



Supplementary Materials for

Reactivation of Latent Working Memories with Transcranial Magnetic Stimulation

Nathan S. Rose^{1*}, Joshua J. LaRocque^{1,2}, Adam C. Riggall^{1,3}, Olivia Gosseries¹, Michael

J. Starrett¹, Emma M. Meyering¹, & Bradley R. Postle^{1,3*}

Corresponding authors: nrose1@nd.edu or postle@wisc.edu

This PDF file includes:

Materials and Methods

Fig. S1. Illustration of the Synaptic Theory of Working Memory

Fig. S2. Wholebrain searchlight conjunction analysis used to select TMS targets

Movie S1. Scalp topography of the TMS-evoked response

Fig. S3. Time frequency representation spectrogram for the whole trial timecourse

Supplemental Results

Fig. S4. Mean proportion of correct recognition memory responses

Fig. S5. Decoding the AMI and UMI from EEG using only Theta, Alpha, or Beta

Fig. S6. Timeseries and source localization of significant current density and scatter

Table S1. Average significant current density and scatter, and the Perturbational Complexity Index, of the TMS-evoked response for AMI, UMI, and absent item

Supplemental Material

Synaptic Theory of Working Memory




Fig. S1. Illustration of the Synaptic Theory of Working Memory. When memory items are presented and retained, their representations are activated. Attentional prioritization results in sustained activation of attended memory items (AMIs), but latent, synaptic-based representations of unattended memory items (UMIs) in “prioritized long-term memory”; an exogenous burst of energy via transcranial magnetic stimulation (TMS) can reactivate the latent, potentiated representations of UMIs.

Methods

Participants

Participants (ages 18-34 yrs) were recruited from the undergraduate and graduate student community of the University of Wisconsin–Madison and were paid for their participation. All had normal or corrected-to-normal vision, no reported history of neurological disease, no other contraindications for MRI or TMS, and all gave written informed consent in accordance with the local Institutional Review Board. Experiment 1 enrolled ten participants who performed both task phases in

the MRI scanner, as well as 14 participants who performed the tasks in a behavioral-only testing session (Fig. S1). Experiment 2 had nine participants for the Phase 1 MRI session and six of these participants participated in all three of the TMS-EEG sessions in Phase 2, with 164 trials in each session (N=18 sessions, 2,952 trials). One participant was excluded due to an incidental finding on the MRI; one participant dropped out of the study after the phase 1 MRI session; and one participant was excluded from analysis due to chance level behavioral performance on Phase 2.

Experiment 3 had six participants who participated in both the Phase 1 MRI session and a Phase 2 TMS-EEG session with 192 trials (N=1,152 trials).

Experiment 4 had 20 participants perform the Phase 2 task with TMS (N=7,200 trials) and 23 control participants that performed the Phase 2 task in a behavioral-only testing session (N=8,280, 360 trials each). Data were excluded from analysis for one control participant due to a programming error, and one TMS and two control participants due to a misunderstanding of the task instructions.

Phase 1 – MRI session

Localizer task. In the first phase of each experiment, participants performed a one-item delayed-recognition task in the MRI scanner that required the short-term retention of a face, a word, or a direction of motion (see Fig. 1), and we applied MVPA to identify category-selective regions for subsequent targeting with TMS. On each trial of the task, a fixation cross appeared for 1.5 sec, followed by target stimulus (1 sec), delay period (7.5 sec), recognition probe stimulus (1 sec), response period and feedback (1 sec), followed by a variable inter-trial interval jittered

between 6 and 12 sec in steps of 2 s with a rectangular distribution. Participants performed four runs of 18 trials each (72 total; 24 per category). Trials of the experimental conditions in fMRI experiment (and all subsequent experiments) were randomly mixed within blocks of trials, with short breaks for rest between blocks.

Participants were to indicate if the probe did or did not match the target stimulus in terms of the features of the face, phonology of the word (rhyme), or direction of the moving dots by pressing button 1 or 2 of the MRI compatible response box, respectively. For the recognition probe, the face was either an exact match or the face with 50-70% shared genetic features morphed using FaceGen Modeller Software; the word was either a rhyme or non-rhyme; the patch (15° in diameter) of coherently moving dots (selected to be off the cardinal axes to minimize verbal recoding) were moving in either the same direction or a direction rotated clockwise or counterclockwise randomly between 5-45°.

Data acquisition and preprocessing. Whole-brain images were acquired with a 3-T MR scanner at the Lane Neuroimaging Laboratory at the University of Wisconsin-Madison. For all subjects, a high-resolution T1-weighted image was acquired with a fast spoiled gradient-recalled-echo sequence (8.132 ms TR, 3.18 ms TE, 12° flip angle, 156 axial slices, 256 X 256 in-plane, 1.0 mm isotropic). A gradient-echo, echo-planar sequence (2 s TR, 25 ms TE) was used to acquire data sensitive to the BOLD signal within a 64 X 64 matrix (39 sagittal slices, 3.5 mm isotropic).

The fMRI data were preprocessed offline using the Analysis of Functional NeuroImages (AFNI) software package (27). All volumes were spatially aligned to

the first volume of each run using a rigid-body realignment and corrected for slice time acquisition. Linear, quadratic, and cubic trends were removed from each run to reduce the influence of scanner drift. For univariate analyses, data were spatially smoothed with a 4 mm FWHM Gaussian kernel and warped to Talairach space (28). For MVPA, data were left in their native space (i.e., unsmoothed) and z-scored separately within each run for each voxel.

fMRI analyses. For univariate analysis, the sample, delay, and probe events within each trial were modeled separately as a boxcar of the corresponding duration. These regressors were convolved with a canonical hemodynamic response function and entered into a general linear model (GLM) for analysis using AFNI. MVPA was performed using the Princeton Multi-Voxel Pattern Analysis (www.pni.princeton.edu/mvpa) toolbox and custom routines in MATLAB. Preprocessed MRI data from individual trial time points were used to train a classifier to distinguish between the three stimulus categories. First, leave-one-trial-out cross-validation with L2 regularized logistic regression was used to classify stimulus category on each TR based on the signal from the top 2000 voxels identified by the omnibus F-test. The λ penalty term was set to 25 (based on prior analyses (11)) in order to reduce the contribution of less informative voxels to classification. Classifier accuracy was based on the proportion of trials in which the classifier chose the correct category label. Separate classifiers were trained and tested on each time point in the trial in the leave-one-trial-out cross-validation scheme. The result of this procedure is a time course of decoding accuracy for the

entire trial, which can reveal the evolution of a memory representation over time including a delay period (Fig. 1B).

Then separate whole-brain searchlight analyses were conducted for each subject (Experiments 2 and 3) to classify both the presence vs. absence of each category and the presence of each category vs. each of the other categories on one post-sample delay-period TR (time = 7 s) with a 3 mm sphere. The purpose of this analysis was to identify category-sensitive regions that could accurately discriminate one category from the others (e.g., face vs. not a face) at a high level (e.g., >75%) yet could not discriminate the other categories from one another (e.g., word vs. motion 40-60%) (Fig. 1C, Fig. S2). Note that we selected TMS targets based on the peak decoding accuracy of non-overlapping areas. For Experiment 2, the target was defined separately for each category for each individual.




Fig. S2. Results of the wholebrain searchlight analysis that was used to select non-overlapping

areas to target with TMS for a representative subject. The color-coding depicts areas that could decode one category from the others (e.g., word vs. not word in red) at >70% accuracy; overlapping areas are indicated according to the color-coded Venn diagram.

The MVPA-defined TMS target for Experiments 3 and 4 was the right precuneus, (MNI 11 -81 44) near the bank of the occipito-parietal sulcus in posterior parietal cortex—the motion-selective area for all subjects in Experiment 2 (>70% wholebrain searchlight classification accuracy, $t > 9.0$, $p < .0001$).

Phase 2 – fMRI (Experiment 1), Simultaneous TMS/EEG (Experiments 2 & 3), or Behavioral only (Experiment 4)

Experimental task

Participants performed a two-item delayed recognition task with two retrocues and two recognition probes per trial. A fixation cross appeared (.5 sec) and then two items from two of the three categories were presented as to-be-remembered items (Exp. 1 = 1 sec; Exps. 2, 3, & 4 = 2 sec), followed by an initial delay period (Exp. 1 = 7 sec; Exp. 2 = 5 sec; Exps. 3 & 4 = 2 sec), cue (.5 sec), a first post-cue delay period (4.5 sec), probe (1 sec), response & feedback (1.5 sec), a second cue (.5 sec), a second post-cue delay period (4.5 sec), probe (1 sec), response & feedback (1.5 sec), and an inter-trial interval randomly jittered between 2 and 4 sec. A single pulse of TMS was delivered during the first post-cue delay period at either 2.5 s after the cue offset (Exps. 2 & 4) or randomly jittered between 2-3 sec after cue offset in steps of 50 ms (Exp. 3) and then again during the second post-cue delay period (Exps 3 & 4, with the same timing parameters as the first post-cue

delay period). All tasks were programmed in Matlab with Psychtoolbox functions to control stimulus presentation and record response accuracy and timing.

TMS targeting and stimulation

TMS was targeted with a Navigated Brain Stimulation (NBS) system that uses infrared-based frameless stereotaxy to coordinate the position of the coil and the participant's head according to the individual's high-resolution MRI overlaid with the TMS target for each category from the MVPA of the Phase 1 localizer task. TMS was delivered with a magnetic stimulator fit with a focal single-pulse, figure-of-eight stimulating coil. NBS allows measurement of the electrical field induced by TMS at the targeted cortex using a model of the participant's head, the coil position, and the estimated distance from the coil to the cortical target. TMS was delivered to the individually defined targets to achieve an estimated intensity at the stimulation target of 90–115 V/m (50–68% of stimulator output, depending on the thickness of the subject's scalp, cortex and depth of the target). The coil was oriented along the sagittal plane to induce an anterior-posterior direction of current, with individual adjustments to minimize artifact. Stimulator intensity, coil position, and coil orientation were held constant for each participant for the duration of each session. To mask the sound of the TMS coil's discharge, participants listened to white noise through earbuds during testing at a titrated volume such that the participant could not detect the click using a staircase thresholding procedure. The stimulation parameters were in accordance with published TMS guidelines.

EEG recording

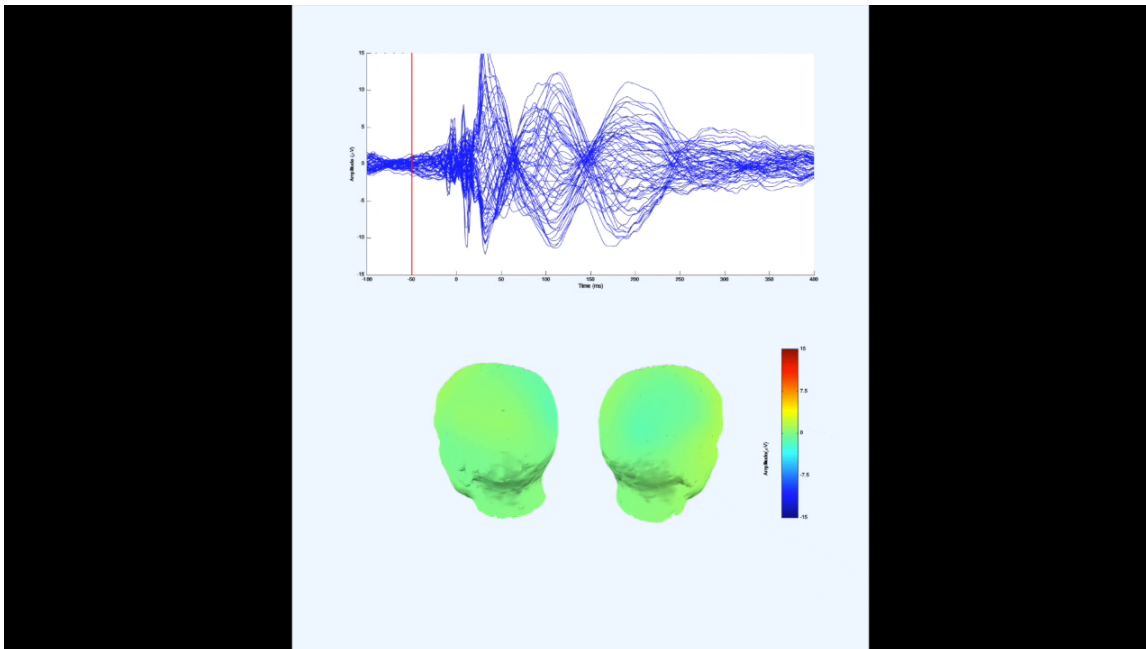
EEG was recorded with a 60-channel cap and TMS-compatible amplifier, equipped with a sample-and-hold circuit that held amplifier output constant from 100 μ s before stimulation to 2 ms after stimulation. Electrode impedance was kept below 5 k Ω . The reference electrode was placed superior to the supraorbital ridge. Eye movements were recorded with two additional electrodes placed near the outer canthi of each eye. The EEG was recorded between 0.1 and 350 Hz at a sampling rate of 1450 Hz with 16-bit resolution.

EEG preprocessing

Data were processed offline using EEGLAB (29) and Fieldtrip (30) toolboxes in MATLAB, downsampled to 500 Hz, and bandpass filtered between 1 and 100 Hz with a notch filter centered at 60 Hz. A spherical spline interpolation was applied to electrodes exhibiting excessive noise. In order to minimize the TMS spike artifact, data were interpolated from 0-30 ms around the TMS pulse for Experiment 2. A TMS pulse was applied on all trial types of the TMS-EEG experiments, which rules out an artifactually-based explanation of a TMS-induced reactivation effect. Nonetheless, in order to minimize the TMS spike artifact, data were interpolated from 0-30 ms around the TMS pulse for Experiment 2. Although this had no effect on the MVPA decoding of power in the binned frequency data, filtering produced some small distortions at the edges of the interpolated window (Movie S1). To better retain the shape and timing of the waveform around the TMS pulse for Experiment 3, a median filter was applied to -5 to 20 ms. Reassuringly, the results were not appreciably

different (31) (for further details on preprocessing steps for TMS-EEG analysis, see <http://nigelrogasch.github.io/TESA/>). Independent components analysis was used to identify and remove components reflecting residual muscle activity, eye movements, blink-related activity, and residual TMS-related artifacts. Finally, the data were rereferenced to the average of all electrodes.

A Morlet wavelet transform was performed on the data using the Fieldtrip toolbox (30). A wavelet at every integer frequency from 2 to 20 Hz and every other integer from 22 to 50 Hz was used with a fixed, Hanning-tapered window of 0.5 sec as in (7). This transform resulted in spectral power values at each of 34 frequencies and 60 channels, sampled every 0.5 sec, for each trial. The spectral time series was smoothed by averaging each value with the two preceding and two subsequent time points, such that each time point reflected the average EEG data from a 2.5-sec window of activity. As in (7), this temporal smoothing procedure was necessary to minimize the noise of the dynamic EEG signal. These data were the features used for all subsequent pattern classification analyses.



Movie S1. The movie shows the scalp topography of the TMS-evoked response for one of the stimulation sites for a representative subject.

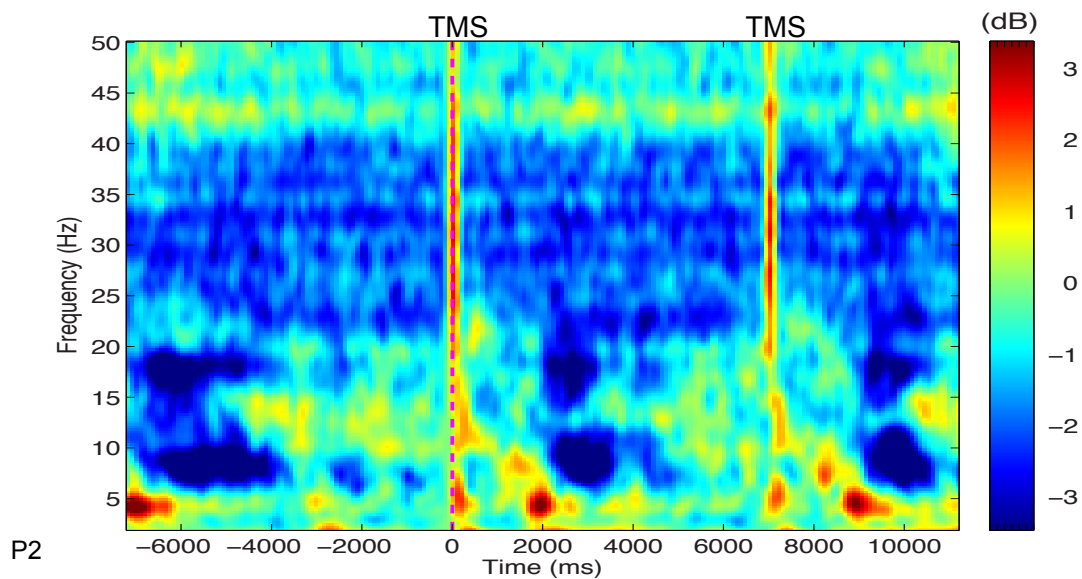


Fig. S3. The time frequency representation spectrogram for the whole trial timecourse from Experiment 3 time-locked to the first TMS pulse in electrode P2. The figure shows event related spectral perturbations associated with the presentation of the stimuli (~ 5 Hz) and administration of TMS across most frequencies. The spectral data depict the features that were used in the EEG decoding analyses.

MVPA decoding

MVPA was performed in MATLAB using the EEG Analysis Toolbox (32) (<https://github.com/mortonne/aperture>) together with the Princeton MVPA toolbox (<https://github.com/PrincetonUniversity/princeton-mvpa-toolbox>). The classification algorithm used for this analysis was L2-regularized logistic regression, with a penalty term set to 1 based on prior research (7). Our analysis scheme trained and tested each 0.5 sec time point separately to construct a decoding time series for the trial. Each exemplar was described by a feature matrix of 34 frequencies by 60 channels. Each feature was z-scored across all trials and time points. The leave-one-out cross-validation scheme trained a classifier on data from 191 trials and then used this classifier to test the one withheld trial. This process was repeated until every trial had been held out for testing.

To measure the classifiers' ability to distinguish trials when a given category was present (and either cued or uncued) from trials in which the category was absent, we used the area under the receiver operating characteristic curve (AUC) as a metric (33, 34). Accuracy was determined by scoring the classifier's chosen category against the actual category, the conventional measure of accuracy. AUC was computed separately for each stimulus condition (i.e., for AMI and for UMI). For the AMI condition, for each stimulus category (i.e., faces, motion, words) we computed the proportion of "hits" (e.g., for faces, the proportion of trials when a face was the AMI and the classifier identified "face" as being present in the signal) vs. "false alarms" (again for faces, the number of trials when a face was not in the memory set yet the classifier identified "face" as being present in the signal). The

analogous procedure was used for the UMI condition. An AUC of greater than 0.5 indicates sensitivity to the category of interest. Statistical significance of sensitivity was assessed with both Bayesian analyses and one-tailed, one-sample t-tests of AUC values averaged across the three categories versus the zero sensitivity value of 0.5, using data from the three sessions for each of the six participants. Classifier evidence was averaged separately across cue-switch and cue-repeat trials for the initially cued, initially uncued, and absent categories (Fig. 3 and Fig. 4, left panel). Note that in Fig. 4, the analysis involved decoding the UMI on the first cue period of the trial and that, after the first memory probe, the UMI would become the AMI in the second cue period on half of the trials (i.e., cue switch trials). For Experiment 3, in which TMS was also administered in the second post-cue delay period, the analysis was repeated to decode the category of the item that was cued second or uncued second (or absent) during the latter half of the trial (Fig. 5, right panel).

Hypotheses regarding the reactivation effect on decoding accuracy were tested with Bayesian analyses to specifically compare evidence favoring the null versus the alternative hypothesis using the formula advocated by (35), with the unit-information scaled-prior conservatively set to the default value. In Experiment 2, for the timepoint immediately following the TMS pulse, classification of the UMI significantly differed from chance, $t(17) = 2.42$, $p = .01$, which resulted in a Bayes Factor (BF) of 3.64, which is considered moderate evidence in favor of the alternative (i.e., against the null) (35). Additionally, the reactivation effect for the TMS targeted category, $t(17)=2.53$, $p=.01$, $BF=4.02$, was larger and more prolonged than for the TMS non-targeted category, $t(17)=2.0$, $p=.03$, $BF=1.72$.

In Experiment 3, for the timepoint immediately following the first TMS pulse, classification of the UMI significantly differed from chance, $t(5) = 4.35$, $p < .005$, resulting in a Bayes Factor of 9.8, which is considered rather strong evidence against the null (35). In contrast, for the timepoint immediately following the second TMS pulse, when the UMI was no longer relevant, classification of the UMI did not differ from chance, $t(5) = 0.2274$ ($p = .42$), resulting in a Bayes Factor of 3.4 in favor of the null. For further confirmation, analyses were also conducted using the Bayes Factor calculations recommended by (36). For the timepoint immediately following the first TMS pulse, with the upper bound of the uniform population distribution set to the maximum classification of the UMI before the first cue (59% classification accuracy), the Bayes Factor was 3,035.84 against the null. In contrast, for the timepoint immediately following the second TMS pulse, with the upper bound of the uniform population distribution set to the maximum reactivation effect of the UMI for all subjects following TMS after the first cue (64% classification accuracy), the Bayes Factor for the timepoint immediately following TMS after the second cue was 3.33 in favor of the null.

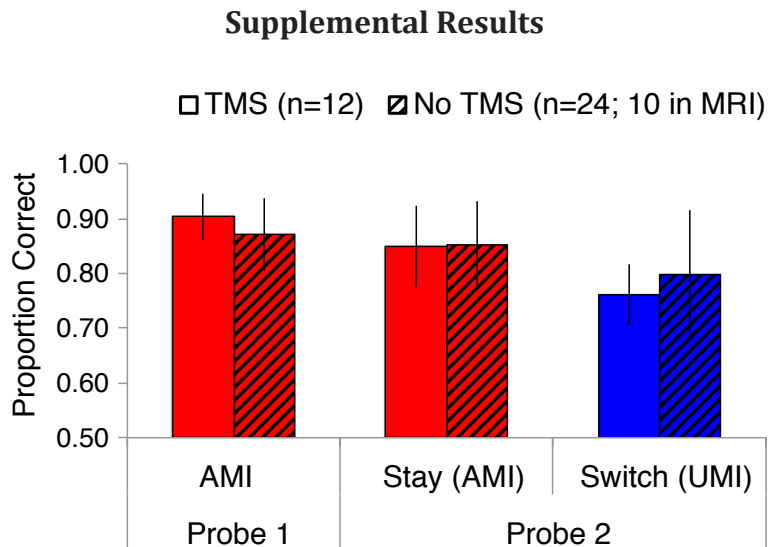


Fig. S4. Mean proportion of correct recognition memory responses on Probe 1 and Probe 2 on the phase 2 task for participants with or without single-pulse TMS during the delay period(s). Relative to Probe 1, recognition accuracy was slightly poorer on probe 2, especially for cue-switch trials (i.e., testing the item that was the UMI during the initial post-cue delay period, but there was no effect of TMS on recognition memory performance $F < 1.0$.

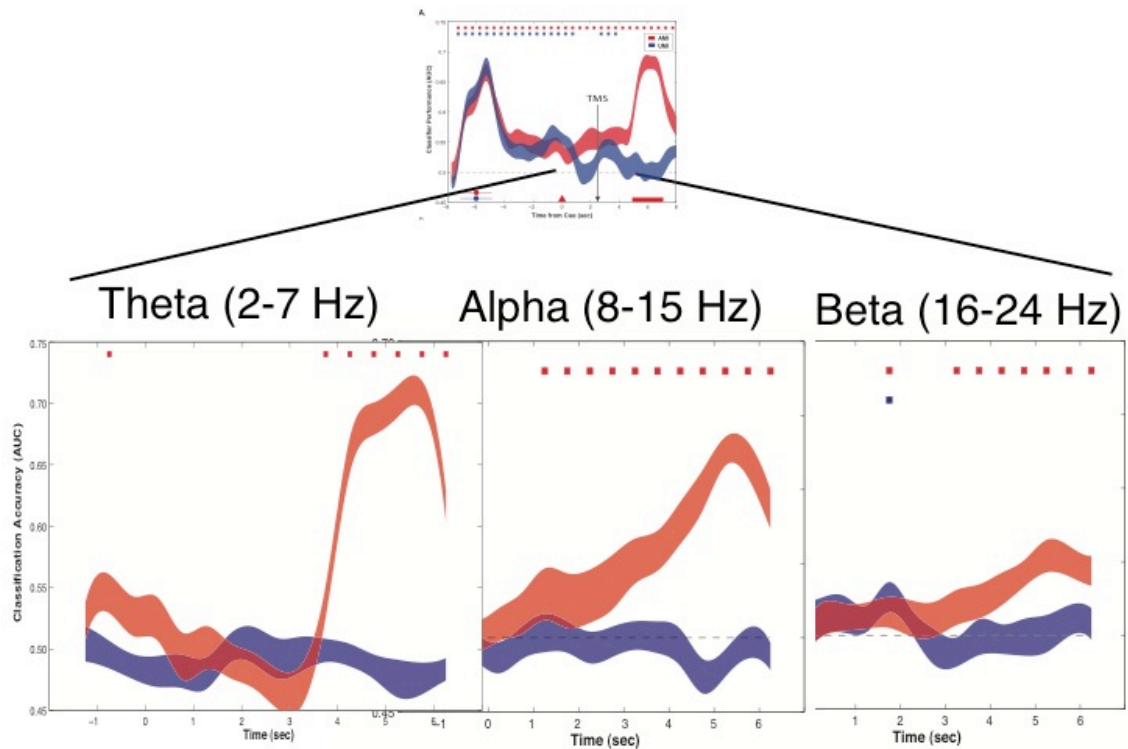


Fig. S5. Features that are important for decoding category specific representations of the AMI and UMI from the EEG: the AMI could be decoded from oscillations in the theta-band during the recognition memory probe period; the AMI could be decoded from alpha-band activity during the delay and memory probe phases; both the AMI and UMI could be decoded from the beta-band following administration of TMS. Thus, oscillations in the beta-band underlie the TMS-induced reactivation effect.

Supplemental Analyses for Experiment 4. To examine the effect of TMS on the ability to reject the UMI nonmatch probes, a mixed, repeated-measures ANOVA on recognition performance on nonmatch probes with stimulus type (UMI, AMInm), probe number (1, 2), and TMS condition was conducted on the data presented in Fig. 5C. There was a significant main effect of stimulus type, $F(1,39)=15.37$, $p<.001$,

as well as a significant interaction between stimulus type and TMS, $F(1,39)=6.87$, $p=.01$, because the UMI nonmatch probes were easier to reject than the AMI nonmatch probes for all contrasts except for Probe 1 with TMS, meaning that there was an increase in false alarms to the UMI probes following TMS on the first recognition memory test when the UMI was still potentially relevant on the trial (Probe 1: TMS UMI=.86, AMInm=.87, $p=.35$; Control UMI=.92, AMInm=.85, $p<.001$; Probe 2: TMS UMI=.85, AMInm=.81, $p<.05$; Control UMI=.87, AMInm=.79, $p<.001$). There was also a significant main effect of probe number because recognition memory performance was better overall for Probe 1 than Probe 2, $F(1,39)=26.93$, $p<.001$.

The critical contrast between Probe 1 UMIs for TMS vs. Controls, $t(39)=2.33$, $p=.01$, was also supported Bayesian analysis: Bayes Factor = 3.50 in favor of the alternative; for Probe 2 UMIs for TMS vs. Controls, $t(39)=0.48$, $p=.32$, the Bayes Factor was 2.24 in favor of the null.

Source Localized Analyses of the TMS-Evoked Response. In an attempt to clarify the anatomical substrates of working memory contents, we conducted source localization analyses of the TMS evoked response on the data from Experiment 2, in which unique TMS targets were utilized for each stimulus category. For each of the targeted stimulus categories, synthetic measures of the brain's response to TMS were computed separately for trials in which the targeted category was represented in each state (attended, unattended, absent). The significant current density (SCD) and significant current scatter (SCS) were computed using the methods detailed in

(37), and the perturbational complexity index was computed using the methods detailed in (38).

First, epochs from -100 to 400 ms post-TMS were source localized using the 3-sphere BERG forward model, modeled using 3,000 meshes derived from each subject's anatomical MRI. The global SCD measure is the sum of the absolute amplitude of all significant TMS-evoked currents (in microamperes per square millimeter) observed between 0 and 400 ms post-TMS in all cortical regions, identified using a nonparametric statistical procedure with $n=70,000$ bootstrapped samples relative to the prestimulus baseline (-100-0) (see (37) for further details). Global SCS is the sum of the geodesic distances (expressed in mm) between the stimulated brain region and significant current sources in distal brain regions. PCI uses the Lempel-Ziv complexity algorithm on the SCD and SCS measures to gauge the amount of non-redundant information in the matrix of significant currents scaled from 0 to 1 (see (38) for further details).

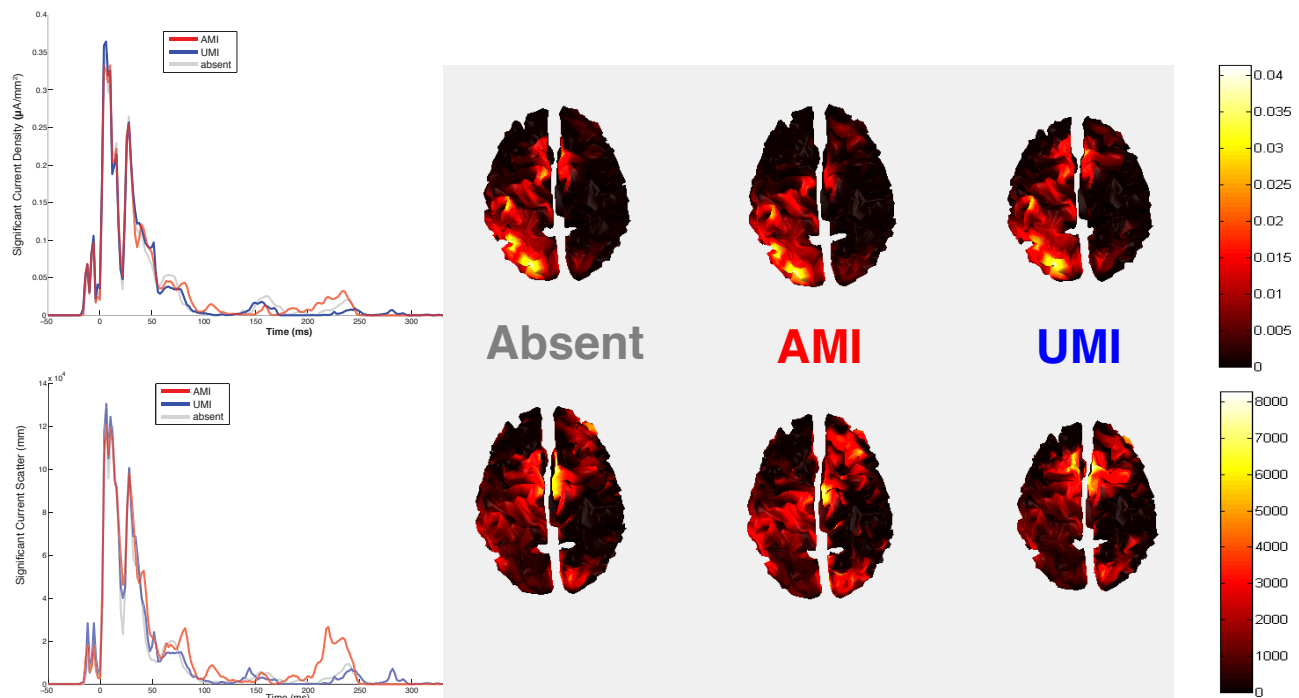


Fig. S6. Timeseries (left panel) and source localization (brain maps in the right panel) of the average significant current density (top row) and significant current scattering (bottom row) of the TMS-evoked response on trials in which the targeted category was the attended memory item (AMI), the unattended memory item (UMI), or absent.

Table S1. Synthetic measures of the TMS-evoked response.

Stimulus	Absent	Attended	Unattended
<i>Significant Current Density ($\mu\text{A}/\text{mm}^2$)</i>			
Motion	6.9 (2.0)	6.0 (1.3)	5.3 (1.2)
Face	5.4 (1.8)	7.9 (2.8)	8.2 (2.7)
Word	7.7 (3.2)	7.5 (2.7)	8.9 (3.0)
Mean	6.7 (1.9)	7.1 (1.6)	7.4 (2.1)
<i>Significant Current Scatter (mm)</i>			
Motion	4.57E+06 (1.22E+06)	3.50E+06 (6.47E+05)	3.22E+06 (7.84E+05)
Face	3.19E+06 (1.01E+06)	3.73E+06 (8.40E+05)	4.03E+06 (9.71E+05)
Word	4.31E+06 (2.01E+06)	4.32E+06 (1.85E+06)	5.17E+06 (2.29E+06)
Mean	4.02E+06 (1.41E+06)	3.85E+06 (1.11E+06)	4.14E+06 (1.35E+06)

*Perturbational Complexity Index (PCI)**

Motion	0.442 (0.036)	0.418 (0.023)	0.405 (0.032)
Face	0.402 (0.041)	0.376 (0.033)	0.378 (0.056)
Word	0.399 (0.050)	0.393 (0.052)	0.425 (0.050)
Mean	0.415 (0.037)	0.396 (0.032)	0.403 (0.040)

These synthetic measures of the TMS-evoked response failed to distinguish between any of the WM states of representation. All main effects and interactions on the global SCD, SCS, and PCI measures failed to reach significance ($ps > .05$). Although these synthetic measures have been successfully used to discriminate between varying states of consciousness, e.g., coma or non-REM sleep vs. wakefulness (38), or fixation vs. the perceptually identical delay period of a WM task (37), it seems unlikely that these global measures of the brain's response to TMS during WM retention could distinguish between different states of accessibility for an individual item that is represented either in an attended state or an unattended state from items that were attended on previous trials. Thus, the MVPA decoding methods were more sensitive to detecting multivariate patterns of brain activity associated with stimulus-specific representations than the synthetic measures of voltage on the scalp in source-localized space. To address more structural hypotheses about the precise anatomical substrates of representations in working memory, future work should more directly probe for the presence of stimulus specific representations in sensory or associative cortices. More invasive methods than those used in the current work may be necessary to sufficiently address such hypotheses.