Abstract representations emerge in human hippocampal neurons during inference behavior

Hristos S. Courellis[1,2,*], Juri Mixha[1,2,*], Araceli R. Cardenas[3], Daniel Kimmel[4], Chrystal M. Reed[5], Taufik A. Valiante[3], C. Daniel Salzman[4], Adam N. Mamelak[1], Stefano Fusi[4,#], Ueli Rutishauser[1,2,5,6,#]

[1] Department of Neurosurgery, Cedars-Sinai Medical Center, Los Angeles, CA, USA

[2] Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, CA, USA

[3] Krembil Research Institute and Division of Neurosurgery, University Health Network (UHN), University of Toronto, Toronto, ON, Canada

[4] Mortimer B. Zuckerman Mind Brain Behavior Institute, Department of Neuroscience, Columbia University, New York, NY, USA

[5] Department of Neurology, Cedars-Sinai Medical Center, Los Angeles, CA, USA

[6] Center for Neural Science and Medicine, Department of Biomedical Sciences, Cedars-Sinai Medical Center, Los Angeles, CA, USA

*=Joint first authors, #=joint senior authors.

Abstract:

Humans have the remarkable cognitive capacity to rapidly adapt to changing environments. Central to this capacity is the ability to form high-level, abstract representations that take advantage of regularities in the world to support generalization[1]. However, little is known about how these representations are encoded in populations of neurons, how they emerge through learning, and how they relate to behavior[2,3]. Here we characterized the representational geometry of populations of neurons (single-units) recorded in the hippocampus, amygdala, medial frontal cortex, and ventral temporal cortex of neurosurgical patients who are performing an inferential reasoning task. We find that only the neural representations formed in the hippocampus simultaneously encode multiple task variables in an abstract, or disentangled, format. This representational geometry is uniquely observed after patients learn to perform inference, and consisted of disentangled directly observable and discovered latent task variables. Interestingly, learning to perform inference by trial and error or through verbal instructions led to the formation of hippocampal representations with similar geometric properties. The observed relation between representational format and inference behavior suggests that abstract/disentangled representational geometries are important for complex cognition.

Introduction:

Humans have a remarkable capacity to make inferences about hidden states that describe their environment[3–5] and use this information to adjust their behavior. One core cognitive function that enables us to perform inference is the construction of abstract representations of the environment[5–7]. Abstraction is a process through which relevant shared structure in the environment is compressed and summarized, while superfluous details are discarded or represented so that they do not interfere with the relevant ones[8,9]. This process often leads to the discovery of latent variables that parsimoniously describe the environment. By performing inference on the value of these variables, frequently from partial information, the appropriate actions for a given context can rapidly be deployed[5,10], thereby generalizing from past experience to novel situations. For example, a latent variable specifying being in a left- or right-driving country can be used by a pedestrian to infer which way to look for oncoming traffic when crossing a road, even in the absence of a sensory cue such as traffic moving in that pedestrian's field of view, and when crossing roads they have never before encountered. Through abstraction, the common, underlying structure of the world is represented in a way that facilitates adaptive behavior.

How are latent variables encoded in the human brain? How are existing representations modified to accommodate them? Does their format matter for behavior? At the level of neuronal activity, the answers to such questions have remained elusive. Prior research has implicated the hippocampus in the implementation of abstraction and inference-related computations, both through neuroimaging in humans during tasks that require abstraction and generalization[10–13], and through neurophysiology in rodents and non-human primates engaged in tasks with abstract spatial and non-spatial components[14–18]. To date, relatively few studies have explored the role that the format, or geometry, of task variable representations plays in shaping computation in the human brain at the single neuron level[19,20], and no study has, to our knowledge, reported the emergence and manipulation of this geometry on the short timescales that would be required for rapid learning in humans.

We recorded the activity of populations of neurons in the brains of awake, behaving epilepsy patients to study the emergence of abstract representations. Patients performed a reversal learning task with two latent contexts, each requiring different responses to the same stimuli. We find that as patients learned to perform inference on the latent context, a representation of context emerged in an abstract format in the hippocampus, being disentangled from simultaneously represented stimulus information (i.e. represented in an approximately orthogonal subspace of the neural activity space). The emergence of the abstract context variable was correlated with an individual's ability to rapidly perform inference on the state of the latent variable context and was absent during error trials. Furthermore, we found that this abstract hippocampal context representation could emerge in two ways: by learning through experience and through verbal instructions informing patients about the latent structure of the task.

Results:

Humans perform inference in a context dependent task.

Patients viewed a sequence of images and indicated for each whether they thought that the associated action was a "left" or "right" button press on a button box (Fig. 1A). Subjects discovered from the feedback provided after each response what the correct response is for a given image. There were two possible fixed mappings ("SRO maps", see Fig. 1C) between each of the four stimuli, the associated correct response (Left/Right button press), and monetary reward given for a correct response (25¢ or 5¢). Which of the two fixed mappings should be used depended on which context a given trial was in (Fig. 1B,C, Context 1 or 2). The two contexts alternated every 15-32 trials. Context was a latent variable that had to be inferred by subjects because no information was provided on the screen on which context was presently active or whether it had changed. Critically, the two stimulus-response maps are systematically related: all stimulus-response pairings are inverted between the two contexts (see Fig. 1C). With this design, an individual performing inference can detect a change in latent context after receiving feedback that their response was incorrect in a single trial and immediately update their stimulus-response associations for the remaining stimuli even though they have not yet been encountered in the new context. We refer to the trials in which a given stimulus is encountered for the first time following a covert context switch as inference trials and to the remaining trials as non-inference trials.

Patients (17 total, see Table 1) completed 42 sessions (180-320 trials/session, 10-16 blocks/session) of the task, typically in pairs of two back-to-back sessions on the same recording day (mean = 2.4 sessions per day, min = 2, max = 4, see Table 1). Novel stimuli were used in every session, thus requiring patients to re-learn the SRO maps through trial and error at the start of every session. Of the 42 sessions, 6 were excluded from analysis due to at-chance performance in non-inference trials (Binomial Test, $p > 0.05$). Performance on non-inference trials was well above chance for the remaining 36 sessions (Fig. 1D). Each of the 36 included sessions was classified as either a "Inference Present" (IP) or "Inference Absent" (IA) session depending on whether the patient performed significantly above chance on the first of the three possible inference trials occurring after context switches (Fig. 1E, trials marked 2-4).

Our task is designed such that by performing inference, patients can respond correctly the first time they see an image in a new context following the initial error trial (seen as significantly below chance performance in Fig. 1E at time point 1). This can be achieved by patients flexibly updating the currently active SRO map immediately after encountering an error, thereby allowing them to perform accurately for the remaining three stimuli that had not yet been seen in the new context. We took accuracy on the first of these three opportunities (the first inference trial) following a context switch as the behavioral signature of successful behavioral inference (timepoint 2 in Fig. 1E, Binomial Test, $p < 0.05$). Block-wise estimates of task performance for IA (Fig. 1F) and IP (Fig. 1G) sessions reveal that during IA sessions, patients exhibit poor inference performance after every context switch throughout the task, although the performance at the end of every block is high. In contrast, during IP sessions, inference performance rapidly rose over the first few blocks and remained high throughout the duration of the session. Note that

within a given session, the two latent contexts had identical stimuli, responses, and outcomes; the only difference was which stimulus was associated with which response and outcome. Correspondingly, subject-level accuracies (Fig. S1B) and reaction times (Fig. S1C) for the two contexts (arbitrarily labeled 1 and 2 across sessions) were not significantly different, indicating that there was no systematic performance bias for one of the two contexts.

When first performing the task, patients were told that they needed to learn arbitrary stimulus-response associations that would change over time, but were not informed about the latent contexts and their related structure. Sessions were recorded in back-to-back pairs with verbal instructions (see Methods) detailing the latent structure of the task provided during the inter-session period (mean length of break = 241 s, range 102-524s), which was considerably shorter than the sessions themselves (mean = 1154 s, range 898-1900s). Importantly, the session following the verbal instructions required re-learning the SRO maps for new stimuli. We next considered whether patients could discover the latent task structure before receiving instructions, and if not, whether verbal instructions successfully shaped behavior. Patients were split into three groups: An "instruction successful" (IS) group, which is composed of patients who did not perform inference during pre-instruction sessions and who did perform inference during post-instruction sessions (5 patients, 10 sessions, Fig. S1D); An "instruction unsuccessful" (IU) group, which were patients who did not perform inference during both the pre-and post-instruction sessions (4 patients, 8 sessions, Fig. S1E); and a "session one interference" (S1) group, which were patients who exhibited inference behavior during both pre- and post-instruction sessions (3 patients, 6 sessions, Fig. S1F). Only patients who performed accurately in non-inference trials in both pre-and post instruction sessions were included in one of these three groups (Fig. S1D-F, "last"; 5 patients excluded, see Table 1). For each subject, no sessions after the first two were considered for this analysis. Thus, patients exhibited a variety of inference behaviors. Below, we contrast the neural representations between these groups of patients to examine how instructions shape neural representations on short timescales.

Electrophysiology and analysis approach

Neural data recorded over the 36 (of 42) included sessions yielded 2694 (of 3124) well isolated single-units, henceforth neurons, distributed across the hippocampus (HPC, 494 neurons), amygdala (AMY, 889 neurons), pre-supplementary motor area (preSMA, 269 neurons), dorsal anterior cingulate cortex (dACC, 310 neurons), ventromedial prefrontal cortex (vmPFC, 463 neurons), and ventral temporal cortex (VTC, 269 neurons) (Fig. 1H,M). Only well isolated neurons as assessed by rigorous spike sorting quality metrics were included (see methods). Action potentials discharged by these neurons were counted during two 1s long trial epochs: during the baseline period (base, -1s to 0s prior to stimulus onset), and during the stimulus period (stim, 0.2s to 1.2s after stimulus onset). For the stimulus period, since patients would sometimes respond before 1.2s (reaction time = 1.08 ± 0.04s over sessions), we determined that 75.15% of all spikes occurred before a response was provided across all recorded neurons, indicating that analyses performed with these spike counts predominantly, but not exclusively, reflect pre-decision processing.

Single neuron responses during the two analysis periods were heterogeneous. During the stimulus period, some neurons exhibited selectivity to one or several of the four variables stimulus identity, response, (predicted) outcome, and context (Fig. 1I-K, example PSTHs of stimulus identity, response, and context coding neurons). Other neurons were modulated by combinations of these variables (Fig. 1L, PSTH of a neuron with conjunctive stimulus and context coding). Across all brain areas, 54% of units (1447/2694) were tuned to task variables, with 26% of units (706/2694) exhibiting only interaction effects, 17% (449/2694) exhibiting only main effects, and 11% (292/2694) exhibiting both when fitting a 3-Way ANOVA for Response, Context, and Outcome (Fig. 1N, RCO column, chance = 135/2694 units, factor significance at $p < 0.05$). When neurons were separated into those recorded from IA and IP sessions, 5-15% of neurons in IA and IP sessions were significantly tuned for each of the main and interaction effects of the 3-Way ANOVA, with significant reductions in the proportion of neurons tuned to Outcome, Response x Context, and Response x Context x Outcome ($p_O = 0.0007$, $p_{RxC} = 0.0395$, $p_{RxCxO} = 0.0048$, two sample z-test) in IP sessions compared to IA sessions (Fig. S1M). Similar analyses conducted on a separate 2-Way ANOVA for Stimulus Identity and Context (Fig. 1N, SC column, Fig. S1N, $p_S = 0.0165$, two sample z-test comparing IA with IP sessions), and for Stimulus Identity and Response (Fig. 1N, SR column, Fig. S1O, $p_S = 0.0287$), revealed a significant decrease in the fraction of neurons tuned to Stimulus Identity in IP compared to IA sessions, again with no significant changes in the proportion of neurons coding for Context, Response, or interactions. These findings indicate diverse tuning to many task variables simultaneously across all brain regions.

Measures of Neural Population Geometry

Given the heterogenous nature of the response pattern at the single neuron level (also see Fig. S1G-L), we adopted a population-level approach. This approach allows us to assess which variables are encoded in the distributed neural patterns, considering also the correlations of the neural responses across multiple conditions. Most importantly, it enables us to examine how these variables are represented, and, in particular, to study the geometry of the neural representation, which has important computational implications. In our task, the geometry of a representation is defined by the arrangement of the eight points in the activity space that represent the experimental conditions. Low dimensional geometries (e.g., when the eight points define a cube) are typically the result of a process of abstraction and confer to a linear readout the ability to generalize. For example, consider a simplified situation with three neurons (the axes) and two stimuli in two contexts, each associated with one of two possible responses (Fig. 2A-C). Imagine that the 4 points (2 per context) are arranged on a relatively low dimensional square (the maximal dimensionality for 4 points is 3), with context encoded along one side and the response along one of the two orthogonal sides (Fig. 2A). Then, a linear decoder for context, trained only on the conditions for which the response is Right, can readily generalize to the Left response conditions (Fig. 2B). This ability to generalize is due to the particular arrangement of the points, which make the coding direction of context for the left responses parallel to the context coding direction for the right responses (Fig. 2C). Also, in this case, context and response are represented in orthogonal subspaces, and hence, they are called disentangled variables[21,22]. In the square example, the ability of a linear decoder to generalize across

conditions (cross condition generalization) also applies to the variable response (i.e. a response decoder trained on context 1 conditions will generalize to context 2 conditions). In previous work[14], we used the cross-condition generalization performance (CCGP) as the defining characteristic of an abstract representation of a variable. Indeed, when context CCGP is high, as in the square example, there is a subspace (the 1D space defined by the context coding direction) in which the representation of context does not depend on the response, or in other words, it is disassociated from the specific instances of the response (Left or Right). In this case we would say that context is represented in an abstract format, or more synthetically that the representation of context is abstract.

Notice that if the 4 points of the example are at random locations in the activity space defining a tetrahedron, the representation is "unstructured" and does not have any of the generalization properties we described. On the other hand, these high dimensional representations allow a linear readout to separate (or shatter) the points in any arbitrary way, and hence confer to the readout the flexibility of implementing any possible task. We refer to the number of ways the points can be separated into two groups by a linear decoder (dichotomies) as shattering dimensionality (SD)[14,23]. Recorded neural representations can have both the generalization properties of the abstract representations, and the flexibility of the high dimensional representations[14].

We characterized and compared the representational geometry between IA and IP sessions for neural pseudopopulations of all recorded neurons in each brain area. For each variable, or dichotomy, we computed the decoding accuracy, which tells us whether the variable is encoded, and the CCGP, which indicates whether the representation of that variable is disentangled from other variables. Both decoding accuracy and CCGP are reported in a cross-validated manner by training and testing decoders on single trials. We complemented this single-trial analysis with a third metric that we call the the parallelism score (PS). PS measures the parallelism of the coding directions of a specific variable. The coding directions are estimated using the average activity for each condition. A high PS strongly indicates that the variable is represented in an abstract format. The PS is also a more direct geometrical measure that focuses on the structure of the representation (the CCGP depends on the geometry but also on the noise and its shape).

To perform the analysis in an unbiased manner, we did not consider only the variables defining the task, but all the variables that correspond to the 35 possible dichotomies of the 8 conditions (see Fig. S2 for illustration of five dichotomies that correspond to specific task variables, Table 2; we refer to these as named dichotomies).

<u>Hippocampal neural geometry correlates with inference behavior</u>.

We first examined the decodability of each balanced dichotomy in different brain areas for sessions where inference was present (IP) and sessions where inference was absent (IA). Following stimulus onset, the hippocampal neural population exhibited a significant increase in average decodability across all balanced dichotomies in IP sessions relative to IA sessions, indicative of an increase in shattering dimensionality (Fig. 2D, IA vs IP $p_{RS} = 2.7x10^{-3}$,

Ranksum over dichotomies). Latent context (Fig. 2D, red, IA vs. IP $p_{RS} = 2.9x10^{-27},\ p_{IA} = 0.12\ , p_{IP} = 5.1x10^{-5}$; p_{IA} and p_{IP} are significance tests vs. chance and p_{RS} a pairwise comparison between IA and IP) and stim pair (Fig. 2D, purple, IA vs. IP, $p_{RS} = 5.0x10^{-27}$, $p_{IA} = 0.015, p_{IP} = 7.9x10^{-7}$) emerged as the most strongly decodable named dichotomies in the IP sessions. The stim pair dichotomy corresponds to the grouping of stimulus identities that elicit the same response (D&B vs. A&C), a relationship that remains the same across both contexts (Fig. 1C, S2). This difference in representation between IA and IP sessions is unique to HPC, as no other recorded region exhibited a significant change in decodability of the variable context (Fig. 2F, red). Rather, in vmPFC stim pair (Fig. 2F, purple, IA vs. IP, $p_{RS} = 6.9x10^{-20}$, $p_{IA} = 0.21, p_{IP} = 0.0097$) and in preSMA response (Fig. 2F, green, IA vs. IP, $p_{RS} = 5.2x10^{-1}\ , p_{IA} = 0.091, p_{IP} = 0.0057$) increased significantly in decodability in IP compared to IA sessions, which is as expected (see discussion). The stim pair variable was also significantly decodable in AMY both during IA and IP, with no significant difference between the two (Fig. 2F, purple, IA vs. IP, $p_{RS} = 0.88, p_{IA} = 0.0016, p_{IP} = 0.0018$).

The expressiveness of the HPC representation can be quantified by the decodability of dichotomies that probe for non-linear interactions of variables in the population. The parity dichotomy (Fig. S2) is only decodable if variables are encoded with a high degree of non-linear interactions in a neural population (see methods). We observed that in the HPC but not in other brain areas, parity decodability increased significantly in IP relative to the IA sessions (Fig. 2D, orange, IA vs. IP, $p_{RS} = 1.5x10^{-21}, p_{IA} = 0.27, p_{IP} = 0.0055$). Generalizing this finding, dividing different dichotomies into increasing levels of "difficulty", with more difficult dichotomies requiring stronger non-linear interactions of task variables, reveals that average shattering dimensionality is highest for the most difficult dichotomies in the HPC (Fig. S7). Together, these findings suggest that non-linearities in the HPC population response in the IP relative to the IA group led to an increase in the number of dichotomies that could be decoded by a linear decoder. Notably, only the hippocampus exhibited significant parity decodability (Fig. 2F, orange, all $p > 0.05$), a significant increase in shattering dimensionality (Fig. 2F), and the emergence of multiple, simultaneously decodable dichotomies between IA and IP sessions.

We next examined the format of the decodable named dichotomies (context, stim pair, parity). During the stim period, CCGP (Fig. 2E, S3D) was significantly elevated for both the context (Fig. 2B, red, IA vs. IP, $p_{RS} = 2.0x10^{-28},\ p_{IA} = 0.51, p_{IP} = 0.02$) and stim pair (Fig. 2E, purple, IA vs. IP, $p_{RS} = 2.0x10^{-28}, p_{IA} = 0.17, p_{IP} = 0.0011$) variables in IP but not in IA sessions. In addition, the PS was significantly larger than expected by chance for the variables context and stim pair in the IP but not IA sessions (Fig. S3E, red, $p_{IA} = 0.55, p_{IP} = 1.4x10^{-15}\ and$ Fig. S3E, purple, $p_{IA} = 0.17, p_{IP} = 1.7x10^{-8}$). CCGP larger than expected by chance was also observed for the stim pair variable in vmPFC during IP and not in IA (Fig. S3A, purple, $p_{IA} = 0.45, p_{IP} = 0.014$), the response variable in preSMA (Fig. S3A, green, $p_{IA} = 0.050, p_{IP} = 0.0010$), and the stim pair variable in AMY during both IA and IP (Fig. S3A, purple, $p_{IA} = 0.050, p_{IP} = 0.039$). Taken together, the increased CCGP and PS values in IP relative to IA sessions indicate that the context and stim pair variables are both simultaneously represented in an abstract format in the HPC in sessions where inference behavior was observed (IP sessions). This means that for example a decoder for context trained only on the conditions in

which the subject responded Left, can readily generalize to all the conditions in which the subject responded Right. These two variables were not represented in an abstract format when inference behavior is absent (IA sessions). While other task variables were also represented in an abstract format in other brain regions, only the hippocampus simultaneously represented these two variables in an abstract format (Fig. S3A,B). These two disentangled variables are thus represented in approximately orthogonal subspaces.

We also conducted a parallel analysis during the pre-stimulus baseline (Fig. 2, inset), analyzing the geometry of persistent representations of the previous trial. We found that context alone was encoded in an abstract format in the HPC only in sessions in which subjects could perform inference (Supplement S.1). This finding indicates that the hippocampal context representation was persistently maintained in an abstract format across trial epochs.

Together, these findings indicate the emergence of an abstract representation of latent context in the HPC that is correlated with patient's ability to perform inference. This context variable is encoded in the hippocampal representation alongside: (i) an abstract, behaviorally-relevant stimulus category variable (stim pair, but see below), and (ii) an increase in parity dichotomy decodability and increased shattering dimensionality from IA to IP sessions. Together, these findings indicate that the ability to perform inference correlates with the hippocampus explicitly encoding a latent context variable (whose state is not directly observable). The resulting representation simultaneously exhibited both features of high dimensionality (parity decodability, increased SD) and low dimensionality (two variables in an abstract format).

<u>Hippocampal representation of context is absent in incorrect trials</u>

We next asked whether the presence of context as an abstract variable was associated with trial-level performance. To examine this question, we compared the decodability and geometry of all dichotomies between correct and incorrect (error) trials in inference present sessions. The first trial of every block was excluded from this analysis due to being necessarily incorrect by design (see Fig. 1E, time point 1). Contrasting correct with incorrect trials, we found that scattering dimensionality was significantly higher in correct trials (Fig. 3A, IP-c vs. IP-i, average SD 0.59 vs 0.54, $p_{RS} = 0.0048$), as was decodability of the parity dichotomy (Fig. 3A, IP-c vs. IP-i, orange, $p_{RS} = 0.029$). Furthermore, the dichotomies context and stimulus pairing were significantly more decodable in correct compared to incorrect trials (Fig. 3A, IP-c vs. IP-i, red, $p_{RS} = 1.2x10^{-20}$, purple, $p_{RS} = 1.0x10^{-9}$). Furthermore, PS for context but not other variables was significantly elevated in correct trials but not in incorrect trials (Fig. 3B, red, IP-c vs. IP-i, $p_{IP-c} = 1.1x10^{-1}$, $p_{IP-i} = 0.083$) in IP sessions. Similarly, the baseline representation of context also showed this effect, being decodable during correct trials but not during incorrect trials (Fig. 3C, red, IP-c vs. IP-i, $p_{RS} = 3.5x10^{-12}, p_{IP-c} = 0.012, p_{IP-i} = 0.47$). Context was present in an abstract format only in correct trials based on the PS for context being significantly larger than chance during correct but not incorrect trials (Fig. 3D, red, IP-c vs. IP-i $p_{IP-c} = 0, p_{IP-i} = 0.94$). The lack of decodability and parallelism of context in the

baseline immediately prior to an incorrect trial indicates that the geometry of the representation at baseline is correlated with correct behavior in the upcoming trial. Together, these findings demonstrate that both the content and format of the hippocampal neural representation is correlated with behavior on a trial-by-trial basis. This effect is present during the stim period in incorrect trials, where shattering dimensionality, context decodability, and context PS significantly decreased. This effect is also present during the baseline period, where reduction in the decodability and parallelism of the context variable is associated with an error on an upcoming trial.

Controls for univariately tuned neurons, seizure-onset zones, and non-inference performance

We performed three sets of control analysis. First, to determine the relationship of our population-level findings to classical (univariate) tuning, we repeated our analyses after removing subsets of neurons from the population. We removed neurons with significant main effects in a (i) a 2x2x2 ANOVA for Response, Context, and Outcome, and (ii) a 4x2 ANOVA for Stimulus Identity and Context (Supplement S.2). In both analyses, context remained significantly decodable and in an abstract format. Also, as expected, in (ii), stimulus pair representations were no longer decodable. These control analyses indicate that the representation of context in the hippocampus we identified did not arise only due to the emergence of classically context tuned neurons. Rather, context was represented by broadly distributed context modulation at the level of the population.

Second, to assess whether our results were influenced by pathology, we repeated our analysis after excluding hippocampal neurons that were located within clinically confirmed seizure onset zones (Supplement S.3). We found no quantitative changes in our results, suggesting that hippocampal pathology did not influence our results.

Lastly, we examined whether our results were sensitive to behavioral accuracy in non-inference trials (Fig. 1E, trial indicated as 'last'). We repeated our analysis in a subset of IA and IP sessions that were chosen such that non-inference trial performance was matched (Supplement S.4). This control analysis revealed no qualitative changes in our results, suggesting that differences in non-inference trial performance cannot explain our results.

Abstraction of stimulus coding across contexts uniquely increases in the hippocampus

Individual HPC neurons in humans prominently encode the identity of visual stimuli[24]. Visually tuned neurons, whose firing rate is strongly modulated by the identity of presented images, are a prominent example of such encoding[25]. We therefore next asked how the variable context, which we show above is encoded in the HPC, interacts with the stimulus identity variable. To examine this question, we conducted geometric analyses over pairs of stimuli (e.g. stimulus A vs B) in the two contexts. This approach is different from the analysis shown so far, in which we considered all 35 possible balanced dichotomies of task conditions. In each balanced dichotomy, different stimulus identities are grouped together (e.g. stimuli A&B vs stimuli C&D), thereby not allowing an examination of the representational geometry for individual stimuli. To examine how stimulus identity encoding relates to context encoding, we instead consider all possible stimulus pairs individually (A vs. B, A vs. C, A vs. D, B vs. C, B vs.

D, and C vs. D) in both contexts. As a comparison, we concurrently performed this analysis on neurons recorded in ventral temporal cortex (VTC). VTC neurons were strongly modulated by stimulus identity (see below)[26,27], but context was not decodable at the level of the balanced dichotomy analysis (basline period: Fig. S8A, red; compare with Fig. 2G; stimulus period: Fig. S8B, red, compare with Fig. 2D).

First, we verified that neurons in both areas encoded the identity of the four stimuli presented. This was the case in both HPC and VTC: 109/494 (22%) of neurons in HPC (Fig. 4A shows an example) and 195/269 (73%) of neurons in VTC (Fig. 4F shows an example) were significantly modulated by stimulus identity following stimulus onset (1x4 ANOVA, $p < 0.05$). Similarly to HPC, at the population level, VTC neurons encoded stimulus identity-related balanced dichotomies in an abstract format (Fig. S8B-D, purple, brown, pink, $p_{IA/IP} < 10^{-10}$). Furthermore, error trial analysis revealed that stimulus-related dichotomies were also decodable during errors in VTC (Fig. S8E, purple, brown, pink, $p_{IP-i} < 10^{-10}$). This finding contrasts with HPC (compare to Fig. 3A, stim pair AC vs. BD dichotomoy) and is consistent with the idea that neurons in VTC veridically represented the stimulus viewed on the screen by the patient during both correct and error trials. These analyses demonstrate that VTC and HPC neurons encode stimulus information, allowing us to next examine how context interacted with stimulus identity coding in these two brain areas.

We next conducted a geometric stimulus-pair analysis to study the interaction of stimulus identity and context coding in the same neural population. The stimulus-pair analysis was designed to detect the presence of simultaneous abstract coding of stimulus identity across contexts (see Fig. S9A-C for illustration) and abstract coding of context across stimuli (see Fig. S9D-F for illustration).

The average stimulus decoding accuracy across all stimulus pairs in the HPC did not differ significantly between IA and IP sessions (0.73 vs. 0.76; Fig. S10C, $p_{RS} = 0.13$, RankSum over stimulus pairs), indicating that the decodability of stimulus information was not different when patients could perform inference vs. when they could not. In contrast, the geometry of the stimulus representation became disentangled from context: both the stimulus CCGP (Fig. 4B, $p_{RS} = 0.041$) and stimulus PS (Fig. 4C, $p_{RS} = 0.040$) were significantly increased in IP compared to IA sessions. This means that a decoder trained to differentiate between stimulus A and B in one context generalized better to the other context in IP compared to IA sessions (and vice-versa for all other possible pairs). This finding suggests that the stimulus responses reorganized with respect to the emerging context variable. Note that context was not decodable in IA sessions as a balanced dichotomy (Fig. 2D, red line). Nevertheless, stimulus decoders did not generalize well between the two contexts in IA sessions. This result indicates that context did modulate stimulus representations in the HPC, but in a way that was entangled with stimulus identity (see below). This effect was specific to the HPC: in VTC, the neural population geometry was unchanged, as indicated by no significant differences in stimulus decodability (Fig. S10D, $p_{RS} = 0.15$), stimulus CCGP (Fig. 4D, $p_{RS} = 0.15$) and stimulus PS (Fig. 4E, $p_{RS} = 0.39$) between IA and IP sessions. In VTC, CCGP was high even in the IA session, indicating that context was not entangled with stimulus identity like in HPC. These analyses demonstrate

that the representational geometry for stimulus identity in HPC becomes significantly more disentangled from context in IP sessions compared to IA sessions in a manner that reflects an abstract format.

Having explored the generalization of stimulus encoding across contexts, we next turned our attention to study the generalization of the context code across stimuli. The presence of abstract coding for one variable (stimulus identity) does not necessarily imply that the other variable is also present in an abstract format, though we do have evidence that this is the case in hippocampus from the CCGP and PS analysis over balanced dichotomies (Fig. 2B, S4E). Context decoding analysis conducted over stimuli (e.g. considering only trials with stimuli A&B shown, decode context) in HPC revealed that context was decodable for many stimuli both during IA and IP sessions, without a significant difference between the two (0.63 vs. 0.67; Fig. S11A, $p_{RS} = 0.065$). However, despite being decodable, context encoding during IA sessions was not in an abstract format for stimulus pairs as indicated by low context CCGP (Fig. 4G, $p_{IA} > 0.10$ for all stim pairs) and low context PS (Fig. 4H, $p_{IA} > 0.17$ for all stim pairs except for AB, where $p_{IA} = 0.033$) values that were not significantly greater than chance. In contrast, in IP sessions, both context CCGP (Fig. 4G, $p_{RS} = 0.012$) and context PS (Fig. 4H, $p_{RS} = 0.015$) increased significantly relative to the IA group. This finding indicates that context emerged as an abstract variable at the level of individual stimulus pairs in the HPC.

We next contrast these findings with VTC. While context was decodable from some stimulus pairs during both IA and IP sessions (Fig. S11B, $p_{IA} \in (0.013, 0.074)$, $p_{IP} \in (0.020, 0.081)$ for all stim pairs), there was no significant change in context decodability between IA to IP sessions (Fig. S11B, $p_{RS} = 0.18$). Rather, there was a significant decrease in context CCGP (Fig. S11C, $p_{RS} = 0.026$) and no significant difference in context PS (Fig. S11D, $p_{RS} = 0.39$) from IA to IP. Together, these findings indicate that in the HPC, the context variable in the IP sessions is in an abstract format because context coding directions become aligned (i.e. parallel) across stimuli. For VTC, on the other hand, the lack of an abstract context representation across stimulus identities in both IA and IP sessions suggests that context, though decodable for individual stimulus pairs, is not organized in an abstract format and does not meaningfully correlate with inference behavior.

In summary, these findings indicate that the emergence of context as an abstract variable in the hippocampus when patients can perform inference is coupled with the reorganization of stimulus representations so they are also more disentangled, thereby forming a jointly abstracted code for stimuli and context. This transformation of the representation is visible directly in the data when projecting the neural representations in 3 dimensions using Multidimensional Scaling (Fig. 4I, S12, Supplementary Video 1). This reorganization occurs without the encoding of additional stimulus information since individual stimuli are equally decodable in the presence or absence of inference behavior. This change in geometry relies on the encoding of context in an abstract format and is unique to hippocampus. In contrast, we found no systematic reorganization of stimulus representations in VTC.

<u>The tuning changes of hippocampal neurons that reflect the geometry changes</u>

We next examined what aspects of neuronal activity changed in the hippocampus to give rise to the abstract neural representations that we observed (i.e. the representations with elevated CCGP). We considered the following non-mutually exclusive possibilities (Fig. 5A-D). (i) The distances between conditions in state space could increase (Fig. 5A vs. Fig. 5B), either as a result of increased firing rate of variable-coding neurons, an increase in the fraction of tuned neurons, or an increase in the depth of tuning of these neurons. (ii) The variance of the population response projected along the coding direction could decrease (Fig. 5C). (iii) Parallelism could increase due to increases in the consistency of firing rate modulation in response to one variable over values of another (Fig. 5D).

We first examined whether mean firing rates across all recorded neurons differed between IA and IP sessions in the HPC. The firing rate across conditions decreased from 3.37±0.13 to 1.36±0.03 Hz, a 60% reduction on average during the stimulus period (Fig. 5E, $p_{RS} = 8.3x10^{-5}$). Firing rates were also reduced during the baseline period (3.29±0.09 to 1.38±0.02 Hz, 58% reduction, Fig. S14G). This firing rate reduction was unique to the HPC, with every other recorded region exhibiting no significant differences or increases in firing rate between IA and IP during the stimulus (Fig. S13C) and baseline periods (Fig. S14H).

The firing rate reduction led to a decrease in the average distance between class centroids (separation) across all dichotomies in IP (5.77 ± 0.22 to 4.17 ± 0.07 Hz, $p_{RS} = 2.9x10^{-8}$, Fig. 5G). However, the centroid distance for a single dichotomy, the context dichotomy, increased from IA to IP (4.3 vs 5.0 Hz, $p_{IA} = 0.87, p_{IP} = 0.076$, $p_{\Delta Dist} = 0.040$, Fig. 5G,H,S13G). In fact, context was the dichotomy with the largest change in distance in firing rate space when comparing the IP and IA conditions (Fig. 5H). This isolated significant rise in context separability was not seen in any of the other recorded areas during the stimulus period (Fig. S13A,B). Similarly, during the baseline period, the centroids for context moved together the fewest in the HPC (5.6 vs 5.0 Hz, $p_{IA} = 0.68, p_{IP} = 0.0007$, $p_{\Delta Dis} = 0.027$, Fig. 5K,L) despite the significant decrease in distance over all dichotomies that was also observed here due to the firing rate reduction (5.85±0.08 to 4.25±0.04 Hz, $p_{RS} = 6.5x10^{-13}$, Fig. 5K).

Next, we assessed changes in the variability of the population response along the coding direction of each dichotomy. The variance along the coding direction of neuronal responses in the HPC decreased for all dichotomies in IP when compared to IA sessions during both the stimulus period (2.51 ± 0.16 vs. 1.53 ± 0.06, $p_{RS} = 6.5x10^{-13}$, Fig. 5I) and the baseline period (2.49 ± 0.09 vs. 1.58 ± 0.02, $p_{RS} = 6.5x10^{-13}$, Fig. S14A,B). However, this decrease could be a simple consequence of the reduction in firing rates under the assumption of Poisson statistics. We conducted a condition-wise Fano-factor analysis to assess whether the variance reduction was beyond that expected for a the reduction in firing rates. This analysis revealed no significant differences in Fano factors between IA and IP sessions during the stimulus period (1.39±0.22 vs 1.36±0.14 , $p_{RS} = 0.99$, Fig. 5F) and the baseline period (1.61±0.26 vs 1.45±0.11, $p_{RS} = 0.19$). Together, these two findings suggest that the decrease in variance along dichotomy coding directions is explained by the decreases in firing rate.

Though the increase in distances between dichotomy centroids for context appears to be a distributed, population-level phenomenon (see Fig. S5A-F), we sought to determine if a signature of this increase could be detected in the tuning of individual neurons (Supplement S.5). We found that the proportion of neurons exhibiting univariate context tuning increased from IA to IP sessions, thus partially explaining the increased representational distance. However, the hippocampal population geometry did not exclusively rely on these neurons because after excluding all neurons with significant univariate coding, we found no qualitative change in the population geometry (Fig. S5).

Finally, we examined the tuning of individual neurons to investigate what gave rise to the increases in parallelism for context across stimuli that we observed (see Fig. S11E-H for examples). A stimulus-tuned neuron also modulated by context could do so consistently across all stimuli (e.g. firing rate increased for all stimuli), or inconsistently (e.g. firing rate increased for some stimuli and decreased for others). Neurons consistently modulated by context would increase the context parallelism. Thus, we quantified the consistency in the direction with which stimulus representations were modulated (firing rate increased or decreased) across contexts for stimulus-identity tuned neurons. Context modulation consistency is computed for each neuron, and can take on values between 0 and 4, with 0 indicating no consistency in modulation and 4 indicating all stimuli exhibit the same firing rate modulation direction between contexts (see Methods for details). We find a significant increase in the consistency of context modulation in HPC from IA to IP sessions (Fig. 5J,11I, 1.8±0.2 vs 2.9±0.3, $p_{RS} = 0.0049$). This effect was specific to HPC: in VTC, this metric decrease significantly (Fig. S11I, 2.6±0.3 vs 1.6±0.2, $p_{RS} = 0.0039$). These findings indicate that, for HPC, context parallelism arises in part due to an increase in the consistency with which the firing rates of stimulus-tuned neurons are modulated by context.

The changes in neural state space responses for HPC are summarized in Fig. 5M, and feature aspects of our previous hypotheses: (i) condition averages for context increase in separation despite relaxing towards the origin (decrease in firing rate), (ii) are accompanied by decreases in variance along the coding direction, and (iii) neurons become increasingly consistent (parallel) in their modulation across stimulus and context dimensions. Together, these changes explain the implementation of the context coding dimension in the hippocampal representation and how it is able to emerge as a simultaneous, linearly encoded variable alongside stimulus identity.

<u>Context representations outside of the Hippocampus.</u>

The only other area of the brain that we examined in which we found a representation of latent context that correlated with inference behavior was in the dACC, but only during the baseline but not the stimulus period (Supplement S.6). Interestingly, context emerged in the task representation through a different implementation strategy, namely an increase in firing rates rather than a decrease, as was observed in HPC.

<u>Verbal instruction induces the representation of context in an abstract format.</u>

In all analyses discussed thus far, we compared sessions in which patients performed inference (IP) with those in which patients did not (IA), without regard to how patients transitioned from the IA to the IP group. We provided verbal instructions detailing the latent task structure following the first session (Fig. 6, inset) to all patients, allowing us to examine whether instructions lead to changes in neural representations. As shown above, patients were divided into three types based on behavior: those who exhibited inference behavior in the very first session (S1 group), those who did so after being given verbal instructions (IS group), and those who did not perform inference even after being provided with verbal instructions (IU group). We next compared the neural representation of context between these three groups of patients.

The instructions provided to these three groups were identical, and all patients acknowledged receipt of the instructions. All included patients exhibited high performance on non-inference trials both before and after being given instructions, indicating that they understood the task and learned the SRO maps. The principal difference between the IS (Fig. 6A, S15A,B) and IU (Fig. 6D, S15C,D) groups is therefore their ability to perform inference following the verbal instructions, with both groups performing the task accurately otherwise. The S1 group, on the other hand, exhibited above-chance inference performance during both pre and post-instruction sessions (Fig. 6F, S15E,F).

In the IS group, context was decodable in the HPC during the stimulus period on correct trials in the session following the verbal instructions (Fig. 6B, $p_{pre-c} = 0.17, p_{post-} = 0.016, p_{RS} = 3.1x10^{-19}$). This representation of context was in an abstract format, as indicated by significant increases in both CCGP (Fig. S16A; $p_{pre} = 0.28, p_{post} = 0.047, p_{RS} = 8.4x10^{-1}$) and PS (Fig. S16B; $p_{pre-c} = 0.023, p_{post-} = 1.2x10^{-6}$). Successful performance in the task was associated with context being represented abstractly in HPC, as both the decodability (Fig. 6B, $p_{post-} = 0.99$, post-c vs. post-i, $p_{RS} = 4.3x10^{-2}$) and PS (Fig. S16B, $p_{post-i} = 1.1x10^{-4}$) of context decreased significantly on incorrect trials in the post-instruction session. Context was also encoded in an abstract format during the baseline period in the same performance dependent manner as context in the stimulus period (Fig. S16C-E). At the single neuron level, this effect can be appreciated by an increase in the proportion of neurons that are significantly linearly tuned to context ($p < 0.05$, one-way ANOVA for context) during both the stimulus (8% (6/75 neurons) vs. 18% (17/93 neurons), $p = 0.027$) and baseline (7% (5/75 neurons) vs. 16% (15/93 neurons), $p = 0.029$) periods in the post-instruction session compared to the pre-instruction session (Fig. 6C shows an example). Thus, the ability of IS group patients to perform inference following instructions was associated with the rapid emergence of an abstract context variable in their hippocampus.

In contrast to the IS group, in patients in the IU group, context was not encoded by HPC neurons neither during the stimulus (Fig. 6E, S16F,G, all $p_{pre/post} > 0.05$) nor the baseline (Fig. S16H-J all $p_{pre/post} > 0.05$) periods in the post-instruction session. Furthermore, there was no significant change in tuning to context at the single-neuron level in the hippocampus of patients in the IU group both during the stimulus period (6% pre vs 6% post, $p = 0.41$) and the baseline period (8% pre vs 5% post, $p = 0.27$). These data indicate that receiving verbal instructions describing the latent context alone is insufficient to generate an abstract context

557 representation in the hippocampus. Instead, context was only represented abstractly in the subset
558 of subjects that performed inference following the instructions.

559 For the S1 patient group, context was decodable already during the pre-instruction
560 session (Fig. 6G, $p_{pre} = 0.014$), and dropped slightly below significance ($p_{post} = 0.17$) during
561 the post-session, likely due to the small number of patients and neurons present in this analysis.
562 The CCGP was not significance (Fig. S16K, $p_{pre} = 0.21$, $p_{post} = 0.31$), but the PS in both the
563 pre and post-sessions was highly significant and near the top of the dichotomy rank order in both
564 cases (Fig. S16L, $p_{pre} = 1.5x10^{-9}$, $p_{post} = 1.7x10^{-6}$). This finding suggests that the context
565 variable these patients learned experientially during the pre-instruction session was in an abstract
566 format as assessed by PS but not CCGP, and that receiving the instructions did not significantly
567 alter the behavior or the neural representation of context. A similar trend was observed with the
568 baseline context representation for these patients (Fig. S16M-O). Note that such discrepancies
569 between CCGP and PS are not unexpected since, when less data is available (fewer neurons,
570 decreased firing rates), single trial measures (Decoding, CCGP) become less sensitive to
571 representational structure than measures that operate on condition averages (PS).

572 We also examined firing rate changes of HPC neurons separately for the IS, IU, and S1
573 groups (Fig. S13D). This analysis revealed significant reductions in firing rate across all
574 conditions from pre- to post-sessions for the IS patients alone (-0.39 ± 0.15 Hz, $p_{IS} =$
575 $1.4x10^{-4}$), confirming that reductions in hippocampal firing rate in the same neurons recorded
576 across adjacent sessions were associated with increases in inference performance. HPC neurons
577 from IU patients exhibited an increase in firing rate (0.10 ± 0.04 Hz, $p_{IU} = 1.2x10^{-4}$) and S1
578 firing rates did not significantly change (0.06 ± 0.08 Hz, $p_{S1} = 0.08$).

579 Lastly, we compared the geometry of the context representations formed by each of these
580 patient groups using the PS (balancing number of neurons, see methods). PS for context
581 increased significantly in the IU group, from levels not different from chance during pre-
582 instructions ($p_{IS,pre} = 0.20, p_{IU,pre} = 0.58$, Fig. 6H) to a level comparable to S1 post-
583 instructions ($p_{IS,post} = 0.0028, p_{S1,post} = 0.0035$, Fig. 6H). PS in the S1 group, on the other
584 hand, did not change significantly and was higher than chance already in the pre-instruction
585 group. This finding suggests that HPC neurons in the S1 group carried an abstract representation
586 of context before receiving high-level instructions, and retained that geometry after receiving
587 instructions. Hippocampal neurons in the IS group, on the other hand, did not encode an abstract
588 representation of context before receiving instructions. However, upon receiving instructions,
589 subjects in the IS group could perform inference, and neurons in their HPC started to encode a
590 task representation whose geometry resembled that of the S1 group. This result indicates that a
591 similar representational geometry can be constructed through either experience or instruction.
592 Lastly, subjects in the IU group could not leverage the information provided in the instructions to
593 perform inference, and accordingly, their hippocampi never encoded an abstract representation
594 of context.

595

Discussion

The ability to perform inference in our task was associated with the hippocampus forming an abstract representation of the environment. This representation encoded latent context and stimulus identity in approximately orthogonal subspaces, was behaviorally relevant on the level of individual trials, and emerged with learning. To implement this representation, the context coding directions for different visual stimuli became more parallel, the distance between contexts in neural state space increased, and the overall variance in firing was reduced due to a reduction in mean firing rates. This representation could emerge quickly, with some patients spontaneously learning the latent task structure during their first session and others exhibiting abstract representations within minutes of receiving verbal instructions explaining the task structure despite having no previous experience before the data shown here was recorded. Abstract representations of context and stimulus identity following stimulus onset were only present in the hippocampus and not in the other brain areas we examined. Together, this data reveals that hippocampal population codes can be restructured by learning and verbal instructions within minutes to support inference in a new task.

How can a neural or biological network efficiently encode multiple variables simultaneously[14,28]? One solution is to encode variables in an abstract format so they can be re-used in novel situations to facilitate generalization and compositionality[21,29–33]. Here, we show that in the human brain, such a disentangled representation emerged as a function of learning to perform inference in our task. The format by which latent context and stimulus identity was represented was predictive of the ability to perform behavioral generalization that relies on contextual inference. Crucially, patients performed well on non-inference trials in all sessions included in the analysis, indicating that they understood the task and successfully learned the stimulus-response associations in both contexts. Therefore, the difference between the inference present and absent sessions was only in whether they performed inference following the covert context switch. For those sessions where patients did not perform inference, there was no systematic relationship between context coding vectors across stimuli. For sessions where patients performed inference, there was alignment of the context coding direction across stimuli (making them parallel), indicating that the context variable had been disentangled from the stimulus identity variable in the hippocampi of these patients. As a result, the two variables became disentangled, thereby allowing for generalization.

Inferential reasoning is thought to rely on cognitive maps, which have been observed in the hippocampus and other parts of the brain[15,34–38] . Cognitive maps are thought to underlie inferential reasoning in various complex cognitive and spatial domains[3,10,34,35,39,40]. However, little is known about how maps for cognitive spaces emerge at the cellular level in the human brain as a function of learning. Here, we show that a cognitive map that organizes stimulus identity and latent context in an ordered manner emerges in the hippocampus. The cognitive map emerges because task states in one context, indexed by stimulus identity, become systematically related to the corresponding task states in the other context through a dedicated context coding direction that is disentangled from stimulus identity. Furthermore, the relational codes between task states (stimuli) in each context are preserved across contexts.

However, hippocampal cognitive maps observed in other studies are often different from those that we observed. Indeed, the encoded variables are often observed to non-linearly interact, which is a signature of high dimensional representations. These representations are believed to be the result of a decorrelation of the neural representations (recoding) that is aimed at maximizing memory capacity[41–43]. This form of pre-processing leads to widely observed response properties, like those of place cells[44]. However, there is some evidence of hippocampal neurons that encode one task variable independently of others[17,45–51]. In these studies, no correspondence was shown between different representational geometries in the hippocampus and differences in behavior. Here, the task representations generated when patients cannot perform inference (but can still perform the task) are systematically different from the abstract hippocampal representations of context and stimulus identity that correlate with inference behavior[14]. Finally, it is important to stress that we also observed an increase in the shattering dimensionality (SD). High SD can be compatible with the low dimensionality of disentangled representations[14,51].

We found that there are other brain regions where stimulus identity is encoded, but does not undergo any major reorganization as a function of learning to perform inference. This is the case in the ventral temporal cortex, a region analogous to macaque IT cortex, in which neurons construct a high-level representation of visual stimuli[52–54]. Some studies conducting unit recordings in this general region in humans show that neurons exhibit strong tuning to stimulus identity [27,55,56]. We similarly find that VTC neurons encode visual stimulus identity. However, these responses were not modulated by latent context in a systematic manner. As a result, despite being decodable for individual stimulus pairs, context was not represented in an abstract format. Rather, in VTC, context was only weakly decodable for a subset of the stimuli, context decodability did not change between IA and IP sessions (Fig. S11B,C), and stimulus identity geometry was not reorganized relative to context in IP sessions (Fig. 4D,E). Our study therefore shows that disentangled context-stimulus representations emerged in the hippocampus, but not in the upstream visually responsive region VTC.

The hippocampus was not unique in its representation of context. A weaker representation of context was also found in the dACC, but only during the baseline period. This finding is in line with work implicating the dACC in the representation of task rules and task sets [57–62]. Following stimulus onset, however, dACC did not contain a representation of latent context (Fig. 2). In contrast, previous studies in tasks with explicitly cued context switches[19,20,63,64] find that neurons in the medial frontal cortex (dACC and preSMA) are tuned to task context following stimulus onset. We hypothesize that this might be due to differences in task demands: in our task, context switches were uncued and had to be inferred from outcomes. It remains an open question to examine whether cued vs. inferred context switches engage different mechanisms of switching between contexts and/or different context encoding schemes in the hippocampus.

The focus of our study was to examine how representations of context, stimulus identity, response, and predicted outcome change as a function of learning in the human brain. In a prior study in macaques[14], the representation of the same variables in very well trained animals was examined in HPC, dlPFC, and ACC in a similar task. Several notable differences exist between

the two studies. First, context was encoded in an abstract format at baseline and was decodable after stimulus onset in all three brain areas examined in the macaques. In contrast, in humans, context is only strongly decodable in the HPC. We hypothesize that the wide-spread encoding of context in the macaque study was due to the extensive training the animals received before recordings commenced. In contrast, our patients had no prior task experience. It is possible that early on during learning, latent context representations are present only in the hippocampus, and are propagated to cortex (dACC) with extensive task experience. This hypothesis is supported by prominent direct and indirect projections from the hippocampus to dACC in primates[65–67], and flexible, context-dependent interactions between medial frontal cortical neurons and hippocampal outputs[19,68]. Second, the human hippocampus exhibited abstract stimulus representations, unlike the abstract response or "choice" representation in the macaques. Notably, the abstract stimulus pair and response dichotomies are constructed such that high CCGP for one will necessarily lead to below-chance CCGP for the other, which was indeed the case for both our study (high stim pair, low response) and the primate study (low stim pair, high response). One potential reason for these differences is a species difference: human HPC neurons are strongly modulated by the identity and semantic category of presented images[19,24,69–71], making it natural to organize representations of context relative to this existing representation. Similarly, representations of choices are not prominent in the human HPC[19]. Another potential reason is a difference in task construction: our task employed semantically identifiable images, whereas the prior experiment with macaques used fractals. Third, unlike macaque HPC, human HPC did not encode predicted outcome. We note that in our task, outcome prediction was not necessary to perform the task because context switches were signaled by the accuracy of the response (correct or incorrect), which was independent of predicted outcome received for a correct response. Furthermore, all possible task-states were uniquely indexable using stimulus identity and context, rendering outcome prediction representation unnecessary for unambiguously defining the current task state. Finally, another possibility is that the reward for the macaques (juice volume) was more motivationally salient than the small monetary reward (25¢ or 5¢) patients received. We hypothesize that these reasons obviated the need for a predictive representation of outcome to complete the task in our patients. It remains an important open question whether representations similar to those seen in macaques emerge in the other brain areas we examined following extensive training. Our data indicates that on short experiential timescales, the human hippocampus generates a representation that encodes the minimum set of variables required to solve the task at hand.

In our study, verbal instructions resulted in changes in hippocampal task representations that correlated with behavioral changes. The emergence of this representation in the session immediately following the instructions in the IS group is correlated with their newfound ability to perform inference and suggests that hippocampal representations can be modified on the timescale of minutes through verbal instructions. This change in representation is qualitatively different from the standard neurophysiological approach of studying the emergence of a "learning set", wherein a low-dimensional representation of abstract task structure emerges slowly over days through trial-and-error learning [46,72,73]. Our finding of similar representational structure in the hippocampus in subjects who learned spontaneously and those who only learned after receiving verbal instructions suggests that both ways of learning can potentially lead to the

same solution in terms of the underlying neural representation. In complex, high-dimensional environments, learning abstract representations through trial and error becomes exponentially costly (the curse of dimensionality), and instructions can be used to steer attention towards previously undiscovered latent structure that can be explicitly represented and utilized for behavior. The process of instruction-dependent restructuring of hippocampal representations is likely cortical-dependent, given the role of the cortex in language comprehension, but the exact mechanism by which this process occurs remains to be explored[30,74]. Our findings suggest that when high-level instructions successfully alter behavior, the underlying neural representation is rapidly modified to resemble one learned through experience.

Acknowledgements: We thank R. Adolphs for advice and support throughout all stages of the project and the members of the labs of R. Adolphs, U. Rutishauser, and M. Meister for discussion. We thank all subjects and their families for their participation and the staff and physisicans of the Cedars-Sinai and Toronto Western Epilepsy Monitoring Units for their support.

Funding: This work was supported by the BRAIN Initiative through the NIH Office of the Director (U01NS117839 to U.R.), NIMH (R01MH110831 to U.R.), the Caltech NIMH Conte Center (P50MH094258 to R.A. and U.R.), the Simons Foundation Collaboration on the Global Brain (to S.F., and U.R.), the Moonshot R&D JPMJMS2294 (to R.A.), and by a merit scholarship from the Josephine De Karman Fellowship Trust (to H.S.C.).

Author contributions: Conceptualization: J.M., U.R., C.D.S., S.F., Task design: J.M. and D.L.K., Data collection: J.M., H.S.C, and A.R.C., Data analysis: H.S.C. and J.M., Writing: H.S.C., U.R., and S.F., Clinical care and experiment facilitation: C.M.R. Surgeries: A.N.M. and T.A.V., Supervision: U.R. and S.F.

Competing interests: None.

Data availability statement: Data will be made available on public repositories such as OSF or DANDI upon acceptance, as we commonly do for our publications.

Code availability statement: Example code to reproduce the results will be made available on public repositories such as Github upon acceptance.

754 **Methods:**

755 **Participants:** The study participants were 17 adult patients who were implanted with depth
756 electrodes for seizure monitoring as part of an evaluation for treatment for drug -resistant
757 epilepsy (see Table 1). 14 were monitored at Cedars-Sinai Medical Center (CSMC) and the other
758 3 were monitored at Toronto Western Hospital (TWH). All patients provided informed consent
759 and volunteered to participate in this study. All research protocols were approved by the
760 institutional review boards of CSMC, TWH, and the California Institute of Technology.

761 **Psychophysical Task and Behavior:** Participants performed a serial reversal learning task.
762 There were two possible static stimulus-response-outcome (SRO) maps, each of which was
763 active in one of the two possible contexts. Context was latent and switches between context were
764 uncued. Each recording session consisted of 280-320 trials grouped into 10-16 blocks of variable
765 size (15-32 trials/block) with block transitions corresponding to a change in the latent context.
766 Each trial consisted of a blank baseline screen, stimulus presentation, speeded response from the
767 participant, followed by feedback after a brief delay (Fig. 1A). Responses were either "left" or
768 "right" in every trial. In each session, stimuli were four unique images, each chosen from a
769 different semantic category (human, macaque, fruit, car). If a patient performed multiple
770 sessions, new images not seen before by the patient were chosen for each session. The task was
771 implemented in MATLAB (The Mathworks, Inc., Natick, MA) using PsychToolbox-3[75]. Images
772 were presented on a laptop positioned in front of the patient and subtended approximately 10
773 degrees of visual arc (300 px^2, 1024x768 screen resolution, 15.6 inch (40 cm) monitor, 50 cm
774 viewing distance). Patients provided responses using a binary response box (RB-844, Cedrus
775 Inc.).

776 Receipt of reward in a given trial was contingent on the accuracy of the response
777 provided. In each trial, either a high or low reward (25¢ or 5¢) was given if the response was
778 correct, and no reward (0¢) if incorrect. Whether a given trial resulted in high or low reward if
779 the response was correct was determined by the fixed SRO map (see Fig. 1C). Stimulus-response
780 associations were constructed such that two out of four images (randomly selected) were
781 assigned one response and the other two images were assigned the other (e.g. human and fruit =
782 left, macaque and car = right). Thus, in each context, each stimulus was uniquely specified by a
783 combination of its correct response (left/right) and reward value (high/low). Crucially, the SRO
784 maps of the two possible contexts were constructed so that they were the opposite of each other
785 from the point of view of the associated response (Fig. 1C). To fully orthogonalize also
786 associated reward, half of the reward values stayed the same and the others switched. This
787 structured relationship of stimuli across contexts led to the full orthogonalization of the response,
788 context, and reward variables (Fig. 1B-C). Crucially, the stimulus-response map inversion across
789 contexts provided the opportunity for patients to perform inferential reasoning about the current
790 state of the SRO map, and therefore the latent context.
791 Since rewards were provided deterministically, participants could switch context upon
792 receiving a single error. Therefore, if patients performed inference, they should be able to
793 respond correctly after receiving a single error. The behavioral signature of inferential reasoning
794 was thus the accuracy in the trials that occurred immediately after the first error trial.

Specifically, we took a participant's performance on the first instance of each of the three remaining stimuli in the new context is to measure a participants inference capabilities.

Patients completed multiple sessions of the task, in each of which new stimuli were chosen. After completion of the first session, the experimenter provided a standardized description of the latent contexts and SRO reversal to the patient (see below). These instructions were given regardless of how well the patient performed in the immediately preceding session. After this brief interlude, the participants completed the task again with a novel set of four stimuli.

Instructions given to patients:

---------- Instruction set 1 (before first session) ----------

In this task, we will show you a series of images, 4 of them in total. Your objective is to learn the correct response for each image (either left or right). In the beginning, you will not know what the correct answer is, so take a guess. The correct answer for an image may occasionally change, so pay close attention. For every correct answer you will receive a reward of either 25 or 5 cents. For an incorrect answer you will receive 0 cents. This is real money that you will receive before you leave the hospital in the form of a gift card to your favorite place (ex. Starbucks). You will have the opportunity to take a break halfway through.

---------- Instruction set 2 (before second session) ----------

You may have noticed that some images have the same correct response and some images have the same reward. Even when the correct response changes, they usually change together. In this experiment, we are going to try a different strategy. Pay attention to which images go together (i.e. have the same correct response and similar reward). This should make it a lot easier to perform the task. To make the task a little more difficult, now the correct response for each image will change a little more frequently.

Behavioral Control: We administered a control version of the task identical to the 'first session' described above to n=49 participants recruited on Amazon Mechanical Turk (MTurk). We then used this data to calibrate the difficulty of the task. A majority (~75%) of the control subjects demonstrated proper inference performance, and the remaining 25% demonstrating slow updating of SROs after a context switch, consistent with a behavioral strategy where each stimulus is updated independently (see Fig. S1A).

Electrophysiology: Electrode Placement and Recording: Extracellular electrophysiological recordings were conducted using microwires embedded within hybrid depth-electrodes (AdTech Medical Inc.). The patients we recruited for this study had electrodes implanted in at least the hippocampus, as well as in addition subsets of amygdala, dACC, pre-SMA, vmPFC, and VTC as determined by clinical needs (see Table 1). Implant locations were often bilateral but some patients only had unilateral implants as indicated by clinical needs. Broadband potentials (0.1Hz – 9kHz) were recorded continuously from every microwire at a sampling rate of 32kHz (ATLAS

833 system, Neuralynx Inc.). All patients included in the study had well isolated single neuron(s) in
834 at least one of the brain areas of interest.

835 Electrode Localization: Electrode localization was conducted using a combination of pre-
836 operative MRI and post-operative CT using standard alignment procedures as previously
837 described[19,63]. Electrode locations were co-registered to the to the MNI152-aligned CIT168
838 probabilistic atlas[76] for standardized location reporting and visualization. Placement of electrodes
839 in gray matter was confirmed through visual inspection of subject-specific CT/MRI alignment,
840 and not through visualization on the atlas.

841 **Spike Detection and Sorting:** Raw electric potentials were filtered with a zero-phase lag filter
842 with a 300Hz-3kHz passband. Spikes were detected and sorted using the OSort software
843 package[77]. All spike sorting outcomes were manually inspected and putative single-units were
844 isolated and used in all subsequent analyses. We evaluated the quality of isolated neurons
845 quantitatively using our standard set of metrics[69,78,79] including proportion of inter-spike interval
846 violations < 3ms, signal-to-noise ratio of the waveform, projection distance between pairs of
847 isolated clusters, and isolation distance of each cluster relative to all other detected spikes.

848 **Selection of Neurons, Trials, and Analysis Periods:** Activity of neurons was considered during
849 two epochs throughout each trial: the baseline period (base), defined as -1s to 0s preceding
850 stimulus onset on each trial, and the stimulus period (stim), defined as 0.2s to 1.2s following
851 stimulus onset on each trial. Spikes were counted for every neuron on every trial during each of
852 these two analysis periods. The resulting firing rate vectors were used for all encoding and
853 decoding analyses. Tests of single-neuron selectivity were conducted using N-way ANOVAs
854 with significance at $P < 0.05$, where N was either 2 for models of stim id (A, B, C, D) and
855 context (1, 2), or 3 for models including outcome (High, Low), response (Left, Right), and
856 context (1, 2). All variables were categorical, and all models were fit with all available
857 interaction terms included.

858 **Population analysis – decoding:** Single-trial population decoding analysis was performed on
859 pseudo-populations of neurons assembled across all neurons recorded across all patients. We
860 pooled across sessions within each anatomically specified recording area as described
861 previously[19,20]. Decoding was conducted using support vector machines (SVM) with a linear
862 kernel and L2 regularization as implement in matlab's fitcsvm function. No hyperparameter
863 optimization was performed. All decoding accuracies are reported for decoding accuracy for
864 individual trials. Decoding accuracy is estimated out-of-sample using 5-fold cross-validation
865 unless otherwise specified (e.g. cross-condition generalization). Many of the decoding analyses
866 in this work consist of grouping sets of distinct task conditions into classes, then training an
867 SVM to discriminate between those two groups of conditions. Neurons included in the analysis
868 were required to have at least K correct trials of every unique condition in order to be included in
869 the analysis (K = 15 trials unless otherwise stated). To construct the pseudopopulation, we then
870 randomly sampled K trials from every unique condition and divided those trials into the groups
871 required for the current decoding analysis for every neuron independently. Randomly sampling
872 correct trials in this way allowed us to destroy noise-correlations that might create locally
873 correlated sub-spaces from neurons recorded in the same area and session[19].

To account for the variance in decoding performance that arose from this random sub-sampling procedure, all reported decoding accuracies are the average resulting from 1000 iterations of sub-sampling and decoder evaluation. A similar trial balancing and sub-sampling procedure was conducted for all analyses that report decoding accuracy on incorrect trials, but with K = 1 trial/condition required as incorrect for the neuron to be included in analysis. Various other analyses conducted throughout this work, including representation geometry measures, centroid distances, and coding direction variances, all rely on this procedure of balanced correct and incorrect trial sub-sampling, and averaging over 1000 iterations of the computed metric to study the relationships between task conditions in an unbiased manner. All reported values have been computed with this approach unless otherwise stated.

Construction of Balanced Dichotomies: Our task has 8 possible states (Fig. 1B). We characterized how neurons represented this task space by assessing how a decoder could differentate between all possible "balanced dichotomies" of these 8 task conditions (Fig. 1B). The set of all possible balanced dichotomies is defined by all possible ways by which the 8 unique conditions can be split into two groups containing 4 of the conditions each (e.g. 4 points in context 1 vs 4 points in context 2 is the context dichotomy). There are 35 possible balanced dichotomies (nchoosek(8,4)/2). Some of the possible balanced dichotomies are easily interpretable because they correspond to variables that were manipulated in the task. We refer to these balanced dichotomies as the "named dichotomies", which are: context, response, outcome, stimulus pair (stim pair), and parity. These dichotomies are shown individually in Fig. S2. The stim pair dichotomy corresponds to the grouping of stimuli for which the response is the same in either context (A&C vs. D&B; see Fig. S2). The parity dichotomy is the balanced dichotomy with the maximal non-linear interaction between the task variables (Fig. S2).

Defining decoding difficulty of dichotomies: We quantify the relative degree of non-linear variable interactions needed by a neural population to classify a given dichotomy using a difficulty metric that rates dichotomies that require proximal task conditions to be placed on opposite sides of the decision boundary as more difficult. Note that proximity of task conditions in task space here is defined with respect to the variables that were manipulated to construct the task space. The conditions corresponding to (Response L, Outcome Low, Context 1) and (Response L, Outcome Low, Context 2) are proximal since their task specifications differ by a single variable (hamming distance 1) whereas (Response L, Outcome Low, Context 1) and (Response R, Outcome High, Context 2) are distal since their task specifications differ by all three variables (hamming distance 3). With this perspective, we can systematically grade the degree of non-linearity required to decode a given dichotomy with high accuracy as a function of the number of adjacent task conditions that are on opposite sides of the classification boundary for that dichotomy. For a set of 8 conditions specified by 3 binary variables, this corresponds to the number of adjacent vertices on the cube defined by the variables that are in opposing classes (See Fig. S7A). We define this number as the "difficulty" for a given dichotomy, and can compute it directly for every one of the 35 balanced dichotomies. The smallest realizable dichotomy difficulty is 4, and corresponds only to named dichotomies that align with the axis of one of the three binary variables used to specify the task space. The largest realizable dichotomy is 12, and this corresponds to the parity dichotomy since the dichotomy difficulty (number of

adjacent conditions with opposing class membership) is maximized in this dichotomy by definition. All remaining dichotomies lie between these two extremes in difficulty, and computing average decoding accuracy over dichotomies of increasing difficulty gives a sensitive readout of the degree of non-linear task variable interaction present in a neural population.

Geometric Analysis of Balanced Dichotomies: We used three measures to quantify the geometric structure of the neural representation[14]: shattering dimensionality (SD), cross-condition generalization performance (CCGP), and parallelism score (PS).

SD is defined as the average decoding accuracy across all balanced dichotomies. It is an index of the expressiveness of a representation, as representations with higher SD allow more dichotomies to be decoded. The content of a representation is assessed by considering which balanced dichotomies are individually decodable better than expected by chance.

CCGP assesses the extent to which training a decoder on one set of conditions generalized to decoding a separate set of conditions. Note that to compute CCGP, all trials from a set of conditions are held out from the training data, which is different from the "leave-one-out" type decoding used to estimate SD. The remaining held-in conditions are used to train the decoder, and performance is then evaluated on the held-out conditions (trial-by-trial performance). The CCGP for a given balanced dichotomy is the average over all possible 16 combinations of held-out conditions on either side of the dichotomy boundary. One of the 4 conditions on each side of the dichotomy are used for testing, whereas the remaining three on each side of the dichotomy are used for training. For each of the 16 possible train/test splits, the decoder is trained on all correct trials from the remaining six conditions, and performance is evaluated on the two held-out conditions.

PS assesses how coding directions for one variable are related to each other across values of other variables in a decoder agnostic manner. The PS is defined for every balanced dichotomy as the cosine of the angle between two coding vectors pointing from conditions in one class to conditions in the other for a given dichotomy. These vectors are computed by selecting four conditions (two on either side of the dichotomy), computing the normalized vector difference between the mean population response for each of the two pairs, then computing the cosine between said coding vectors.This procedure is repeated for all possible pairs of coding vectors, and the average over all cosines is reported. Since the correct way of "pairing" conditions on either side of the dichotomy is not known a-priori, we compute the cosine average for all possible configurations of pairing conditions on either side of the dichotomy, then report the PS as the maximum average cosine value over configurations.

Null distribution for geometric measures: We used two approaches to construct null distributions for significance testing of the geometric measures SD, CCGP, and PS.

For the SD and decoding accuracy of individual dichotomies, the null distribution was constructed by shuffling trial labels between the two classes on either side of each dichotomy prior to training and testing the decoder. After shuffling the order of the trial labels, the identical procedures for training and testing were employed. This way of constructing the null distribution destroys the information content of the neural population while preserving single-neuron properties such as mean firing rate and variance.

For the CCGP and PS, we employed a geometric null distribution[14]. Prior to training, we randomly swapped the responses of pairs of neurons within a given condition. For example, for one task condition, all of neuron 1's responses are assigned to neuron 2 and all of neuron 2's responses are assigned to neuron 1, for another task condition, all of neuron 1's responses are assigned to neuron 3, etc…). This way of randomly shuffling entire condition responses leads to the situation where neural population response statistics by-condition are held constant, but the systematic cross-condition relationships that exist for a given neuron are destroyed. This way of shuffling creates a maximally high dimensional representation, thereby establishing a conservative null distribution for the geometric measures CCGP and PS.

Analysis of Incorrect Trials: For determining decoding accuracy for trials in which subjects provided an incorrect response ("error trials"), decoders were trained and evaluated out of sample on all correct trials in IA and IP sessions (denoted as "IA-c" and "IP-c" trials respectively). The accuracy of the decoder was then evaluated on the left out incorrect trials in the IP sessions (denoted as "IP-i"" trials) that were balanced by task condition. Neurons from sessions without at least one incorrect trial for each of the 8 conditions were excluded. We did not estimate CCGP separately for correct and incorrect trials. The PS was estimated using only correct trials for IP-c and IA-c. For IP-i, parallelism was computed using one coding vector (difference between two conditions) from correct trials and one coding vector from incorrect trials. All other aspects of the PS calculation remained as described earlier. The very first trial after a context switch was excluded from analysis (it was incorrect but by design, as the subject cannot know when a context switch occurred).

Stimulus Identity Geometry Analysis (Fig. 4): We repeated the geometric analysis described above for subsets of trials to examine specifically how the two variables context and stimulus interact with each other. To do so, we considered each possible pair of stimuli (AB, AC, AD, BC, BD, CD) separately. For each stimulus pair, we then examine the ability to decode and the structure of the underlying representation for two variables: stimulus identity (see Table 3) and context (see Table 4).

For stimulus identity, what is decoded is whether the stimulus identity is the first or second possible identity in each pair (i.e. "A vs. B" for the AB pair). Stimulus CCGP (Fig. 4B,D) is calculated by training a decoder to decide "A vs. B" in context 1 and testing the decoder in context 2 and vice-versa (the CCGP is the average between these two decoders). Stimulus PS (Fig. 4C,E) is the angle between the two coding vectors "A vs. B" in context 1 and 2.

For context, decoding accuracy is estimated by training two decoders to decide "Context 1 vs. Context 2" for each of the two stimuli in a stimulus pair. The reported decoding accuracy is the average between these two decoders (Fig. S11A,B). For example, for the stimulus pair AB, one such decoder each is trained for all "A" trials and all "B" trials. Context CCGP (Fig. 4G, S11C) is calculated by training a decoder to differentiate between Context 1 and 2 based on the trials in the first identity of the pair, and tested in the second pair and vice-versa. The reported Context CCGP value for a given stimulus pair is the average between the two. Similarly, context PS (Fig. 4H, S11D) is the angle between the two coding vectors Context 1 vs. Context 2 estimated separately for the first and second stimulus in a pair.

Distance/Variance Analysis (Fig. 5): We computed a series of metrics to quantify aspects of the population response that changed between IA and IP sessions. We used (i) the firing rate, (ii) distance in neural state space between classes for balanced dichotomies and stimulus dichotomies (dichotomy distance), (iii) variance of neural spiking projected along the coding directions for those dichotomies (coding direction variance), and (iv) the condition-wise fano factor.

Firing rate (Fig. 5E) was the mean firing rate averaged across all neurons during the stimulus period, reported separately for correct trials of every unique task condition. Values reported during the baseline (Fig. S14G,H) are computed with an identical procedure using firing rates from before 1s prior to stimulus onset.

Dichotomy distance (Fig. 5G,H,K,L) was defined as the Euclidean distance in neural state space between the centroids of the two classes on either side of the decision boundary for that dichotomy. Centroids were computed by constructing the average response vector for each class using a balanced number of correct trials from every condition included in each class through a resampling procedure (described below). Null distributions reported for dichotomy distances are geometric null distributions.

Coding direction variance (Fig. 5I) was computed for a given balanced dichotomy by projecting individual held-out trials onto the coding vector of the decoder trained to differentiate between the two groups of the balanced dichotomy being evaluated. The coding direction was estimated by training a linear decoder on all trials except eight (one from each condition either side of the dichotomy). The vector of weights estimated by the decoder (one for each neuron) was normalized to unit magnitude to estimate the coding vector. The projection of the left out trial onto this coding vector was than calculated using the dot product. This process was repeated 1000 times, generating a distribution of single trial projections onto the coding vector for each dichotomy. The variance of the distribution of 1000 projected data point was then computed and reported as the variance for a given balanced dichotomy (Fig. 5I).

The condition-wise Fano factor (Fig. 5F) was computed separately for each neuron. We used all correct trials for a given balanced dichotomy to estimate the mean firing rate and standard deviation and then took the ratio between the two to calculate the Fano factor for each neuron. Reported fano factors are the average of all fano factors across all neurons from that area/behavioral condition. Fano factors are computed by-condition since grouping trials across conditions could lead to task variable coding (signal) contaminating the fano-factor measurement, which should ideally only reflect trial-by-trial variation around the mean for approximately poisson-distributed firing rates.

The context-modulation consistency (Fig. 5J) was also computed separately for each neuron. Context modulation consistency is the tendency for a neuron's firing rate to shift consistently (increase or decrease) to encode context across stimuli. For each neuron, it was computed by deteriming the sign of the difference (+/-) between the mean firing rate for a given stimulus between the two contexts, and summing the number of stimuli that exhibit the same modulation (either increase or decrease) across the two contexts. This consistency can take on values between 0 (increase in firing rate to encode context for half of the stimuli, decrease in firing rate for the other half) and 4 (either increase or decrease in firing rate for all four stimuli).

Bootstrap Re-sampled Estimation of Measures and Null Distributions: All the measures described in the preceding sections were estimated using a trial and neuron-based re-sampling wmethod. This resampling strategy was used to assure that every measure reported is comparable between a set of conditions by assuring that the same number of neurons and data points are used to train and test classifiers. Metrics were re-computed 1000 times with resampling and all null distributions were computed with 1000 iterations of shuffling and re-computing. Plotted boundaries of null distributions correspond to the 5th and 95th percentiles as sestimated from the 1000 repetitions.

A single iteration of the re-sampling estimation procedure proceeds as follows. For all analyses that involved a comparison of a metric between two behavioral conditions (IA vs IP or pre vs. post instruction), the same number of neurons was included in both conditions by on a region by region basis. For a neuron to be included, at least 15 correct trials for each of the 8 unique task conditions had to exist (120 correct trials total). Across patients, the number of correct trials per condition varied: min = 10.9 ± 1.3 trials/condition, mean = 25.0 ± 0.6 trials/condition, max = 39.6 ± 1.2 trials/condition (mean ± s.e.m.). After identifying the neurons that met this inclusion criteria, an equal number were randomly sampled from both behavioral conditions. The number of considered neurons was set to the number of neurons available in the smallest group.

When constructing feature matrices for decoding, 15 trials were randomly selected from each unique condition that was included in the given analysis. Trial order was shuffled independently for every neuron within condition to destroy potential noise correlations between neurons that were simultaneously recorded. For decoding and SD, out-of-sample accuracy was estimated with 5-fold cross validation. For generalization analyses (CCGP), all trials were used in training since performance is evaluated on entirely held-out conditions. For vector-based measures (dichotomy distance, variance, PS), all trials in relevant conditions were used to compute condition centroids. In the case of variance estimation, all trials except one on either side of the dichotomy boundary were used to learn the coding axis, then the held-out trials were projected onto the coding axis. As previously stated, these procedures were repeated 1000 times with independent random seeds to ensure independent random sampling of neurons and trials across iterations.

Figure Captions:

Figure 1. Task, behavior, recording locations, and single-neuron tuning. (A) The task consisted of variable-length blocks (15-32 trials) that alternated between two latent contexts (red and blue). Context changes (red arrows) were covert. Trials consisted of a pre-stimulus baseline followed by stimulus presentation during which patients executed the associated response (left or right button press) in a speeded manner. After button press, the stimulus was replaced with a fixation cross, followed by the outcome (either high/low reward or incorrect) was presented after a fixed 0.5s delay. **(B-C)** Illustration of the task structure. Each stimulus (A-D) is associated with a single correct response and results in either a high or low reward if the correct response is given. All stimulus-response-reward relationships are inverted between context 1 (blue) and 2 (orange). This visualization is reflective of the disentangled structure of the task variables, and does not necessarily reflect how neurons will organize their responses in neural state-space to each of these conditions. **(C)** Example images (left) associated with the stimuli A-D. Note: stimuli A and B are masked due to copyright. These associations are randomized for every session. **(D)** Patients exhibited high accuracy on non-inference (baseline) trials. Each dot corresponds to the average non-inference trial performance over a single session. Black dashed line indicates chance. Only sessions where patients exhibited above-chance accuracy on non-inference trials are shown (36/42 sessions, $p < 0.05$, Binomial Test on all non-inference trials). **(E)** Task performance split by whether inference was present (IP) or absent (IA) on the first inference trial following context switches in a given session. Sessions where inference performance was significantly above chance (22 sessions, $p < 0.05$, Binomial Test on inference trial 1) were deemed "inference present" (IP, blue), and those where inference performance was not above chance (14 sessions, $p > 0.05$, Binomial Test on inference trial 1) were considered "inference absent" (IA, red). Plot shows performance on the last trial before the context switch, the first trial after the context switch, and for the remaining three inference trials averaged over all trials in each session (mean ± s.e.m. across sessions). Dashed line marks chance. Black box indicates inference trial 1. **(F-G)** Task performance as a function of time in the task for the **(F)** inference absence and **(G)** inference present groups. Shown is the accuracy for the last non-inference trial before a switch (black) and the first inference trial after a switch (gray). Accuracy is shown block-by-block averaged over a 3-block window (mean ± s.e.m. across sessions). **(H)** Electrode locations. Each dot corresponds to a single microwire-bundle. Locations are shown on the same hemisphere (right) for visualization purposes only. Shown are pre-Supplementary Motor Area (preSMA, purple), dorsal Anterior Cingulate Cortex (dACC, blue), ventromedial Prefrontal Cortex (vmPFC, green), Hippocampus (HPC, red), Amygdala (AMY, yellow), and Ventral Temporal Cortex (VTC, teal). **(I-L)** Example neurons tuned to task variables. **(I)** Hippocampal neuron that encodes stimulus identity. Raster trials are reordered based on stimulus identity, and sorted by reaction time therein (black curves). Stimulus onset occurs at time 0. Black points above PSTH indicate times where 1-way ANOVA over the plotted task variables was significant ($p < 0.05$). **(J-L)** PSTH of three other example neurons that encode response **(J)**, context **(K)**, and mixtures of stimulus id and context **(L)**. **(M)** Number of single units recorded across all brain areas (3124 neurons recorded in total). **(N)** Number of single units across all brain areas exhibiting significant Main effects or interaction effects (n-way ANOVA with interactions, $p < 0.05$, see methods) to at least one of the principal task variables (R = Response, C = Context, O = Outcome, S = Stimulus ID) or to combinations of variables. A unit is linearly tuned if it has at least one significant main effect, and non-linearly tuned if it has at least one significant interaction term in the ANOVA model.

Figure 2. Emergence of multiple abstract variables in hippocampus supports inference.

(A) Simplified example of a neural state space where each axis is the firing rate of one neuron. Points correspond to the response of the neurons to different task states, i.e. two stimuli (red and blue) that elicit two responses (R and L) in two contexts. Note the coding vectors for response and context (black arrows) are not aligned with the axes as each neuron might respond to mixtures of variables. The axes of this state space differ from those shown in Fig. 1B, with the latter being defined by experimenter-selected variables rather than neural firing rates.

(B) Example of cross-condition generalization. A decoder is trained to classify context only on response "R" conditions (green) and is evaluated on its ability to decode context on response "L" conditions (purple). If context is represented in an abstract format (i.e. disentangled from response), then the decoder should generalize to the held-out response condition, yielding a high cross-condition generalization performance (CCGP) for context.

(C) Example of context parallelism. Coding vectors for context (gray arrows) are parallel, indicating that the coding direction for context is identical for different responses, and thus that context and response are disentangled. Details for computing the parallelism score (PS) of a balanced dichotomy with 8 conditions are provided in the methods.

(D-F) During the stimulus presentation period, context (red) and stimulus pair (purple) become decodable in the IP group in the HPC. Context is encoded in an abstract format. Decoding accuracy **(D)** and CCGP **(E)** are shown for all 35 balanced dichotomies during the stimulus period (0.2 to 1.2 s following stimulus onset, see inset). A subset of the dichotomies are named (color code) because they represent task specific variables (see Fig. S2). Swarm plots for decoding accuracy and CCGP are light circles and dark circles respectively. Gray bars denote the 5th-95th percentile of the shuffle-null distribution for decoding accuracy and geometric null distribution for CCGP. Stars denote named dichotomies that are above chance in IP and are significantly different from their corresponding IA value ($p_{RS} < 0.05/35$, Ranksum Test, Bonferroni corrected for multiple comparisons across all dichotomies).

(F) Identical analysis to **(D)** showing decodability of balanced dichotomies from neurons recorded in other brain regions (except VTC, which is shown in Fig. 4 and S8).

(G-H) Same as for **(D-F)**, but for spikes counted during the baseline period prior to stimulus onset. Context (red) becomes decodable in IP sessions and is in an abstract format. Trials are labeled according to the current trial. Decoding accuracy **(G)** and CCGP **(H)** computed in HPC for all balanced dichotomies with spikes counted during the pre-stimulus baseline period (-1 to 0s prior to stimulus onset, see inset). All plotting conventions identical to those in **(D-F)**, except Baseline analysis is conducted with task variables from previous trial.

Note: all reported geometric measures, decoding accuracies, or angles are the average of 1000 runs with condition-wise trial resampling as described in the methods. All null distributions are constructed from 1000 iterations of shuffled trial-resampling using either trial-label shuffling (shuffle null) or random rotations designed to destroy low-dimensional structure (geometric null). Also, neuron counts are balanced between IA and IP for every brain area to ensure that dimensionally-sensitive values (e.g. vector angles, decoding accuracies, etc..) are directly comparable. See methods for details.

1163 **Figure 3. Geometry of hippocampal context representations is compromised during error trials.**

1164 **(A-B)** Following stimulus onset, context (red) is not decodable (A) and not in an abstract format (B) in
1165 incorrect trials occurring during IP sessions. Decoding accuracy and parallelism scores are estimated
1166 separately during correct and incorrect trials in IP sessions (IP-c and IP-i, respectively) and in correct
1167 trials only during IA sessions (IA-c). Stars denote named dichotomies that are above chance in the IP-c
1168 trials and are significantly different from their corresponding IA-c value ($p < 0.05/35$, Ranksum Test,
1169 Bonferroni multiple comparison corrected across all dichotomies). **(C-D)** Same as **(A-B)**, but estimated
1170 for spikes counted during the baseline period and task states from previous trial. Note: IP-c values here
1171 are slightly lower than the IP values reported in Fig. 2 because additional neurons were removed by
1172 subsampling to equalize the number of neurons used for both correct and incorrect trials.

1173

Figure 4. Stimulus representations become structured around context with inference in HPC but not VTC.

(A-C) Responses in HPC following stimulus onset carry information about stimulus identity. **(A)** Example neuron in the HPC. Trials are reordered by stimulus identity, and are sorted by reaction time therein (black line).

(B,C) Stimulus geometry across contexts, with geometric analysis conducted over pairs of stimuli in each context. Data points shown correspond to different stimulus pairs (color coded, see lower right for legend). Significance of differences is tested using RankSum comparing IA and IP over all stimulus pairs (* indicates $p < 0.05$, n.s. otherwise). All other conventions identical to those in Fig. 2,3.

(B) CCGP ($p_{RS} = 0.041$) and **(C)** parallelism score ($p_{RS} = 0.040$) for stimulus coding across contexts significantly increased in IP compared to IA.

(D-F) Same as **(A-C)**, but for VTC.

(D) CCGP ($p_{RS} = 0.15$) and **(E)** parallelism score ($p_{RS} = 0.39$) for stimulus coding across contexts does not differ significantly between IA and IP.

(G-H) Context encoding across stimulus pairs for HPC. Plotting conventions are identical to those in panels (B-E). **(G)** CCGP for context across stimuli ($p_{RS} = 0.012$) and **(H)** parallelism score for context coding vectors between pairs of stimuli ($p_{RS} = 0.015$) both significantly increase from IA to IP.

(I) Changes in neural geometry in HPC. MDS of condition-averaged responses of all recorded HPC neurons shown for IA (left) and IP (right). Colored points are average population vector responses to a stimulus (point color) in each context (plane color). Stimuli of the same identity in either context are connected by a line of the same color. Black arrows represent the coding direction between two different stimuli in each context. Comparing left with right, note that context emerges as a decodable variable (red and blue planes separate), and both stimulus coding (black arrows) and context coding (colored lines) are parallel reflecting the abstract format of both variables.

Note: only three stimuli are shown here since the full stimulus ID space is 3D (4 unique, unrelated stimuli) and context is encoded in its own independent dimension, thus requiring 4 dimensions to fully visualize the neural geometry of the hippocampus during inference. All 2D projections of stimulus pairs are shown in Fig. S12.

Figure 5. Firing rate properties underlying the observed changes of population-level hippocampal neural geometry.

(A-D) Illustration of different hypothesized firing rate pattern changes that could give rise to the observed population level geometry changes. The four different hypotheses are illustrated with two stimuli (C and D), each present in two different contexts (blue and red). Condition responses are defined by the firing rates of two hypothetical neurons. The solid-colored points represent the condition average for each stimulus and the larger shaded circles represent the trial-by-trial variation of the two neurons for each context. Gray arrows signify changes that have occurred in IP plots **(B-D)** relative to the IA plot **(A)**.

These response of the neurons to the stimuli during IA **(A)** can be shaped by **(B)** increasing the distance between the context centroids, **(C)** decreasing the variance along the coding direction in the absence of changes in distance, or **(D)** straightening the neural responses without changing distances/variances so that the geometry becomes more orthogonalized.

(E) Changes in hippocampal firing rate from IA to IP. Points correspond to the average firing rate over neurons in the HPC for each of 8 unique task conditions, and are colored according to stimulus identity in that condition (e.g. task condition $C1_L^-$ describes: stimulus C, context 1, outcome $-$, response L). Neuronal firing rates were lower during IP compared to IA sessions ($p_{RS} = 8.3x10^{-5}$, RankSum over conditions).

(F) Same as **(E)**, but for condition-wise fano factors. Fano factor (FF) here is computed as the ratio of the condition-wise variance and the condition-averaged firing rate, computed by neuron and averaged over neurons. Points correspond to average FF over all hippocampal neurons. There was no significant difference (RankSum over conditions between IA and IP ($p_{RS} = 0.99$)).

(G) Population distances between centroids for all 35 balanced dichotomies. The colored connected points represent distances for the named dichotomies indicated in the legend to the right. Gray bars indicate the 5th-95th percentile of the geometric null distribution. Across all dichotomies, distances decreases from IA to IP ($p_{RS} = 2.9x10^{-8}$, RankSum over dichotomies).

(H) Context alone is the only dichotomy whose distance significantly increases from IA to IP **(**red, $p_{\Delta Dist} = 0.040$**)**. The null distribution shown H is the distribution of differences between the IP and IA null distributions shown in.

(I) Average variance projected along the coding direction decreased on average between IA and IP sessions ($p_{RS} = 6.5x10^{-13}$). Variance was computed in a cross-validated manner (see methods) resulting in a distribution of trial-by-trial population activity along each dichotomy coding direction from which the coding variance was computed. The 5th-95th percentile of the geometric null distribution was also used here for the null distribution.

(J) Change in the consistency of context-modulation for stimuli averaged over all neurons in HPC. Greater context modulation consistency for individual neurons results in greater parallelism score for context at the population level. HPC neurons on average exhibit a significant increase in context modulation consistency between IA and IP sessions ($p_{RS} = 0.0039$) during the stimulus period.

(K,L) Same as **(G,H)**, but for spike counts during the baseline period and grouping trials by task state of the previous trial. Distance was significantly reduced across all dichotomies (**K**, $p_{RS} = 6.4x10^{-1}$, RankSum over dichotomies) and context alone exhibits a distance reduction that is smaller than would be expected by chance (**L,** red, $p_{\Delta Dist} = 0.027$)

(M) Illustration of implementational changes to neural state space using the conventions introduced in **(A-D)**. We find that, when comparing IA with IP sessions, that (i) context dichotomy distance increased (indicated by the increased distance between the red and blue shaded circles), (ii) variance decreased due to a reduction in firing rate (indicated by decreased shaded circle radius and movement towards the origin of state space), and (iii) an increase in the consistency of stimulus modulation across contexts (indicated by lines becoming parallel).

Figure 6. Abstract hippocampal representation of context is present following successful verbal instructions about latent context.

(A,D,F) Behavioral performance shown separately for pre-instruction and post-instruction sessions (pre and post). **(A)** Inference performance for the instruction-successful (IS) patient group (5 patients, 10 sessions). **(C)** Example hippocampal neuron with univariate context encoding in the post (bottom) but not the pre instruction (top) session. (one-way ANOVA, $p_{pre} = 0.40, p_{post} = 0.010$). **(D)** Inference performance for the instruction-unsuccessful (IU) patient group (4 patients, 8 sessions). **(F)** Inference performance for the session one inference (S1) patient group (3 patients, 6 sessions).

(B,E,G) Encoding of context in the stimulus period in pre-and post-instruction sessions during pre-instruction correct trials (pre-c), post-instruction correct trials (post-c) and post-instruction incorrect trials (post-i). The first trial following a switch is excluded from this analysis. **(E, inset)** Schematic of the recording procedure, showing a pre- and post-instruction session shaded in gray (30 min duration), with a 4 minute inter-session break (mean duration = 241 s, range 102-524s) during which instructions detailing task structure were provided. **(B)** Context emerges as significantly decodable post-instructions but not pre-instructions in the IS group in a task-performance-dependent manner ($p_{pre-c} = 0.17$, $p_{post} = 0.016, p_{RS} = 3.1x10^{-19}, p_{post-} = 0.99$). **(E)** Context is not significantly decodable in the IU group neither pre- nor post-instructions ($p_{pre} = 0.44$, $p_{post} = 0.42$). **(G)** Context is decodable both pre- and post-instructions in the S1 group ($p_{pre} = 0.014, p_{post} = 0.17$).

(H) Summary of changes in PS for context for all three subject groups (S1 red, IS blue, IU green). Neuron counts are sub-sampled to match across all groups so that PS values are directly comparable. Significant increase in context PS from pre to post is indicated for the IS group ($p_{IS,pre} = 0.20, p_{IS,post} = 0.0028$), but not for the IU ($p_{IU,pre/post} < 0.5$) and S1 ($p_{S1,pre/post} < 0.005$) groups.

Bibliography

1. Tolman, E. C. Cognitive maps in rats and men. *Psychological Review* **55**, 189–208 (1948).

2. Chung, S. & Abbott, L. F. Neural population geometry: An approach for understanding biological and artificial neural networks. *Current Opinion in Neurobiology* **70**, 137–144 (2021).

3. Whittington, J. C. R., McCaffary, D., Bakermans, J. J. W. & Behrens, T. E. J. How to build a cognitive map. *Nat Neurosci* **25**, 1257–1272 (2022).

4. Tenenbaum, J. B., Kemp, C., Griffiths, T. L. & Goodman, N. D. How to Grow a Mind: Statistics, Structure, and Abstraction. *Science* **331**, 1279–1285 (2011).

5. Kemp, C. & Tenenbaum, J. B. Structured statistical models of inductive reasoning. *Psychological Review* **116**, 20–58 (2009).

6. McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S. & Smith, L. B. Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences* **14**, 348–356 (2010).

7. Griffiths, T. L., Chater, N., Kemp, C., Perfors, A. & Tenenbaum, J. B. Probabilistic models of cognition: exploring representations and inductive biases. *Trends in Cognitive Sciences* **14**, 357–364 (2010).

8. Ho, M. K., Abel, D., Griffiths, T. L. & Littman, M. L. The value of abstraction. *Current Opinion in Behavioral Sciences* **29**, 111–116 (2019).

9. Konidaris, G. On the necessity of abstraction. *Current Opinion in Behavioral Sciences* **29**, 1–7 (2019).

10. Vaidya, A. R., Jones, H. M., Castillo, J. & Badre, D. Neural representation of abstract task structure during generalization. *eLife* **10**, e63226 (2021).

11. Gershman, S. J. & Niv, Y. Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology* **20**, 251–256 (2010).

12. Robert, S., Arno, V. & J.d, M. M. Distinct hippocampal and cortical contributions in the representation of hierarchies. *eLife* **12**, (2023).

13. Constantinescu, A. O., O'Reilly, J. X. & Behrens, T. E. J. Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**, 1464–1468 (2016).

14. Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S. & Salzman, C. D. The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell* **183**, 954-967.e21 (2020).

15. Knudsen, E. B. & Wallis, J. D. Hippocampal neurons construct a map of an abstract value space. *Cell* **184**, 4640-4650.e10 (2021).

16. Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature* **543**, 719–722 (2017).

17. Nieh, E. H., Schottdorf, M., Freeman, N. W., Low, R. J., Lewallen, S., Koay, S. A., Pinto, L., Gauthier, J. L., Brody, C. D. & Tank, D. W. Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80–84 (2021).

18. Courellis, H. S., Nummela, S. U., Metke, M., Diehl, G. W., Bussell, R., Cauwenberghs, G. & Miller, C. T. Spatial encoding in primate hippocampus during free navigation. *PLOS Biology* **17**, e3000546 (2019).

19. Minxha, J., Adolphs, R., Fusi, S., Mamelak, A. N. & Rutishauser, U. Flexible recruitment of memory-based choice representations by human medial-frontal cortex. *Science* **368**, eaba3313 (2020).

20. Fu, Z., Beam, D., Chung, J. M., Reed, C. M., Mamelak, A. N., Adolphs, R. & Rutishauser, U. The geometry of domain-general performance monitoring in the human medial frontal cortex. *Science* **376**, eabm9922 (2022).

21. Higgins, I., Chang, L., Langston, V., Hassabis, D., Summerfield, C., Tsao, D. & Botvinick, M. Unsupervised deep learning identifies semantic disentanglement in single inferotemporal face patch neurons. *Nat Commun* **12**, 6456 (2021).

22. Higgins, I., Amos, D., Pfau, D., Racaniere, S., Matthey, L., Rezende, D. & Lerchner, A. Towards a Definition of Disentangled Representations. Preprint at https://doi.org/10.48550/arXiv.1812.02230 (2018).

23. Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K. & Fusi, S. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).

24. Fried, I., Rutishauser, U., Cerf, M. & Kreiman, G. *Single Neuron Studies of the Human Brain: Probing Cognition*. (MIT Press, 2014).

25. Kreiman, G., Koch, C. & Fried, I. Category-specific visual responses of single neurons in the human medial temporal lobe. *Nat Neurosci* **3**, 946–953 (2000).

26. Gross, C. G. Single neuron studies of inferior temporal cortex. *Neuropsychologia* **46**, 841–852 (2008).

27. Decramer, T., Premereur, E., Zhu, Q., Van Paesschen, W., van Loon, J., Vanduffel, W., Taubert, J., Janssen, P. & Theys, T. Single-Unit Recordings Reveal the Selectivity of a Human Face Area. *J Neurosci* **41**, 9340–9349 (2021).

28. Bengio, Y., Courville, A. & Vincent, P. Representation Learning: A Review and New Perspectives. *arXiv:1206.5538 [cs]* (2014).

29. Lake, B. M., Ullman, T. D., Tenenbaum, J. B. & Gershman, S. J. Building machines that learn and think like people. *Behavioral and Brain Sciences* **40**, e253 (2017).

30. Ito, T., Klinger, T., Schultz, D. H., Murray, J. D., Cole, M. W. & Rigotti, M. Compositional generalization through abstract representations in human and artificial neural networks. Preprint at https://doi.org/10.48550/arXiv.2209.07431 (2022).

31. Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T. & Wang, X.-J. Task representations in neural networks trained to perform many cognitive tasks. *Nat Neurosci* **22**, 297–306 (2019).

32. Johnston, W. J. & Fusi, S. Abstract representations emerge naturally in neural networks trained to perform multiple tasks. *Nat Commun* **14**, 1040 (2023).

33. Muhle-Karbe, P. S., Sheahan, H., Pezzulo, G., Spiers, H. J., Chien, S., Schuck, N. W. & Summerfield, C. Goal-seeking compresses neural codes for space in the human hippocampus and orbitofrontal cortex. *Neuron* (2023) doi:10.1016/j.neuron.2023.08.021.

34. Epstein, R. A., Patai, E. Z., Julian, J. B. & Spiers, H. J. The cognitive map in humans: spatial navigation and beyond. *Nat Neurosci* **20**, 1504–1513 (2017).

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

36. O'Keefe, J. & Nadel, L. *The Hippocampus as a Cognitive Map*. *Oxford University Press: Oxford, UK. (1978)* (Oxford University Press, 1978).

37. Wilson, R. C., Takahashi, Y. K., Schoenbaum, G. & Niv, Y. Orbitofrontal Cortex as a Cognitive Map of Task Space. *Neuron* **81**, 267–279 (2014).

38. Igarashi, K. M., Lee, J. Y. & Jun, H. Reconciling neuronal representations of schema, abstract task structure, and categorization under cognitive maps in the entorhinal-hippocampal-frontal circuits. *Current Opinion in Neurobiology* **77**, 102641 (2022).

39. Vaidya, A. R. & Badre, D. Abstract task representations for inference and control. *Trends in Cognitive Sciences* **26**, 484–498 (2022).

40. Morton, N. W., Schlichting, M. L. & Preston, A. R. Representations of common event structure in medial temporal lobe and frontoparietal cortex support efficient inference. *Proceedings of the National Academy of Sciences* **117**, 29338–29345 (2020).

41. Marr, D. Simple memory: a theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* **262**, 23–81 (1971).

42. McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* **102**, 419–457 (1995).

43. Gluck, M. A. & Myers, C. E. Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus* **3**, 491–516 (1993).

44. Benna, M. K. & Fusi, S. Place cells may simply be memory cells: Memory compression leads to spatial tuning and history dependence. *Proceedings of the National Academy of Sciences* **118**, e2018422118 (2021).

45. Tang, W., Shin, J. D. & Jadhav, S. P. Geometric transformation of cognitive maps for generalization across hippocampal-prefrontal circuits. *Cell Reports* **42**, 112246 (2023).

46. Samborska, V., Butler, J. L., Walton, M. E., Behrens, T. E. J. & Akam, T. Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems. *Nat Neurosci* **25**, 1314–1326 (2022).

47. Wood, E. R., Dudchenko, P. A., Robitsek, R. J. & Eichenbaum, H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–633 (2000).

48. Grieves, R. M., Wood, E. R. & Dudchenko, P. A. Place cells on a maze encode routes rather than destinations. *eLife* **5**, e15986 (2016).

49. Frank, L. M., Brown, E. N. & Wilson, M. Trajectory Encoding in the Hippocampus and Entorhinal Cortex. *Neuron* **27**, 169–178 (2000).

50. Sun, C., Yang, W., Martin, J. & Tonegawa, S. Hippocampal neurons represent events as transferable units of experience. *Nat Neurosci* **23**, 651–663 (2020).

51. Boyle, L. M., Posani, L., Irfan, S., Siegelbaum, S. A. & Fusi, S. Tuned geometries of hippocampal representations meet the demands of social memory. 2022.01.24.477361 Preprint at https://doi.org/10.1101/2022.01.24.477361 (2023).

52. Bao, P., She, L., McGill, M. & Tsao, D. Y. A map of object space in primate inferotemporal cortex. *Nature* **583**, 103–108 (2020).

53. Hesse, J. K. & Tsao, D. Y. The macaque face patch system: a turtle's underbelly for the brain. *Nat Rev Neurosci* **21**, 695–716 (2020).

54. Tanaka, K. Inferotemporal Cortex and Object Vision. *Annual Review of Neuroscience* **19**, 109–139 (1996).

55. Khuvis, S., Yeagle, E. M., Norman, Y., Grossman, S., Malach, R. & Mehta, A. D. Face-Selective Units in Human Ventral Temporal Cortex Reactivate during Free Recall. *J. Neurosci.* **41**, 3386–3399 (2021).

56. Axelrod, V., Rozier, C., Malkinson, T. S., Lehongre, K., Adam, C., Lambrecq, V., Navarro, V. & Naccache, L. Face-selective neurons in the vicinity of the human fusiform face area. *Neurology* **92**, 197–198 (2019).

57. Duncan, J. & Owen, A. M. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci* **23**, 475–483 (2000).

58. Dosenbach, N. U. F., Visscher, K. M., Palmer, E. D., Miezin, F. M., Wenger, K. K., Kang, H. C., Burgund, E. D., Grimes, A. L., Schlaggar, B. L. & Petersen, S. E. A Core System for the Implementation of Task Sets. *Neuron* **50**, 799–812 (2006).

59. Ma, L., Chan, J. L., Johnston, K., Lomber, S. G. & Everling, S. Macaque anterior cingulate cortex deactivation impairs performance and alters lateral prefrontal oscillatory activities in a rule-switching task. *PLOS Biology* **17**, e3000045 (2019).

60. Monosov, I. E., Haber, S. N., Leuthardt, E. C. & Jezzini, A. Anterior cingulate cortex and the control of dynamic behavior in primates. *Curr Biol* **30**, R1442–R1454 (2020).

61. Economides, M., Guitart-Masip, M., Kurth-Nelson, Z. & Dolan, R. J. Anterior Cingulate Cortex Instigates Adaptive Switches in Choice by Integrating Immediate and Delayed Components of Value in Ventromedial Prefrontal Cortex. *J Neurosci* **34**, 3340–3349 (2014).

1421 62. Cohen, Y., Schneidman, E. & Paz, R. The geometry of neuronal representations during rule
1422 learning reveals complementary roles of cingulate cortex and putamen. *Neuron* **109**, 839-851.e9 (2021).

1423 63. Fu, Z., Wu, D.-A. J., Ross, I., Chung, J. M., Mamelak, A. N., Adolphs, R. & Rutishauser, U. Single-
1424 Neuron Correlates of Error Monitoring and Post-Error Adjustments in Human Medial Frontal Cortex.
1425 *Neuron* **101**, 165-177.e5 (2019).

1426 64. Wang, S., Mamelak, A. N., Adolphs, R. & Rutishauser, U. Abstract goal representation in visual
1427 search by neurons in the human pre-supplementary motor area. *Brain* **142**, 3530–3549 (2019).

1428 65. Insausti, R. & Muñoz, M. Cortical projections of the non-entorhinal hippocampal formation in
1429 the cynomolgus monkey (Macaca fascicularis). *Eur J Neurosci* **14**, 435–451 (2001).

1430 66. Aggleton, J. P., Wright, N. F., Rosene, D. L. & Saunders, R. C. Complementary Patterns of Direct
1431 Amygdala and Hippocampal Projections to the Macaque Prefrontal Cortex. *Cereb Cortex* **25**, 4351–4373
1432 (2015).

1433 67. Wang, J., John, Y. & Barbas, H. Pathways for Contextual Memory: The Primate Hippocampal
1434 Pathway to Anterior Cingulate Cortex. *Cereb Cortex* **31**, 1807–1826 (2020).

1435 68. Spellman, T., Rigotti, M., Ahmari, S. E., Fusi, S., Gogos, J. A. & Gordon, J. A. Hippocampal–
1436 prefrontal input supports spatial encoding in working memory. *Nature* **522**, 309–314 (2015).

1437 69. Kamiński, J., Sullivan, S., Chung, J. M., Ross, I. B., Mamelak, A. N. & Rutishauser, U. Persistently
1438 active neurons in human medial frontal and medial temporal lobe support working memory. *Nat*
1439 *Neurosci* **20**, 590–601 (2017).

1440 70. Rutishauser, U., Schuman, E. M. & Mamelak, A. N. Activity of human hippocampal and amygdala
1441 neurons during retrieval of declarative memories. *Proceedings of the National Academy of Sciences* **105**,
1442 329–334 (2008).

1443 71. Rutishauser, U., Ye, S., Koroma, M., Tudusciuc, O., Ross, I. B., Chung, J. M. & Mamelak, A. N.
1444 Representation of retrieval confidence by single neurons in the human medial temporal lobe. *Nat*
1445 *Neurosci* **18**, 1041–1050 (2015).

1446 72. Zhou, J., Jia, C., Montesinos-Cartagena, M., Gardner, M. P. H., Zong, W. & Schoenbaum, G.
1447 Evolving schema representations in orbitofrontal ensembles during learning. *Nature* **590**, 606–611
1448 (2021).

1449 73. Zhou, J., Montesinos-Cartagena, M., Wikenheiser, A. M., Gardner, M. P. H., Niv, Y. &
1450 Schoenbaum, G. Complementary Task Structure Representations in Hippocampus and Orbitofrontal
1451 Cortex during an Odor Sequence Task. *Current Biology* **29**, 3402-3409.e3 (2019).

1452 74. Cole, M. W., Laurent, P. & Stocco, A. Rapid instructed task learning: a new window into the
1453 human brain's unique capacity for flexible cognitive control. *Cogn Affect Behav Neurosci* **13**, 1–22
1454 (2013).

1455 75. Brainard, D. H. The Psychophysics Toolbox. *Spatial Vision* **10**, 433–436 (1997).

1456 76. Tyszka, J. M. & Pauli, W. M. In vivo delineation of subdivisions of the human amygdaloid
1457 complex in a high-resolution group template. *Human Brain Mapping* **37**, 3979–3998 (2016).

1458 77. Rutishauser, U., Schuman, E. M. & Mamelak, A. N. Online detection and sorting of extracellularly
1459 recorded action potentials in human medial temporal lobe recordings, in vivo. *J Neurosci Methods* **154**,
1460 204–224 (2006).

1461 78. Hill, D. N., Mehta, S. B. & Kleinfeld, D. Quality Metrics to Accompany Spike Sorting of
1462 Extracellular Signals. *J. Neurosci.* **31**, 8699–8705 (2011).

1463 79. Courellis, H., Nummela, S., Miller, C. & Cauwenberghs, G. A computational framework for
1464 effective isolation of single-unit activity from in-vivo electrophysiological recording. in *2017 IEEE*
1465 *Biomedical Circuits and Systems Conference (BioCAS)* 1–4 (2017). doi:10.1109/BIOCAS.2017.8325164.

1466

Figure 1

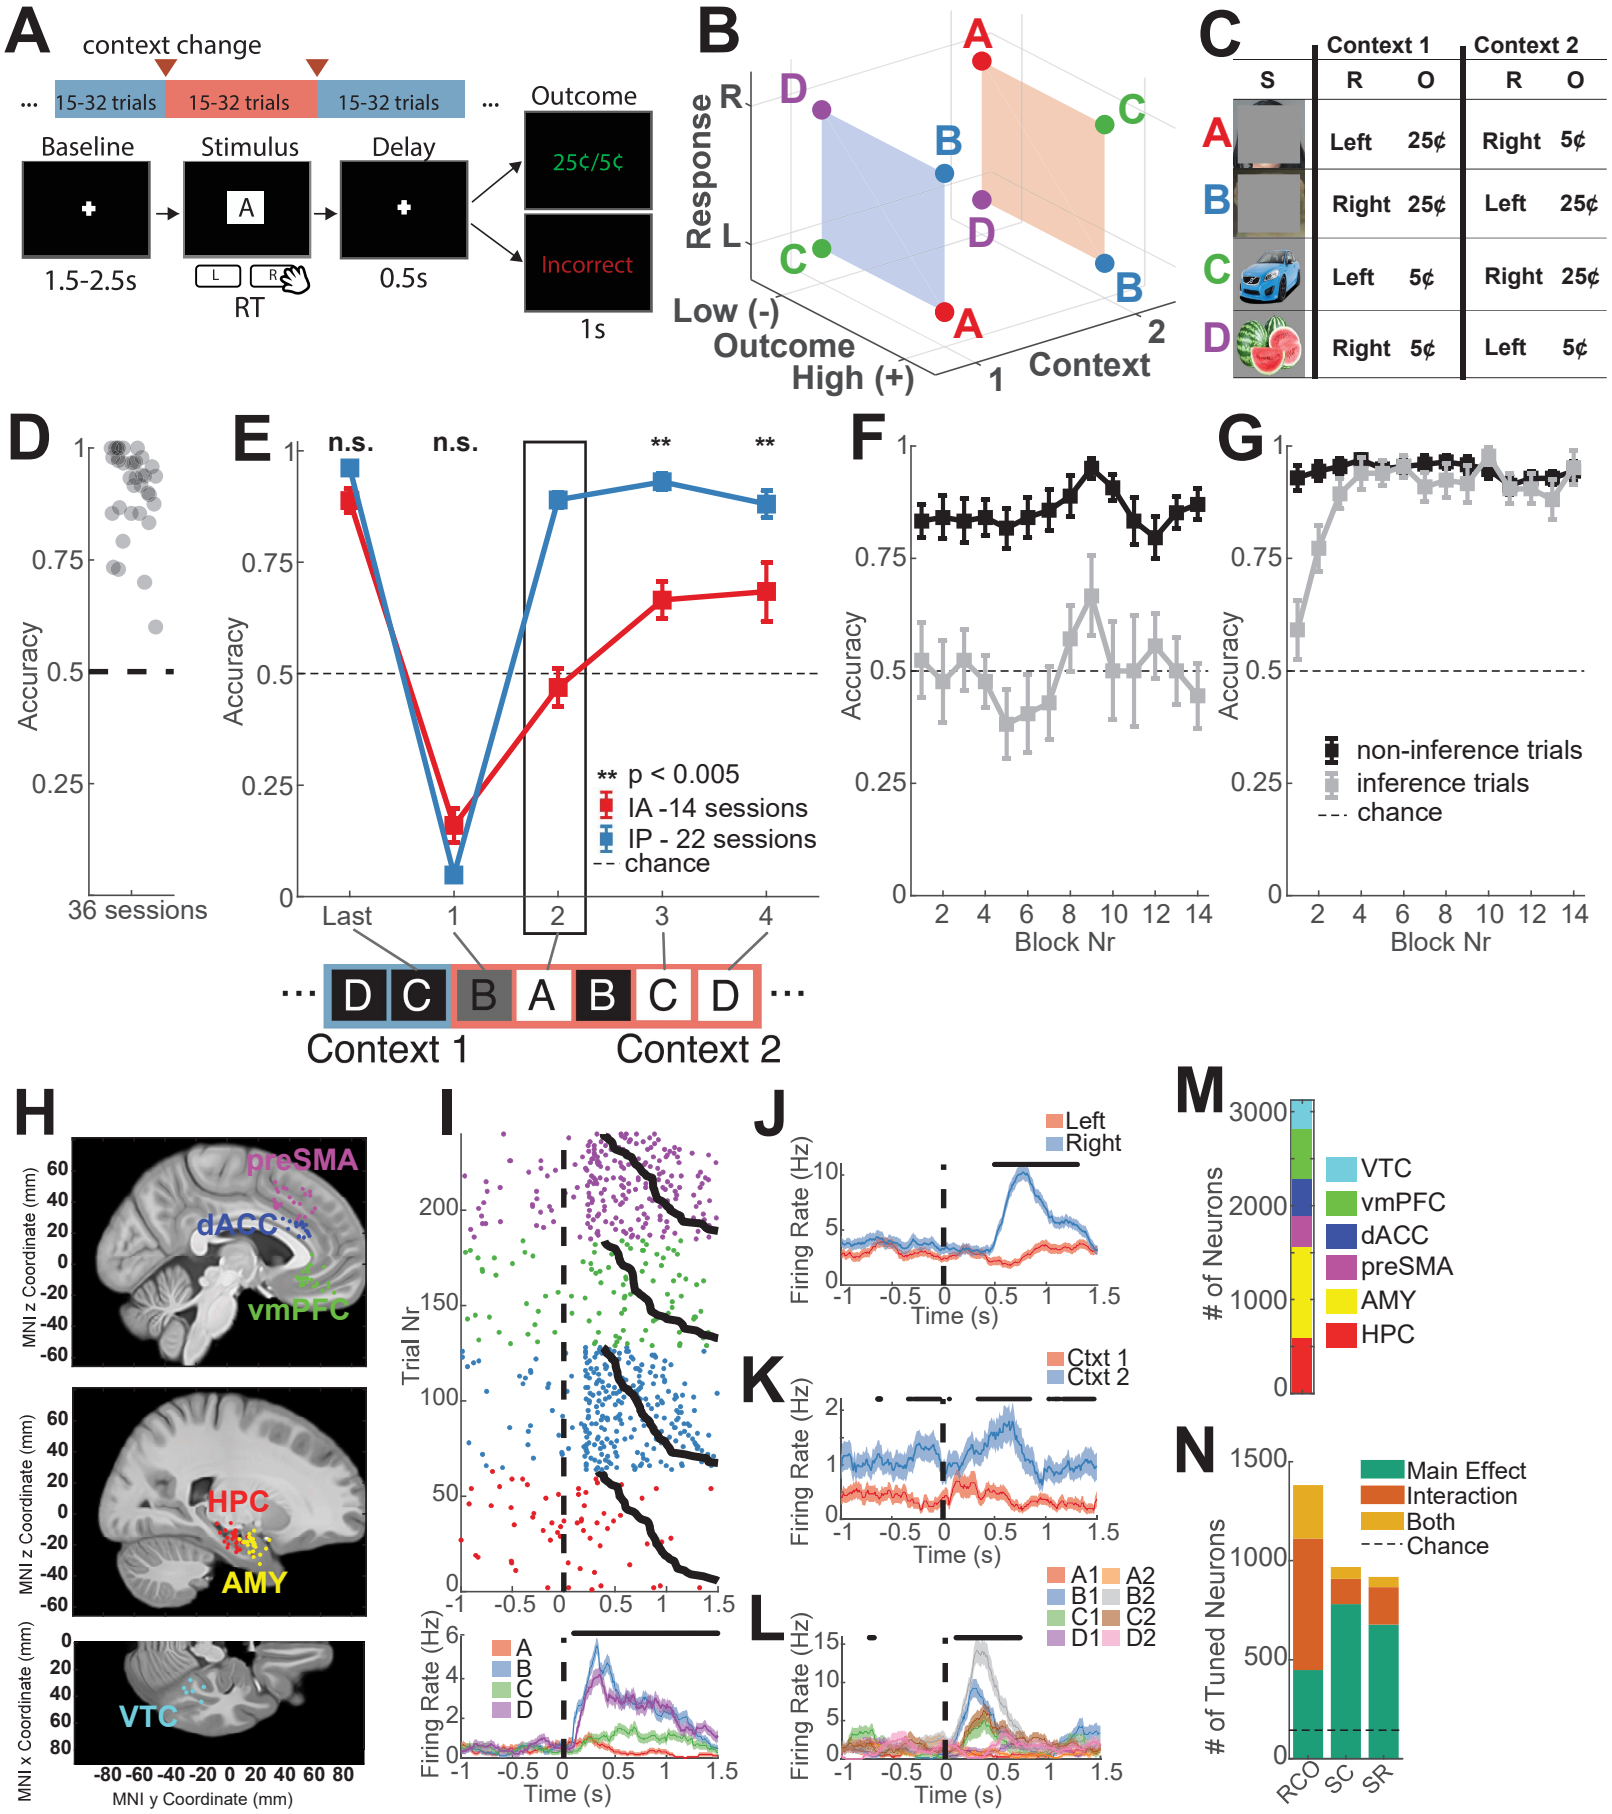
A
context change
15-32 trials
15-32 trials
15-32 trials
Baseline
Stimulus
Delay
Outcome
25¢/5¢
Incorrect
1.5-2.5s
RT
0.5s
1s
B
Response
R
L
Low (-)
Outcome
High (+)
Context
1
2
C
Context 1
Context 2
S
R
O
R
O
A
Left
25¢
Right
5¢
B
Right
25¢
Left
25¢
C
Left
5¢
Right
25¢
D
Right
5¢
Left
5¢
D
Accuracy
36 sessions
E
n.s.
n.s.
**
**
Accuracy
** p < 0.005
IA -14 sessions
IP - 22 sessions
chance
Last
1
2
3
4
Context 1
Context 2
F
Accuracy
Block Nr
G
Accuracy
non-inference trials
inference trials
chance
Block Nr
H
preSMA
dACC
vmPFC
HPC
AMY
VTC
MNI z Coordinate (mm)
MNI x Coordinate (mm)
MNI y Coordinate (mm)
I
Trial Nr
Firing Rate (Hz)
Time (s)
A
B
C
D
J
Left
Right
Firing Rate (Hz)
Time (s)
K
Ctxt 1
Ctxt 2
Firing Rate (Hz)
Time (s)
L
A1
A2
B1
B2
C1
C2
D1
D2
Firing Rate (Hz)
Time (s)
M
# of Neurons
VTC
vmPFC
dACC
preSMA
AMY
HPC
N
# of Tuned Neurons
Main Effect
Interaction
Both
Chance
RCO
SC
SR

Figure 2

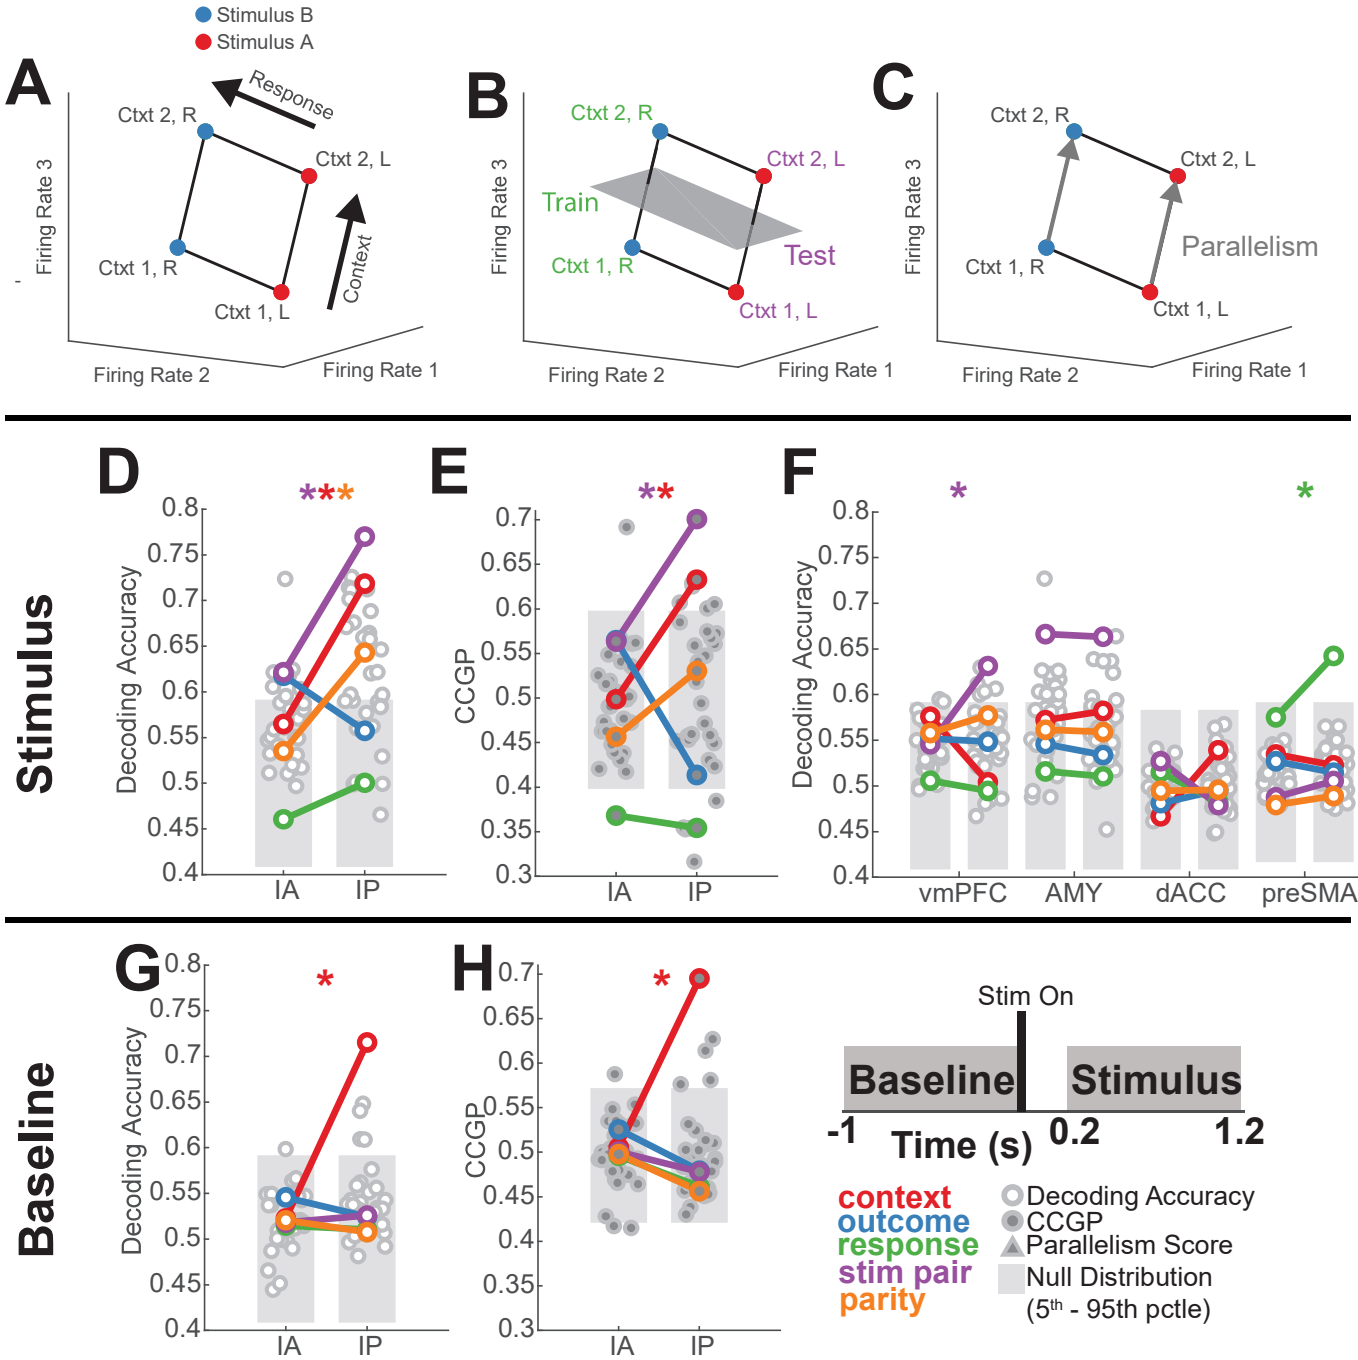
Stimulus B
Stimulus A
A
Response
Context
Ctxt 2, R
Ctxt 2, L
Ctxt 1, R
Ctxt 1, L
Firing Rate 3
Firing Rate 2
Firing Rate 1
B
Train
Test
C
Parallelism
Stimulus
Baseline
D
E
F
G
H
Decoding Accuracy
CCGP
IA
IP
vmPFC
AMY
dACC
preSMA
Stim On
Baseline
Stimulus
-1
0.2
1.2
Time (s)
context
outcome
response
stim pair
parity
Decoding Accuracy
CCGP
Parallelism Score
Null Distribution
(5th - 95th pctle)

Figure 3

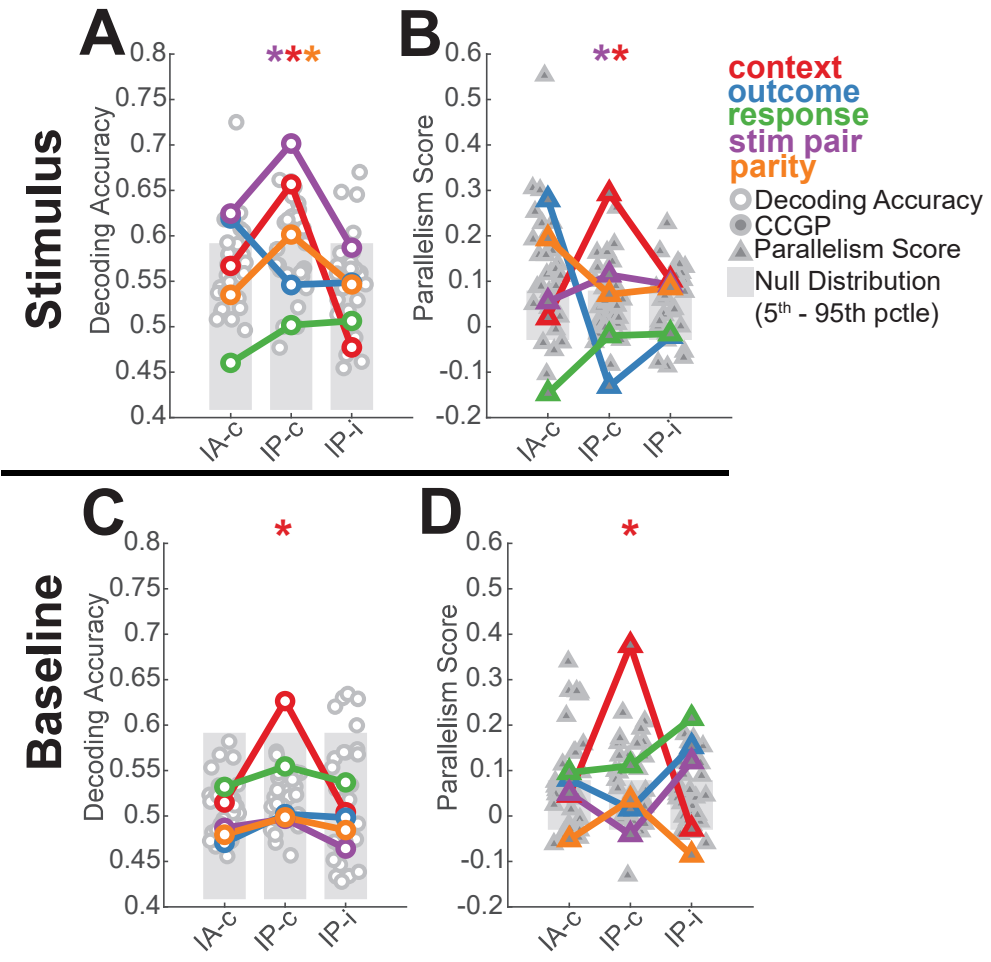
A
B
C
D
Stimulus
Baseline
Decoding Accuracy
Parallelism Score
IA-c
IP-c
IP-i
context
outcome
response
stim pair
parity
Decoding Accuracy
CCGP
Parallelism Score
Null Distribution
(5th - 95th pctle)

Figure 4

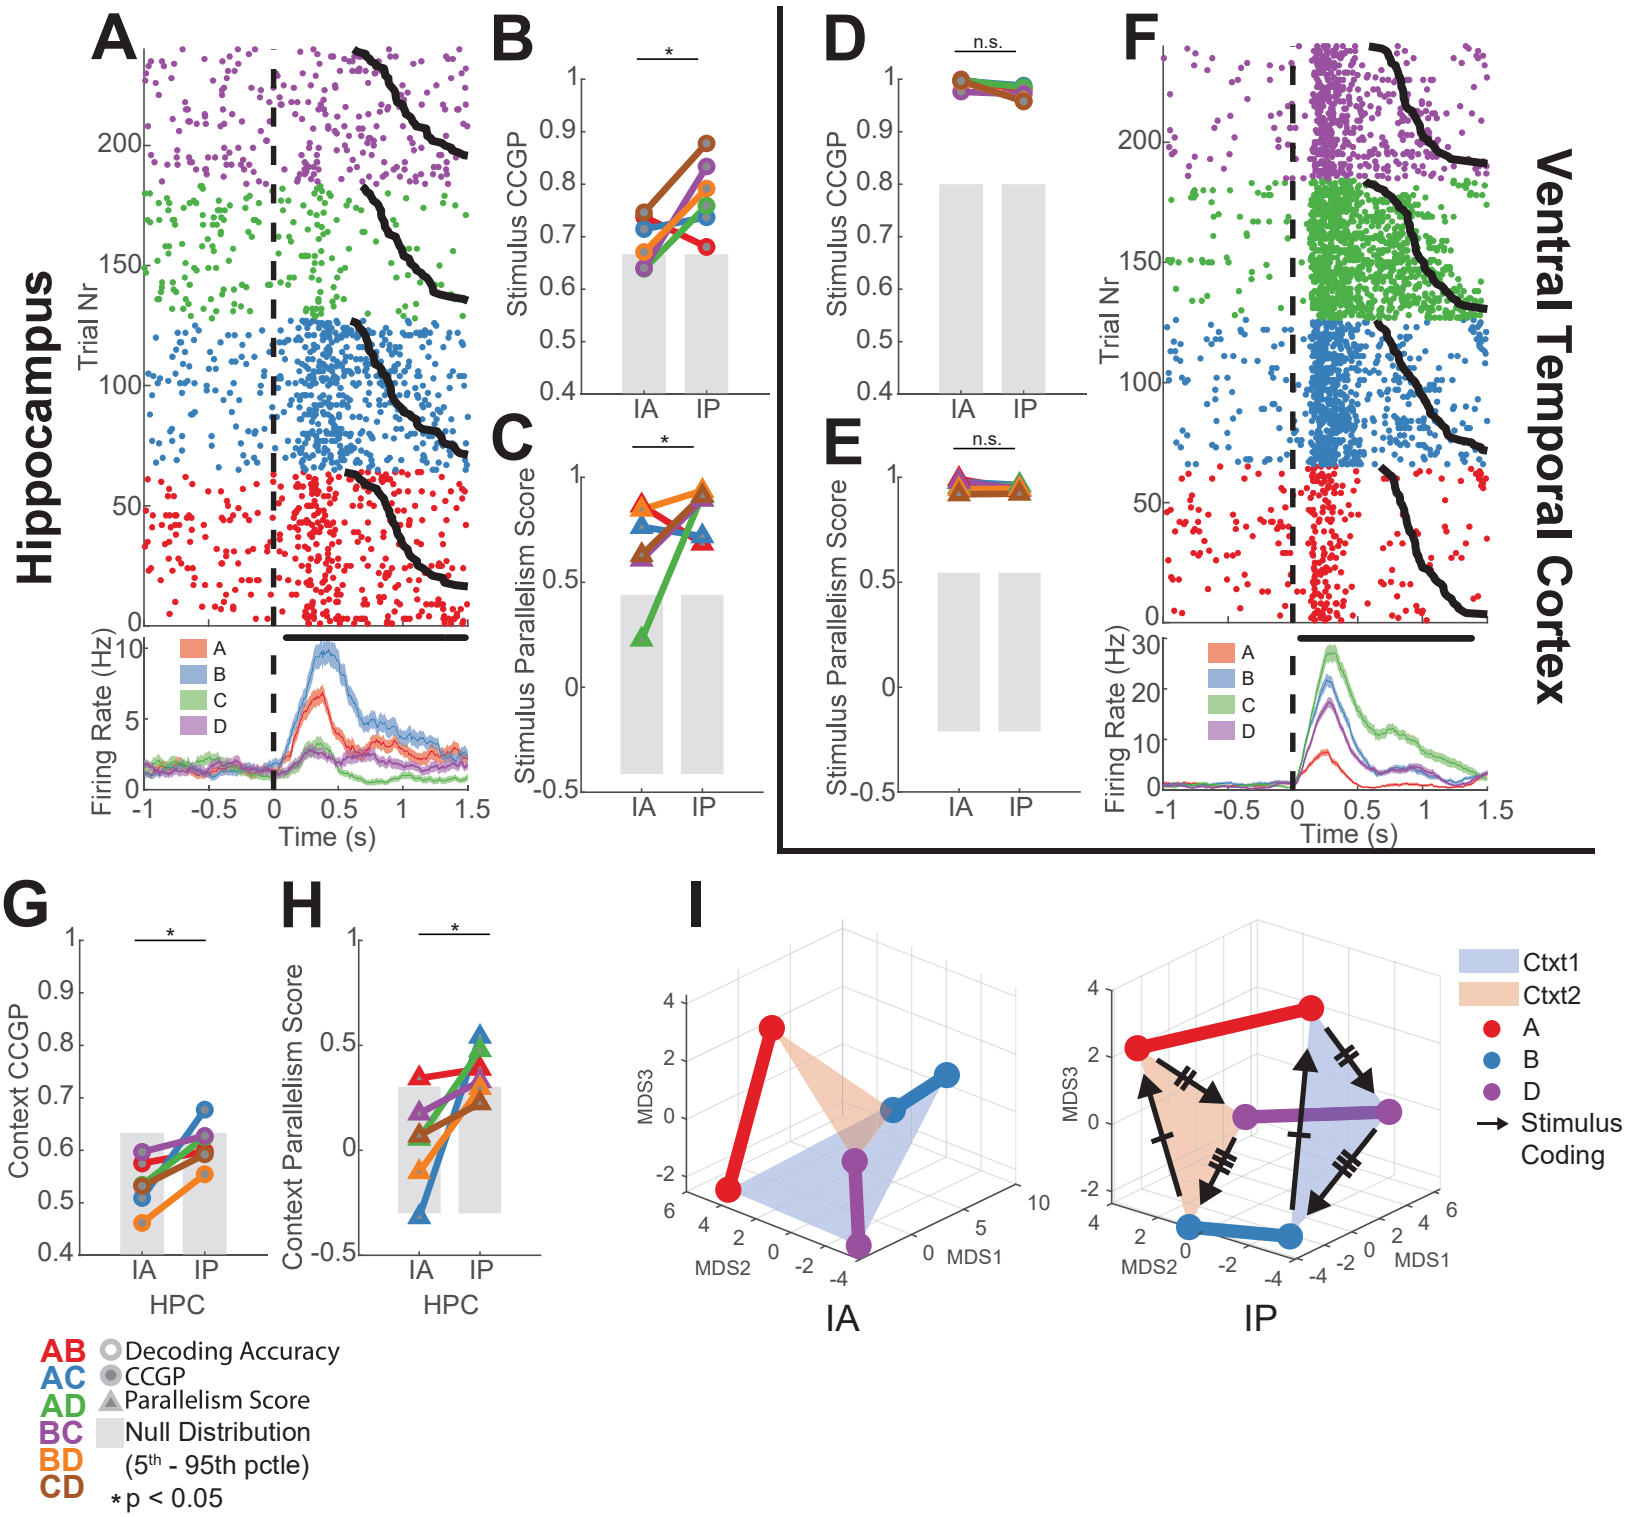
Hippocampus
Ventral Temporal Cortex
A
B
C
D
E
F
G
H
I
Trial Nr
Firing Rate (Hz)
Time (s)
A
B
C
D
Stimulus CCGP
Stimulus Parallelism Score
Context CCGP
Context Parallelism Score
IA
IP
HPC
n.s.
*
MDS1
MDS2
MDS3
Ctxt1
Ctxt2
Stimulus Coding
AB
AC
AD
BC
BD
CD
Decoding Accuracy
CCGP
Parallelism Score
Null Distribution
(5th - 95th pctle)
* p < 0.05

Figure 5

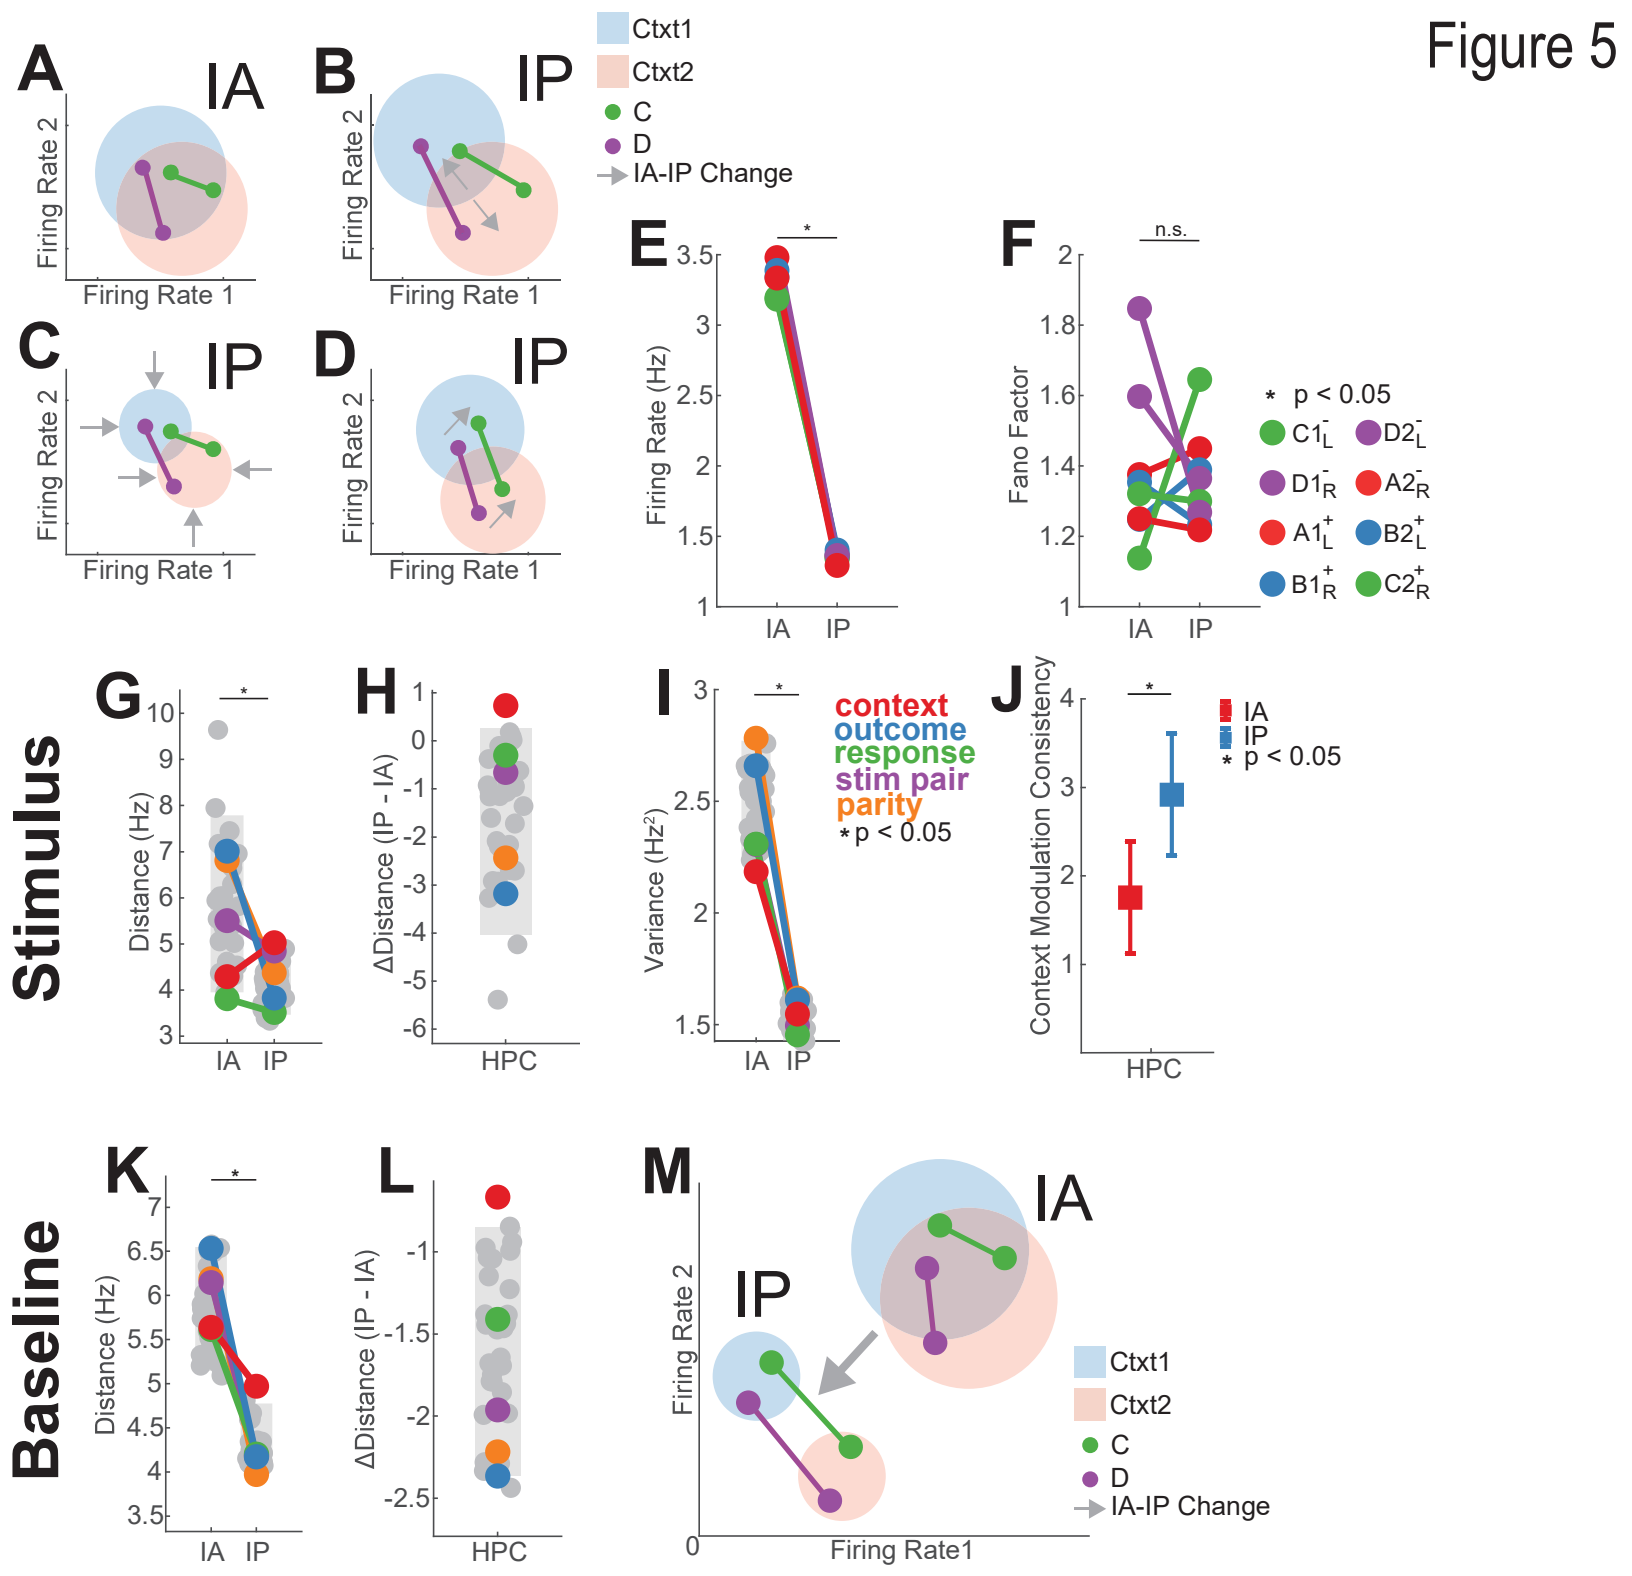
A
IA
B
IP
C
IP
D
IP
Ctxt1
Ctxt2
C
D
IA-IP Change
Firing Rate 2
Firing Rate 1
E
Firing Rate (Hz)
IA
IP
F
Fano Factor
n.s.
* p < 0.05
Stimulus
G
Distance (Hz)
H
ΔDistance (IP - IA)
HPC
I
Variance (Hz²)
context
outcome
response
stim pair
parity
* p < 0.05
J
Context Modulation Consistency
IA
IP
* p < 0.05
Baseline
K
Distance (Hz)
L
ΔDistance (IP - IA)
M
Firing Rate2
Firing Rate1
Ctxt1
Ctxt2
C
D
IA-IP Change

Figure 6

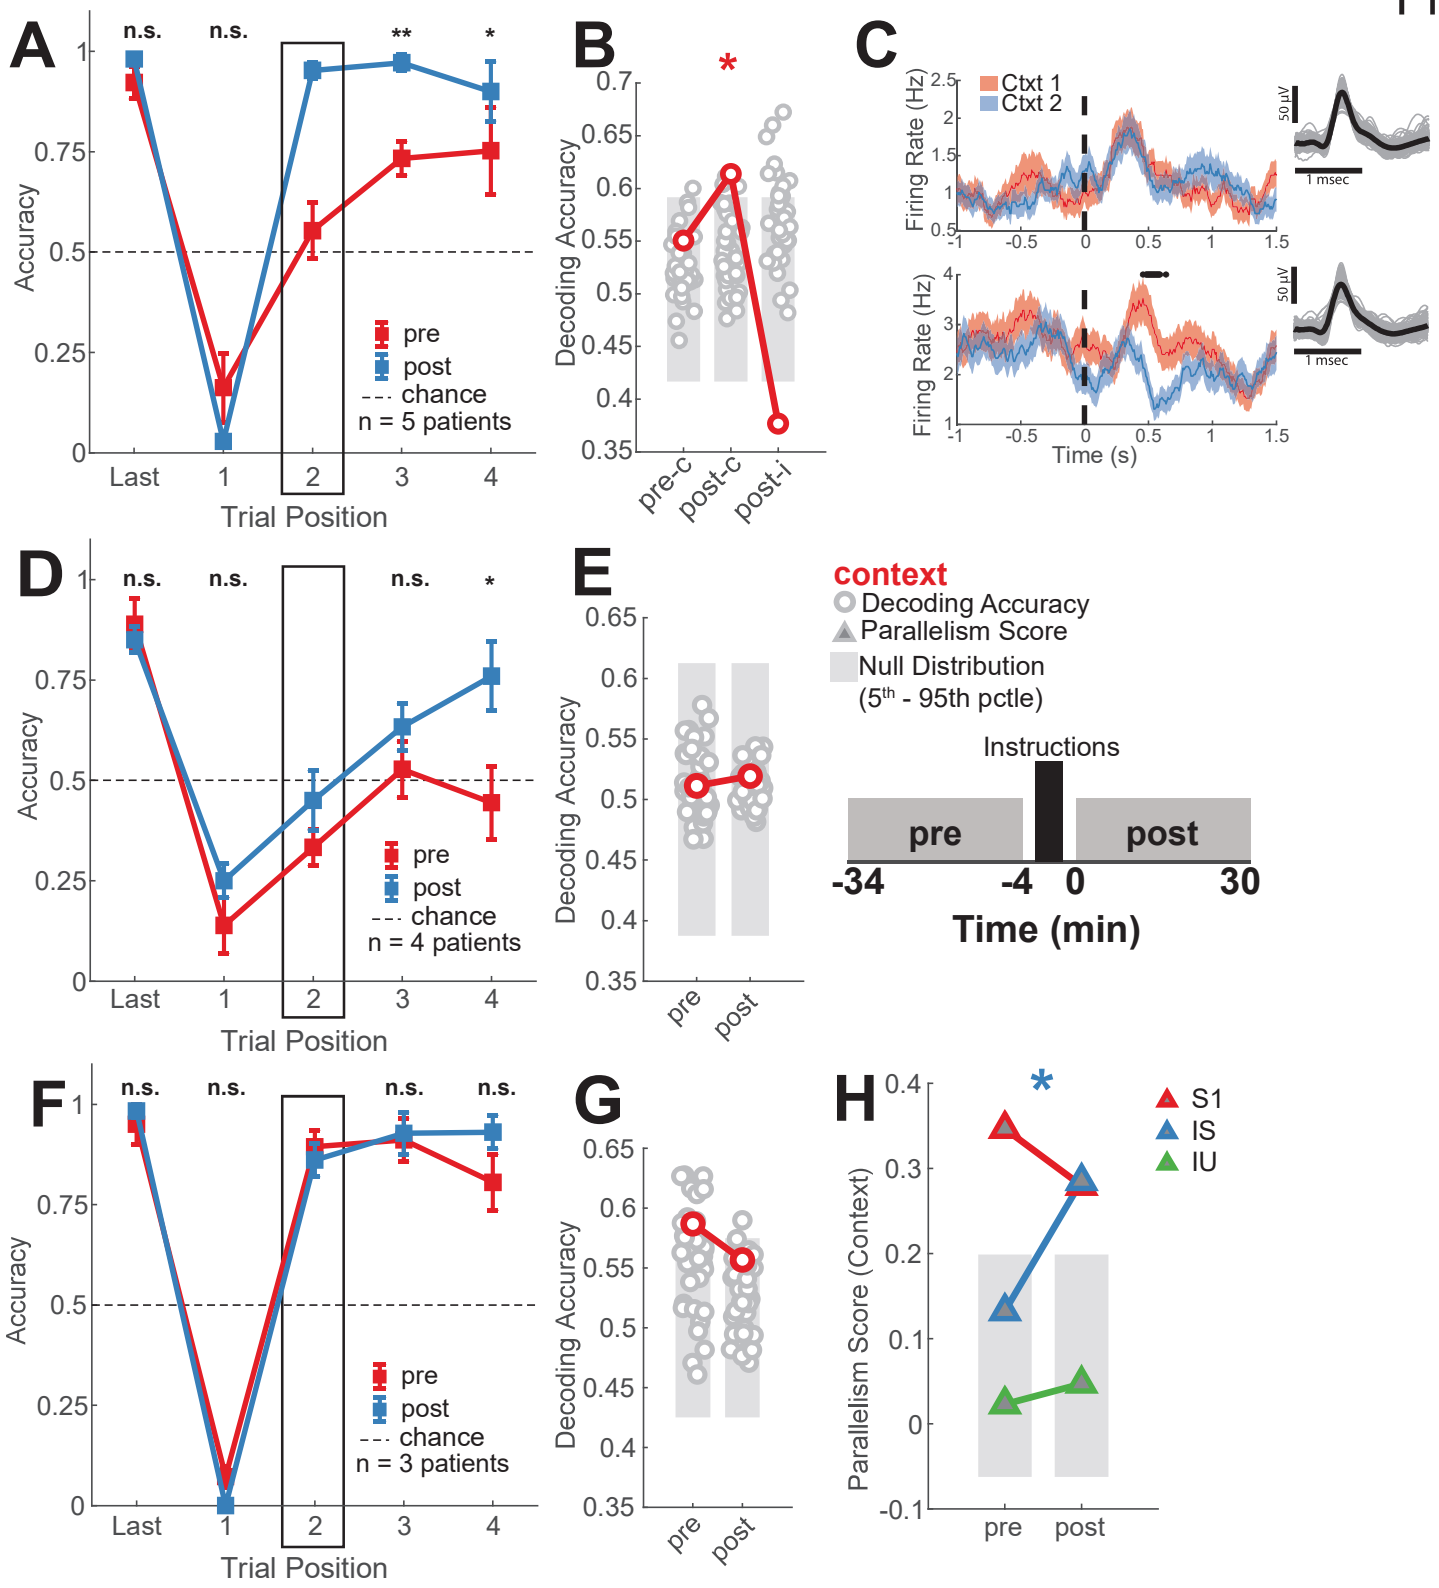
A
n.s.
n.s.
**
*
Accuracy
pre
post
chance
n = 5 patients
Last
1
2
3
4
Trial Position
B
*
Decoding Accuracy
pre-c
post-c
post-i
C
Ctxt 1
Ctxt 2
Firing Rate (Hz)
50 µV
1 msec
Time (s)
D
n.s.
n.s.
n.s.
*
Accuracy
pre
post
chance
n = 4 patients
Trial Position
E
Decoding Accuracy
pre
post
context
Decoding Accuracy
Parallelism Score
Null Distribution
(5th - 95th pctle)
Instructions
pre
post
-34
-4
0
30
Time (min)
F
n.s.
n.s.
n.s.
n.s.
Accuracy
pre
post
chance
n = 3 patients
Trial Position
G
Decoding Accuracy
pre
post
H
*
S1
IS
IU
Parallelism Score (Context)
pre
post