Supporting Information

Manning et al. 10.1073/pnas.1015174108

SI Materials and Methods

Intracranial Recordings. ***Note on intracranial recordings.*** Despite the high data quality provided by human intracranial recordings, there are several factors one should consider when interpreting the results of any intracranial study of human epilepsy patients, including those we report in the present study. First, whereas in animal studies, electrodes are placed according to researchers' needs, the placements of implanted human electrodes are determined solely by clinical teams with the goal of localizing the seizure focus to ensure the best possible outcome for the patient. For some patients, this means that the brain areas most relevant to a particular research question may receive little or no electrode coverage. To obtain adequate coverage of all relevant brain areas, we have analyzed data from many patients (Table S1). A second concern is that medications or recent seizures might change the electrophysiological properties of the brain. For this reason, we refrained from collecting data while the patients were on high dosages of pain medications or antiepileptic drugs, or during the 6-h period following any clinically significant seizure. A third issue is that the brain is known to rewire itself to compensate for damage, including damage caused by epilepsy (1), which could lead to cognitive remapping. Although we cannot control for cognitive remapping that may have occurred in individual participants, we have averaged our anatomical analyses over many patients; thus, results attributable to remapping in one patient will average out in the population analyses. A fourth concern is that severe epilepsy can lead to cognitive impairment. To address this issue, we have analyzed data only from patients with scores on the Wechsler Intelligence and Wechsler Memory Scales within 1.5 SDs of the mean for their age group.

Recording methods. Subdural grids or depth electrodes (Ad-Tech, Inc.) were implanted by neurosurgical teams solely for clinical purposes. The locations of the electrodes were determined by means of coregistered postoperative computed tomography and preoperative MRI scans, or from postoperative MRI scans, by an indirect stereotactic technique and converted into Montreal Neurological Institute coordinates. ECoG signals were recorded referentially using a Telefactor, Bio-Logic, XLTek, Neurofile, or Nicolet electroencephalographic digital video-EEG system. Depending on the amplifier, signals were sampled at 200, 256, 500, 512, or 1,024 Hz. Several hospitals applied band-pass filters to the recorded signals before writing to disk (Table S2). Where applicable, frequencies outside of the filtered range were excluded from further analysis. Data were subsequently notch-filtered with a Butterworth filter with zero phase distortion at 50 or 60 Hz to eliminate electrical line and equipment noise. ECoG signals and behavioral events were aligned using synchronization pulses sent from the testing computer (mean precision <4 ms).

Analysis Methods. ***Quantifying the contiguity effect.*** Fig. 4*C* depicts an analysis relating the neural reinstatement effect to the recall behavior of the participants. Specifically, we show that participants showing stronger neural reinstatement effects tend to exhibit a stronger contiguity effect (whereby neighboring list items tend to be recalled successively). The contiguity effect is measured using the temporal clustering score, an analysis technique described previously (2). The temporal clustering score is calculated as follows.

For each recall transition, we create a distribution of temporal distances between the just-recalled word and the set of words that have not yet been recalled. These distances are simply the absolute value of the difference between the serial position of the just-recalled word and the set of words that have not yet been recalled. A percentile score is generated by comparing the temporal distance value corresponding to the next item in the recall sequence with the rest of the distribution. Specifically, we calculate the proportion of the possible distances that the observed value is less than, because strong temporal clustering will cause observed lags to be smaller than average. As is often the case, when there is a tie, we score this as the percentile falling halfway between the two items. If the participant always chose the closest temporal associate (which is only possible for pure serial recall in the forward or backward direction), the temporal clustering score would yield a value of 1 (because there would never be an opportunity for a tie). A value of 0.5 indicates no effect of temporal clustering. Each patient was assigned a temporal clustering score by taking the average of the percentile scores across all observed recall transitions.

Quantifying the primacy and recency effects. The primacy and recency effects refer to an enhancement in memory for early and late list items, respectively, compared with memory for intermediate list items (3, 4). The number of items that show a boost in memorability attributable to primacy or recency is relatively invariant to changes in list length; the primacy effect generally affects the first few items, whereas the recency effect generally affects the last six or so items (4). To measure the strength of the primacy effect, we labeled the first three serial positions on each list as primacy positions and the last six serial positions as recency positions. The remaining positions were labeled as intermediate list positions (i.e., items 4–9 for 15-word lists, items 4–14 for 20-word lists). We then measured the strength of the primacy effect for each participant by dividing his or her mean probability of recalling items from primacy positions by his or her mean probability of recalling items from intermediate list positions. The main text reports that the neural signature of context reinstatement (t value) is not correlated with the strength of the primacy effect ($r = 0.13$, $P = 0.42$).

We also performed an analogous analysis to test whether the neural signature of context reinstatement was influenced by the factors underlying the recency effect. We measured the strength of the recency effect for each participant by dividing his or her mean probability of recalling items from recency positions by his or her mean probability of recalling items from intermediate list positions. The neural signature of context reinstatement is not correlated with the strength of the recency effect ($r = 0.13$, $P = 0.40$). Mean serial position curves for participants showing strong (top 50%) and weak (bottom 50%) neural signatures of context reinstatement (by t value) are shown in Fig. S2. As shown in the figure, the primacy effect, recency effect, and overall probability of recall are roughly conserved across the two groups of participants.

Simulations. We conducted three neural network simulations (Fig. 3 and Fig. S1) that predict the expected outcome of our test for context reinstatement under various model assumptions. As described in the main text, the autocorrelated noise model has neural activity evolve randomly over time, irrespective of what is happening in the experiment. The content reinstatement model has each neuron represent a different word; a neuron is activated if its associated word is presented or recalled. The context reinstatement model also has each neuron represent a different word. We simulate context reinstatement by activating not only the neuron associated with the word being recalled, but also

other neurons that were active at the time the recalled word was studied.

For all three simulations, we define an activity vector, $\mathbf{f}$, that defines the pattern of activation across the network. Each neuron in the network takes on a value between 0 (inactive) and 1 (maximally active). Let $\mathbf{f}_i$ denote the state of $\mathbf{f}$ after the ith experimental event (i.e., study presentation, distracting task, recall). Our main analysis entails selecting autocorrelated components of neural activity as the candidate context representation (*Results*). After this feature selection, the feature vectors we analyze are autocorrelated, a property we need to take into account in our simulations. In particular,

$$\mathbf{f}_i = \rho_i \mathbf{f}_{i-1} + \beta \mathbf{w}_i,$$

where β is a constant; ρ_i is a function of $\mathbf{f}_{i-1}$,$\mathbf{w}_i$, and β (with $0 \leq \rho_i$, $\beta \leq 1$); and $\mathbf{w}_i$ is the pattern of neural activity specifically evoked by the ith experimental event [details are presented by Polyn and Kahana (5)]. In this way, the neural activity measured after a given experimental event (e.g., presentation of the fifth list item) is a recency-weighted blend of the activity evoked by previous experimental events (e.g., activity evoked by presentations of items 5, 4, 3, 2, and 1). We initialize $\mathbf{f}_0$ by setting the activation of the first neuron to 1 and the activations of the other neurons to 0. We then simulate different experimental events by adjusting $\mathbf{w}_i$ according to the particular rules of each model. We ensure that $\mathbf{f}_i$ is always of unit length by setting

$$\rho_i = \sqrt{1 + \beta^2 \left[(\mathbf{f}_{i-1} \cdot \mathbf{w}_i)^2 - 1 \right]} - \beta (\mathbf{f}_{i-1} \cdot \mathbf{w}_i).$$

For the autocorrelated noise model, each $\mathbf{w}_i$ is set to a vector of 0's, plus a 1 in a single random position. In this way, each $\mathbf{w}_i$ activates one of the neurons in the network at random. As shown in Fig. S1*A*, for $\beta < 0.5$, similarity between $\mathbf{f}_i$ during presentation and $\mathbf{f}_j$ during recall increases as a function of i. This is because, by definition, an autocorrelated signal measured at times t and $t + \Delta$ becomes more similar as $\Delta \to 0$. For $\beta > 0.5$, similarity as a function of lag flattens out, because as $\mathbf{f}_i$ is dominated by $\mathbf{w}_i$, the average similarity between $\mathbf{f}_i$ and $\mathbf{f}_j$ approaches the expected similarity between two independent draws of $\mathbf{w}_i$.

For the content reinstatement model, $\mathbf{w}$ is set differently depending on the type of experimental event. In this model, each neuron is assigned a different word or distractor. During presentation of study items or distractors, $\mathbf{w}_i$ is set to a vector of all 0's except for a 1 in the position of the neuron representing the item or distractor being presented. During recall of the jth presented item, we set $\mathbf{w}_i = \mathbf{w}_j$. As shown in Fig. S1*B*, for $\beta < 0.5$, similarity increases as a function of lag. Because β is small, $\mathbf{f}_i$ is dominated by $\mathbf{f}_{i-1}$ rather than $\mathbf{w}_i$. Because the specifics of the experimental event contribute only minimally to $\mathbf{f}$, the simulation approximates the autocorrelated noise simulation. For $0.5 < \beta < 1$, neural similarity is roughly constant as a function of lag for negative lags but decreases as a function of lag for positive lags. This is because the pattern of activation during the ith presentation will only contain traces of $\mathbf{w}_j$ if $i > j$. Finally, for $\beta = 1$, similarity is 1 when lag = 0 and is 0 everywhere else. This is attributable to the fact that when $\beta = 1$, $\mathbf{f}_i = \mathbf{w}_i$; thus, the neural activity evoked by the ith item will be present only during its presentation or recall.

The context reinstatement model is identical to the content reinstatement model during the presentation of study items and distractors. To simulate context reinstatement during recall of the jth presented item, we set $\mathbf{w}_i = \mathbf{f}_j$ (recall that $\mathbf{f}_j$ will contain a recency-weighted average of the activations associated with the previously presented items). As shown in Fig. S1*C*, for $\beta < 0.5$, similarity increases as a function of lag, just as in the other simulations. Importantly, for $0.5 < \beta < 1$, neural similarity decreases with absolute lag in both the positive and negative directions, as seen in the neural data (Fig. 4*A*). Finally, as in the content reinstatement simulation, for $\beta = 1$, similarity is 1 when lag = 0 and is 0 everywhere else.

These simulations show that regardless of the precise rate at which neural activity evolves over time, the simplest model consistent with our neural results (Fig. 4*A*) is one in which the temporal context in which an item is studied is reinstated when the item is recalled. Although we have not ruled out every possible model that does not include some form of context reinstatement, neither autocorrelated noise (Fig. S1*A*) nor content reinstatement alone (Fig. S1*B*) can account for the neural signature of context reinstatement we observed in our ECoG recordings.

Neural symmetry vs. behavioral asymmetry. The neural data (Fig. 4*A*) show that the decrease in neural similarity with absolute lag falls off symmetrically in the forward (positive) and backward (negative) directions. A natural question, then, concerns why the behavioral data exhibit a clear forward asymmetry in the conditional response probability as a function of lag (Fig. 4*B*). In particular, if the neural signature of context reinstatement we observe is truly related to participants' behavior (as implied in Fig. 4*C*), why is the neural signature of context reinstatement symmetrical, whereas the contiguity effect is forward asymmetrical?

Consistent with the neural data, our simulations show that context reinstatement, per se, implies a symmetrical decrease in neural similarity with lag (Fig. 3*F*). Thus, the forward asymmetry in the behavioral data must arise as the result of some additional process that is not captured by our neural analysis. One possibility is that, in addition to reinstating the recalled item's context, the representation of the recalled item itself receives an additional "boost." As described above, reinstating the representation of an item (without its associated context) implies a decrease in neural similarity as a function of lag in the forward direction, but not in the backward direction (Fig. 3*E*). In this way, the behavioral data might reflect both context and content reinstatement [e.g., figure 6 in the article by Howard and Kahana (6)]. However, because we examine only autocorrelated components of neural activity, our neural analysis is (intentionally) biased toward examining neural features related to context rather than neural features related to item representations. An interesting question for future studies will be to clarify the extent to which the context and item representations overlap.

Selecting autocorrelated features. Context-based theories of memory posit the existence of a gradually changing pattern of neural activity that becomes associated with each studied word during study and is reinstated during recall. To identify candidate components of the context representation for a given recording session, we selected autocorrelated PCA-derived features of the neural representation (Fig. 1*C*) as follows. Separately for each feature x, we computed the Pearson's lag 1 autocorrelation coefficient (r) and associated P value for the values of x within each list. We then combined the autocorrelation coefficients across lists into a summary autocorrelation measure, $\bar{r}$:

$$\bar{r} = F^{-1}\left(\sum_{i=1}^{L} F(r_i) \right),$$

where r_i is the Pearson's lag 1 autocorrelation coefficient for the values of x measured during list i and $F()$ is the Fisher z-prime transformation:

$$F(r) = \frac{ln(1+r) - ln(1-r)}{2},$$

and $F^{-1}()$ is the inverse of $F()$:

$$F^{-1}(z) = \frac{e^{2z}-1}{e^{2z}+1}.$$

In this way, if r_i has large positive values across all lists, $\bar{r}$ will have a large positive value. Similarly, if r_i is negative across all lists, $\bar{r}$ will have a large negative value. If r_i is sometimes positive and sometimes negative (with approximately equal probability), $\bar{r}$ will take on a value near zero. (Note that $-1 \le r_i, \bar{r} \le 1$.)

We also obtained a P value, $\bar{p}$, associated with $\bar{r}$ by applying the inverse Normal transformation to the P values associated with the Pearson's lag 1 autocorrelation coefficients for each list. We then summed across the transformed P values and evaluated the cumulative normal distribution function at this sum to obtain $\bar{p}$. We selected features with $\bar{r} > 0$ and $\bar{p} < 0.1$ for further analysis (*Results*).

Identifying time interval of the recall event. Our main analysis (Fig. 4*A*) compares the neural activity elicited by a studied word with the neural activity elicited by a word's retrieval during the recall period. We restrict our analysis of the study period to ECoG activity beginning 200 ms after the appearance of a word and ending when the word disappears from the screen. Here, the 200-ms delay was meant to account for the lag between the word's appearance on-screen and the processing of the word by the participant.

To search for the optimal time interval for the recall event, we tested for context reinstatement while varying both the duration and onset of the time interval for the recall event. We tested time intervals ranging in duration from 100 to 1,000 ms (in increments of 100 ms) and onsets ranging from −1000 to 0 ms (in increments of 100 ms) relative to the time the participant began his or her vocalized recall. This analysis indicates that the context reinstatement effect is strongest for the recall interval ranging from −600 to 200 ms relative to vocalization.

To account for the possibility that different brain regions reinstate context at different times relative to vocalization, we repeated this optimization analysis separately for each region of interest. The best time interval for the temporal lobe was from −400 to −300 ms (Fig. 5*B*). The time interval that gave the strongest frontal lobe effect was from −900 to −400 ms; however, the frontal effect was not statistically reliable (*Results*).

Additional details of selected features. In addition to asking whether specific brain regions contribute to the representation of context (Fig. 5), a natural question is whether the principal components comprising the feature vectors tend to weight particular oscillatory components of ECoG activity more heavily than others. Because PCA performs a linear mapping from the n-dimensional space of the original set of activity vectors onto the m-dimensional PCA space (where $m \le n$), we can use the PCA coefficients to perform the inverse mapping of the feature vectors back onto the original n-dimensional space. The PCA coefficients tell us how much each of the elements in the original principal components vectors contributes to each of the principal components in the feature vectors. This allowed us to determine the degree to which each oscillatory component recorded from each electrode contributes to each element of the feature vectors. For a given frequency band, we assessed the degree to which that frequency band contributed to the feature vectors across all study and recall events by examining the distribution of PCA coefficients assigned to that frequency band across all participants. An analysis of PCA coefficients across frequency bands revealed no significant differences among frequency bands [repeated measures ANOVA: $F(4,37) = 0.57$, $P = 0.69$]. This finding suggests that the selected features are composed of oscillatory activity at a broad range of frequencies (Fig. S3*A*).

1. Ribak CE, Dashtipour K (2002) Neuroplasticity in the damaged dentate gyrus of the epileptic brain. *Prog Brain Res* 136:319–328.
2. Polyn SM, Norman KA, Kahana MJ (2009) A context maintenance and retrieval model of organizational processes in free recall. *Psychol Rev* 116:129–156.
3. Deese J, Kaufman RA (1957) Serial effects in recall of unorganized and sequentially organized verbal material. *J Exp Psychol* 54:180–187.
4. Murdock BB (1962) The serial position effect of free recall. *J Exp Psychol* 64: 482–488.
5. Polyn SM, Kahana MJ (2008) Memory search and the neural representation of context. *Trends Cogn Sci* 12:24–30.
6. Howard MW, Kahana MJ (2002) A distributed representation of temporal context. *J Math Psychol* 46:269–299.

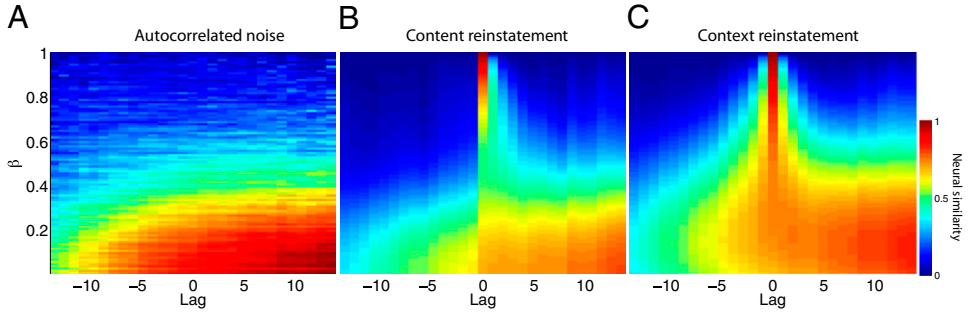

Fig. S1. Simulated neural similarity as a function of lag and drift rate (β), given no content or context information in the neural recordings (*A;* Fig. 3 *A* and *D*), content reinstatement without context reinstatement (*B;* Fig. 3 *B* and *E*), and context reinstatement (*C;* Fig. 3 *C* and *F*). Similarity is computed as the normalized dot product between the simulated feature vector after the recall of the *i*th word and the feature vector corresponding to presentation of word *i* + lag. The first dimension (initialized to 1 before the start of the simulation) was ignored for the similarity calculations. Simulation results in Fig. 3 used β = 0.7 [this choice was motivated by previously reported simulation results (2)].

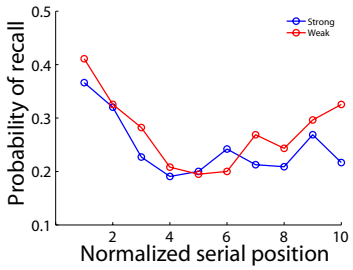

Fig. S2. Serial position curves. The average serial position curve for participants exhibiting the top 50% strongest neural signatures of context reinstatement (by *t* value) is shown in blue. The average serial position curve for participants exhibiting the bottom 50% strongest neural signatures of context reinstatement is shown in red. The serial positions have been normalized such that all participants, regardless of whether they studied 15- or 20-word lists, are shown in the same curves. Normalized serial positions 1–3 (i.e., primacy positions) correspond to absolute serial positions 1–3. Normalized serial position 4 (i.e., intermediate position) reflects the average probabilities of recall for items in serial positions 4–9 (15-word lists) or 4–14 (20-word lists). Normalized serial positions 5–10 (i.e., recency positions) correspond to the last 6 list items.

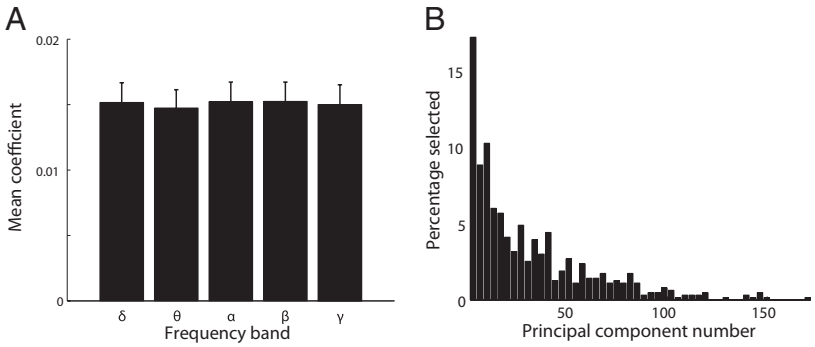

Fig. S3. Selected features. (*A*) Mean contributions of each frequency band to selected features. Error bars indicate 95% confidence intervals (1). (*B*) Percentage of selected features by principal component number. Smaller principal component numbers explain a higher proportion of the variance in the raw neural data.

1. Loftus GR, Masson MEJ (1994) Using confidence intervals in within-subject designs. *Psychon Bull Rev* 1:476–490.

Table S1. Patient and task information

ID	HOSP	AGE, y	SEX	HAND	ELC	FEA	L LEN	SES	LST	REC	REP	PLI	ELI
1	BW	33	F	R	64	3	20	1	15	55	2	26	12
2	BW	51	F	R	40	5	20	1	15	66	2	8	1
3	BW	32	M	R	32	2	15	3	39	212	5	9	4
4	BW	40	M	R	96	5.5	15	2	20	82	5	38	28
5	BW	44	M	R	16	1	15	2	20	58	1	12	19
6	BW	27	M	R	64	1.5	15	2	20	76	49	12	3
7	BW	38	M	R	104	9.33	15	3	30	136	3	30	15
8	CH	13	F	R	64	9	20	1	12	59	0	2	1
9	CH	12	F	R	104	19	20	1	15	39	0	0	2
10	CH	15	M	L	128	15.67	20	3	30	90	1	12	8
11	CH	17	M	R	64	4	20	3	45	178	20	27	17
12	CH	15	M	R	123	13	20	1	15	86	3	6	3
13	CH	11	M	R	104	0	20	2	30	104	3	2	3
14	CH	14	F	R	72	0	20	1	15	104	2	7	5
15	CH	8	F	R	86	6.5	20	2	30	159	5	18	10
16	CH	17	M	R	84	12	20	1	14	30	2	9	12
17	CH	17	M	L	124	10.5	20	4	60	116	2	104	13
18	CH	20	F	R	128	8.5	15	2	24	114	2	9	4
19	CH	14	M	R	94	6.67	15	3	30	94	0	18	13
20	CH	17	M	L	80	7	15	2	20	14	0	10	22
21	CH	19	F	R	125	8.5	15	2	17	47	2	10	1
22	CH	16	M	R	156	13	15	1	16	76	1	4	1
23	CH	12	M	L	83	5	15	2	20	52	9	12	25
24	CH	13	M	R	72	4.75	15	4	40	200	3	4	2
25	FR	33	M	R	98	8	20	1	9	43	0	0	0
26	FR	25	M	R	85	21	20	1	9	45	18	1	2
27	FR	31	M	L	56	4	20	1	9	24	0	4	0
28	FR	41	F	R	63	9	20	1	7	23	0	9	6
29	FR	34	F	L	40	3	20	1	7	38	3	5	3
30	FR	45	F	L	100	10	20	1	8	27	0	20	3
31	FR	46	F	L	14	0	20	1	1	4	3	0	5
32	FR	20	M	R	84	4	20	1	15	42	6	6	2
33	FR	53	F	L	41	5	20	1	15	49	26	21	21
34	FR	50	M	R	68	4	20	2	30	116	5	36	15
35	FR	28	M	L	112	11	20	1	15	37	1	1	1
36	FR	30	F	R	60	7	20	1	15	67	9	10	2
37	FR	37	F	L	30	6	15	1	20	65	84	73	42
38	FR	18	M	L	30	2	15	1	20	121	7	12	13
39	FR	23	M	L	58	3	15	4	56	281	98	32	29
40	FR	21	M	L	93	4	15	1	10	49	2	10	3
41	FR	28	F	R	86	7	15	1	10	36	4	4	6
42	FR	35	F	L	122	3	15	2	20	54	0	6	6
43	FR	37	F	L	52	2.5	15	4	35	161	38	29	28
44	FR	19	M	L	74	2	15	2	30	148	19	14	35
45	FR	41	F	R	30	3	15	1	15	15	0	3	38
46	FR	21	F	R	64	6	15	1	15	50	0	3	7
47	FR	43	F	R	56	0	15	1	15	23	5	11	59
48	FR	19	M	R	30	3	15	2	25	120	2	5	27
49	FR	21	M	R	70	5	15	5	53	408	41	8	194
50	FR	35	F	R	62	6	15	1	15	44	24	16	114
51	FR	25	M	R	84	5	15	2	30	145	5	3	96
52	FR	47	M	L	82	4	15	1	4	13	1	3	35
53	FR	45	F	R	88	4	15	1	10	43	11	3	0
54	TJ	25	M	R	62	0.67	15	3	48	232	3	6	1
55	TJ	40	F	R	94	3.25	15	4	64	164	8	54	43
56	TJ	39	M	L	56	2	15	1	16	53	1	8	20
57	TJ	34	F	R	111	4.3	15	10	154	513	7	110	24
58	TJ	44	M	R	125	7	15	1	13	31	1	7	6
59	UP	38	M	R	62	4.75	15	4	40	135	3	68	24
60	UP	30	M	R	86	3	15	2	20	54	5	24	21
61	UP	43	M	R	66	2.33	15	3	18	31	22	12	33
62	UP	36	M	R	88	5.75	15	4	40	70	6	114	50
63	UP	25	M	R	62	3	15	4	40	135	2	1	2

Table S1. Cont.

ID	HOSP	AGE, y	SEX	HAND	ELC	FEA	L LEN	SES	LST	REC	REP	PLI	ELI
64	UP	18	F	R	76	7.33	15	3	30	104	5	6	3
65	UP	27	F	R	48	2	15	2	32	104	2	43	20
66	UP	55	F	L	80	1.5	15	2	32	81	11	61	24
67	UP	18	M	A	100	2.33	15	3	48	253	7	8	3
68	UP	38	F	R	86	6	15	1	16	48	14	3	73
69	UP	40	M	R	58	5.75	15	4	64	304	1	14	9

This table provides the hospital (HOSP) at which each patient's data were collected (Table S2), as well as each patient's age (AGE), sex (SEX), handedness or language mapping (HAND), number of implanted electrodes (ELC), and mean number of features selected for analysis across all sessions for that patient (FEA). Information about the task includes the list length (L LEN) used for each participant, number of testing sessions (SES), and number of lists each participant encountered across all sessions (LST). Performance information includes the total number of correct recalls across all lists (REC), total number of repeated recalls (REP), and total number of incorrect recalls, which include recalls of previously presented items [prior list intrusions (PLI)] and recalls of items that were never presented [extralist intrusions (ELI)]. In total, the 69 patients contributed 5,299 electrodes and 739 selected features, studying 29,030 items presented in 1,790 lists. A, ambidexterous; BW, Brigham & Women's Hospital (Boston, MA); CH, Children's Hospital (Boston, MA); F, female; FR, University Hospital of Freiburg (Freiburg, Germany); L, left; M, male; R, right; TJ, Thomas Jefferson University Hospital (Philadelphia, PA); UP, Hospital of the University of Pennsylvania (Philadelphia, PA).

Table S2. Bandpass filters used by our collaborating hospitals

Hospital	HOSP	Lower	Upper
Brigham & Women's Hospital, Boston, MA	BW	0.5 Hz	60 Hz
Children's Hospital Boston, Boston, MA	CH	0.3 Hz	50 Hz
University Hospital of Freiburg, Freiburg, Germany	FR	0.1 Hz	100 Hz
Thomas Jefferson University Hospital, Philadelphia, PA	TJ	$-\infty$	∞
Hospital of the University of Pennsylvania, Philadelphia, PA	UP	$-\infty$	∞

Hospital codes (HOSP) are referenced in Table S1. The Lower and Upper columns denote the lower and upper limits of the bandpass filters, respectively. Frequencies outside of the band-passed range were excluded from further analysis.