

Reactivation strength during cued recall is modulated by graph distance within cognitive maps

Simon Kern^{*1,2,3}, Juliane Nagel^{1,2,3}, Martin F. Gerchen^{1,7,8}, Cagatay Guersoy^{1,2,3}, Andreas Meyer-Lindenberg^{2,8}, Peter Kirsch^{1,7,8} Raymond J. Dolan^{4,5}, Steffen Gais⁶, Gordon B. Feld^{*1,2,3,7}

1 Clinical Psychology, Central Institute of Mental Health, Medical Faculty Mannheim, University of Heidelberg, Mannheim, Germany

2 Psychiatry and Psychotherapy, Central Institute of Mental Health, Medical Faculty Mannheim, University of Heidelberg, Mannheim, Germany

3 Addiction Behavior and Addiction Medicine, Central Institute of Mental Health, Medical Faculty Mannheim, University of Heidelberg, Mannheim, Germany

4 Max Planck UCL Centre for Computational Psychiatry and Ageing Research, London, UK

5 Wellcome Centre for Human Neuroimaging, University College London, London, UK

6 Institute of Medical Psychology and Behavioral Neurobiology, Eberhard-Karls-University Tübingen, Germany

7 Department of Psychology, Ruprecht Karl University of Heidelberg, Germany

8 Bernstein Center for Computational Neuroscience Heidelberg/Mannheim, Mannheim, Germany

* Correspondence: Email: gordon.feld@zi-mannheim.de and simon.kern@zi-mannheim.de

Central Institute of Mental Health, J5, 68159 Mannheim, Germany, Tel. +49 621 1703 6540, Fax +49 621 1703 6505.

Abstract

Declarative memory retrieval is thought to involve reinstatement of the neuronal activity patterns elicited and encoded during a prior learning episode. Recently, it has been suggested that two mechanisms operate during reinstatement, dependent on task demands: individual memory items can be reactivated simultaneously as a clustered occurrence or, alternatively, replayed sequentially as temporally separate instances. In the current study, participants learned associations between images that were embedded in a directed graph network and retained over a brief 8-minute consolidation period. During a subsequent cued recall session, participants retrieved the learned information while undergoing magnetoencephalographic (MEG) recording. Using a trained stimulus decoder, we found evidence for clustered reactivation of learned material. Reactivation strength of individual items during clustered reactivation decreased as a function of increasing graph distance, an ordering present solely for successful retrieval but not with retrieval failure. In line with previous research, we found evidence that sequential replay was dependent on retrieval performance and limited to low performers. The results provide further evidence for the existence of different performance-dependent retrieval mechanisms suggesting graded clustered reactivation as a plausible mechanism to search within abstract cognitive maps.

Introduction

Memory relies on three distinct stages: encoding (learning), consolidation (strengthening and transforming) and retrieval (reinstating) of information. New episodic memories are learned by encoding a representation, thought to be realized in a specific spatio-temporal neuronal firing pattern in hippocampal and neocortical networks (Frank et al., 2000; Preston & Eichenbaum, 2013). These firing patterns are reactivated during later rest or sleep, sometimes in fast sequential sequences, a process linked to memory consolidation (Born & Wilhelm, 2012; Feld & Born, 2017). Similarly, during retrieval, the same firing patterns seen during encoding are replayed in a manner that predicts retrieval success (Carr et al., 2011; Foster, 2017). Even though replay has been studied most intensely with respect to the hippocampus, a replay of memory traces in temporal succession is suggested as a general mechanism for planning, consolidation, and retrieval (Buhry et al., 2011). While a rich body of evidence exists in rodents (Ambrose et al., 2016; Chen & Wilson, 2023; Foster & Knierim, 2012; Ólafsdóttir et al., 2018), the contributions of replay to memory storage and retrieval in humans are only beginning to be understood (Brunec & Momennejad, 2022; Eichenlaub et al., 2020; Fuentemilla et al., 2010; Wimmer et al., 2020).

Heretofore, one obstacle has been the difficulty in measuring sequential replay or general network reactivation in humans (NB here we follow the definition of Genzel et al., 2020, where *reactivation* is used as an umbrella term for any form of reoccurrence of a previously encoded neural pattern related to information-encoding, and *replay* refers to reactivation events with a temporally sequential nature). The most straightforward method would be to use intracranial electroencephalography (iEEG), though this is generally only feasible within individuals undergoing treatment for epilepsy (Axmacher et al., 2008; Engel et al., 2005; Staresina et al., 2015; Zhang et al., 2015). Another approach is to use functional magnetic resonance imaging (Schuck & Niv, 2019; Wittkuhn & Schuck, 2021) though the latter is burdened by the challenge posed by the sluggishness of the hemodynamic response. Researchers have recently started to leverage the spatio-temporal precision of magnetoencephalography (MEG), in combination with machine learning based brain decoding techniques, to reveal sequential human replay in humans across a range of settings that includes memory, planning and inference (Eldar et al., 2018; Kurth-Nelson et al., 2016; Liu et al., 2019; Liu, Mattar, et al., 2021; McFadyen et al., 2023; Nour et al., 2021; Wimmer et al., 2020, 2023; Wise et al., 2021). Many of the latter studies deploy a novel statistical analysis technique, temporally delayed linear modeling (Liu, Dolan, et al., 2021). TDLM, and its variants, has enabled identification of sequential replay for previously learned material during resting state (Liu et al., 2019; Liu, Mattar, et al., 2021), during planning of upcoming behavioral output (Eldar et al., 2020; Kurth-Nelson et al., 2016; McFadyen et al., 2023; Wise et al., 2021) and retrieval (Wimmer et al., 2020).

In relation to memory, Wimmer et al. (2020) reported sequential reactivation of episodic content after a single initial exposure during cued recall one day post-encoding. Specifically, they showed participants eight short, narrated stories, each consisting of four different visual story anchor

elements taken from six different categories (faces, buildings, body parts, objects, animals, and cars) and a unique ending element. During a next day recall session, participants were shown two story elements and asked whether both elements were part of the same story, and whether the second appeared before or after the first. During the retrieval test they showed that stories were replayed in reverse order to the prompt (i.e., when prompting element 3 and element 5, successful retrieval would traverse element 5 through 4 and arrive at element 3). However, this effect was only found in those with regular performance, while in high performers there was no evidence for temporal succession. Instead, high performers simultaneously reactivated all related story elements at in a clustered manner.

In memory research, declarative tasks often avail of item lists or paired associates (Barnett et al., 2016; Cho et al., 2020; Feld et al., 2013; Kolibius et al., 2021; Roux et al., 2022; Schönauer et al., 2014; Stadler et al., 1999, 1999). When studying sequential replay, the task structure must have a linear element (Liu et al., 2019; Liu, Mattar, et al., 2021; Wimmer et al., 2020; Wise et al., 2021) and such linearity is a defining feature of episodic memory (Tulving, 1993). By contrast, semantic memory is rarely organized linearly and instead involves complex and interconnected knowledge networks or cognitive maps (Behrens et al., 2018) motivating researchers to investigate how memory works when organized into a complex graph structure (Eldar et al., 2020; G. Feld et al., 2021; Garvert et al., 2017; Schapiro et al., 2013; for an overview see Momennejad, 2020). However, little is currently known regarding replay involvement in consolidation and retrieval processes for information embedded in graph structures.

We examined the relationship of graph learning to reactivation and replay in a task where participants learned a directed, cyclic graph, represented by ten connected images. Eight nodes had exactly one direct predecessor and successor node, two *hub nodes*, each had two direct predecessors and successors (See Figure 2B). The task was arranged such that participants could not rely on simple pair mappings but needed to learn the context of each edge. Additionally, the graph-structure was never shown to the participant as a ‘birds-eye-view’, encouraging implicit learning of the underlying structure. Following a retention period, consisting of eight minutes eyes closed resting state, participants then completed a cued recall task, which is the focus of the current analysis.

Methods

Participants

We recruited thirty participants (15 men and 15 women) between 19 and 32 years old (mean age 24.7 years). Inclusion criteria were right-handedness, no claustrophobic tendencies, no current or previous diagnosed mental disorder, non-smoker, fluency in German or English, age between 18 and 35 and normal or corrected-to-normal vision. Caffeine intake was requested to be restricted for four hours before the experiment. Participants were recruited through the institute's website and mailing list and various local Facebook groups. A single participant was excluded due to a corrupted data file and replaced with another participant. We acquired written informed consent from all participants, including consent to share anonymized raw and processed data in an open access online repository. The study was approved by the ethics committee of the Medical Faculty Mannheim of Heidelberg University (ID: 2020-609). While we had preregistered the study design and an analysis approach for the resting state data (<https://aspredicted.org/kx9xh.pdf>, #68915) here we report analyses of the retrieval period. The current analysis conceptually replicates the analyses and hypotheses of Wimmer et al. (2020) focusing on the retrieval period albeit in a much more complex and therefore naturalistic paradigm and are therefore, despite not being preregistered, mainly of confirmatory nature. We wish to maintain transparency by acknowledging that the findings from the preregistered analysis concerning the resting state data, are being prepared for publication as part of a distinct submission.

Procedure

Participants came to the laboratory for a single study session of approximately 2.5 hours. After filling out a questionnaire about their general health, their vigilance state (Stanford Sleepiness Scale, Hoddes et al., 1973) and mood (PANAS, Watson et al., 1988), participants performed five separate tasks while in the MEG scanner. First, an eight-minute eyes-closed resting state was recorded. This was followed by a localizer task (~30 minutes), in which all 10 items were presented 50 times in pseudo-randomized order, using auditory and visual stimuli. Next, participants learned a sequence of the 10 visual items embedded into a graph structure until they achieved 80% accuracy or reached a maximum of six blocks (7 to 20 minutes). Following this, we recorded another eight-minutes eyes-closed resting state to allow for initial consolidation and, finally, a cued recall testing session (four minutes). For an overview see Figure 1.

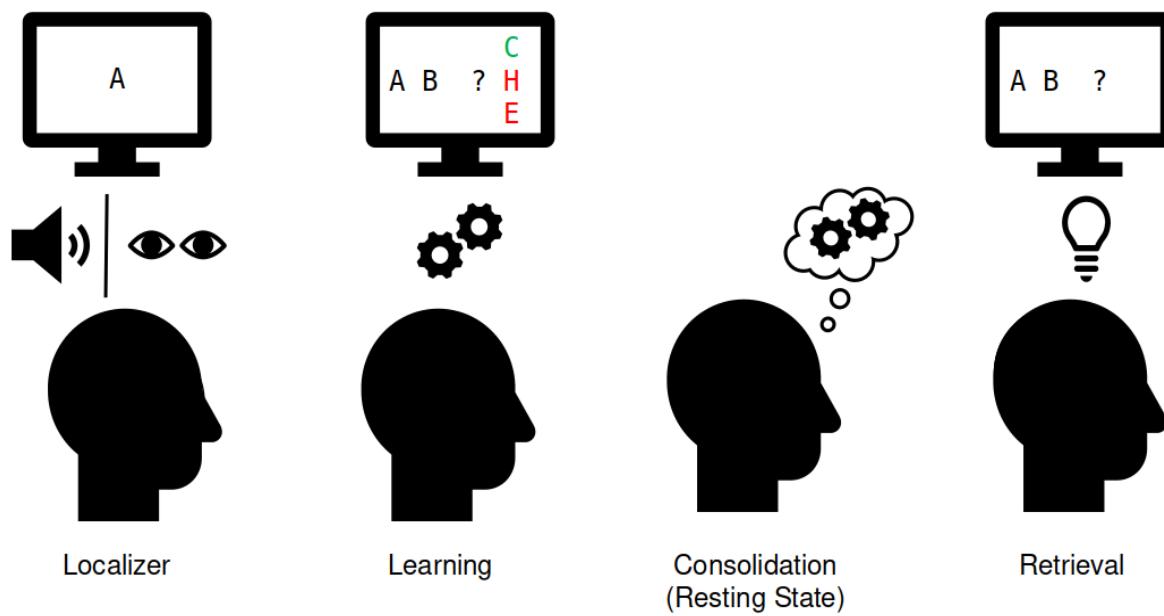


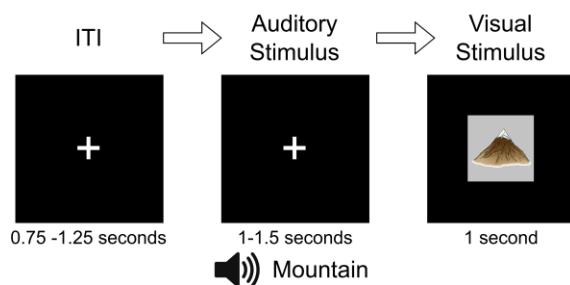
Figure 1: Experiment procedure in the MEG. Localizer task: The ten individual items were repeatedly presented to the participant auditorily and visually to extract multisensory activity patterns. Learning: Participants learned pseudo-randomly generated triplets of the ten items by trial and error. These triplets were determined by an underlying graph structure. Participants were unaware of the exact structure and graph layout. Consolidation: Eight minutes of resting state activity were recorded. Retrieval: Participants' recall was tested by cueing triplets from a sequence. The letters in the pictograms are placeholders for individual images.

Stimulus material

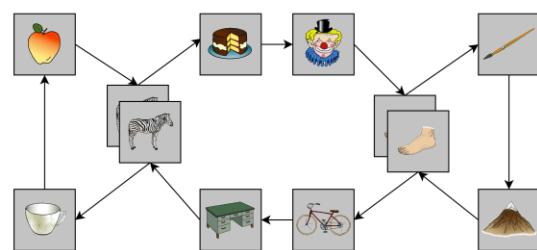
Visual stimuli were taken from the colored version (Rossion & Pourtois, 2001) of the Snodgrass & Vanderwart (1980) stimulus dataset. To increase brain pattern discriminability, images were chosen with a focus on diversity of color, shape and category (see Figure 2B) and for having short descriptive words (one or two syllables) both in German and English. Auditory stimuli were created using the Google text-to-speech API, availing of the default male voice (*Ssm/VoiceGender.NEUTRAL*) with the image description labels, either in German or English, based on participants language preference. Auditory stimulus length ranged from 0.66 to 0.95 seconds.

Task description

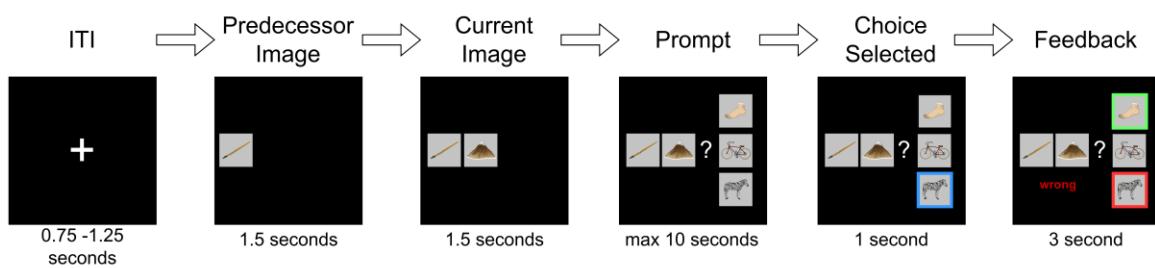
A Localizer Trial



B Sequence



C Learning Trial



D Retrieval Trial

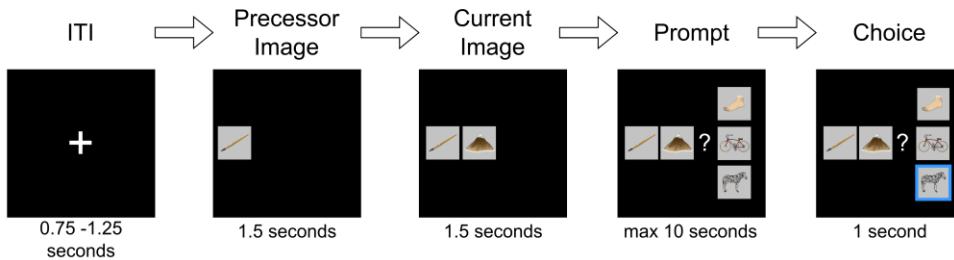


Figure 2: Task structure: A) During the localizer task, a word describing the stimulus was played back via headphones and the item then shown to the participant. In 4% of trials, the audio and visual cue did not match and in this case, participants were instructed to press a button (attention check). B) Graph layout of the task. Two elements could appear in two different triplets. The graph was directed such that each tuple had exactly one successor (e.g., apple → zebra could only be followed by cake and not mug), but individual items could have different successors (zebra alone could be followed by mug or cake). Participants never saw the illustrated birds-eye-view. C) During learning, one node was randomly chosen as the current node in each trial. First, its predecessor node was shown, then the current node was shown, then the participant was given the choice of three items and must choose the correct node that followed the displayed cue tuple. Feedback was then provided to the participant. This process was repeated until the participant reached 80% accuracy for any block or reached a maximum of six blocks of learning. D) The retrieval testing followed the same structure as the learning task, except that no feedback was given.

Localizer task

In the localizer task, the ten graph items were shown to the participants repeatedly in a pseudo-random order where a DeBruijn-sequence (DeBruijn, 1946) was used to ensure that the number of transitions between any two stimuli was equal. Two runs of the localizer were performed per participant, comprising 250 trials with 25 item repetitions. Each trial started with a fixation cross followed by an inter-trial interval of 0.75 to 1.25 seconds. Next, to encourage a multi-sensory and neural representation, the name of the to-be-shown image was played through in-ear head-phones (maximum 0.95 seconds) followed 1.25 to 1.75 seconds later by the corresponding stimulus image, shown for 1.0 second. As an attention check, in ~4% of the trials the auditory stimulus did not match the image and participants were instructed to press a button as fast as possible to indicate detection of an incongruent auditory-visual pair. A short break of maximum 30 seconds was scheduled every 80 trials. Between the two parts of the localizer task, another short break was allowed. Stimulus order was randomized and balanced between subjects. To familiarize the participant with the task, a short exemplar of the localizer task with dummy images was shown beforehand. All subsequent analyses were performed using the visual stimulus onset as a point of reference.

Graph-Learning

The exact same images deployed in the localizer task were randomly assigned to nodes of the graph, as shown in Figure 2B. Participants were instructed to learn a randomized sequence of elements with the goal of reaching 80% performance within six blocks of learning. During each block, participants were presented with each of the twelve edges of the graph exactly once, in a balanced, pseudo-randomized order. After a fixation cross of 3.0 seconds a first image (predecessor) was shown on the left of the screen. After 1.5 seconds, the second image (current image) appeared in the middle of the screen. After another 1.5 seconds, three possible choices were displayed in vertical order to the right of the two other images. One of the three choice options was the correct successor of the cued edge. Of the two distractor stimuli, one was chosen from a distal location on the graph (five to eight positions away from the current item), and one was chosen from a close location (two to four positions away from the current item). Neither of the latter were directly connected to any of the other elements onscreen. Participants were given a controller with three buttons to indicate their answer. The correct item was then highlighted for 3.0 seconds, and the participant's performance was indicated ("correct" or "wrong") (see Figure 2C). No audio was played during learning. The participant was instructed to learn the sequence transitions by trial-and-error, and instructed that there was no semantic connection between the items (i.e., that the sequence did not follow any specific logic related to image content). Participants completed a minimum of two, and a maximum of six blocks of learning. To prevent ceiling effects, learning was discontinued if a participant reached 80% accuracy during any block. To familiarize participants with the task, a short example with dummy images was shown before the learning task. Triplets were shown in a random order and choices were displayed in a pseudo-random position that ensured the on-screen position of the correct item could never be at the same position for more than three consecutive trials. Distractor choices were

balanced such that exposure to each individual item was approximately equal.

Resting State

After graph learning, participants completed a resting state session of eight minutes. Here, they were instructed to close their eyes and “to not think of anything particular”. The resting state data are not reported here.

Retrieval Test

After the resting state, we presented subjects with a single testing session block, which followed the exact layout of the learning task with the exception that no feedback was provided as to whether choices were correct or incorrect (Figure 2D).

MEG Acquisition and Pre-Processing

MEG was recorded in a passively shielded room with a MEGIN TRIUX (MEGIN Oy, Helsinki, Finland) with 306 sensors (204 planar gradiometers and 102 magnetometers) at 1000 Hz with a 0.1–330 Hz band-pass acquisition filter at the ZIPP facility of the Central Institute for Mental Health in Mannheim, Germany. Before each recording, empty room measurements made sure that no ill-functioning sensors were present. Head movement was recorded using five head positioning coils. Bipolar vertical and horizontal electrooculography (EOG) as well as electrocardiography (ECG) was recorded. After recording, the MEGIN proprietary MaxFilter algorithm (version 2.2.14) was run using temporally extended signal space separation (tSSS) and movement correction with the MaxFilter default parameters (Taulu & Simola, 2006, raw data buffer length of 10 s, and a subspace correlation limit of .98. Bad channels were automatically detected at a detection limit of 7; none had to be excluded. The head movement correction algorithm used 200 ms windows and steps of 10 ms. The HPI coil fit accept limits were set at an error of 5 mm and a g-value of .98). Using the head movement correction algorithm, the signals were virtually re-positioned to the mean head position during the initial localizer task to ensure compatibility of sensor-level analysis across the recording blocks. The systematic trigger delay of our presentation system was measured and visual stimuli appeared consistently 19 milliseconds after their trigger value was written to the stimulus channel, however, to keep consistency with previous studies that do not report trigger delay, timings in this publication are reported uncorrected (i.e., ‘as is’, not corrected for this delay).

Data were pre-processed using Python-MNE (version 1.1, Gramfort, 2013). Data were down-sampled to 100 Hz using the MNE function ‘*resample*’ (with default settings, which applies an anti-aliasing filter before resampling with a brick-wall filter at the Nyquist frequency in the frequency domain) and ICA applied using the ‘*picard*’ algorithm (Ablin et al., 2018) on a 1 Hz high-pass filtered copy of the signal using 50 components. As recommended, ICA was set to ignore segments that were marked as bad by *Autoreject* (Jas et al., 2017) on two-second segments. Components belonging to EOG or ECG and muscle artifacts were identified and removed automatically using MNE functions ‘*find_bads_eog*’, ‘*find_bads_ecg*’ and ‘*find_bads_emg*’, using the EOG and ECG as

reference signals. Finally, to reduce noise and drift, data were filtered with a high-pass filter of 0.5 Hz using the MNE filter default settings (hamming window FIR filter, -6 dB cutoff at 0.25 Hz, 53 dB stop-band attenuation, filter length 6.6 seconds).

Trials for the localizer task were created from -0.1 to 0.5 seconds relative to visual stimulus onset to train the decoders and for the retrieval task, from 0 to 1.5 seconds after onset of the second visual cue image. No baseline correction was applied. To detect artifacts, *Autoreject* was applied using the default settings, which repaired segments by interpolation in case artifacts were present in only a limited number of channels and rejected trials otherwise (see Supplement 1). Finally, to improve numerical stability, signals were re-scaled to similar ranges by multiplying values from gradiometers by $1e^{10}$ and from magnetometers by $2e^{11}$. These values were chosen empirically by matching histograms for both channel types. As outlier values can have significant influence on the computations, after re-scaling, values that were still above 1 or below -1 were “cutoff” and transformed to smaller values by multiplying with $1e^{-2}$. Anonymised and maxfiltered raw data are openly available at Zenodo (<https://doi.org/10.5281/zenodo.8001755>), code is made public on GitHub (<https://github.com/CIMH-Clinical-Psychology/DeSMRRest-clustered-reactivation>).

Decoding framework and training

In line with previous investigations (Kurth-Nelson et al., 2016; Liu et al., 2019; Wimmer et al., 2020) we applied Lasso regularized logistic regression on sensor-level data of localizer trials using the Python package Scikit-Learn (Pedregosa et al., 2011). Decoders were trained separately for each participant, and each stimulus, using *liblinear* as a solver with 1000 maximum iterations and a L1 regularization of C=6. This value was determined based on giving the best average cross-validated peak accuracy across all participants when searching within the parameter space of C = 1 to 20 in steps of 0.5 using the same approach as outlined below (note that Scikit-Learn shows stronger regularization with lower C values, opposite to e.g., MATLAB). To circumvent class imbalance due to trials removed by *Autoreject*, localizer trials were stratified such that they contained an equal number of trials from each stimulus presentation by randomly removing trials from over-represented classes. Using a cross validation schema (leaving one trial out for each stimulus per fold, i.e., 10 trials left out per fold), for each participant the decoding accuracy was determined across time (**Figure 3A**). During cross validation, for each fold, decoders were trained on data of each 10 milliseconds time step and tested on left out data from the same time step. Therefore, decoding accuracy reflects the separability of the stimulus classes by the sensor values for each time step independently.

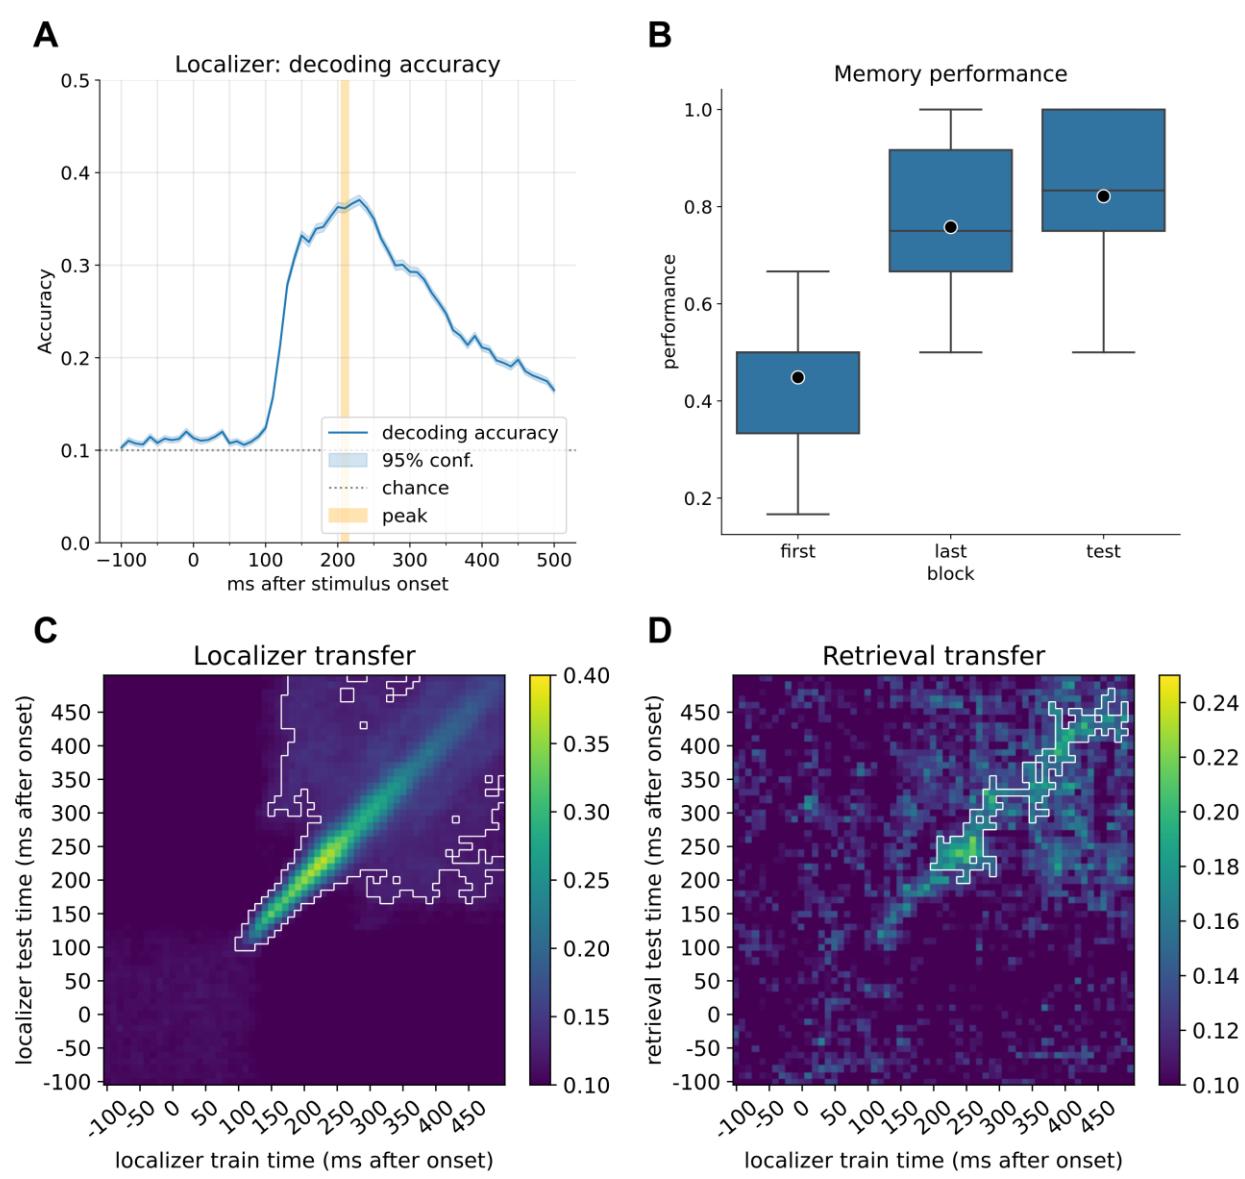


Figure 3: A) Decoding accuracy of the currently displayed item during the localizer task for participants with a decoding accuracy higher than 30% ($n=21$). The mean peak time point across all participants corresponded to 210 ms, with an average decoding peak decoding accuracy of 42% ($n=21$). Note that the displayed graph combines accuracies across participants, where peak values were computed on an individual level and then averaged. Therefore, the indicated individual mean peak does not match the average at a group level. B) Memory performance of participants after completing the first block of learning, the last block (block 2 to 6, depending on speed of learning), and the test performance. C) Classifier transfer within the localizer when trained and tested at different time points determined by cross validation. D) Classifier transfer from the localizer session to the retrieval session when trained at different time points during training and tested at different time points during cue presentation of the first (predecessor) image cue during retrieval. For B and C: Within the white outline, classification was significantly above chance level (cluster permutation testing, $\alpha<0.05$).

For each participant, a final set of decoders (i.e., 10 decoders per participant, for each stimulus one decoder) were trained at 210 milliseconds after stimulus onset, a time point reflecting the average peak decoding time point computed for all participants (for individual decoding accuracy plots see

Supplement 3). For the final decoders, data from before the auditory stimulus onset was added as a negative class with a ratio of 1:2. Adding null data allows decoders to report low probabilities for all classes simultaneously in absence of a matching signal and reduces false positives while retaining relative probabilities between true classes. Together with use of a sparsity constraint on the logistic regression coefficients, this increases the sensitivity of sequence detection by reducing spatial correlations of decoder weights (see also Liu, Dolan, et al., 2021). For a visualization of relevant sensor positions see Supplement 5.

Decoders were then applied to trials of the test session, starting from the stimulus onset of the second sequence cue (“current image”) to just prior to onset of the selection prompt (1.5 seconds). For each trial, this resulted in ten probability vectors across the trial, one for each item, in steps of 10 milliseconds. These probabilities indicate the similarity of the current sensor-level activity to the activity pattern elicited by exposure to the stimulus and can therefore be used as a proxy for detecting active representations, akin to a representational pattern analysis approach (RSA, Grootswagers et al., 2017). As a sanity check, we confirmed that we could decode the currently on-screen image by applying the final trained decoders to the first image shown during test (predecessor stimulus, see Figure 3D).

Sequential replay analysis

To test whether individual items were reactivated in sequence at a particular time lag, we applied temporally delayed linear modeling (TDLM, Liu, Dolan, et al., 2021) on the time span after the stimulus onset of the sequence cue (“current image”). In brief, this method approximates a time lagged cross-correlation of the reactivation strength in the context of a particular transition pattern, quantifying the strength of a certain activity transition pattern distributed in time.

Using a linear model, we first estimate evidence for sequential activation of the decoded item representations at different time lags. For each item i at each time lag Δt up to 250 milliseconds we estimated a linear model of form:

$$Y_i = Y(\Delta t) \times \beta_i(\Delta t)$$

where Y_i contains the decoded probability output of the classifier of item i and $Y(\Delta t)$ is simply Y time lagged by Δt . When solving this equation for $\beta_i(\Delta t)$ we can estimate the predictive strength of $Y(\Delta t)$ for the occurrence of Y_i at each time lag Δt . Calculated for each stimulus i , we then create an empirical transition matrix $T_e(\Delta t)$ that indexes evidence for a transition of any item j to item i at time lag Δt (i.e., a 10x10 transition matrix per time lag, each column j contains the predictive strength of j for each item i at time lag Δt). These matrices are then combined with a ground truth transition matrix T (encoding the valid sequence transitions of interest) by taking the Frobenius inner product. This returns a single value $Z_{\Delta t}$ for each time lag, indicating how strongly the detected transitions in

the empirical data follow the expected task transitions, which we term “sequenceness”. Using different transition matrices to depict forward (T_f) and backward (T_b) replay, we quantified evidence for replay at different time lags for each trial separately. This process is applied to each trial individually, and resulting sequenceness values are averaged to provide a final sequenceness value per participant for each time lag Δt . To test for statistical significance, we create a baseline distribution by permuting the rows of the transition matrix 1000 times (creating transition matrices with random transitions; identity-based permutation, Liu, Dolan, et al., 2021) and calculate sequenceness across all time lags for each permutation. The null distribution is then constructed by taking the peak sequenceness across all time lags for each permutation.

Differential reactivation analysis

To test for clustered, non-sequential reactivation, we adopted a similar approach to that of Wimmer et al. (2020). As decoders were trained independently for each stimulus, all decoders reacted to presentation of any visual stimulus to some extent. By using differences in reactivation between stimuli, this aggregated approach allowed us to examine more closely whether near items are more strongly activated than distant items, thereby quantifying non-sequential reactivation with more sensitivity. For each trial, the mean probability of the two items following the current on-screen item were contrasted with the mean probability of all items further away by subtraction. The two items currently displayed on-screen (i.e., predecessor and current image) were excluded. As only few trials were available for this analysis per participant, the raw probabilities were noisy. Therefore, to address this we applied a Gaussian smoothing kernel ($\sigma = 1$) to the probability vectors across the time dimension. By shuffling the stimulus labels 1000 times, we constructed an empirical permutation distribution to determine at which time points the differential reactivation of close items was significantly above chance ($\alpha = 0.05$).

Graph reactivation analysis

To detect whether reactivation strength was modulated by underlying graph structure, we compared the raw reactivation strength of all items by distance on the directed graph. First, we calculated a time point of interest by computing the peak probability estimate of decoders across all trials. Then, for each participant, for each trial we sorted all nodes based on their distance to the current on-screen item on the directed graph. Again, we smoothed probability values with a Gaussian kernel ($\sigma = 1$) and ignored the predecessor on-screen item. Following this, we evaluated the sorted decoder probabilities at the previously determined peak time point. Using a repeated measures ANOVA on the mean probability values per distance per participant, we then estimated whether reactivation strength was modulated by graph distance.

Exclusions

Replay analysis relies on a successive detection of stimuli where the chance of detection exponentially decreases with each step (e.g., detecting two successive stimuli with a chance of 30%

leaves a 9% chance of detecting the replay event). Therefore, all participants with a peak decoding accuracy of below 30% were excluded from the analysis (nine participants), a threshold set before starting the analysis process. Additionally, as successful learning was necessary for the paradigm, we ensured all remaining participants had a retrieval performance of at least 50% (see Supplement 2).

Results

Behavioral

All but one participant learned the sequence of ten images embedded into the directed graph with partial overlap (Supplement 3). On average, participants needed 5 blocks of learning (range 2 to 6, see Supplement 4) and manifested a memory performance of 76% during their last block of learning (range: 50% to 100%). After eight minutes of rest, retrieval performance improved marginally to a mean of 82% ($t=-2.053$, $p=0.053$, effect size $r=0.22$, Figure 3). Note that since the last learning block included feedback, this marginal increase cannot necessarily be attributed to consolidation processes.

Decoder training

We first confirmed we could decode brain activity elicited by the ten items using a cross-validation approach. Indeed, decoders were able to separate the items presented during the localizer task (see Figure 3A) well, with an average peak decoding accuracy of ~42% across all participants (range: 32% to 57%, chance level: 10%, excluding participants with peak accuracy < 30%, for all participants see Supplement 3). We calculated the time point of the mean peak accuracy for each participant separately and subsequently used the average best time point, across all included participants, at 206 milliseconds (rounded to 210 milliseconds) for training of our final decoders. This value is very close in range to the time point found in previous studies (Kurth-Nelson et al., 2016; Liu et al., 2019; Liu, Mattar, et al., 2021; Wimmer et al., 2020). The decoders also transferred well to stimulus presentation during the retrieval trials and could effectively decode the current prompted image cue with above chance significance (cluster permutation test, see Figure 3D).

Sequential forward replay in subjects with lower memory performance

Next, we assessed whether there was evidence for sequential replay of the learned sequences during cued recall. Using TDLM we assessed whether decoded reactivation probabilities followed a sequential temporal pattern, in line with transitions on the directed graph. Here we focused on all allowed graph transitions and analyzed the entire time window, of 1500 milliseconds, after onset of the test cue ("current image"). We found positive sequenceness across all time lags for forward sequenceness, with a significant increase at around 40-50 milliseconds for forward sequenceness (Figure 4A). As discussed in Liu, Dolan et al. (2021), correction for multiple comparisons for this

sequenceness measure across time is non-trivial and the maximum of all permutations represents a highly conservative statistic. Due to this complexity, we additionally report the 95% percentile of sequenceness maxima across time per permutation. Nevertheless, as we did not have a pre-defined time lag of interest beforehand, and to mitigate multiple-comparisons, we additionally computed the mean sequenceness across all computed time lags for each participant (similar as previously proposed in the context of a sliding window approach in Wise et al., 2021). This measure can help reveal an overall tendency for replay of task states that is invariant to a specific time lag. Our results show that across all participants, there is a significant increase in task-related forward sequential reactivation of states ($p=0.027$, two-sided permutation test with 1000 permutations; 95% of permutation maxima reached at 40-50 ms, Figure 4B). Following up on this, in a second analysis, we asked whether mean sequential replay was associated with memory performance and found a significant negative correlation between retrieval performance and forward replay (forward: $r=-0.46$, $p=0.031$; backward: $r=-0.13$, $p=0.56$, see Figure 4C). In line with previous results (Wimmer et al., 2020) low-performing participants had higher forward sequenceness when compared to high-performing participants, whose mean sequenceness tended towards zero.

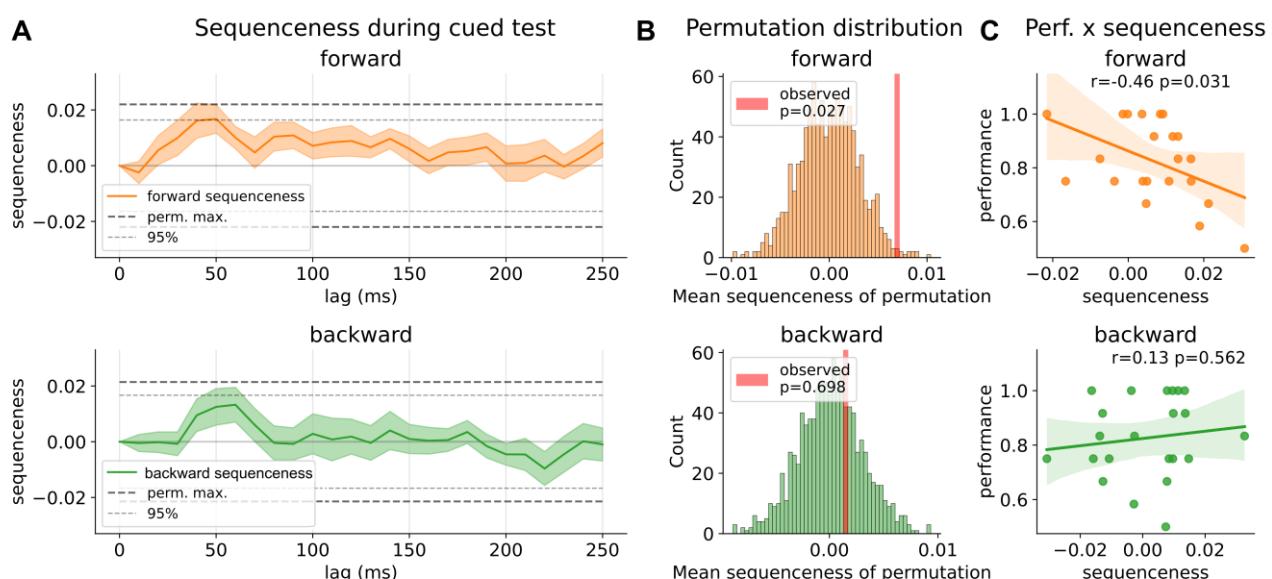


Figure 4: A) Strength of forward and backward sequenceness across different time lags up to 250 ms during the 1500 ms window after cue onset. Two significance thresholds are shown: Conservative threshold of the maximum of 1000 permutations of classification labels across all time lags and the 95% percentiles (see Methods section for details). B) Permutation distribution of mean sequenceness values across 1000 state permutations. Observed mean sequenceness is indicated with a red line. C) Association between memory performance and mean sequenceness value computed across all trials, and time lags, for each participant.

Closer nodes show stronger reactivation than distant nodes

Next, in a complementary analysis, we asked whether a non-sequential clustered reactivation of items occurs after onset of a cue image (as shown previously for high performers in Wimmer et al., 2020). We compared reactivation strength of the two items following the cue image with all items with a distance of more than two steps, subtracting the mean decoded reactivation probabilities from each other. Using this differential reactivation, we found that near items were significantly reactivated compared to items further away within a time window of 220 ms to 260 ms after cue onset (Figure 5A, $p<0.05$, permutation test with 10000 shuffles).

To further explore the relation of reactivation strength and graph distances, we analyzed the mean reactivation strength by item distance at peak classifier probabilities and found reactivation strength significantly related to graph distance (repeated measures ANOVA, $F(4, 80)=2.98$, $p=0.023$ Figure 5B). When subdividing trials into correct and incorrect responses, we found that this relationship was only significant for trials where a participant successfully retrieved the currently prompted sequence excerpt (repeated measures ANOVA, $F(4, 80)=5.0$, $p=0.001$ for correctly answered trials, Figure 5C). For incorrect trials we found no evidence for this relationship ($F(4, 48)=1.45$ $p=0.230$ for incorrectly answered trials), albeit we found no interaction between distance and response type ($F(4, 48)=1.8$, $p=0.13$). Note, that the last two analysis are based on $n = 14$ since 7 participants had no incorrect trials.

Questionnaire results

Participants were concentrated and alert as indicated by the Stanford Sleepiness Scale ($M = 2.3$, $SD = 0.6$, range 1-3). Participants' summed positive affect score was on average 33.2 ($SD = 4.5$), their summed negative affect score was on average 12.2 ($SD = 1.9$) (PANAS).

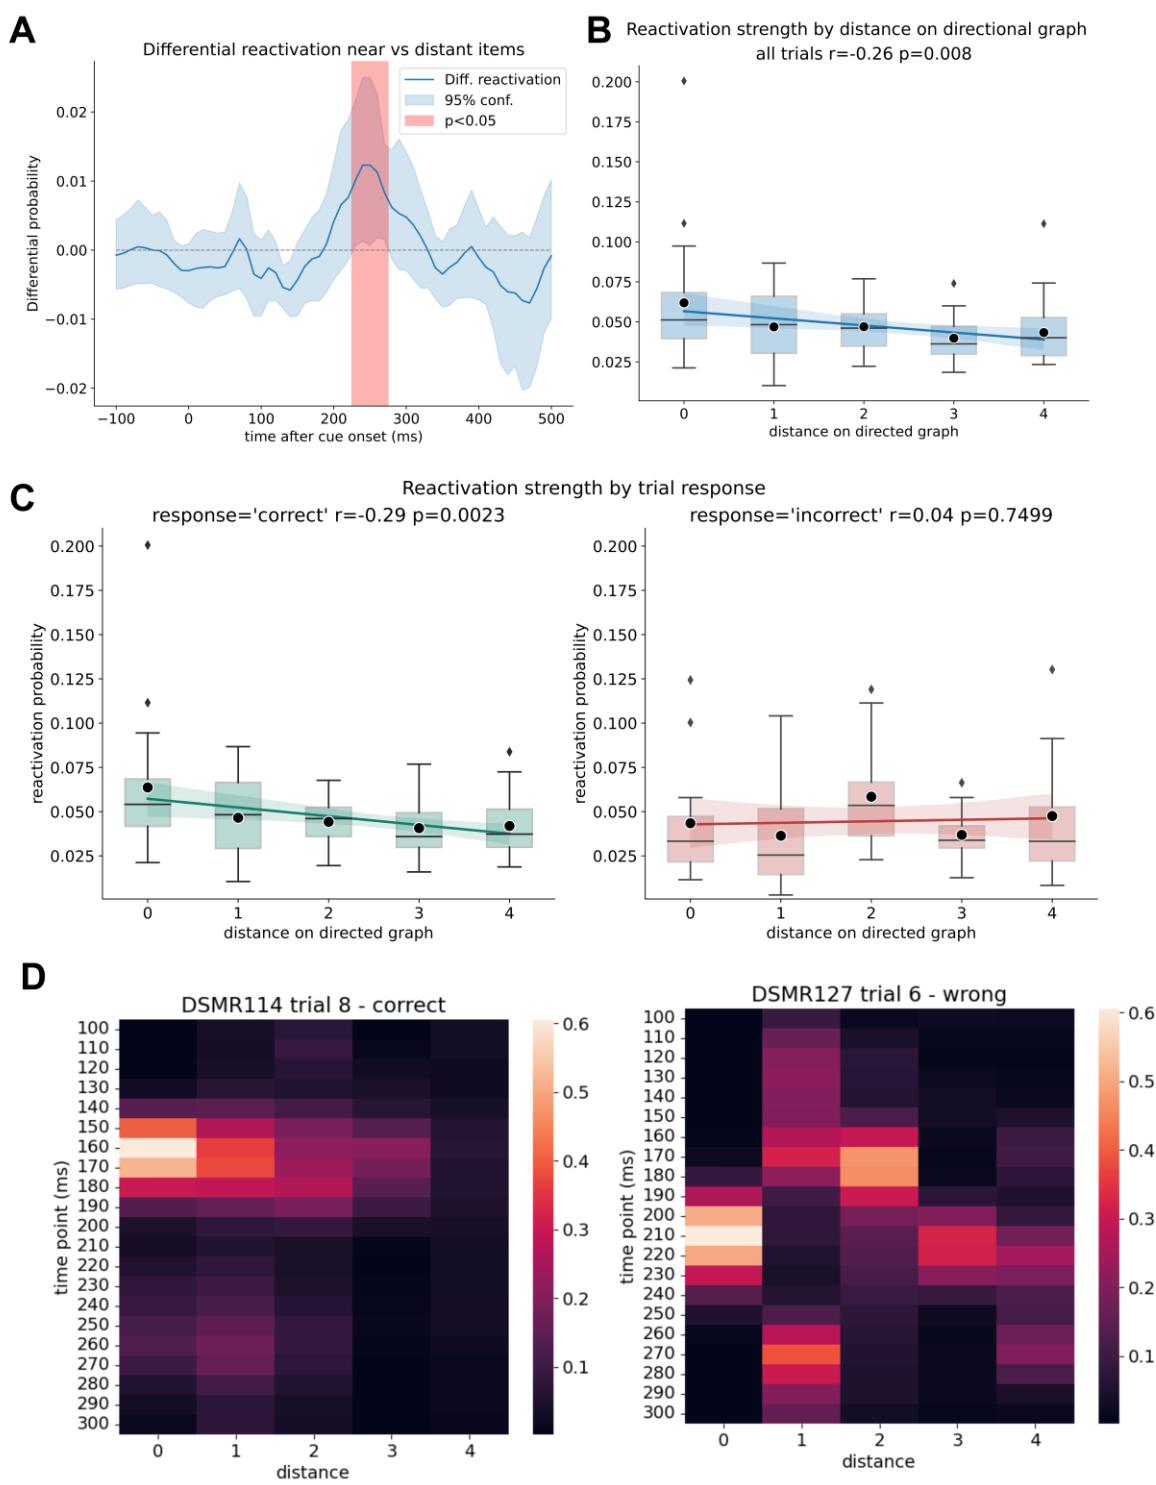


Figure 5: Clustered Reactivation: A) Differential reactivation probability between off-screen items that were up to two steps ahead of the current stimulus cue vs. distant items that were more than two steps away on the graph for trials with correct answers. Between 220 and 260 ms the next items are simultaneously reactivated significantly more than items that are further away ($p<0.05$; permutation test with 10000 shuffles). B) Reactivation strength of items after retrieval cue onset by distance of items to the currently on-screen stimulus. A significant negative correlation between distance on a directional graph and reactivation strength can be seen ($p=0.008$). C) Same as B, but subdivided into trials in which participants answered correctly (left) and in which participants did not know the correct answer (right). A correlation between reactivation strength and distance can only be seen in case of successful retrieval (but see also limitations for a discussion of the low trial and participant number in this sub-analysis). Mean probability values are marked by black dots. D) Example activations of a successful retrieval (left) and a failed retrieval (right), sorted by distance to current cue. Colors indicate probability estimates of the decoders. Indicated time points in D) are after onset of the current image cue.

Discussion

We combined a graph-based learning task with machine learning to study neuronal events linked to memory retrieval. Participants learned triplets of associated images by trial and error, where these were components of a simple directed graph with ten nodes and twelve edges. Using machine learning decoding of simultaneously recorded MEG data we asked what brain processes are linked to retrieval of this learned information, and how this related to the underlying graph structure. We found evidence that graph items are retrieved by a simultaneous, clustered, reactivation of items and that the associated reactivation strength relates to graph distances.

Memory retrieval is thought to involve reinstatement of previously evoked item-related neural activity patterns (Danker & Anderson, 2010; Johnson & Rugg, 2007; Staresina et al., 2012). Both spatial and abstract information is purported to be encoded into cognitive maps within the hippocampus and related structures (Behrens et al., 2018; Bellmund et al., 2018; Epstein et al., 2017; Garvert et al., 2017; O'Keefe & Nadel, 1979; Peer et al., 2021). While, for example, spatial distance within cognitive maps is encoded within hippocampal firing patterns (Theves et al., 2019), it is unclear how competing, abstract, candidate representations are accessed during retrieval (Kerrén et al., 2018, 2022; Spiers, 2020). Two separate mechanisms seem plausible. First, depth-first search might enable inferences in not yet fully consolidated cognitive maps by sequential replay of potential candidates (Mattar & Daw, 2018; Nyberg et al., 2022); second, breadth-first search could be deployed involving simultaneous activation of candidates when these are sufficiently consolidated within maps that support non-interfering co-reactivation of competing representations (Mattar & Lengyel, 2022), or when exhaustive replay would be too expensive computationally. Indeed, consistent with this, Wimmer et al., (2020) showed that for regular memory performance, sequential and temporally spaced reactivation of items seems to 'piece together' individual elements. This is contrasted with high performers who showed a clustered, simultaneous, reactivation profile. We replicate this clustered reactivation and show that its strength reflects distance on a graph structure. This complements previous findings of graded pattern similarity during memory search representing distance within the search space (Manning et al., 2011; Tarder-Stoll et al., 2023). As this effect was evident only for correct choices the finding points to its importance for task performance.

In line with Wimmer et al. (2020), we found that the strength of replay was linked to weaker memory performance. This suggests that the expression of sequential replay or simultaneous reactivation depends on the stability of an underlying memory trace. However, we acknowledge that it remains unclear which factors enable recruitment of either of these mechanisms. A crucial step in consolidation encompasses an integration of memory representations into existing networks (Dudai et al., 2015; Sekeres et al., 2017). In Wimmer et al. (2020), participants had little exposure to the learning material and replay was measured after a substantial retention period that included sleep, where the latter is considered to strengthen and transform memories via repeated replay

(Diekelmann & Born, 2010; Feld & Born, 2017). This contrasts with the current task design, which involved several blocks of learning and retrieval and only a relatively brief period of consolidation.

Intriguingly, it has been speculated that retrieval practice may elicit the same transformation of memory traces as offline replay (Antony et al., 2017). Following this reasoning, it is possible that both consolidation during sleep and repeated practice have similar effects on the transformation of memories, and consequently the mechanisms that support their subsequent retrieval. This possibility is especially interesting in the light of retrieval practice enhancing memory performance more than is the case for restudy (McDermott, 2021) and is also in line with evidence that replay during rest prioritizes weakly learned memories (Schapiro et al., 2018). It is known that retrieval practice reduces the pattern similarity of competing memory traces in the hippocampus (Hulbert & Norman, 2015) and, as in the case of our graph-based task, may enable clustered reactivation since differences in timing of reactivation are no longer required to distinguish correct from incorrect items. Therefore, we speculate that clustered reactivation may be a physiological correlate of retrieval facilitated either by repeated retrieval testing-based learning (as in our study) or by sleep dependent memory consolidation (as in Wimmer et al., 2020). This implies that there may be a switch from sequential replay to clustered reactivation corresponding to when learned material can be accessed simultaneously without interference. This suggestion could be systematically investigated by, for example, manipulating retrieval practice, retention interval, and the difficulty of a graph-based task. While we closely follow the analysis approach taken in Wimmer et al., (2020), we did not explicitly preregister the confirmatory analysis of the retrieval data as such. We do acknowledge that only a somewhat limited number of trials were available for analysis, impacting especially the analysis of incorrect answers. In addition, the number of low-performing participants was low in our study which would render a performance-dependent sub-analysis underpowered.

In conclusion, the reported findings support a role of clustered reactivation mechanism for well-learned items during memory retrieval. When interconnected semantic information is retrieved, the retrieval process seems to resemble a breadth-first search, with items sorted by neural activation strength. Additionally, we find that sequential replay relates to low memory performance. The likely coexistence of two types of retrieval process, recruited dependent on the participants' learning experience, is an important direction for future research. Using more complex memory tasks, such as explicitly learned associations of graph networks, should enable a more systematic study of this process. Finally, we suggest that accessing information embedded in a knowledge network may benefit from recruitment of either process, replay or reactivation, on the fly.

Data availability

MaxFiltered and anonymized MEG raw data as well as behavioural results are available at Zenodo (<https://doi.org/10.5281/zenodo.8001755>).

Code availability

The code of the analysis as well as the experiment paradigm and the stimulus material is available at <https://github.com/CIMH-Clinical-Psychology/DeSMRRest-clustered-reactivation>.

Supplement figures

For supplementary figures of this draft, see after references.

Acknowledgment

This research was supported by an Emmy-Noether research grant awarded to GBF by the DFG (FE1617/2-1) and a project grant by the DGSM as well as a doctoral scholarship of the German Academic Scholarship Foundation, both awarded to SK.

References

Ablin, P., Cardoso, J.-F., & Gramfort, A. (2018). Faster Independent Component Analysis by Preconditioning With Hessian Approximations. *IEEE Transactions on Signal Processing*, 66(15), 4040–4049. <https://doi.org/10.1109/TSP.2018.2844203>

Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron*, 91(5), 1124–1136. <https://doi.org/10.1016/j.neuron.2016.07.047>

Antony, J. W., Ferreira, C. S., Norman, K. A., & Wimber, M. (2017). Retrieval as a Fast Route to Memory Consolidation. *Trends in Cognitive Sciences*, 21(8), 573–576. <https://doi.org/10.1016/j.tics.2017.05.001>

Axmacher, N., Elger, C. E., & Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. *Brain*, 131(7), 1806–1817. <https://doi.org/10.1093/brain/awn103>

Barnett, J. H., Blackwell, A. D., Sahakian, B. J., & Robbins, T. W. (2016). The Paired Associates Learning (PAL) Test: 30 Years of CANTAB Translational Neuroscience from Laboratory to Bedside in Dementia Research. In T. W. Robbins & B. J. Sahakian (Eds.), *Translational Neuropsychopharmacology* (pp. 449–474). Springer International Publishing. https://doi.org/10.1007/7854_2015_5001

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, 100(2), 490–509. <https://doi.org/10.1016/j.neuron.2018.10.002>

Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415), eaat6766. <https://doi.org/10.1126/science.aat6766>

Born, J., & Wilhelm, I. (2012). System consolidation of memory during sleep. *Psychological*

Research, 76(2), 192–203. <https://doi.org/10.1007/s00426-011-0335-6>

Brunec, I. K., & Momennejad, I. (2022). Predictive Representations in Hippocampal and Prefrontal Hierarchies. *The Journal of Neuroscience*, 42(2), 299–312.
<https://doi.org/10.1523/JNEUROSCI.1327-21.2021>

Buhry, L., Azizi, A. H., & Cheng, S. (2011). Reactivation, replay, and preplay: How it might all fit together. *Neural Plasticity*, 2011. <https://doi.org/10.1155/2011/203462>

Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: A potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14(2), 147–153.
<https://doi.org/10.1038/nn.2732>

Chen, Z. S., & Wilson, M. A. (2023). How our understanding of memory replay evolves. *Journal of Neurophysiology*, 129(3), 552–580. <https://doi.org/10.1152/jn.00454.2022>

Cho, K. W., Tse, C.-S., & Chan, Y.-L. (2020). Normative data for Chinese-English paired associates. *Behavior Research Methods*, 52(1), 440–445. <https://doi.org/10.3758/s13428-019-01240-2>

Danker, J. F., & Anderson, J. R. (2010). The ghosts of brain states past: Remembering reactivates the brain regions engaged during encoding. *Psychological Bulletin*, 136, 87–102.
<https://doi.org/10.1037/a0017937>

DeBruijn, N. G. (1946). A combinatorial problem. *Proceedings of the Section of Sciences of the Koninklijke Nederlandse Akademie van Wetenschappen Te Amsterdam*, 49(7), 758–764.

Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, 11(2), 114–126. <https://doi.org/10.1038/nrn2762>

Dudai, Y., Karni, A., & Born, J. (2015). The Consolidation and Transformation of Memory. *Neuron*, 88(1), 20–32. <https://doi.org/10.1016/j.neuron.2015.09.004>

Eichenlaub, J.-B., Jarosiewicz, B., Saab, J., Franco, B., Kelemen, J., Halgren, E., Hochberg, L. R., & Cash, S. S. (2020). Replay of Learned Neural Firing Sequences during Rest in Human Motor Cortex. *Cell Reports*, 31(5), 107581. <https://doi.org/10.1016/j.celrep.2020.107581>

Eldar, E., Bae, G. J., Kurth-Nelson, Z., Dayan, P., & Dolan, R. J. (2018). Magnetoencephalography decoding reveals structural differences within integrative decision processes. *Nature Human Behaviour*, 2(9), 670–681. <https://doi.org/10.1038/s41562-018-0423-3>

Eldar, E., Lièvre, G., Dayan, P., & Dolan, R. J. (2020). The roles of online and offline replay in planning. *ELife*, 9, e56911. <https://doi.org/10/gms5sp>

Engel, A. K., Moll, C. K. E., Fried, I., & Ojemann, G. A. (2005). Invasive recordings from the human brain: Clinical insights and beyond. *Nature Reviews Neuroscience*, 6(1), Article 1.
<https://doi.org/10.1038/nrn1585>

Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: Spatial navigation and beyond. *Nature Neuroscience*, 20(11), 1504–1513.
<https://doi.org/10.1038/nn.4656>

Feld, G. B., & Born, J. (2017). Sculpting memory during sleep: Concurrent consolidation and

forgetting. *Current Opinion in Neurobiology*, 44, 20–27. <https://doi.org/10/gbrvqm>

Feld, G. B., Lange, T., Gais, S., & Born, J. (2013). Sleep-Dependent Declarative Memory Consolidation—Unaffected after Blocking NMDA or AMPA Receptors but Enhanced by NMDA Coagonist D-Cycloserine. *Neuropsychopharmacology*, 38(13), Article 13.
<https://doi.org/10.1038/npp.2013.179>

Feld, G., Bernard, M., Rawson, A., & Spiers, H. (2021). *Learning graph networks: Sleep targets highly connected global and local nodes for consolidation* [Preprint]. Neuroscience.
<https://doi.org/10.1101/2021.08.04.455038>

Foster, D. J. (2017). Replay Comes of Age. *Annual Review of Neuroscience*, 40(1), 581–602.
<https://doi.org/10.1146/annurev-neuro-072116-031538>

Foster, D. J., & Knierim, J. J. (2012). Sequence learning and the role of the hippocampus in rodent navigation. *Current Opinion in Neurobiology*, 22(2), 294–300.
<https://doi.org/10.1016/j.conb.2011.12.005>

Frank, L. M., Brown, E. N., & Wilson, M. (2000). Trajectory Encoding in the Hippocampus and Entorhinal Cortex. *Neuron*, 27(1), 169–178. [https://doi.org/10.1016/S0896-6273\(00\)00018-0](https://doi.org/10.1016/S0896-6273(00)00018-0)

Fuentemilla, L., Penny, W. D., Cashdollar, N., Bunzeck, N., & Düzel, E. (2010). Theta-Coupled Periodic Replay in Working Memory. *Current Biology*, 20(7), 606–612.
<https://doi.org/10.1016/j.cub.2010.01.057>

Garvert, M. M., Dolan, R. J., & Behrens, T. E. (2017). A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *ELife*, 6, e17086. <https://doi.org/10/f95g5z>

Gramfort, A. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7.
<https://doi.org/10.3389/fnins.2013.00267>

Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *Journal of Cognitive Neuroscience*, 29(4), 677–697.
https://doi.org/10.1162/jocn_a_01068

Hardwicke, T. E., & Wagenmakers, E.-J. (2023). Reducing bias, increasing transparency and calibrating confidence with preregistration. *Nature Human Behaviour*, 7(1), Article 1.
<https://doi.org/10.1038/s41562-022-01497-2>

Hoddes, E., Zarcone, V., Smythe, H., Phillips, R., & Dement, W. C. (1973). Quantification of Sleepiness: A New Approach. *Psychophysiology*, 10(4), 431–436. <https://doi.org/10.1111/j.1469-8986.1973.tb00801.x>

Hulbert, J. C., & Norman, K. A. (2015). Neural Differentiation Tracks Improved Recall of Competing Memories Following Interleaved Study and Retrieval Practice. *Cerebral Cortex*, 25(10), 3994–4008. <https://doi.org/10.1093/cercor/bhu284>

Jas, M., Engemann, D. A., Bekhti, Y., Raimondo, F., & Gramfort, A. (2017). Autoreject: Automated artifact rejection for MEG and EEG data. *NeuroImage*, 159, 417–429.
<https://doi.org/10.1016/j.neuroimage.2017.06.030>

Johnson, J. D., & Rugg, M. D. (2007). Recollection and the reinstatement of encoding-related cortical activity. *Cerebral Cortex (New York, N.Y.: 1991)*, 17(11), 2507–2515.

<https://doi.org/10.1093/cercor/bhl156>

Kerrén, C., Linde-Domingo, J., Hanslmayr, S., & Wimber, M. (2018). An Optimal Oscillatory Phase for Pattern Reactivation during Memory Retrieval. *Current Biology*, 28(21), 3383-3392.e6.

<https://doi.org/10.1016/j.cub.2018.08.065>

Kerrén, C., van Bree, S., Griffiths, B. J., & Wimber, M. (2022). Phase separation of competing memories along the human hippocampal theta rhythm. *eLife*, 11, e80633.

<https://doi.org/10.7554/eLife.80633>

Kolibus, L. D., Born, J., & Feld, G. B. (2021). Vast Amounts of Encoded Items Nullify but Do Not Reverse the Effect of Sleep on Declarative Memory. *Frontiers in Psychology*, 11, 607070.

<https://doi.org/10.3389/fpsyg.2020.607070>

Kurth-Nelson, Z., Economides, M., Dolan, R. J., & Dayan, P. (2016). Fast Sequences of Non-spatial State Representations in Humans. *Neuron*, 91(1), 194–204.

<https://doi.org/10.1016/j.neuron.2016.05.028>

Liu, Y., Dolan, R. J., Higgins, C., Penagos, H., Woolrich, M. W., Ólafsdóttir, H. F., Barry, C., Kurth-Nelson, Z., & Behrens, T. E. (2021). Temporally delayed linear modelling (TDLM) measures replay in both animals and humans. *eLife*, 10, e66917. <https://doi.org/10/gkf6sk>

Liu, Y., Dolan, R. J., Kurth-Nelson, Z., & Behrens, T. E. J. (2019). Human Replay Spontaneously Reorganizes Experience. *Cell*, 178(3), 640-652.e14. <https://doi.org/10.1016/j.cell.2019.06.012>

Liu, Y., Mattar, M. G., Behrens, T. E. J., Daw, N. D., & Dolan, R. J. (2021). Experience replay is associated with efficient nonlocal learning. *Science*, 372(6544), eabf1357.

<https://doi.org/10.1126/science.abf1357>

Manning, J. R., Polyn, S. M., Baltuch, G. H., Litt, B., & Kahana, M. J. (2011). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences*, 108(31), 12893–12897. <https://doi.org/10.1073/pnas.1015174108>

Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21(11), 1609–1617. <https://doi.org/10.1038/s41593-018-0232-z>

Mattar, M. G., & Lengyel, M. (2022). Planning in the brain. *Neuron*, 110(6), 914–934.

<https://doi.org/10.1016/j.neuron.2021.12.018>

McDermott, K. B. (2021). Practicing Retrieval Facilitates Learning. *Annual Review of Psychology*, 72(1), 609–633. <https://doi.org/10.1146/annurev-psych-010419-051019>

McFadyen, J., Liu, Y., & Dolan, R. J. (2023). Differential replay of reward and punishment paths predicts approach and avoidance. *Nature Neuroscience*, 26(4), Article 4.

<https://doi.org/10.1038/s41593-023-01287-7>

Momennejad, I. (2020). Learning Structures: Predictive Representations, Replay, and Generalization. *Current Opinion in Behavioral Sciences*, 32, 155–166.

<https://doi.org/10.1016/j.cobeha.2020.02.017>

Nour, M. M., Liu, Y., Arumuham, A., Kurth-Nelson, Z., & Dolan, R. J. (2021). Impaired neural replay of inferred relationships in schizophrenia. *Cell*, 184(16), 4315-4328.e17. <https://doi.org/10/gmcq6d>

Nyberg, N., Duvelle, É., Barry, C., & Spiers, H. J. (2022). Spatial goal coding in the hippocampal formation. *Neuron*, 110(3), 394–422. <https://doi.org/10.1016/j.neuron.2021.12.012>

O'Keefe, J., & Nadel, L. (1979). Précis of O'Keefe & Nadel's The hippocampus as a cognitive map. *Behavioral and Brain Sciences*, 2(4), 487–494. <https://doi.org/10.1017/S0140525X00063949>

Ólafsdóttir, H. F., Bush, D., & Barry, C. (2018). The Role of Hippocampal Replay in Memory and Planning. *Current Biology*, 28(1), R37–R50. <https://doi.org/10.1016/j.cub.2017.10.073>

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(85), 2825–2830.

Peer, M., Brunec, I. K., Newcombe, N. S., & Epstein, R. A. (2021). Structuring Knowledge with Cognitive Maps and Cognitive Graphs. *Trends in Cognitive Sciences*, 25(1), 37–54. <https://doi.org/10.1016/j.tics.2020.10.004>

Preston, A. R., & Eichenbaum, H. (2013). Interplay of Hippocampus and Prefrontal Cortex in Memory. *Current Biology*, 23(17), R764–R773. <https://doi.org/10.1016/j.cub.2013.05.041>

Rossion, B., & Pourtois, G. (2001). Revisiting snodgrass and Vanderwart's object database: Color and texture improve object recognition. *Journal of Vision*, 1(3), 413. <https://doi.org/10.1167/1.3.413>

Roux, F., Parish, G., Chelvarajah, R., Rollings, D. T., Sawlani, V., Gollwitzer, S., Kreiselmeyer, G., Wimber, M., Self, M. W., & Hanslmayr, S. (2022). Oscillations support short latency co-firing of neurons during human episodic memory formation. *eLife*, 44. <https://doi.org/10.7554/eLife.78109>

Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C., & Norman, K. A. (2018). Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nature Communications*, 9(1), Article 1. <https://doi.org/10.1038/s41467-018-06213-1>

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, 16(4), Article 4. <https://doi.org/10.1038/nn.3331>

Schönauer, M., Pawlizki, A., Köck, C., & Gais, S. (2014). Exploring the Effect of Sleep and Reduced Interference on Different Forms of Declarative Memory. *Sleep*, 37(12), 1995–2007. <https://doi.org/10.5665/sleep.4258>

Schuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. *Science*, 364(6447). <https://doi.org/10.1126/science.aaw5181>

Sekeres, M. J., Moscovitch, M., & Winocur, G. (2017). Mechanisms of Memory Consolidation and Transformation. In N. Axmacher & B. Rasch (Eds.), *Cognitive Neuroscience of Memory Consolidation* (pp. 17–44). Springer International Publishing. https://doi.org/10.1007/978-3-319-45066-7_2

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name

agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 174–215. <https://doi.org/10.1037/0278-7393.6.2.174>

Spiers, H. J. (2020). The Hippocampal Cognitive Map: One Space or Many? *Trends in Cognitive Sciences*, 24(3), 168–170. <https://doi.org/10.1016/j.tics.2019.12.013>

Stadler, M. A., Roediger, H. L., & McDermott, K. B. (1999). Norms for word lists that create false memories. *Memory & Cognition*, 27(3), 494–500. <https://doi.org/10.3758/BF03211543>

Staresina, B. P., Bergmann, T. O., Bonnefond, M., van der Meij, R., Jensen, O., Deuker, L., Elger, C. E., Axmacher, N., & Fell, J. (2015). Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nature Neuroscience*, 18(11), Article 11. <https://doi.org/10.1038/nn.4119>

Staresina, B. P., Henson, R. N. A., Kriegeskorte, N., & Alink, A. (2012). Episodic Reinstatement in the Medial Temporal Lobe. *Journal of Neuroscience*, 32(50), 18150–18156. <https://doi.org/10.1523/JNEUROSCI.4156-12.2012>

Tarder-Stoll, H., Baldassano, C., & Aly, M. (2023). *The brain hierarchically represents the past and future during multistep anticipation* (p. 2023.07.24.550399). bioRxiv. <https://doi.org/10.1101/2023.07.24.550399>

Taulu, S., & Simola, J. (2006). Spatiotemporal signal space separation method for rejecting nearby interference in MEG measurements. *Physics in Medicine and Biology*, 51(7), 1759–1768. <https://doi.org/10/b5bjqq>

Theves, S., Fernandez, G., & Doeller, C. F. (2019). The Hippocampus Encodes Distances in Multidimensional Feature Space. *Current Biology*, 29(7), 1226-1231.e3. <https://doi.org/10.1016/j.cub.2019.02.035>

Tulving, E. (1993). What Is Episodic Memory? *Current Directions in Psychological Science*, 2(3), 67–70. <https://doi.org/10.1111/1467-8721.ep10770899>

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54(6), 1063–1070. <https://doi.org/10.1037/0022-3514.54.6.1063>

Wimmer, G. E., Liu, Y., McNamee, D. C., & Dolan, R. J. (2023). Distinct replay signatures for prospective decision-making and memory preservation. *Proceedings of the National Academy of Sciences*, 120(6), e2205211120. <https://doi.org/10.1073/pnas.2205211120>

Wimmer, G. E., Liu, Y., Vehar, N., Behrens, T. E. J., & Dolan, R. J. (2020). Episodic memory retrieval success is associated with rapid replay of episode content. *Nature Neuroscience*, 23(8), Article 8. <https://doi.org/10.1038/s41593-020-0649-z>

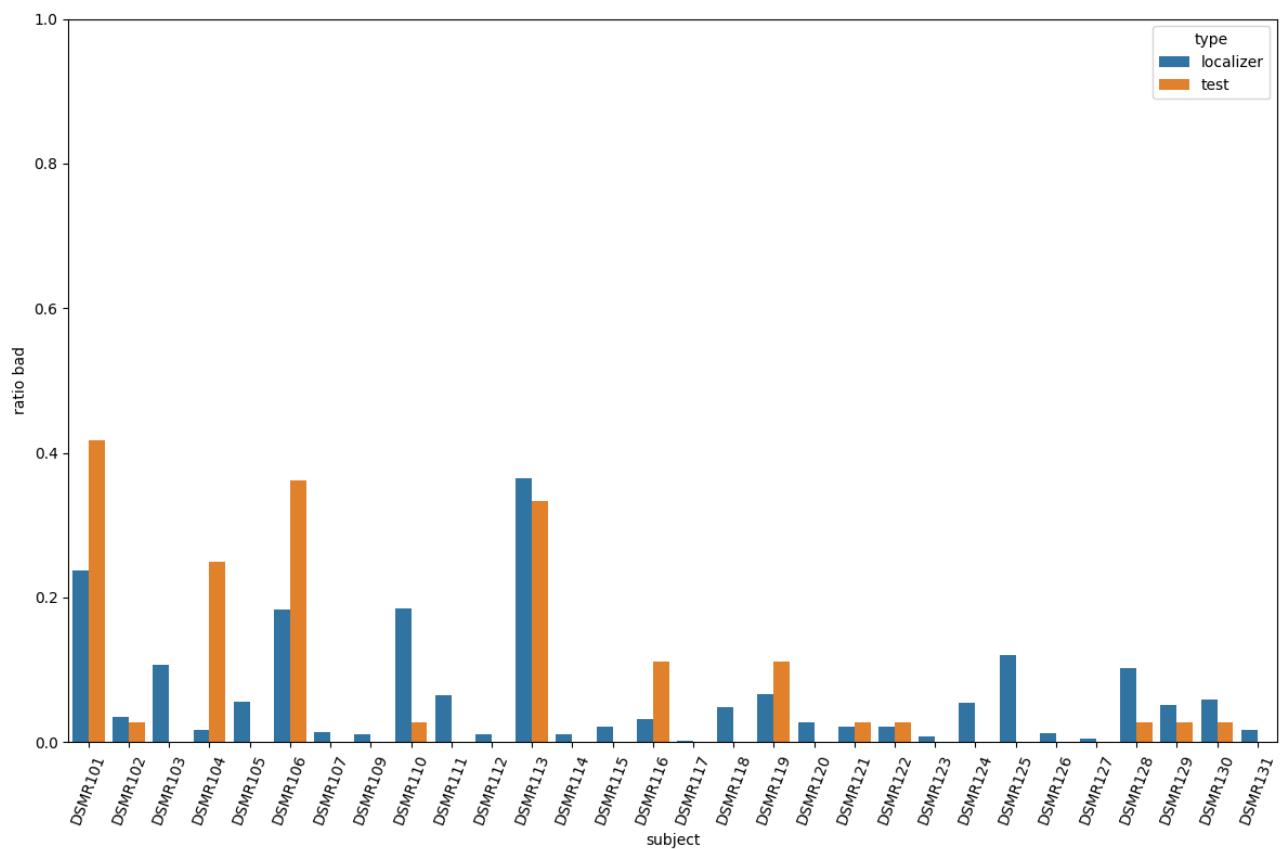
Wise, T., Liu, Y., Chowdhury, F., & Dolan, R. J. (2021). Model-based aversive learning in humans is supported by preferential task state reactivation. *Science Advances*, 7(31), eabf9616. <https://doi.org/10.1126/sciadv.abf9616>

Wittkuhn, L., & Schuck, N. W. (2021). Dynamics of fMRI patterns reflect sub-second activation sequences and reveal replay in human visual cortex. *Nature Communications*, 12(1), 1795.

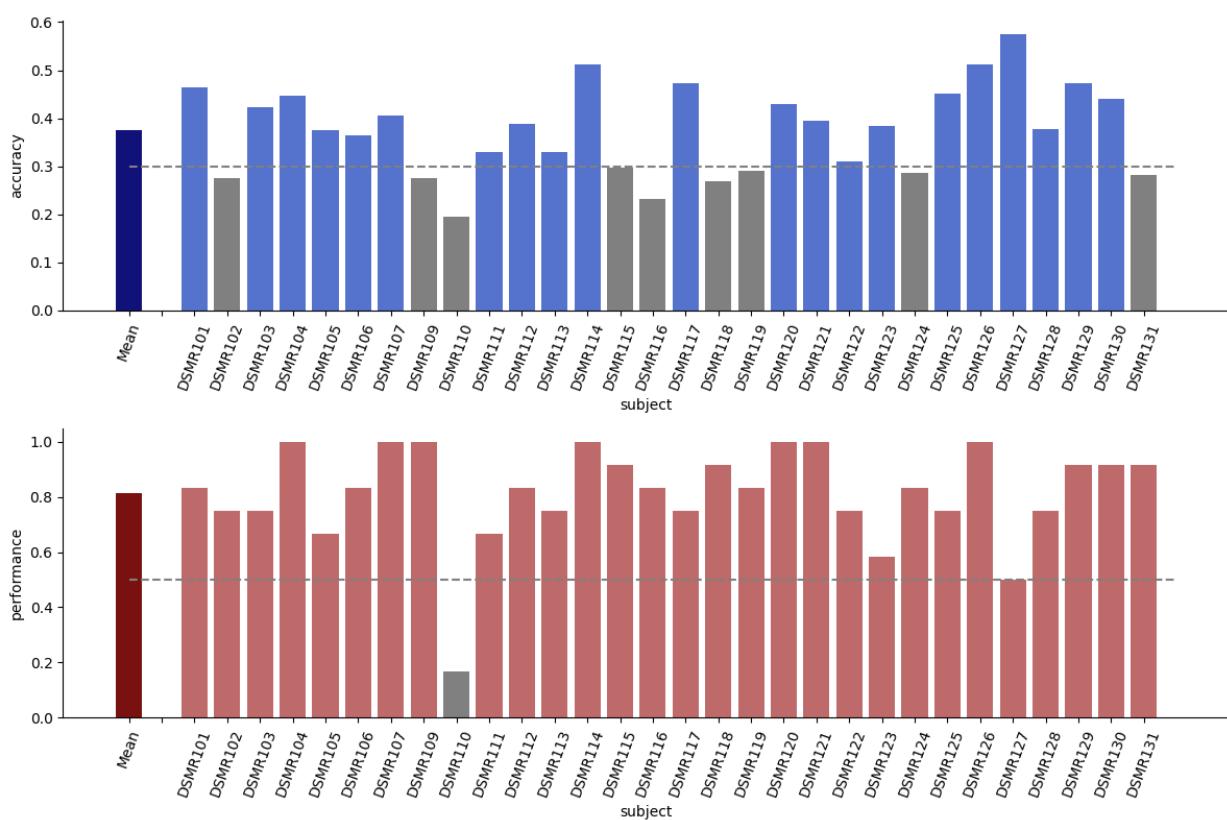
<https://doi.org/10/gjhf54>

Zhang, H., Fell, J., Staresina, B. P., Weber, B., Elger, C. E., & Axmacher, N. (2015). Gamma Power Reductions Accompany Stimulus-Specific Representations of Dynamic Events. *Current Biology*, 25(5), 635–640. <https://doi.org/10.1016/j.cub.2015.01.011>

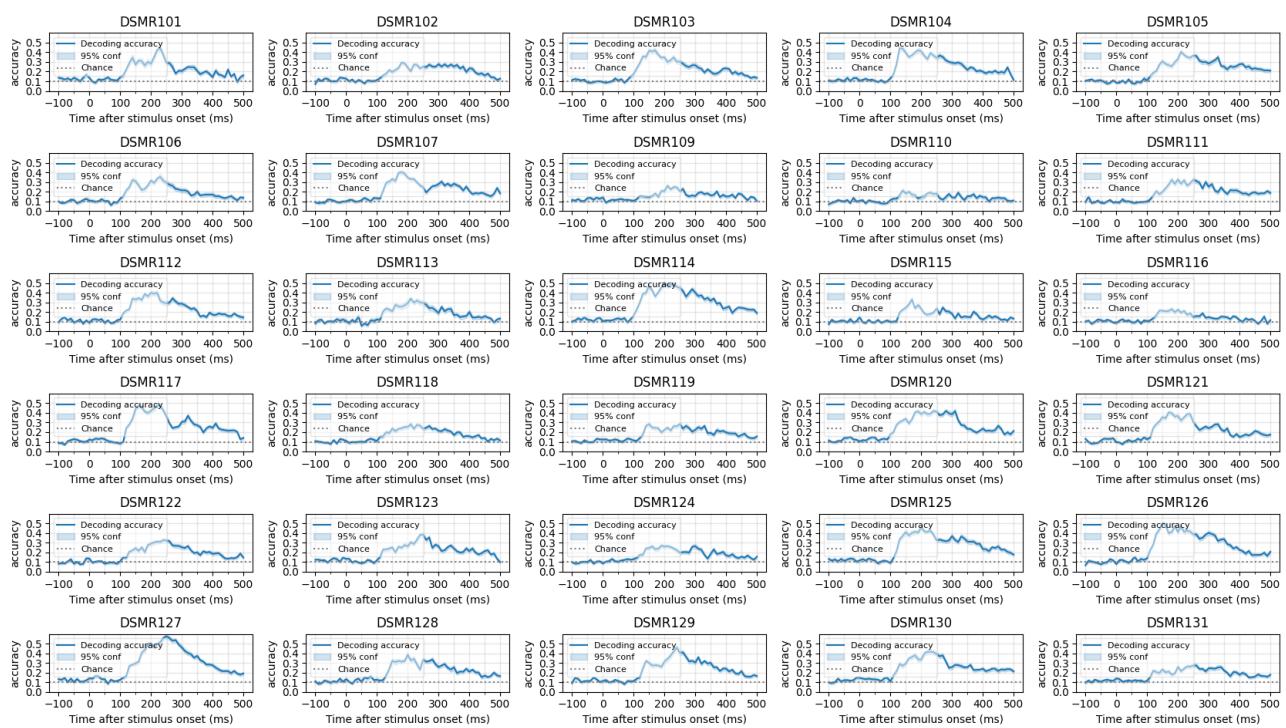
Supplement



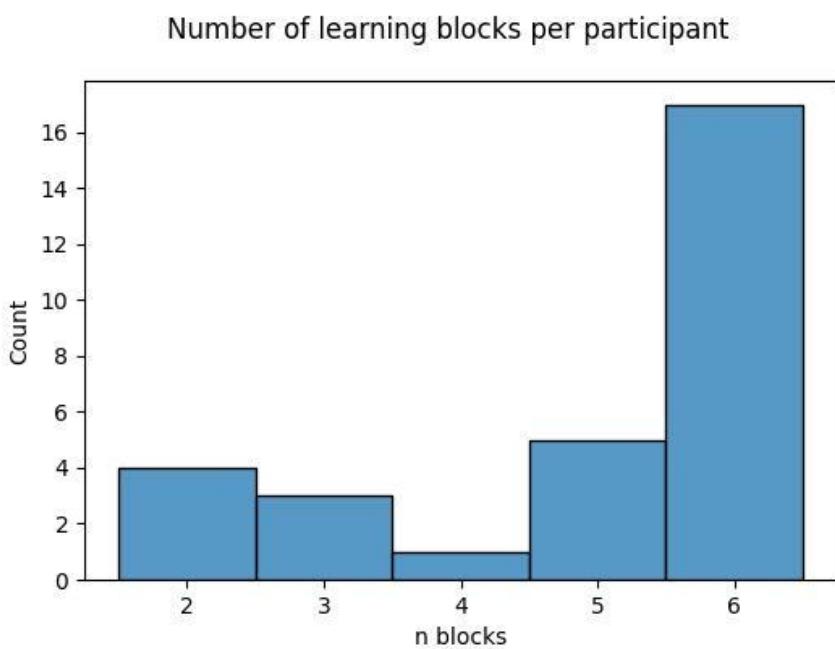
Supplement 1: Percentage of rejected trials for each participant. Artifacts were detected automatically by AutoReject. If possible, channels were interpolated for the affected time span, else the trial was rejected.



Supplement 2: Excluded participants based on decoding accuracy and memory performance during testing.



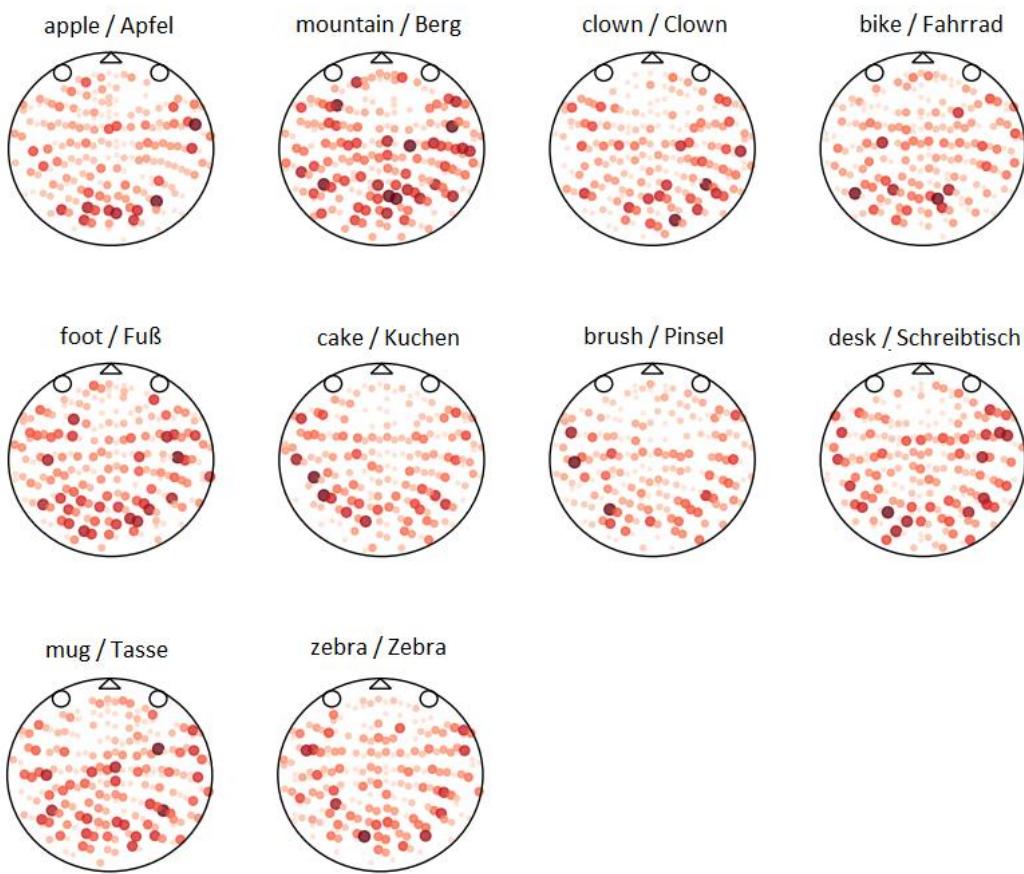
Supplement 3: Decoding accuracy across time determined by a leave-one-out cross-validation per participant.



Supplement 4: Number of learning blocks that each participant completed.

Learning was stopped if participants reached at least 80% memory performance in a block or if they reached 6 blocks.

Sensor distribution for different images



Supplement 5: Percentage of sensors relevant for each image across all participants (beta weight of sensor location unequal to zero). Larger/darker dots indicate more participants' decoders' used information from this sensor. The largest dot indicates that this sensor was used for all participants for this image for this image. The smallest/lightest dot indicates that almost no participant's decoder used information from this sensor.