), D_{self} contains the top ten self PCs and σ_{prey} is the i^{th} singular value of C_{prey}. The numerator quantifies how much prey variance lies in the self subspace; the denominator normalizes this by the maximum prey variance captured by a ten-dimensional subspace (that is, the prey variance captured by the top ten prey-PCs). Thus, the alignment index varies from 0 (subspace orthogonality) to 1 (subspace collinearity).

To determine whether the observed alignment was lower than expected given dimensionality and the overall covariance structure, we estimated the covariance of the combined self and prey activity and sampled 1,000 random ten-dimensional orthonormal bases from this covariance. Alignment indices were recomputed for each random subspace to form a shuffled baseline. One-sided *p* values were defined as the fraction of random subspaces whose alignment index was smaller than the observed value, testing whether the neural subspaces were more orthogonal than expected by chance.

Identifying orthogonal subspaces

To identify two mutually orthogonal subspaces that best capture self and chosen prey (or any other agents or gaze pair) variance, we adapted the orthogonal subspace optimization methods developed by Elsayed et al.[29] Starting from the *Z*-scored population activity matrices for self and prey (M_{self} and M_{prey}) we computed the associated covariances (C_{self} and C_{prey}). We then searched for two orthonormal bases $Q_{self} \in \mathbb{R}^{N \times d_{self}}$ and $Q_{prey} \in \mathbb{R}^{N \times d_{prey}}$ that maximized the sum of normalized variance captured in each subspace:

$$\left[\widehat{Q}_{self}, \widehat{Q}_{prey}\right] = argmax_{[Q_{self},\, Q_{prey}]} \frac{1}{2}\left(\frac{tr\left(Q_{self}^{T} C_{self} Q_{self}\right)}{\sum_{i=1}^{d_{self}} \sigma_{self}(i)} + \frac{tr\left(Q_{prey}^{T} C_{prey} Q_{prey}\right)}{\sum_{i=1}^{d_{prey}} \sigma_{prey}(i)}\right)$$

$$\text{Subject to } Q_{self}^{T} Q_{prey} = 0,\ Q_{self}^{T} Q_{self} = I,\ Q_{prey}^{T} Q_{prey} = I$$

$\sigma_{self}(i)$ and $\sigma_{prey}(i)$ denote the singular values of C_{self} and C_{prey}. The dimensionality *d* was selected based on the proportion of variance retained under the orthogonality constraint. Specifically, we retained the minimum number of dimensions required to explain at

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

least 35% of the variance, consistent with previous studies.[29,36] Because orthogonality constraints limit the variance that can be simultaneously captured across subspaces, the retained variance is expected to be lower than in unconstrained dimensionality-reduction approaches such as standard PCA. Importantly, dimensionality selection was performed independently of downstream decoding analyses. The optimization was performed on the Stiefel manifold using a manifold-optimization toolbox.[55] This method simultaneously identifies the two subspaces, ensures complete orthogonality, and normalizes for differences in variance and dimensionality between agents.[29,56]

Linear transformation across subspaces

To test whether the inferred orthogonal subspaces for each agent were related by a consistent linear transformation, we projected the population activity onto each orthogonalized subspace. For each pair of agents and gaze (described here for self vs chosen prey), this projection yielded low-dimensional representation for self and prey ($X_{self} \in \mathbb{R}^{d_{self} \times C}$ and $X_{prey} \in \mathbb{R}^{d_{prey} \times C}$, with C number of samples (spatial bins conditions). We then fit a linear decoder with least squares:

$$X_{prey} = WX_{self}$$

where $W \in \mathbb{R}^{d_{prey} \times d_{self}}$ contains the regression weights mapping the activity in the self subspace onto the activity of the prey subspace. The quality of fit was quantified as:

$$R^2 = 1 - \frac{\left\|X_{prey} - WX_{self}\right\|_F^2}{\left\|X_{prey}\right\|_F^2}$$

R^2 were computed on the full data and using leave-one-out cross-validation across spatial bins. As a control, we generated a shuffled distribution by independently permuting the entries within each row of X_{self} and refitting the decoder (1000 shuffles). The p value was the fraction of shuffled CV R^2 values greater than or equal to the observed CV R^2 and tested whether the linear mapping between self and prey subspaces exceeded chance.

Cross-correlation generalization performance

Data preprocessing

Because CCGP requires many independent population samples rather than a single trial averaged map per bin, we segmented the continuous neural recordings into multiple 40 s chunks to produce pseudo population firing rate vectors for each spatial bin and each chunk. For each patient, we concatenated spike trains and the x and y coordinates of each spatial variable across trials. We then divided the concatenated data into 40 s non overlapping chunks and computed firing rates for each of the 36 spatial bins within each chunk using a 6×6 grid. This resulted in a population activity matrix (neurons x 36 bins) for each chunk and each spatial variable.

Defining classes for CCGP

Because our goal was to test generalization across agents and gaze, we first had to define two categorical conditions that were comparable across self, both prey, and gaze.[7] Since the x-coordinate is shared across spatial variables and can be discretized consistently across all spatial maps, we used left vs right spatial position on the screen as the decoding variables.

For each 40 s chunk derived population matrix, we identified the 24 (12 per side) spatial bins whose x-coordinates were the most leftward and the most rightward. For each side, among these 12 bins, we then selected the 6 bins with the highest occupancy. Notably, choosing 6 random bins from the left and right set yielded comparable CCGP, indicating that our results were not driven by a specific bin selection rule (data not shown). Each selected bin contributed one population activity vector (neurons × 1), and these vectors were stacked across bins, classes, and chunks. This produced a design matrix with a number of rows equal to 6 (bins) x 2 (left/right side) x number of chunks in a session, and a number of columns equal to the number of neurons ($D \in \mathbb{R}^{(n_{bins} \times n_{classes} \times n_{chunk}) \times n_{neurons}}$). This procedure yielded four matrices, one for self, one for chosen prey, one for unchosen prey, and one for gaze. Each matrix was paired with a binary class label vector indicating whether each population vector belonged to the left or right side of the screen. These matrices and label vectors were the inputs to the CCGP decoding analyses.

CCGP computation

For ease of description, we focus on the self vs. chosen prey generalization, though the same logic was applied to all agents and gaze pairs. The goal was to determine whether a left/right decoder learned from self-related activity could generalize to the population activity of the chosen prey. On each repetition (50 total), the self design matrix D was split into a training set (70%) and a test set (30%), preserving the balance between left and right bins. A linear support vector machine (SVM; MATLAB's fitcsvm with 'KernelFunction', 'linear') was trained on the Z-scored self training data; chosen prey test data were normalized using the mean and standard deviation computed from the self training set only. The trained decoder was then applied, without retraining, to the chosen prey test data (same left/right structure as the self test split). The CCGP for chosen prey was defined as the mean classification accuracy across repetitions.

Shuffle baseline

To determine whether the observed CCGP exceeded chance, we constructed a null distribution in which left/right labels were independently permuted across samples within each context (self, chosen prey, unchosen prey, gaze). For each shuffle (500 total), we

Please cite this article in press as: Chericoni et al., Neural geometry in the human hippocampus enables generalization across spatial position and gaze, Neuron (2026), https://doi.org/10.1016/j.neuron.2026.08.004

reran the full 50 repetition procedure, yielding a distribution of shuffled CCGP values. The *p* value was computed as the proportion of shuffled CCGP values that exceeded (or, for metrics below chance, were lower than) the observed CCGP.

Decoding axes visualization in neuron space

To compare how similarly different agents separated left/right position in neural space, we extracted one linear decoding axis per context (self, chosen prey, unchosen prey and gaze) and measured the angles between them. For each context, we started from the corresponding CCGP design matrix *D* and the corresponding left/right labels, *Z*-scored firing rates across samples (per neuron) and trained a linear SVM (MATLAB fitcsvm, 'KernelFunction','linear') to discriminate left/right. The resulting SVM weight vector (beta coefficients) defines a single decoding axis in neuron space for that context. We repeated this procedure separately for each agent and normalized each axis to unit length. We then quantified the similarity between any two decoders axes by computing the cosine of the angle between their weight vectors and converted these cosines to angles in degrees. These angles report the geometric separation between agents' left/right decision boundaries in the shared neuronal population space.

To visualize how self, chosen prey, unchosen prey and gaze encode left/right position in the shared neuronal population, we computed the cosine similarity between every pair of axes, yielding a 4×4 cosine matrix. We converted these cosines to pairwise distances using $D_{ij} = \sqrt{2(1 - \cos(\theta_{ij}))}$ and obtained a cosine distance matrix. We then applied multidimensional scaling (MDS) to this matrix and plotted the first two dimensions. The resulting 2D embedding provides a visualization of the relative geometry of the four decoding axes while preserving their pairwise angular relationships as closely as possible.