

# CA3 size predicts the precision of memory recall

Martin J. Chadwick<sup>1</sup>, Heidi M. Bonnici, and Eleanor A. Maguire<sup>1</sup>

Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, London WC1N 3BG, United Kingdom

Edited by Marcus E. Raichle, Washington University in St. Louis, St. Louis, MO, and approved June 12, 2014 (received for review October 18, 2013)

**There is enduring interest in why some of us have clearer memories than others, given the substantial individual variation that exists in retrieval ability and the precision with which we can differentiate past experiences. Here we report novel evidence showing that variation in the size of human hippocampal subfield CA3 predicted the amount of neural interference between episodic memories within CA3, which in turn predicted how much retrieval confusion occurred between past memories. This effect was not apparent in other hippocampal subfields. This shows that subtle individual differences in subjective mnemonic experience can be accurately gauged from measurable variations in the anatomy and neural coding of hippocampal region CA3. Moreover, this mechanism may be relevant for understanding memory muddles in aging and pathological states.**

hippocampus | pattern separation | pattern completion | fMRI | decoding

Our memories often contain overlapping elements, because they tend to feature the same people and places that form the cornerstones of our lives. Nevertheless, we are generally able to recall many of these past experiences as distinct episodes, although we are not all equally adept at doing so. There is substantial individual variation in retrieval ability and the precision with which we can differentiate past events (1, 2). This is most acute as we age and in conditions such as dementia, where confusion about the past is often evident (2). There is keen interest, therefore, in elucidating the neural mechanisms that allow us to recollect numerous life experiences despite a high degree of intermemory similarity.

We know little about how this is achieved in humans, but theoretical models propose that computations within hippocampal subfields facilitate the efficient storage and retrieval of similar memories (3–7). When we experience an event, pattern separation leads to the formation of a distinct neural representation within region CA3 (8–11). At retrieval, a previously stored memory representation within CA3 can be reactivated through the process of pattern completion (12, 13). However, when episodes are highly similar, the CA3 neuronal representations may not be completely distinct, leading to partial overlap (14). It is therefore not clear precisely what the limits of CA3 pattern separation might be. Here we directly tested the capacity of human CA3 to maintain distinct episodic representations in the presence of overlapping elements. We further investigated whether variation in this ability provides an explanatory account of individual differences in the precision of episodic memory retrieval.

## Results

We combined high-resolution functional MRI (14) (fMRI) with an ultra-high resolution structural MRI scanning protocol that permitted the separate identification of CA1, CA3, dentate gyrus (DG), and subiculum (15, 16). Stimuli were created by filming two brief action events against a green-screen background. Each event was then superimposed onto the same two spatial contexts, creating four movie clips that included every combination of the two events and the two contexts (17) (Fig. 1). Each participant viewed the four movies before scanning and then vividly recalled each one numerous times during fMRI. Because the four episodes completely overlapped with one another in terms of their constituent elements, any successful differentiation of the four

memories from patterns of activation would suggest the presence of individual episodic memory traces (note that an alternative possibility based on the presence of separate event and spatial context components is considered and ruled out below). Using this paradigm we could also assess whether the recall of an episode led to the coactivation of overlapping episodes through CA3 pattern completion. Such a result would suggest that the episodic representations were not entirely distinct, but overlapped with one another owing to a partial failure of pattern separation. This design therefore allowed us to test for the presence of episodic information within CA3, as well as any possible overlap between the episodic representations.

Data were analyzed using a model-based decoding approach (18–20) called multivariate Bayes (21–23) (MVB) in which a hierarchical Bayesian model is used to discover the pattern of voxels that best explains a given target variable from the experimental design (in this case, the activity profile associated with recalling a specific episodic memory). For each model, this voxel pattern was used to produce a predicted information time course, and model fit was assessed by correlating the predicted time course against the target variable. Statistical significance was assessed using a nonparametric permutation approach (24). This involved shuffling the memory labels, which provided a baseline level of information taking into account general retrieval processes, thereby ensuring that any significant results could only be due to the presence of information about each individual memory (*Methods*).

We first assessed whether information about each of the four individual episodes was present within each of the four hippocampal subfields. Using a Bonferroni-corrected threshold of  $P < 0.0125$ , we found that only CA3 contained episodic information [ $t(14) = 3.34, P = 0.0049$ ], whereas the other three subfields did not [CA1:  $t(14) = 2.02, P = 0.063$ ; DG:  $t(14) = -0.53, P = 0.6$ ; subiculum:  $t(14) = 2.09, P = 0.056$ ]. This result is consistent with

## Significance

How does the brain allow us to recall numerous life experiences despite there often being a high degree of similarity between memories? This is a key question in neuroscience. Moreover, there is also keen interest in understanding why some people are able to recall memories with greater clarity than other people. In this study, we identified a specific brain region, CA3, an area within a structure called the hippocampus, and a mechanism within it that helps to explain individual differences in recollection. These findings have relevance for all of us in elucidating memory muddles in general, in aging, and possibly also in conditions such as dementia, where confusion about the past is often evident.

Author contributions: M.J.C. and E.A.M. designed research; M.J.C. performed research; H.M.B. contributed new reagents/analytic tools; M.J.C. and E.A.M. analyzed data; and M.J.C. and E.A.M. wrote the paper.

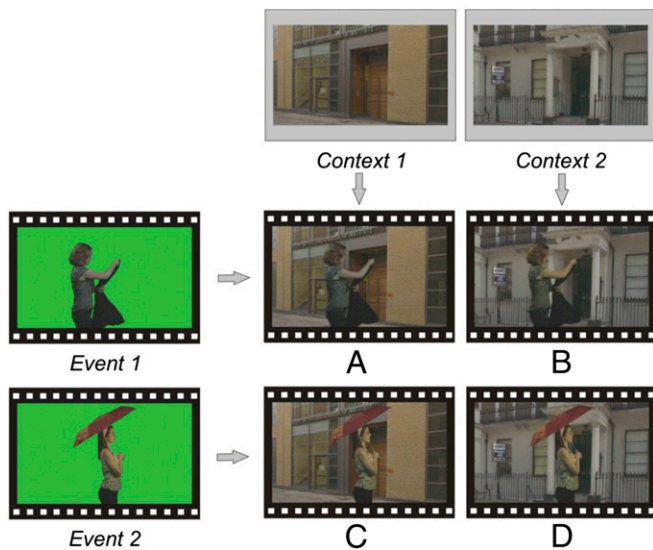
The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

<sup>1</sup>To whom correspondence may be addressed. Email: martin.chadwick@ucl.ac.uk or e.maguire@ucl.ac.uk.

This article contains supporting information online at [www.pnas.org/lookup/suppl/doi:10.1073/pnas.1319641111/-DCSupplemental](http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1319641111/-DCSupplemental).



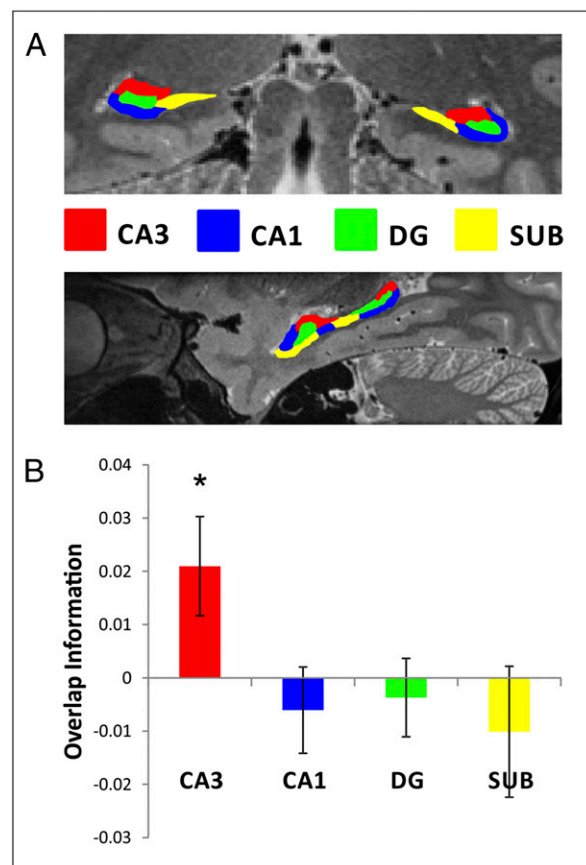
**Fig. 1.** The movie stimuli. A set of overlapping episodes was created by filming two brief action events against a green-screen background (Left). Each event was then superimposed onto the same two spatial contexts (Upper), creating four movie clip “episodes” (A–D) that included every combination of the two events and contexts. Participants recalled memories of these movies during fMRI scanning.

the theoretical role of CA3 as the recipient of pattern-separated episodic information, which can then be retrieved at recall via pattern completion. Furthermore, it demonstrates that such information is still present and detectable despite the high degree of overlap between the four episodes. In a further MVB analysis we focused specifically on the cue period (Methods). This failed to find any significant cue-related information in any of the subfields (all  $P > 0.05$ ). Furthermore, a direct comparison of the information contained within subfield CA3 revealed significantly greater information during memory retrieval than during cue presentation [one-tailed  $t$  test  $t(15) = 1.90, P = 0.039$ ]. Thus, our effects seemed to be driven by episodic memory retrieval, rather than visual information related to the cues. After scanning, participants completed a debriefing questionnaire assessing eight factors relating to each memory (e.g., vividness). None of these factors differed significantly across the four episodes, demonstrating that they could not explain the decoding results (Table S1).

We next asked whether the overlapping episodes within CA3 were coded as entirely distinct neural representations or whether they partially overlapped with one another, leading to coactivation of similar episodes during retrieval. Because each MVB model estimates the pattern of voxel activity, we were able to assess memory overlap by correlating the MVB voxel patterns for each overlapping pair of memories. Pattern correlations for each pair of nonoverlapping memories were also calculated to provide a baseline measure of similarity. Fig. 2 displays the average overlap information after subtracting the baseline overlap for each subfield. Any information that is significantly greater than zero can only be due to the presence of overlapping mnemonic representations within that subfield. We found that CA3 contained a significant amount of neural overlap information [ $t(14) = 2.18, P = 0.047$ ]. Given that only CA3 was found to contain episodic information in the initial analysis, it was no surprise that there was no evidence for the presence of overlap information in the other subfields [CA1:  $t(14) = -0.66, P = 0.52$ ; DG:  $t(14) = -1.28, P = 0.22$ ; subiculum:  $t(14) = 0.16, P = 0.88$ ]. Thus, it seems that CA3 displays a significant degree of coactivation between overlapping episodic representations during recall. We suggest that this result may be due to a partial failure

of pattern separation in the presence of highly overlapping episodes. Such a failure would lead to the creation of partially overlapping CA3 neural representations, which are then coactivated during pattern completion. Notably, the specificity of this effect within CA3 cannot be explained by the number of voxels within this region compared with the other subfields, because CA1 was the largest subfield tested. Thus, if the effect were simply driven by number of voxels, we would have found a result within CA1 rather than CA3 (see Supporting Information for further consideration of this point).

So far the results suggest that individual episodic representations are present within CA3, but that the retrieval of an episode also partially coactivates overlapping memory traces. This suggests a competitive pattern completion process, whereby the desired memory trace is activated alongside the undesired, overlapping memory traces. If this process is indeed competitive, we would expect an increase in interfering memory traces to lead to a decrease in the desired memory trace through a process of inhibitory interference. We made use of a degree of intersubject variance in neural information for both the recalled episodic memory (episodic information) and the overlapping memories (overlap information) within CA3 to test this prediction and discovered a negative correlation between the two [ $r(13) = -0.63, P = 0.012$ ]. This was not the case in any other subfield



**Fig. 2.** (A) Coronal (Upper) and sagittal (Lower) views of the subfields in the hippocampi of an example participant. (B) The amount of overlapping memory coactivation within each subfield during episodic retrieval. This is measured as the average correlation between the MVB voxel patterns for overlapping memories, minus the correlation for nonoverlapping memories (baseline). The group mean is shown with SE bars. Only CA3 displayed a significant degree of coactivation, demonstrating that the overlapping episodes were not represented by completely distinct neuronal representations.

(all  $P > 0.25$ ), which is not surprising given that we did not find evidence for episodic information in any region other than CA3. We therefore find evidence for a competitive pattern completion process taking place within CA3. Furthermore, the amount of neural interference seems to vary across individuals.

One key question is whether these results genuinely reflect the dynamic competition between individual episodic representations within CA3, as we have suggested. In theory, our finding of distinct episodic information with partial overlap between overlapping memories could be explained instead by the presence of separate representations of event content and spatial context within CA3 without requiring the presence of distinct episodic representations. In this case, the retrieval of a given memory would activate the relevant event and context representations within CA3, rather than any individual episodic memory trace. However, the correlation data allow us to adjudicate between these two alternatives. If CA3 contained only separate event and spatial representations then a participant who has stronger representations should have a higher information measure for both the episodic and overlap measures, because both of these should be based directly on the strength of the underlying component representations. In other words, if this were the case, we ought to see a positive correlation across participants between unique and overlap information. However, if CA3 contains distinct episodic representations that are coactivated through a competitive process of pattern completion, then a participant with a strong episodic representation should show a weak overlap representation, because these overlapping episodes are suppressed. In other words, we should see a negative correlation between episodic and overlap information. Importantly, therefore, these two rival explanations make opposite predictions about the expected direction of the correlation. To formally demonstrate the validity of these two predictions, we created two sets of simulated fMRI data (25) based on the two different explanatory models and ran the same MVB analyses that were applied to our real data. These analyses confirmed that the presence of separate event and spatial context information should lead to a positive correlation between the strength of episodic and overlap information, whereas the presence of distinct episodic representations with competitive coactivation should lead to a negative correlation (see *Supporting Information* for full details of these simulation analyses). The fact that we found a significant negative correlation between episodic and overlap information in our real dataset therefore strongly supports the idea that we have detected individual episodic representations within CA3.

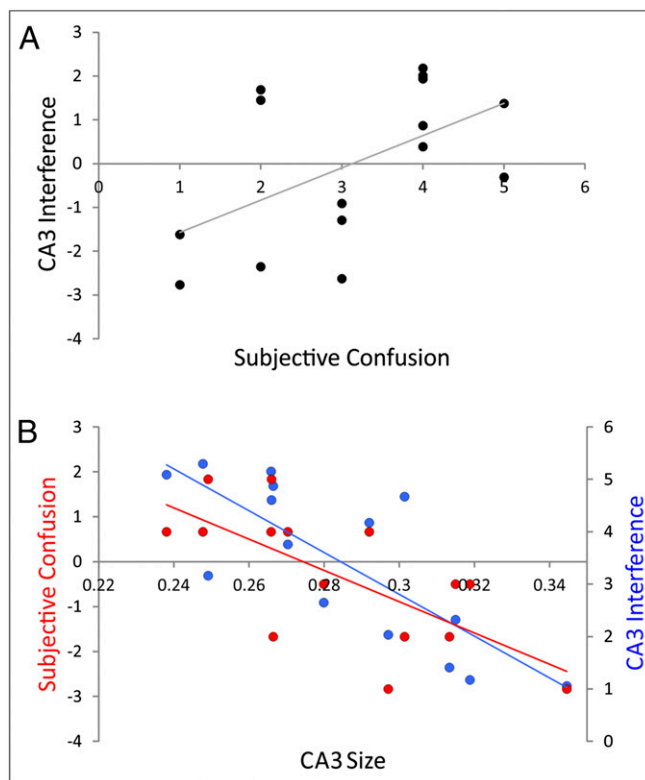
Another question prompted by this result is whether the neural interference taking place within CA3 actually influenced retrieval. Participants could recall the correct memory on each trial, suggesting that the CA3 overlap did not preclude memory retrieval. However, further analysis revealed a more complex story. In the postscan session, in addition to assessing possible confounds using the eight ratings mentioned earlier (Table S1) we also asked participants to provide a rating of how aware they were of the similarities across the four episodes during retrieval (from 1 to 5). We used this rating as a proxy for the degree of subjective confusion between the overlapping episodes (see *Supporting Information* for consideration of the relationship between subjective confusion and the other behavioral variables). We operationalized CA3 neural interference as the relative weighting of neural overlap information compared with episodic information for each participant (*Methods*), because this measure accurately reflects individual variation in the tension between the retrieval of the desired versus competing memories. We found a significant positive correlation between the degree of CA3 neural interference and subjective confusion across individuals [ $r(13) = 0.53, P = 0.04$ ; see also *Supporting Information*]. This indicates that measurable neural processes occurring within CA3, specifically the degree of neural

interference between episodes, can have a significant impact on retrieval processes (Fig. 3A).

We next asked whether either of these variables might be predicted by differences in the underlying physical substrate of subfield CA3. We found that the size of CA3 (adjusting for total hippocampal volume) correlated negatively with both subjective confusion [ $r(13) = 0.70, P = 0.0037$ ] and with CA3 neural interference [ $r(13) = -0.79, P = 0.0005$ ]. In other words, subjects with a larger CA3 showed a reduction in both CA3 neural interference and subjective confusion (Fig. 3B). Thus, there seems to be a tight relationship between CA3 size, CA3 processing, and subjective levels of retrieval confusion. To explore the relationship between these three variables further, we conducted a mediation analysis (26, 27). This determines whether or not a predictor variable (CA3 size) has a causal effect on another variable (subjective confusion) via a third, mediator variable (CA3 neural interference). We found a significant effect of the mediation pathway (using 10,000 bootstrap samples,  $P = 0.001$ ), which demonstrates that these relationships were not independent. These findings provide further evidence that variation in CA3 size can directly influence the amount of neural interference within CA3, which then influences a person's subjective memory experience, determining how much retrieval confusion occurs.

## Discussion

This study provides two novel insights into hippocampal processing. First, we found evidence that the recall of highly similar episodes leads to overlapping neural representations selectively within hippocampal CA3. This result is consistent with the idea



**Fig. 3.** (A) We found a positive correlation between the degree of subjective confusion and CA3 neural Interference. Subjective confusion ratings were on a scale of 1–5 (5 indicating high levels of confusion). (B) Both subjective confusion (displayed in red) and CA3 neural interference (shown in blue) correlated negatively with the relative size of CA3 (measured as a ratio of total hippocampal volume to control for total hippocampal volume).

that high levels of episodic similarity may elicit a partial failure of pattern separation, leading to the creation of overlapping representations within CA3. During retrieval, these overlapping patterns are then partially coactivated through pattern completion, which can lead to mnemonic confusion. Second, we found evidence for a neural mechanism underlying individual variation in retrieval precision whereby CA3 size directly affects the efficacy with which overlapping memories are differentiated, which in turn influences how we experience our memories. Put together, this set of results provides a striking demonstration that subtle individual differences in subjective mnemonic experience can be accurately predicted from measurable differences in the anatomy and neural coding of hippocampal region CA3.

These results relate to previous work on mnemonic interference. Several studies have demonstrated that competing associative memories are represented within visual cortex (28–30), and our results now demonstrate that such interference can also take place within hippocampal CA3. This suggests a possible explanatory mechanism underlying cortical interference effects, whereby cortical interference is the output of associative retrieval processes taking place within the hippocampus. This would be fully consistent with the idea that cortical reinstatement of memory traces follows pattern completion within the hippocampus (3). Specifically, our results suggest that a partial failure of pattern separation occurred, leading to the presence of overlapping representations within CA3. At retrieval, the presence of this representational overlap leads to a competitive pattern completion process, and concomitant mnemonic interference. Although we did not set out to test cortical representations in this study, our findings nevertheless suggest the possibility that cortical interference observed in previous studies may reflect the output of competitive processes taking place within CA3.

We also found a clear relationship between hippocampal structure and both neural function and subjective memory. There is a long history of studying the relationship between hippocampal structure and cognition (e.g., refs. 31–33), and recent studies have also begun to investigate the structure of specific hippocampal subfields (e.g., refs. 34–36). However, none of these studies focused on how individual variation in neural anatomy may relate to the computations taking place within the hippocampal subfields. To our knowledge, only one previous study has investigated this issue, and that was in aging participants in the context of being unable to separate CA3 from DG, and using simple objects as stimuli (37). They found that white matter integrity in the CA3/DG region correlated with both behavioral pattern separation and a neural measure of pattern separation within the same region. This suggests that the degree of memory decline with aging may depend on these hippocampal computations, which in turn may be associated with the integrity of white matter. Here we provide the first evidence, to our knowledge, of substantial individual variation in, specifically, CA3 anatomy and functioning on the one hand and complex episodic memory on the other in healthy, young participants. This three-way relationship places individual variation in complex episodic memory within the framework of computational theory for the first time to our knowledge and is an advance over previous studies that have simply shown correlations between anatomy and episodic memory with no accompanying mechanism to explain these correlations. Furthermore, as discussed above, our results provide a direct link between CA3 variation and individual differences in mnemonic interference. Put together, these results offer an important bridge between the neural computations of the hippocampus and mnemonic interference.

The evidence presented here supports the conclusion that there is a tight correspondence between CA3 anatomy, pattern separation and completion, and the subjective experience of episodic memory recall. We speculate that a larger CA3 may

promote a decrease in retrieval confusion via an increased number of CA3 neurons, or enhanced lateral connectivity within CA3, either of which could precipitate more efficient pattern separation. We suggest that the precise neural mechanisms underlying normal variation in episodic memory are now tractable and may prove fruitful in elucidating why some people have clearer memories than others, and why confusion about past experiences often characterizes normal aging (37) and pathological states (36, 38), issues that continue to provoke wide interest.

## Methods

**Participants and Experimental Design.** Fifteen healthy, right-handed participants (eight female) took part in the experiment (mean age 21.17 y, SD 2.18 y, range 18–25 y). All had normal or corrected-to-normal vision and gave informed written consent to participation in accordance with the University College London research ethics committee. The participants and experimental design have been reported previously (17) in a study that was focused on a different set of questions that did not concern the hippocampal subfields and that did not involve the high-resolution structural MRI scans or any of the data analysis methods deployed here. Full details of the stimuli and experimental design are provided in [Supporting Information](#).

**Image Acquisition.** High-resolution (1.5-mm<sup>3</sup>) fMRI scans were acquired using a 3T Magnetom Allegra head only MRI scanner (Siemens Healthcare), taking a partial volume of 35 slices focused on the temporal lobes. High-resolution T2-weighted structural images (0.52 × 0.52 × 0.5 mm<sup>3</sup>) were acquired on a 3T whole-body MRI scanner (Magnetom TIM Trio; Siemens Healthcare) also in a partial volume focused on the temporal lobes (39). To ensure optimal data quality, images were reconstructed online and underwent online quality assurance (40). See [Supporting Information](#) for further details of the MRI sequences and [Fig. S1](#) for an example of the high-resolution T2-weighted structural images.

**Segmentation of the Hippocampal Subfields.** Manual segmentation of the hippocampal subfields was conducted using the ITK-SNAP software package (41) and a recently devised subfield segmentation protocol (15, 16). It took on average 1 d to segment the subfields of one hippocampus. Crucially, the T2-weighted images acquired in this study, along with the new segmentation protocol, permitted the separate delineation of DG and CA3 (as well as CA1 and the subiculum). Some of the hippocampi from the current study (four left and three right) were segmented by a second trained operator to assess interrater reliability using the Dice metric (42). The Dice metric provides a measure of overlap that is normalized between 0 and 1, where 1 indicates perfect overlap. It is calculated as  $2|A \cap B| / (|A| + |B|)$ , where  $A \cap B$  represents the number of voxels that overlap between the two segmentations and  $A$  and  $B$  represent the number of nonoverlapping voxels in each of the two segmentations. It is therefore effectively a measure of the ratio of overlapping to nonoverlapping voxels. Dice scores for each subfield were as follows: CA1 0.78, CA3 0.68, DG 0.75, and subiculum 0.58. These scores are similar to those reported by other methods, indicating that the segmentations were reliable. Further consideration of the subfield segmentation is provided in [Supporting Information](#).

**fMRI Preprocessing.** All neuroimaging analyses were conducted using SPM8 ([www.fil.ion.ucl.ac.uk/spm](http://www.fil.ion.ucl.ac.uk/spm)). The first six functional volumes were discarded to allow for T1 equilibration. The remaining functional volumes were spatially realigned to the first image of the series, and distortion corrections were applied based on the field maps using the unwarp routines in SPM (43, 44). Each participant's whole-brain MT FLASH structural scan was then coregistered to a mean image of their realigned, distortion-corrected functional scans. Following this, the high-resolution T2-weighted structural average was coregistered to the MT FLASH structural scan, bringing all images into alignment (this coregistration was performed before the segmentation of the subfields). Functional data were left unsmoothed for the decoding analyses so that information present across patterns of voxels across the small subfields could be detected. All data were analyzed in the native space of each participant, using subject-specific regions of interest.

**Decoding Analysis.** We used a Bayesian model-based decoding method, MVB, for all decoding analyses (21–23). An MVB model maps multivariate voxel responses to a psychological target variable (e.g., individual memories) using a hierarchical approach known as parametric empirical Bayes. This modeling process therefore attempts to find the pattern of voxel responses that

best explains the target variable, which in this case is the activity profile associated with a single episodic memory. MVB uses the same design matrix of experimental variables used in a conventional SPM analysis. When a decoding contrast is specified, a target variable  $X$  is derived from this contrast, after removing confounds. The multivariate voxel activity can be used to produce a predictor variable  $Y$ . For each MVB model, the model fit was assessed by measuring the (Fisher-transformed) correlation between predictor variable  $Y$  and target variable  $X$ . It is possible to specify priors on the pattern of voxel weights in an MVB design, and in this case we used a sparse prior, because the distribution of episodic representations is expected to be sparse (3–5). The number of greedy search steps in the MVB modeling was set to three for all analyses for computational efficiency and to guard against the possibility of overfitting.

**Decoding Individual Memories.** The first analysis examined whether there was any information about each of the four episodes present within each subfield, despite the high degree of overlap across episodes. To do this, we first set up an appropriate SPM design matrix. We created a single regressor for each individual memory, where recall trials for that memory were modeled with a boxcar function covering the length of the recall period. Movement parameters were included as regressors of no interest. Following this, for each subfield separately we fitted four MVB models, one to each individual memory regressor (i.e., one model for memory A, one regressor for memory B, and so on; Fig. 1). The model fit correlation score was averaged across these four models, creating a single summary measure of information about episodic information. For each subfield we tested for significant differences between the hemispheres and found none. We therefore averaged the information measures across the hemispheres, and all analyses reported here are based on these pooled measures.

To determine whether this information was statistically significant, we used a permutation approach to derive an estimate of the null information. Each memory recall event was randomly assigned to one of four regressors, this time ignoring the memory label (e.g., memory A). The four regressors were matched to the original regressors in terms of the number of memories included in each. An MVB model was fitted to each of these regressors, and the model fit correlation score was averaged across the four. Because each regressor contains a mixture of the four memory types, this new model fit correlation can only reflect information about general retrieval processes rather than information about each individual memory. For each participant and each subfield, this permutation process was repeated 10 times, and the model fit correlation was averaged across these 10 permutations to derive a single estimate of the null information. If a region contains information about each individual memory, then the empirical model fit should be greater than the permuted null model fit. To determine statistical significance at the group level we therefore used a paired  $t$  test to compare the empirical model fit against the null. For the visual cue control analysis, the method was identical to that described above, with the exception that the regressors for each “memory” now modeled the onset of each visual cue as a single event, rather than a boxcar covering the retrieval period.

**Overlapping Neural Representations.** Our key question was whether evidence existed for partially overlapping neural representations, particularly within CA3. If so, this would indicate that pattern separation may have partially failed because of the high similarity across the episodes. Owing to the nature of the design, each memory had another memory that shared spatial background, another memory that shared event content, and one that had no shared elements. For example, memory A shared the background with memory C and shared the event content with memory B (Fig. 1). To investigate the presence of memory overlap during episodic recall, we looked at the relationship between the MVB voxel patterns. If two memories have overlapping neural representations, then the MVB patterns for those two memories should correlate, whereas the MVB patterns for two non-overlapping memories should not. To test this, for each participant and subfield we calculated the Fisher-transformed Pearson correlation between the MVB voxel patterns of each pair of overlapping episodes. These were averaged to create a single summary statistic of the overlap similarity. We then did the same for each pair of nonoverlapping memories. These two sets of similarity scores were compared in each subfield using a paired  $t$  test.

To ensure that this analysis was not confounded by differences in correlation between the psychological variable time courses themselves, we extracted the time course of each memory after convolution with the canonical HRF. We then measured the correlation between the overlapping and nonoverlapping memories. Crucially, the overlapping memories are not significantly more correlated than the nonoverlapping memories. Indeed, we found a significant difference in the opposite direction [ $t(14) = 3.1, P = 0.008$ ], with significantly greater correlation between the nonoverlapping memories. Thus, if anything, this should bias our data in the opposite direction to the results we actually found, suggesting that our results are robust even in the presence of this added noise.

**CA3 Neural Interference.** Given the negative correlation between episodic and overlap information within CA3 (*Results*), we argue that the most appropriate measure of neural interference for each participant is not simply the neural overlap score (as described above). Instead, it should take into account the relative amount of overlap with respect to the amount of unique information present within CA3. We therefore operationalized CA3 neural interference as the relative weighting of overlap information compared with distinct episodic information (the model fit correlation from the individual memory MVB models) for each individual. To calculate this, we first calculated the z-score for each participant for each CA3 variable separately, to normalize the measures. We then subtracted the episodic from the overlap measure for each individual. The resulting score provides a measure of the relative impact of the neural overlap for each individual, which we refer to as CA3 neural interference. A more positive neural interference score indicates a relatively higher impact of neural overlap within CA3.

**ACKNOWLEDGMENTS.** We thank Karl Friston for methodological advice and Demis Hassabis and Catherine Sebastian for useful discussions. E.A.M. is supported by the Wellcome Trust.

- LePort AKR, et al. (2012) Behavioral and neuroanatomical investigation of Highly Superior Autobiographical Memory (HSAM). *Neurobiol Learn Mem* 98(1):78–92.
- Nyberg L, Lövdén M, Riklund K, Lindenberg U, Bäckman L (2012) Memory aging and brain maintenance. *Trends Cogn Sci* 16(5):292–305.
- Marr D (1971) Simple memory: A theory for archicortex. *Philos Trans R Soc Lond B Biol Sci* 262(841):23–81.
- Treves A, Rolls ET (1994) Computational analysis of the role of the hippocampus in memory. *Hippocampus* 4(3):374–391.
- Rolls ET (2010) A computational theory of episodic memory formation in the hippocampus. *Behav Brain Res* 215(2):180–196.
- McClelland JL, McNaughton BL, O'Reilly RC (1995) 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(3):419–457.
- O'Reilly RC, Bhattacharyya R, Howard MD, Ketz N (2011) Complementary Learning Systems. *Cogn Sci*, 10.1111/j.1551-6709.2011.01214.x.
- Leutgeb S, Leutgeb JK, Treves A, Moser M-B, Moser EI (2004) Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305(5688):1295–1298.
- Leutgeb JK, Leutgeb S, Moser M-B, Moser EI (2007) Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* 315(5814):961–966.
- Bakker A, Kirwan CB, Miller M, Stark CEL (2008) Pattern separation in the human hippocampal CA3 and dentate gyrus. *Science* 319(5870):1640–1642.
- Lacy JW, Yassa MA, Stark SM, Muftuler LT, Stark CEL (2011) Distinct pattern separation related transfer functions in human CA3/dentate and CA1 revealed using high-resolution fMRI and variable mnemonic similarity. *Learn Mem* 18(1):15–18.
- Lee I, Yoganarasimha D, Rao G, Knierim JJ (2004) Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature* 430(6998):456–459.
- Vazdarjanova A, Guzowski JF (2004) Differences in hippocampal neuronal population responses to modifications of an environmental context: Evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. *J Neurosci* 24(29):6489–6496.
- Carr VA, Rissman J, Wagner AD (2010) Imaging the human medial temporal lobe with high-resolution fMRI. *Neuron* 65(3):298–308.
- Bonnici HM, et al. (2012) Multi-voxel pattern analysis in human hippocampal subfields. *Front Hum Neurosci* 6:290.
- Bonnici HM, Chadwick MJ, Maguire EA (2013) Representations of recent and remote autobiographical memories in hippocampal subfields. *Hippocampus* 23(10):849–854.
- Chadwick MJ, Hassabis D, Maguire EA (2011) Decoding overlapping memories in the medial temporal lobes using high-resolution fMRI. *Learn Mem* 18(12):742–746.
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10(9):424–430.
- Haynes J-D, Rees G (2006) Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7(7):523–534.
- Chadwick MJ, Bonnici HM, Maguire EA (2012) Decoding information in the human hippocampus: A user's guide. *Neuropsychologia* 50(13):3107–3121.
- Friston K, et al. (2008) Bayesian decoding of brain images. *Neuroimage* 39(1):181–205.
- Morcom AM, Friston KJ (2012) Decoding episodic memory in ageing: A Bayesian analysis of activity patterns predicting memory. *Neuroimage* 59(2):1772–1782.
- FitzGerald THB, Friston KJ, Dolan RJ (2012) Action-specific value signals in reward-related regions of the human brain. *J Neurosci* 32(46):16417–1623a.
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Hum Brain Mapp* 15(1):1–25.
- Welvaert M, Durnez J, Moerkerke B, Verdoelae G, Rosseel Y (2011) NeuRosim: An R package for generating fMRI data. *J Stat Softw* 44(10):1–18.

