Neural fatigue influences memory encoding in the human hippocampus

Lynn J. Lohnas[a,1], Lila Davachi[a,b,2], Michael J. Kahana[c]
[a]Department of Psychology, [b]Center for Neural Science, New York University;
6 Washington Place, 8th floor, New York, NY 10003
[c]Department of Psychology, University of Pennsylvania;
433 South University Avenue, Philadelphia, PA 19104

[1]Present address: Department of Psychology, Syracuse University, 430 Huntington Hall, Syracuse, NY 13244
[2]Present address: Department of Psychology, Columbia University, 371 Schermerhorn Hall, 1190 Amsterdam Avenue, New York, NY 10027

Corresponding author:
Lynn J. Lohnas
Department of Psychology
430 Huntington Hall
Syracuse, NY 13244
ljlohnas@syr.edu

Abstract

Here we examine the variability underlying successful memory encoding. Successful encoding of successive study items may fatigue encoding resources, thus decreasing the ability to encode subsequent items (Tulving and Rosenbaum, 2006); alternatively, successful encoding may be persistent, leading to more successful encoding (Kahana, Aggarwal, and Phan, 2018). Analyzing intracranial electroencephalographic activity while subjects studied lists of words for subsequent free recall, we examined high-frequency activity (HFA) in hippocampus and dorsolateral prefrontal cortex (DLPFC), as HFA was greater for subsequently recalled than non-recalled items in these regions. We compared non-recalled items with good encoding history (i.e. one of the two preceding items was recalled) with non-recalled items with poor encoding history (i.e. neither prior item was recalled). In the hippocampus, good encoding history led to reduced HFA. By contrast, in DLPFC good encoding history led to enhanced HFA. The hippocampal results appear consistent with neural fatigue, whereas the DLPFC results appear consistent with persistent encoding states.

Keywords: memory; episodic memory; iEEG; temporal lobe; encoding

1. Introduction

The ability to measure physiological activity in the human brain as people study and subsequently attempt to retrieve memoranda has uncovered a diverse network of regions whose activity predicts encoding or retrieval success. To derive biomarkers of successful encoding, researchers have compared physiological activity recorded during the study of items that are subsequently remembered to activity recorded during the study of items that are subsequently forgotten. This subsequent memory analysis has revealed increased hemodynamic and electrophysiological activity in a core network of brain regions, including the hippocampus and medial temporal lobe structures as well as in dorsolateral prefrontal cortex (DLPFC; Brewer et al., 1998; Davachi, 2006; Diana et al., 2007; Hanslmayr and Staudigl, 2013; Kim, 2011; Paller and Wagner, 2002; Rugg et al., 2012; Sederberg et al., 2007; Wagner et al., 1998). These signals are thought to reflect variability in the goodness of memory encoding arising not only from item properties but also from endogenous variation in the neurocognitive processes underlying successful memory storage (Kahana et al., 2018).

Variability is a ubiquitous feature of any complex dynamical system. However, the nature of this variability and the mechanisms that give rise to it can be characterized and explained, and such a characterization would provide further insight into the underlying system. For example, variability could be described by a stochastic process, with goodness of memory encoding starting at some value and rising or falling unpredictably over the course of a list. Such a process may by described by the degree to which it is persistent, with goodness of memory processes either drifting slowly around some mean value (i.e., having a high autocorrelation), or jumping unpredictably from item to item. In this

formulation, the rise or fall in memory efficiency is modeled as a random shock that is independent of the recent history of encoding efficiency.

An alternative account of variability in encoding goodness comes from the idea of "neural fatigue". According to this view, the neural processes that support effective encoding cannot be sustained indefinitely but rather depend on a latent resource that depletes during sustained periods of high activity and replenishes over periods of low activity. Tulving and Rosenbaum (2006) advanced this idea as an explanation for the well-known "law of primacy" which describes better memory for the initial items experienced within a given context (Murdock, 1962; Spurgeon et al., 2014). Neural fatigue predicts that the encoding of a given item should be influenced by the past history of successful or unsuccessful encoding. A biomarker of the neural activity that is required to sustain good encoding should thus decline following a period of sustained successful encoding. Behaviorally, this may also lead to items being less likely to be successfully encoded.

In this study we attempt to distinguish between neural correlates of encoding goodness that are transitory, as predicted by the standard autoregressive model, and those that specifically reflect changes associated with the recent past history of encoding success, as predicted by a neural fatigue model.

2. Materials and methods

2.1 Subjects

We present novel analyses of intracranial electrophysiological recordings taken from patients undergoing invasive monitoring as treatment for drug-resistant epilepsy. Recordings from subdural grids and intraparenchymal depth electrodes were taken from patients who volunteered to participate in memory studies during their 1-3 week

hospitalization. The clinical team determined the placement of electrodes entirely for purposes of seizure localization. Both the electrophysiological and the behavioral data were collected in a multi-center study from 2000-2017. Although portions of this dataset have been reported on previously (e.g., Burke et al., 2014; Long et al., 2017; Merkow et al., 2015), all of the analyses and results described here are novel. Data used in this report may be freely obtained from the cognitive electrophysiology portal at the University of Pennsylvania (http://memory.psych.upenn.edu/Electrophysiological_Data).

For the purpose of the present study we selected patients who had electrodes in at least one of our two regions of interest (ROIs): hippocampus and DLPFC. The clinical team determined the placement of electrodes entirely for purposes of seizure localization. In total, data were included from 223 subjects, each of whom contributed 1–8 testing sessions (99 contributed 1 session). Of these subjects, 131 had electrodes in hippocampus and 163 had electrodes in DLPFC. Our research protocol was approved by the institutional review boards at the University of Pennsylvania and our collaborating hospitals, and was carried out in accordance with the provisions of the World Medical Association Declaration of Helsinki. Informed consent was obtained from patients (or their guardians, in the case of teenage subjects). Our motivation for the sample size was simply to use every subject available from this multicenter study, producing large sample sizes for both subjects with hippocampal electrodes (N=131) and DLPFC electrodes (N=163).

2.2 Procedure

Subjects studied lists of common nouns chosen at random and without replacement from a pool of high-frequency words. Following an orienting stimulus (+ sign) the computer displayed each list item in capital letters for 1600 ms, followed by an 800 to 1200

ms blank ISI. The variation in the duration of the ISI served to decorrelate the physiological responses from successive word presentations (Sederberg et al., 2007). Following the final list item, subjects were shown a series of arithmetic problems of the form A+B+C=?, where A, B, and C were randomly chosen digits in the set {1,...,9}. Subjects responded by typing the answer on a computer keyboard. Immediate feedback was given in the form of a high-pitched tone for correct entries and a low-pitched tone for incorrect answers. After performing this distractor task for 20 s a row of asterisks accompanied by a 300 ms tone signaled the start of the recall period. Subjects were given 45 s to recall items aloud from the current list in any order. Fourteen subjects received 20-item lists; 145 received 12-items lists; 64 received 15-item lists. Vocal responses were digitally recorded and scored for analysis following each session (Solway et al., 2010).

2.3 Electrophysiology

Intracranial electroencephalography (EEG) activity was recorded from subdural arrays (grids of 3 mm diameter contacts spaced 1 cm apart) or depth probes with 1 mm collar electrodes spaced 8 mm apart. EEG recordings were sampled at 256 – 1KHz depending on the clinical center. Physiological and behavioral measures were synchronized using optically isolated pulses received from the testing computer on an additional recording channel. Signals were converted to a bipolar montage by differencing the signals between each pair of immediately adjacent electrodes on grids and depth electrodes (Burke et al., 2013). The resulting bipolar signals were treated as new virtual electrodes and are referred to as such in the remainder of the text.

Images with the electrode placements on individual patient brains were created by coregistering a computed tomography (CT) scan with a preoperative magnetic resonance

image (MRI). The electrodes were manually identified using the postoperative CT scans. These images were then normalized to a standardized brain in MNI space (Maldjian et al., Burdette, 2003). The MNI coordinates were then transformed to Tailarach space (Lancaster et al., 2000), and Tailarach coordinates were used to determine the side of each electrode as well as the Brodmann areas of cortical electrodes. Cortical electrodes in Brodmann areas 9 and 46 defined the DLPFC region of interest. For depth electrodes implanted in the temporal lobe, an experienced clinician reviewing CT scans and MRIs labeled anatomic locations.

Electrodes were positioned by clinical teams to identify seizure foci and functional brain regions to guide potential resective surgery. As a result, most electrodes were usually placed in temporal cortex, but many electrodes were also placed in the hippocampi and frontal, occipital, and parietal cortices. Because the clinical procedure of identifying seizure foci entails placing electrodes in any region that is potentially epileptogenic, the majority of recordings come from brain regions outside the area that is eventually determined to be involved in seizures (Jacobs and Kahana, 2010).

2.4 Data analysis

To eliminate electrical line and equipment noise, data were notch-filtered on-line at 50 or 60 Hz with a Butterworth filter with zero phase distortion. Power was calculated during each word presentation (0-1600 ms post-onset) with a 1000 ms buffer on either side. The Morlet wavelet transform with a wave number of 6 was used to compute spectral power as a function of time. Following previous studies, we defined wavelets at 46 log-spaced frequencies between 2-100 Hz, but here only report results for HFA, 44-100 Hz (Long et al., 2014; Long and Kahana, 2015). After calculating power values at each

electrode and frequency value, we then log transformed the power values. Next, we Z-scored the values across all events and timepoints within a session, separately at each frequency and electrode. Finally, we took the mean across time points and frequency values contained within the HFA range.

To ensure reliable statistics from each subject, we required that each subject contribute at least 10 observations per session for each of the four behavioral conditions of interest: good encoding versus poor encoding (i.e. whether an item was successfully recalled), each considered with good versus poor encoding history (i.e. whether at least one of the prior two items was recalled). In addition, any lists with no remembered items were excluded from analysis. These criteria led to the exclusion of 19 subjects. Requiring a minimum number of observations per session also led to the exclusion of individual sessions for individual subjects: 33 subjects had 1 session excluded, 6 had 2 sessions excluded, 3 had 3 sessions excluded, and 1 subject had 6 sessions excluded. In the behavioral analysis, we include the same sessions that were included for subjects, but nonetheless include recall from lists with no recalls, to highlight that memory performance was not at floor.

To determine significance for a pairwise comparison, for each subject and electrode, a t-statistic was generated through an unpaired t-test comparing the Z-scored HFA power between the two conditions. These t-statistics were averaged across electrodes within an ROI, creating a single t-statistic for each subject and ROI. The distribution of subject t-statistics was compared to zero using an unpaired t-test. To determine significance between pairwise comparisons, we extracted the distribution of t-statistics from each comparison, then compared them using a paired t-test.

3. Results

According to a standard view of variability in encoding efficacy, periods of good encoding will tend to be persistent, with encoding states that lead to successful subsequent recall tending to be followed by further good encoding states. In contrast, according to a neural fatigue account, extended periods of good encoding will tend to deplete neural resources thus making subsequent epochs more likely to be poorly encoded. To evaluate these hypotheses, we examined HFA in the local-field potential across two regions of interest within the core verbal memory network: hippocampus and DLPFC. We chose to examine HFA because of its strong correlation with the firing rates of individual neurons in both human and non-human animals (Hirabayashi et al., 2014; Manning et al., 2009). Furthermore, hippocampus and prefrontal cortical HFA have been established as a biomarker of good memory encoding in humans (Long et al., 2014, Long and Kahana, 2015).

Here we compared HFA as subjects encoded lists of items for subsequent free recall (Figure 1a,b). We found that HFA was significantly greater during encoding of items that were subsequently recalled than items that were not recalled (Figure 1c), both in the hippocampus ($t(130) = 4.74$, $p < .001$) and in DLPFC ($t(162) = 4.10$, $p < .001$). We thus interpret greater HFA in these regions as a biomarker of good memory encoding.

Having established HFA in hippocampus and DLFPC as biomarkers of good memory encoding, we next sought to characterize how these biomarkers vary over time during

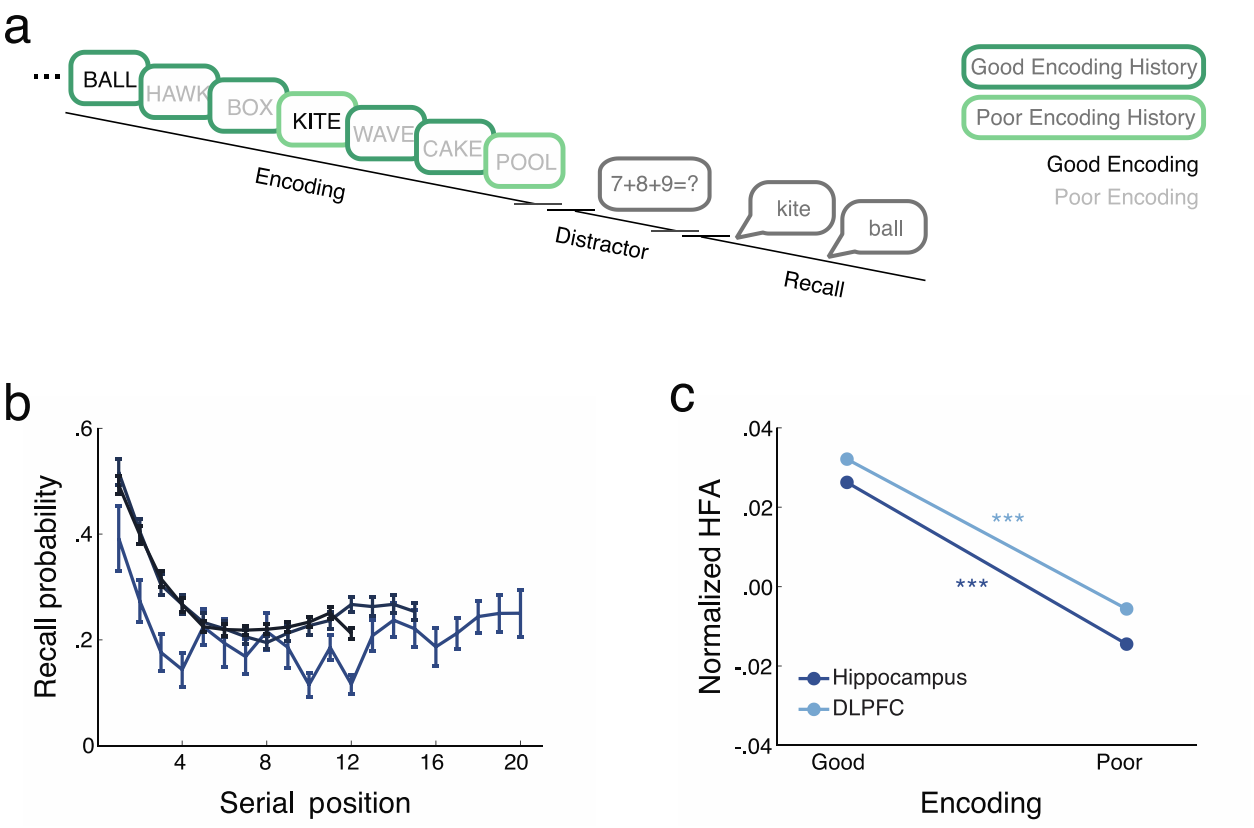

Figure 1. Experiment design, behavioral performance, and high frequency activity (HFA) as a biomarker of good encoding. [*Two columns*]

A. Subjects studied and then performed delayed free recall of lists of common nouns. Each item was retroactively labeled according to two variables: 1) successful encoding, based on whether the item was subsequently recalled (black for good and light gray for poor); 2) successful recent encoding history, based on whether either of the two preceding items were recalled (dark green for good encoding history and light green for poor encoding history).

B. Probability of recall as a function of serial position, divided based on the subject's list-length: 12 (N=145), 15 (N=64), 20 (n=14).

C. HFA in hippocampus and dorsolateral prefrontal cortex (DLPFC) HFA is a biomarker of good encoding, being significantly greater for items that were subsequently recalled (good encoding) in comparison to items that were subsequently not recalled (poor encoding). Asterisks indicate significance between groups, with the color corresponding to the region referenced in the figure legend (***p < .001).

encoding. For each item we defined two variables: whether it was subsequently recalled and whether at least one of the prior two words was subsequently recalled (Figure 1a). The former measure is a surrogate for successful current encoding, while the latter measure is a surrogate for successful encoding history. These two variables enable us to distinguish between an autocorrelated account and a neural fatigue account of encoding success, as these hypotheses make a different prediction regarding items with good encoding history yet poor encoding: According to the neural fatigue account, this reflects a depletion in neural resources, thus preventing encoding of the current item, an autocorrelated account would assume that this simply reflects fluctuations in good or poor encoding states. Thus, the neural fatigue account predicts that a good encoding history should reduce HFA for non-recalled items. In contrast, a strong autocorrelated process predicts that HFA for non-recalled items should be greater with a good encoding history, as these items with good encoding should have higher HFA. Figure 2 presents representative signals recorded during individual word lists. In these example lists, filled markers denote subsequently remembered items whereas open markers denote subsequently forgotten items; larger markers denote good encoding history whereas smaller markers denote poor encoding history. In the hippocampal electrode examples, when there is a transition from a good to poor encoding state (such that an item has poor current encoding but is preceded by a good encoding history), this coincides with a sharp decrease in HFA (green arrows). This is consistent with a neural fatigue account, whereby a good encoding state, reflected both by successful subsequent recall and greater HFA, cannot be maintained for extended periods of time. In contrast, fluctuations in DLPFC HFA from good to poor encoding states are not as drastic, and thus are more consistent with an autocorrelated process.

Going beyond single examples, we next quantified the extent of the changes in HFA for non-recalled items as a function of their encoding history. In the hippocampus, HFA for non-recalled items was significantly lower for items with a good than a poor encoding history (Figure 3; $t(130) = 3.18$, $p = .002$). This is consistent with the neural fatigue hypothesis prediction that a recent history of successful encoding should deplete cognitive resources, thus leading the current item to be in a poor encoding state. As such, the item will not be recalled, and hippocampal HFA will be lower. In contrast, HFA in DLPFC trended towards being significantly greater for good vs. poor encoding history ($t(162) = 1.90$, $p = .059$). Whether this difference is interpreted as reliable, it is nonetheless most consistent with an autocorrelated account. If the autocorrelation was strong, then the prior encoding state should persist into the current encoding state. As a result, even if an item were not recalled, it would nonetheless exhibit greater HFA if it were preceded by a good encoding history than a poor encoding history. In the case of weaker autocorrelation, where the random shocks are highly variable, encoding history may have little influence on HFA of the current item, and thus HFA would not differ with encoding history. The results in DLPFC may be consistent with a weaker form of the autocorrelated account, but regardless is qualitatively inconsistent with the neural fatigue account.

To further query HFA differences in encoding history between hippocampus and DLPFC, we asked whether there was an interaction between these regions. As a conservative estimate of this interaction, we only considered those subjects who had electrodes in both hippocampus and DLFPC. In this subset of 71 subjects, the main effects by region were still present albeit weaker, with greater hippocampal HFA for items with a poor than good encoding history ($t(70)=3.34$, $p = .001$), and in DLPFC there was a

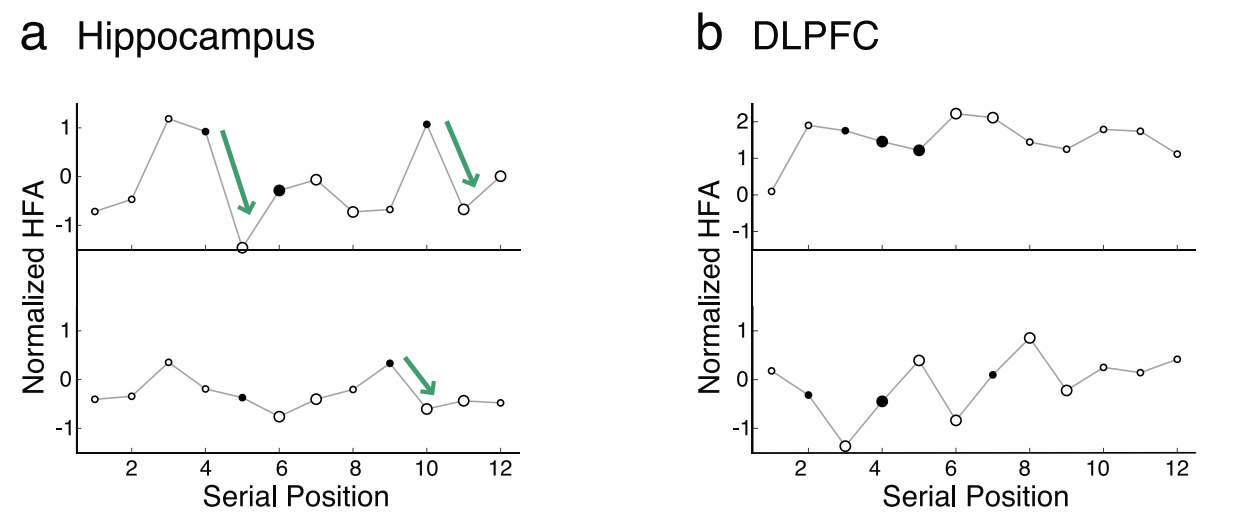

Figure 2. Modulations of high-frequency activity (HFA) by encoding state. [*Two columns*]
A. Examples of HFA in single hippocampal electrodes exhibiting activity consistent with a neural fatigue process in a single list. The top panel is from an electrode from left hippocampus, and the bottom panel is from an electrode in right hippocampus.
B. Examples of HFA in a single DLPFC electrodes exhibiting activity consistent with an autocorrelated process in a single list are shown in (d). The top panel is an electrode from right Brodmann Area (BA) 9, and the bottom panel is an electrode from left BA9.

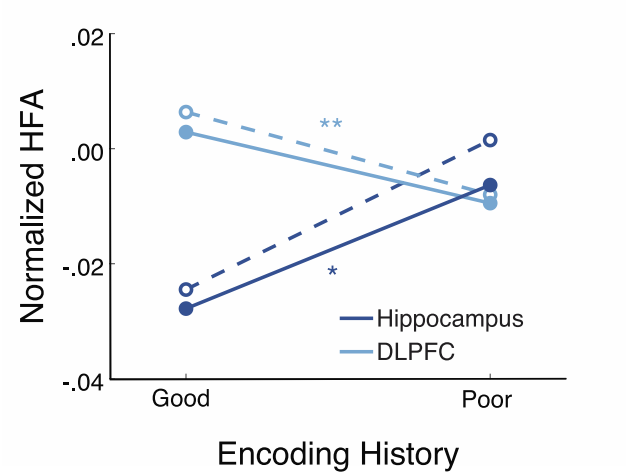

Figure 3. Modulations of high-frequency activity (HFA) by encoding history for nonrecalled items. [*One column*]
In the hippocampus, HFA is significantly greater for items with poor encoding history than good encoding history (N=131). The dorsolateral prefrontal cortex (DLPFC) shows a trend toward greater HFA for items with good than poor encoding history (N=163). Solid lines and filled circles indicate values across all subjects who had electrodes in either region. Dashed lines and open circles indicate values in the same regions when calculated for the subset of subjects who had electrodes in both regions (N=71). Asterisks indicate significance between conditions for each brain region, for all subjects, with the color corresponding to the region referenced in the figure legend (*p<.06, **p < .005).

trend towards greater HFA for items with a good encoding history ($t(70) = 1.52$, $p=.13$). Critically, there was a significant interaction between regions ($t(70)=3.31$, $p = .001$), due to greater HFA for nonrecalled items with a poor encoding history in hippocampus but not in DLPFC. Taken together, these results suggest that hippocampal HFA is more consistent with the neural fatigue account and DLPFC is more consistent with autocorrelated encoding states.

4. Discussion

Our ability to form new memories varies over time, with periods of good encoding being interrupted by periods of poor encoding. By recording neural activity during the learning process we can identify biomarkers of this variability. Research using both non-invasive and invasive measures of human brain activity have identified reliable biomarkers of good memory encoding (Davachi, 2006; Paller and Wagner, 2002; Sederberg et al., 2003; Wagner et al., 1998). Intracranial EEG studies in particular have shown that in both hippocampus and DLPFC, HFA increases during the encoding of words that will be subsequently remembered as compared with those that are forgotten (Long et al., 2014, Long and Kahana, 2015). Although there are several possible mechanisms that could give rise to these dynamics, we sought to test the hypothesis that sustained successful memory encoding will induce fatigue in the core memory network, and that this fatigue will appear as a subsequent drop in the good encoding biomarkers. An alternative to this neural fatigue hypothesis is the idea that goodness of encoding varies stochastically, being equally likely to rise or fall independent of the recent history of memory encoding. This prediction would arise if the biomarker followed a standard autoregressive process.

To test the neural fatigue hypothesis, we analyzed hippocampal (N=131) and DLPFC

(N=163) HFA measured as neurosurgical patients studied lists of items for a subsequent recall test. We classified an item as having good encoding history if either one or both of the two preceding items was successfully recalled. In general, we expected to observe higher HFA for recalled than for non-recalled items. If encoding resources are limited then HFA is more likely to decline than it is to rise following a sustained period of good encoding, and this would result in a failure to recall the subsequent item. As such, a neural fatigue account would predict that non-recalled items with good encoding history would have lower HFA than those with a poor encoding history. We found that the hippocampal HFA dynamics were consistent with this hypothesis, such that HFA for non-recalled items was significantly greater for items with a poor than a good encoding history. In contrast, HFA in the DLPFC was not consistent with this hypothesis, but rather was more consistent with an autocorrelated account: non-recalled items with a poor encoding history, like the items that preceded them, had lower HFA in comparison to items with a good encoding history.

We considered predictions of a neural fatigue account outlined in Tulving and Rosenbaum (2006). According to this account, the neural ensemble engaged during an item's first presentation loses its efficacy when an item or components of that item are repeated. Consistent with the neural fatigue hypothesis, Serruya et al. (2014) found that HFA in both DLPFC and lateral temporal cortex reliably decreased over the course of list presentation, demonstrating a strong correlation with the primacy effect. Merzagora et al. (2014) examined HFA during the encoding of repeated items in a Sternberg short-term item recognition task. They found that HFA during the second occurrence of an item was markedly attenuated, consistent with our findings that decreases in HFA may reflect a depletion of resources to devote to a particular stimulus.

Tulving's neural fatigue hypothesis also provides an explanation for the build-up and release of proactive interference exhibited in lists of items from the same semantic category (Loess, 1967; Wickens, 1970); items drawn from the same category share many stimulus features, thus imposing similar demands on encoding and fatiguing encoding processes. Although the current study did not involve item or category repetition, all of the items within a list share contextual features representing their overlapping temporal attributes and encoding task context (Manning et al., 2011; Polyn et al., 2009). Further, subjects may automatically segment subsequences of items into meaningful "chunks" or events (Clewett and Davachi, 2017; Farrell, 2012; Zacks et al., 2001). Such structure, whether endogenously created or exogenously imposed, has been shown to modulate both encoding success (Ezzyat and Davachi, 2011; Heusser et al., 2018; Kurby and Zacks, 2008; Speer and Zacks, 2005) as well as hippocampal activity (DuBrow and Davachi, 2014; Ezzyat and Davachi, 2011, 2014). These findings provide support for our results of modulations of hippocampal activity by encoding success.

Our findings are also consistent with evidence suggesting that, during successful encoding, hippocampal activity is more likely to be in a good attentional state (Aly and Turk-Browne, 2016; Uncapher and Rugg, 2009). In a complementary way, several studies have found that hippocampal activity is associated with guiding attention more efficiently based on previously encoded information (Chun et al., 2011; Goldfarb et al., 2016; Summerfield et al., 2006). Taken together, this suggests that hippocampal HFA not only reflects the current encoding state but also reflects the influence of prior encoding states on the current encoding state.

Measures of neural activity in the core neural memory network provide a window

into the cognitive processes underlying successful encoding. By recording intracranial activity during memory encoding for subsequent free recall, we found that brain regions established as being modulated by encoding success are also modulated by a recent history of encoding success. Although both DLPFC and hippocampus exhibit standard increases in HFA that mark successful encoding, the hippocampus exhibits a unique signature consistent with Tulving's neural fatigue hypothesis. By defining an item's recent encoding history as at least one of the prior two items being recalled, hippocampal HFA of non-recalled items with a good encoding history was significantly lower than HFA for non-recalled items with a poor encoding history. According to the neural fatigue account, sustained periods of successful encoding and greater hippocampal HFA eventually fatigue local neural networks, thus leading to poor encoding and lower hippocampal HFA. These results provide the first neural evidence that the biomarkers of good memory encoding exhibit fatigue following successful encoding of prior items.

Acknowledgments

We thank Blackrock Microsystems for providing neural recording and stimulation equipment. This work was supported by the DARPA Restoring Active Memory (RAM) program [Cooperative Agreement N66001-14-2-4032]; and National Institutes of Health [grant numbers MH106266, MH55687]. The views, opinions, and/or findings contained in this material are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

Declarations of interest: none.

Competing interests

The authors have no competing interests to declare.

References

Aly, M., Turk-Browne, N. B. (2016). Attention promotes episodic encoding by stabilizing hippocampal representations. Proceedings of the National Academy of Sciences.

Brewer, J. B., Zhao, Z., Desmond, J. E., Glover, G. H., Gabrieli, J. D. (1998). Making memories: Brain activity that predicts how well visual experience will be remembered. Science, 281, 1185-7.

Burke, J. F., Long, N. M., Zaghloul, K. A., Sharan, A. D., Sperling, M. R., Kahana, M. J. (2014). Human intracranial high-frequency activity maps episodic memory formation in space and time. NeuroImage, 85 , 834–843.

Burke, J. F., Zaghloul, K. A., Jacobs, J., Williams, R. B., Sperling, M. R., Sharan, A. D., Kahana, M. J. (2013). Synchronous and asynchronous theta and gamma activity during episodic memory formation. Journal of Neuroscience, 33, 292–304.

Chun, M. M., Golomb, J. D., Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. Annual Review of Psychology, 62, 73-101.

Clewett, D., Davachi L. (2017). The ebb and flow of experience determines the temporal structure of memory. Current Opinion in Behavioral Sciences, 17, 186-193.

Cohen, J. (1988). Statistical power analysis for the behavioral sciences (2nd ed.) Hillsdale, NJ: Earlbaum.

Davachi, L. (2006). Item, context and relational episodic encoding in humans. Current Opinion in Neurobiology, 16, 693—700.

Diana, R. A., Yonelinas, A. P., Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: a three-component model. Trends in Cognitive Sciences, 11, 379–386.

DuBrow, S., Davachi, L. (2014). Temporal memory is shaped by encoding stability and intervening item reactivation. The Journal of Neuroscience, 34, 13998–14005.

Ezzyat, Y., Davachi, L. (2011). What constitutes an episode in episodic memory? Psychological Science, 22, 243-252.

Ezzyat, Y., Davachi, L. (2014). Similarity breeds proximity: Pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. Neuron, 81, 1179–1189.

Farrell, S. (2012). Temporal clustering and sequencing in short-term memory and episodic memory. Psychological Review, 119, 223-71

Goldfarb, E. V., Chun, M. M., Phelps, E. A. (2016). Memory-guided attention: Independent contributions of the hippocampus and striatum. Neuron, 89, 1-8.

Heusser, A. C., Ezzyat, Y., Shiff, I., Davachi, L. (2018). Perceptual boundaries cause trade-offs between local and boundary processing and across-trial associative binding. Journal of Experimental Psychology: Learning, Memory, and Cognition. 44(7), 1075-1090.

Hirabayashi, T., Tamura, K., Takeuchi, D., Takeda, M., Koyano, K., Miyashita, Y. (2014). Distinct neuronal interactions in anterior inferotemporal areas of macaque monkeys during retrieval of object association memory. Journal of Neuroscience, 34, 9377-9388.

Jacobs, J., Kahana, M. J. (2010). Direct brain recordings fuel advances in cognitive electrophysiology. Trends in Cognitive Sciences, 14(4), 162-171.

Kahana, M. J., Aggarwal, E. V., Phan, T. D. (2018). The variability puzzle in human memory. Journal of Experimental Psychology: Learning, Memory and Cognition (44)12, 1857-1863.

Kim, H. (2011). Neural activity that predicts subsequent memory and forgetting: a meta-analysis of 74 fMRI studies. NeuroImage, 54, 2446–2461.

Kurby, C. A., Zacks, J. M. (2008). Segmentation in the perception and memory of events. Trends in Cognitive Sciences, 12.

Hanslmayr, S., Staudigl, T. (2013). How brain oscillations form memories – A processing based perspective on oscillatory subsequent memory effects. Neuroimage, 85, 648-655.

Lancaster, J. L., Woldorff, M. G., Parsons, L. M., Liotti, M., Freitas, C. S., Rainey, L., Kochunov, P. V., Nickerson, D., Mikiten, S. A., Fox, P. T. (2000). Automated Talairach atlas labels for functional brain mapping. Hum Brain Mapp, 10, 120–131.

Loess, H. (1967). Short-term memory, word class, and sequence of items. Journal of Experimental Psychology, 74, 556-561.

Long, N. M., Burke, J. F., Kahana, M. J. (2014). Subsequent memory effect in intracranial and scalp EEG. NeuroImage, 84, 488–494.

Long, N. M., Kahana, M. J. (2015). Successful memory formation is driven by contextual encoding in the core memory network. NeuroImage, 119, 332-337.

Long, N. M., Sperling, M. R., Worrell, G. A., Davis, K. A., Gross, R. E., Lega, B. C., Jobst, B. A., Sheth, S. A., Zaghloul, K., Stein, J. M., Kahana, M. J. (2017). Contextually mediated spontaneous retrieval is specific to the hippocampus. Current Biology, 27, 1-6.

Maldjian, J. A., Laurienti, P. J., Kraft, R. A., Burdette, J. H. (2003). An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. Neuroimage, 19, 1233–1239.

Manning, J. R., Jacobs, J., Fried, I., Kahana, M. J. (2009). Broadband shifts in local field

potential power spectra are correlated with single-neuron spiking in humans. Journal of Neuroscience, 29, 13613–13620.

Manning, J. R., Polyn, S. M., Baltuch, G., Litt, B., Kahana, M. J. (2011). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. Proceedings of the National Academy of Sciences, USA, 108, 12893–12897.

Merkow, M. B., Burke, J. F., Kahana, M. J. (2015). The human hippocampus contributes to both the recollection and familiarity components of recognition memory. Proceedings of the National Academy of Sciences, 112, 14378–14383.

Merzagora, A. R., Coffey, T. J., Sperling, M. R., Sharan, A., Litt, B., Baltuch, G., Jacobs, J. (2014). Repeated stimuli elicit diminished high-gamma electrocorticographic responses. NeuroImage, 85, 844–852.

Murdock, B. B., Jr. (1962). The serial position effect of free recall. Journal of Experimental Psychology, 64 (5), 482-488.

Paller, K. A., Wagner, A. D. (2002). Observing the transformation of experience into memory. Trends in Cognitive Sciences, 6, 93-102.

Polyn, S. M., Norman, K. A., Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. Psychological Review, 116, 129-56.

Rugg, M. D., Vilberg, K. L., Mattson, J. T., Yu, S. S., Johnson, J. D., Suzuki, M. (2012). Item memory, context memory and the hippocampus: fMRI evidence. Neuropsychologia, 50, 3070-3079.

Sederberg, P. B., Kahana, M. J., Howard, M. W., Donner, E. J., Madsen, J. R. (2003). Theta and gamma oscillations during encoding predict subsequent recall. Journal of Neuroscience, 23, 10809–10814.

Sederberg, P. B., Schulze-Bonhage, A., Madsen, J. R., Bromfield, E. B., McCarthy, D. C., Brandt, A., . . . Kahana, M. J. (2007). Hippocampal and neocortical gamma oscillations predict memory formation in humans. Cerebral Cortex, 17, 1190–1196.

Serruya, M. D., Sederberg, P. B., Kahana, M. J. (2014). Power shifts track serial position and modulate encoding in human episodic memory. Cerebral Cortex , 24 , 403–413.

Solway, A., Geller, A. S., Sederberg, P. B., Kahana, M. J. (2010). Pyparse: A semi-automated system for scoring spoken recall data. Behavior Research Methods, 42, 141-147.

Speer, N. K., Zacks, J. M. (2005). Temporal changes as event boundaries: Processing and memory consequences of narrative time shifts. Journal of Memory and Language, 53, 125–140.

Spurgeon, J., Ward, G., Matthews, W. J. (2014). Why do participants initiate free recall of short lists of words with the first list item? Toward a general episodic memory explanation. Journal of Experimental Psychology: Learning, Memory, and Cognition, 40(6), 1551-1567.

Summerfield, J. J., Lepsien, J., Gitelman, D. R., Mesulam, M. M., Nobre, A. C. (2006). Orienting attention based on long-term memory experience. Neuron, 49, 905-916.

Tulving, E., Rosenbaum, R. S., (2006). What do explanations of the distinctiveness effect need to explain? R.R. Hunt, J.B. Worthen (Eds.), Distinctiveness and memory, Oxford University Press, New York (2006), pp. 407-423

Uncapher, M. R., Rugg, M. D. (2009). Selecting for memory? the influence of selective attention on the mnemonic binding of contextual information. Journal of Neuroscience, 29, 8270-8279.

Wagner, A. D., Schacter, D. L., Rotte, M., Koutstaal, W., Maril, A., Dale, A. M., Koutstaal, W.,

Maril, A., Dale, A. M., Rosen, B. R., Buckner, R. L. (1998). Building memories: remembering and forgetting of verbal experiences as predicted by brain activity. Science, 281, 1188-1191.

Wickens, D. D. (1970). Encoding categories of words: An empirical approach to meaning. Psychological Review, 77, 1-15.

Zacks, J. M., Braver, T. S., Sheridan, M. A., Donaldson, D. I., Snyder, A. Z., Ollinger, J. M., Buckner, R. L., Raichle, M. E. (2001). Human brain activity time-locked to perceptual event boundaries. Nature, 4, 651-655.