

A Gaussian process model of human electrocorticographic data

Lucy L. W. Owen¹, Tudor A. Muntianu¹, Andrew C. Heusser^{1,2},
Patrick Daly³, Katherine Scangos³, and Jeremy R. Manning^{1*}

¹Department of Psychological and Brain Sciences, Dartmouth College,
Hanover, NH 03755, USA

²Akili Interactive,
Boston, MA 02110, USA

³Weill Institute for Neurosciences, University of California, San Francisco,
San Francisco, CA 94121, USA

Abstract

We present a model-based method for inferring full-brain neural activity at millimeter-scale spatial resolutions and millisecond-scale temporal resolutions using standard human intracranial recordings. Our approach makes the simplifying assumptions that different people's brains exhibit similar correlational structure, and that activity and correlation patterns vary smoothly over space. One can then ask, for an arbitrary individual's brain: given recordings from a limited set of locations in that individual's brain, along with the observed spatial correlations learned from other people's recordings, how much can be inferred about ongoing activity at *other* locations throughout that individual's brain? We show that our approach generalizes across people and tasks, thereby providing a person- and task-general means of inferring high spatiotemporal resolution full-brain neural dynamics from standard low-density intracranial recordings.

Keywords: Electrocorticography (ECoG), intracranial electroencephalography (iEEG), local field potential (LFP), epilepsy, maximum likelihood estimation, Gaussian process regression

Introduction

Modern human brain recording techniques are fraught with compromise (Sejnowski et al. 2014). Commonly used approaches include functional magnetic resonance imaging (fMRI), scalp electroencephalography (EEG), and magnetoencephalography (MEG). For each of these techniques, neuroscientists and electrophysiologists must choose to optimize spatial resolution at the cost of temporal resolution (e.g., as in fMRI) or temporal resolution at the cost of spatial

25 resolution (e.g., as in EEG and MEG). A less widely used approach (due to requiring work
26 with neurosurgical patients) is to record from electrodes implanted directly onto the cortical
27 surface (electrocorticography; ECoG) or into deep brain structures (intracranial EEG; iEEG).
28 However, these intracranial approaches also require compromise: the high spatiotemporal
29 resolution of intracranial recordings comes at the cost of substantially reduced brain coverage,
30 since safety considerations limit the number of electrodes one may implant in a given patient's
31 brain. Further, the locations of implanted electrodes are determined by clinical, rather than
32 research, needs.

33 An increasingly popular approach is to improve the effective spatial resolution of MEG or
34 scalp EEG data by using a geometric approach called *beamforming* to solve the biomagnetic or
35 bioelectrical inverse problem (Sarvas 1987). This approach entails using detailed brain conduc-
36 tance models (often informed by high spatial resolution anatomical MRI images) along with
37 the known sensor placements (localized precisely in 3D space) to reconstruct brain signals orig-
38 inating from theoretical point sources deep in the brain (and far from the sensors). Traditional
39 beamforming approaches must overcome two obstacles. First, the inverse problem beamform-
40 ing seeks to solve has infinitely many solutions. Researchers have made progress towards
41 constraining the solution space by assuming that signal-generating sources are localized on a
42 regularly spaced grid spanning the brain and that individual sources are small relative to their
43 distances to the sensors (Baillet et al. 2001; Hillebrand et al. 2005; Snyder 1991). The second,
44 and in some ways much more serious, obstacle is that the magnetic fields produced by external
45 (noise) sources are substantially stronger than those produced by the neuronal changes being
46 sought (i.e., at deep structures, as measured by sensors at the scalp). This means that obtaining
47 adequate signal quality often requires averaging the measured responses over tens to hundreds
48 of responses or trials (e.g., see review by Hillebrand et al. 2005).

49 Another approach to obtaining high spatiotemporal resolution neural data has been to col-

50 lect fMRI and EEG data simultaneously. Simultaneous fMRI-EEG has the potential to balance
51 the high spatial resolution of fMRI with the high temporal resolution of scalp EEG, thereby,
52 in theory, providing the best of both worlds. In practice, however, the signal quality of both
53 recordings suffers substantially when the two techniques are applied simultaneously (e.g., see
54 review by Huster et al. 2012). In addition, the experimental designs that are ideally suited to
55 each technique individually are somewhat at odds. For example, fMRI experiments often lock
56 stimulus presentation events to the regularly spaced image acquisition time (TR), which max-
57 imizes the number of post-stimulus samples. By contrast, EEG experiments typically employ
58 jittered stimulus presentation times to maximize the experimentalist’s ability to distinguish
59 electrical brain activity from external noise sources such as from 60 Hz alternating current
60 power sources.

61 The current “gold standard” for precisely localizing signals and sampling at high temporal
62 resolution is to take (ECoG or iEEG) recordings from implanted electrodes (but from a limited
63 set of locations in any given brain). This begs the following question: what can we infer about
64 the activity exhibited by the rest of a person’s brain, given what we learn from the limited
65 intracranial recordings we have from their brain and additional recordings taken from *other*
66 people’s brains? Here we develop an approach, which we call *SuperEEG*¹, based on Gaussian
67 process regression (Rasmussen 2006). SuperEEG entails using data from multiple people to
68 estimate activity patterns at arbitrary locations in each person’s brain (i.e., independent of their
69 electrode placements). We test our SuperEEG approach using two large datasets of intracranial
70 recordings (Ezzyat et al. 2017, 2018; Horak et al. 2017; Kragel et al. 2017; Kucewicz et al. 2017,
71 2018; Lin et al. 2017; Manning et al. 2011, 2012; Sederberg et al. 2003, 2007a,b; Solomon et al.
72 2018; Weidemann et al. 2019). We show that the SuperEEG algorithm recovers signals well

¹The term “SuperEEG” was coined by Robert J. Sawyer in his popular science fiction novel *The Terminal Experiment* (Sawyer 1995). SuperEEG is a fictional technology that measures ongoing neural activity throughout the entire living human brain at arbitrarily high spatiotemporal resolution.

73 from electrodes that were held out of the training dataset. We also examine the factors that
74 influence how accurately activity may be estimated (recovered), which may have implications
75 for electrode design and placement in neurosurgical applications.

76 **Approach**

77 The SuperEEG approach to inferring high temporal resolution full-brain activity patterns is
78 outlined and summarized in Figure 1. We describe (in this section) and evaluate (in *Results*) our
79 approach using two large previously collected datasets comprising multi-session intracranial
80 recordings. Dataset 1 comprises multi-session recordings taken from 6876 electrodes implanted
81 in the brains of 88 epilepsy patients (Manning et al. 2011, 2012; Sederberg et al. 2003, 2007a,b).
82 Each recording session lasted from 0.2–3 h (total recording time: 0.3–14.2 h; Fig. S6E). During
83 each recording session, the patients participated in a free recall list learning task, which lasted for
84 up to approximately 1 h. In addition, the recordings included “buffer” time (the length varied
85 by patient) before and after each experimental session, during which the patients went about
86 their regular hospital activities (confined to their hospital room, and primarily in bed). These
87 additional activities included interactions with medical staff and family, watching television,
88 reading, and other similar activities. For the purposes of the Dataset 1 analyses presented here,
89 we aggregated all data across each recording session, including recordings taken during the
90 main experimental task as well as during non-experimental time. We used Dataset 1 to develop
91 our main SuperEEG approach, and to examine the extent to which SuperEEG might be able
92 to generate task-general predictions. Dataset 2 comprised multi-session recordings from 14860
93 electrodes implanted in the brains of 131 epilepsy patients (Ezzyat et al. 2017, 2018; Horak
94 et al. 2017; Kragel et al. 2017; Kucewicz et al. 2017, 2018; Lin et al. 2017; Solomon et al. 2018;
95 Weidemann et al. 2019). Each recording session lasted from 0.4–2.2 h (total recording time:
96 0.4–6.6 h; Fig. S6K). Whereas Dataset 1 included recordings taken as the patients participated

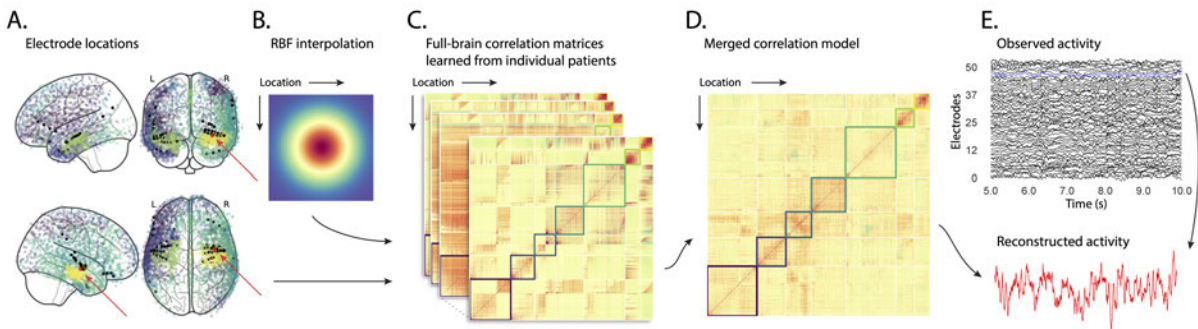


Figure 1: **Methods overview.** **A. Electrode locations.** Each dot reflects the location of a single electrode implanted in the brain of a Dataset 1 patient. A held-out recording location from one patient is indicated in red, and the patient’s remaining electrodes are indicated in black. The electrodes from the remaining patients are colored by k -means cluster (computed using the full-brain correlation model shown in Panel D). **B. Radial basis function kernel.** Each electrode contributed by the patient (black) weights on the full set of locations under consideration (all dots in Panel A, defined as \bar{R} in the text). The weights fall off with positional distance (in MNI152 space) according to an RBF. **C. Per-patient correlation matrices.** After computing the pairwise correlations between the recordings from each patient’s electrodes, we use RBF-weighted averages to estimate correlations between all locations in \bar{R} . We obtain an estimated full-brain correlation matrix using each patient’s data. **D. Merged correlation model.** We combine the per-patient correlation matrices (Panel C) to obtain a single full-brain correlation model that captures information contributed by every patient. Here we have sorted the rows and columns to reflect k -means clustering labels (using $k=7$; Yeo et al. 2011), whereby we grouped locations based on their correlations with the rest of the brain (i.e., rows of the matrix displayed in the panel). The boundaries denote the cluster groups. The rows and columns of Panel C have been sorted using the Panel D-derived cluster labels. **E. Reconstructing activity throughout the brain.** Given the observed recordings from the given patient (shown in black; held-out recording is shown in blue), along with a full-brain correlation model (Panel D), we use Equation 12 to reconstruct the most probable activity at the held-out location (red).

97 in a variety of activities, Dataset 2 included recordings taken as each patient performed each of
98 two specific experimental memory tasks: a random word list free recall task (Experiment 1) and
99 a categorized word list free recall task (Experiment 2). We used Dataset 2 to further examine the
100 ability of SuperEEG to generalize its predictions within versus across tasks. Figure S6 provides
101 additional information about both datasets.

102 We first applied fourth order Butterworth notch filters to remove 60 Hz (± 0.5 Hz) line
103 noise from every recording (from every electrode). Next, we downsampled the recordings

104 (regardless of the original samplerate) to 250 Hz. This downsampling step served to both
105 normalize for differences in sampling rates across patients and to ease the computational burden
106 of our subsequent analyses. We then excluded any electrodes that showed putative epileptiform
107 activity. Specifically, we excluded from further analysis any electrode that exhibited a maximum
108 kurtosis of 10 or greater across all of that patient's recording sessions. We also excluded any
109 patients with fewer than 2 electrodes that passed this criteria, as the SuperEEG algorithm
110 requires measuring correlations between 2 or more electrodes from each patient. For Dataset
111 1, this yielded clean recordings from 4168 electrodes implanted throughout the brains of 67
112 patients (Fig. 1A, colored dots); for Dataset 2, this yielded clean recordings from 5023 electrodes
113 implanted in 78 patients. Each individual patient contributed electrodes from a limited set
114 of brain locations, which we localized in a common space (MNI152; Grabner et al. 2006); an
115 example Dataset 1 patient's 54 electrodes that survived the kurtosis thresholding procedure are
116 highlighted in black and red (Fig. 1A).

The recording from a given electrode is maximally informative about the activity of the neural tissue immediately surrounding its recording surface. However, brain regions that are distant from the recording surface of the electrode also contribute to the recording, albeit (*ceteris paribus*) to a much lesser extent. One mechanism underlying these contributions is volume conduction. The precise rate of falloff due to volume conduction (i.e., how much a small volume of brain tissue at location x contributes to the recording from an electrode at location η) depends on the size of the recording surface, the electrode's impedance, and the conductance profile of the volume of brain between x and η . As an approximation of this intuition, we place a Gaussian radial basis function (RBF) at the location η of each electrode's recording surface (Fig. 1B). We use the values of the RBF at any brain location x as a rough estimate of how much structures

around x contributed to the recording from location η :

$$\text{rbf}(x|\eta, \lambda) = \exp\left\{-\frac{\|x - \eta\|^2}{\lambda}\right\}, \quad (1)$$

117 where the width variable λ is a parameter of the algorithm (which may in principle be set
118 according to location-specific tissue conductance profiles) that governs the level of spatial
119 smoothing. In choosing λ for the analyses presented here, we sought to maximize spatial
120 resolution (which implies a small value of λ) while also maximizing the algorithm's ability
121 to generalize to any location throughout the brain, including those without dense electrode
122 coverage (which implies a large value of λ). Here we set $\lambda = 20$, guided in part by our prior
123 related work (Manning et al. 2014, 2018), and in part by examining the brain coverage with
124 non-zero weights achieved by placing RBFs at each electrode location in Dataset 1 and taking
125 the sum (across all electrodes) at each voxel in a 4 mm³ MNI brain. (We then held λ fixed for
126 our analyses of Dataset 2.) We note that this value could in theory be further optimized, e.g.,
127 using cross validation or a formal model (e.g., Manning et al. 2018).

128 A second mechanism whereby a given region x can contribute to the recording at η is
129 through (direct or indirect) anatomical connections between structures near x and η . Although
130 anatomical and functional correlations can differ markedly (e.g., Adachi et al. 2012; Goñi et al.
131 2014; Honey et al. 2009), we use temporal correlations in the data to estimate these anatomical
132 connections (Becker et al. 2018). Let \bar{R} be the set of locations at which we wish to estimate local
133 field potentials, and let $R_s \subseteq \bar{R}$ be set of locations at which we observe local field potentials from
134 patient s (excluding the electrodes that did not pass the kurtosis test described above). In the
135 analyses below we define $\bar{R} = \cup_{s=1}^S R_s$. We can calculate the expected inter-electrode correlation
136 matrix for patient s , where $C_{s,k}(i, j)$ is the correlation between the time series of voltages for
137 electrodes i and j from subject s during session k , using:

$$\bar{C}_s = r \left(\frac{1}{n} \left(\sum_{k=1}^n z(C_{s,k}) \right) \right), \text{ where} \quad (2)$$

$$z(r) = \frac{\log(1+r) - \log(1-r)}{2} \text{ is the Fisher } z\text{-transformation and} \quad (3)$$

$$z^{-1}(z) = r(z) = \frac{\exp(2z) - 1}{\exp(2z) + 1} \text{ is its inverse.} \quad (4)$$

138 Next, we use Equation 1 to construct a number of to-be-estimated locations by number of
 139 patient electrode locations weight matrix, W_s . Specifically, W_s approximates how informative
 140 the recordings at each location in R_s are in reconstructing activity at each location in \bar{R} , where
 141 the contributions fall off with an RBF according to the distances between the corresponding
 142 locations:

$$W_s(i, j) = \text{rbf}(i|j, \lambda). \quad (5)$$

143 Given this weight matrix, W_s , and the observed inter-electrode correlation matrix for patient
 144 s , \bar{C}_s , we can estimate the correlation matrix for all locations in \bar{R} (\hat{C}_s ; Fig. 1C) using:

$$\hat{N}_s(x, y) = \sum_{i=1}^{|R_s|} \sum_{j=1}^{i-1} W(x, i) \cdot W(y, j) \cdot z(\bar{C}_s(i, j)) \quad (6)$$

$$\hat{D}_s(x, y) = \sum_{i=1}^{|R_s|} \sum_{j=1}^{i-1} W(x, i) \cdot W(y, j). \quad (7)$$

$$\hat{C}_s = r \left(\frac{\hat{N}_s}{\hat{D}_s} \right). \quad (8)$$

After estimating the numerator (\hat{N}_s) and denominator (\hat{D}_s) placeholders for each \hat{C}_s , we aggregate these estimates across the S patients to obtain a single expected full-brain correlation matrix (\hat{K} ; Fig. 1D):

$$\hat{K} = r \left(\frac{\sum_{s=1}^S \hat{N}_s}{\sum_{s=1}^S \hat{D}_s} \right). \quad (9)$$

145 Intuitively, the numerators capture the general structures of the patient-specific estimates of full-
 146 brain correlations, and the denominators account for which locations were near the implanted
 147 electrodes in each patient. To obtain \hat{K} , we compute a weighted average across the estimated
 148 patient-specific full-brain correlation matrices, where patients with observed electrodes near a
 149 particular set of locations in \hat{K} contribute more to the estimate.

150 Having used the multi-patient data to estimate a full-brain correlation matrix at the set
 151 of locations in \bar{R} that we wish to know about, we next use \hat{K} to estimate activity patterns
 152 everywhere in \bar{R} , given observations at only a subset of locations in \bar{R} (Fig. 1E).

153 Let α_s be the set of indices of patient s 's electrode locations in \bar{R} (i.e., the locations in R_s),
 154 and let β_s be the set of indices of all other locations in \bar{R} . In other words, β_s reflects the locations
 155 in \bar{R} where we did not observe a recording for patient s (these are the recording locations we
 156 will want to fill in using SuperEEG). We can sub-divide \hat{K} as follows:

$$\hat{K}_{\beta_s, \alpha_s} = \hat{K}(\beta_s, \alpha_s), \text{ and} \quad (10)$$

$$\hat{K}_{\alpha_s, \alpha_s} = \hat{K}(\alpha_s, \alpha_s). \quad (11)$$

157 Here $\hat{K}_{\beta_s, \alpha_s}$ represents the correlations between the “unknown” activity at the locations indexed
 158 by β_s and the observed activity at the locations indexed by α_s , and $\hat{K}_{\alpha_s, \alpha_s}$ represents the
 159 correlations between the observed recordings (at the locations indexed by α_s).

160 Let Y_{s,k,α_s} be the number-of-timepoints (T) by $|\alpha_s|$ matrix of (observed) voltages from the
 161 electrodes in α_s during session k from patient s . Then we can estimate the voltage from patient
 162 s 's k^{th} session at the locations in β_s as follows (Rasmussen 2006):

$$\hat{Y}_{s,k,\beta_s} = ((\hat{K}_{\beta_s, \alpha_s} \cdot \hat{K}_{\alpha_s, \alpha_s}^{-1}) \cdot Y_{s,k,\alpha_s}^T)^T. \quad (12)$$

163 This equation is the foundation of the SuperEEG algorithm. Whereas we observe recordings

164 only at the locations indexed by α_s , Equation 12 allows us to estimate the recordings at all loca-
165 tions indexed by β_s , which we can define *a priori* to include any locations we wish, throughout
166 the brain. This yields estimates of the time-varying voltages at *every* location in \bar{R} , provided that
167 we define \bar{R} in advance to include the union of all of the locations in R_s and all of the locations
168 at which we wish to estimate recordings (i.e., a timeseries of voltages).

169 We designed our approach to be agnostic to electrode impedances, as electrodes that do not
170 exist do not have impedances. Therefore our algorithm recovers voltages in standard deviation
171 (z-scored) units rather than attempting to recover absolute voltages. (This property reflects the
172 fact that $\hat{K}_{\beta_s, \alpha_s}$ and $\hat{K}_{\alpha_s, \alpha_s}$ are correlation matrices rather than covariance matrices.) Also, we
173 note that Equation 12 requires computing a T by T matrix, which can become computationally
174 expensive when T is very large (e.g., for the Dataset 1 patient with the longest recording time,
175 $T = 12,786,750$; also see Fig. S6, Panels E and K). However, because Equation 12 is time
176 invariant, we may compute Y_{s,k,β_s} in a piecewise manner by filling in Y_{s,k,β_s} one row at a time
177 (using the corresponding samples from Y_{s,k,α_s}).

178 The SuperEEG algorithm described above and in Figure 1 allows us to estimate, up to a
179 constant scaling factor, local field potentials (LFPs) for each patient at all arbitrarily chosen
180 locations in the set \bar{R} , *even if we did not record that patient's brain at all of those locations*. We next
181 turn to an evaluation of the accuracy of those estimates.

182 Results

183 We used a cross-validation approach to test the accuracy with which the SuperEEG algorithm
184 reconstructs activity throughout the brain. For each patient in turn, we estimated full-brain
185 correlation matrices (Eqn. 9) using data from all of the *other* patients. This step ensured that the
186 data we were reconstructing could not also be used to estimate the between-location correlations
187 that drove the reconstructions via Equation 12 (otherwise the analysis would be circular). For

188 that held-out patient, we held out each electrode in turn. We used Equation 12 to reconstruct
189 activity at the held-out electrode location, using the correlation matrix learned from all other
190 patients' data as \hat{K} , and using activity recorded from the other electrodes from the held-out
191 patient as Y_{s,k,α_s} . (For analyses examining the stability of our estimates of \hat{K} across time
192 and patients, see Figs. S7 and S8, respectively). We then asked: how closely did each of the
193 SuperEEG-estimated recordings at those electrodes match the observed recordings from those
194 electrodes (i.e., how closely did the estimated \hat{Y}_{s,k,β_s} match the observed Y_{s,k,β_s})?

195 We used this general approach to quantify the algorithm's performance across the full
196 dataset. For each held-out electrode, from each held-out patient in turn, we computed the
197 average correlation (across recording sessions) between the SuperEEG-reconstructed voltage
198 traces and the observed voltage traces from that electrode. For this analysis we set \bar{R} to be the
199 union of all electrode locations across all patients. This yielded a single correlation coefficient
200 for each electrode location in \bar{R} , reflecting how well the SuperEEG algorithm was able to recover
201 the recording at that location by incorporating data across patients (black histogram in Fig. 2A,
202 map in Fig. 2C). The observed distribution of correlations was centered well above zero (mean:
203 $r = 0.51$; t -test comparing mean of distribution of z -transformed average patient correlation
204 coefficients to 0: $t(66) = 23.55, p < 10^{-10}$), indicating that the SuperEEG algorithm recovers
205 held-out activity patterns substantially better than random guessing.

206 Next, we compared the quality of these across-participant reconstructions (i.e., computed
207 using a correlation model learned from other patients' data) to reconstructions generated using
208 a correlation model trained using the in-patient's data. In other words, for this within-patient
209 benchmark analysis we estimated \hat{C}_s (Eqn. 8) for each patient in turn, using recordings from all
210 of that patient's electrodes except at the location we were reconstructing. These within-patient
211 reconstructions serve as an estimate of how well data from all of the other electrodes from that
212 single patient may be used to estimate held-out data from the same patient. This allows us to

213 ask how much information about the activity at a given electrode might be inferred through
214 (a) volume conductance or other sources of “leakage” from activity patterns measured from
215 the patient’s other electrodes and (b) across-electrode correlations learned from that single
216 patient. As shown in Figure 2A (gray histogram), the distribution of within-patient correlations
217 was centered well above zero (mean: $r = 0.32$; t -test comparing mean of distribution of z -
218 transformed average patient correlation coefficients to 0: $t(66) = 15.16, p < 10^{-10}$). However, the
219 across-patient correlations were substantially higher (t -test comparing average z -transformed
220 within versus across patient electrode correlations: $t(66) = 9.17, p < 10^{-10}$). This is an especially
221 conservative test, given that the across-patient SuperEEG reconstructions exclude (from the
222 correlation matrix estimates) all data from the patient whose data is being reconstructed. We
223 repeated each of these analyses on a second independent dataset and found similar results
224 (Fig. 2B, D; within versus across reconstruction accuracy: $t(77) = 11.25, p < 10^{-10}$). We also
225 replicated this result separately for each of the two experiments from Dataset 2 (Fig. S3). This
226 overall finding, that reconstructions of held-out data using correlation models learned from
227 *other* patient’s data yield higher reconstruction accuracy than correlation models learned from
228 the patient whose data is being reconstructed, has two important implications. First, it implies
229 that distant electrodes provide additional predictive power to the data reconstructions beyond
230 the information contained solely in nearby electrodes. This follows from the fact that each
231 patient’s grid, strip, and depth electrodes are implanted in a unique set of locations, so for any
232 given electrode the closest electrodes in the full dataset tend to come from the same patient.
233 Second, it implies that the spatial correlations learned using the SuperEEG algorithm are, to
234 some extent, similar across people.

235 The recordings we analyzed from Dataset 1 comprised data collected as the patients per-
236 formed a variety of (largely idiosyncratic) tasks throughout each day’s recording session. That
237 we observed reliable reconstructions across patients suggests that the spatial correlations de-

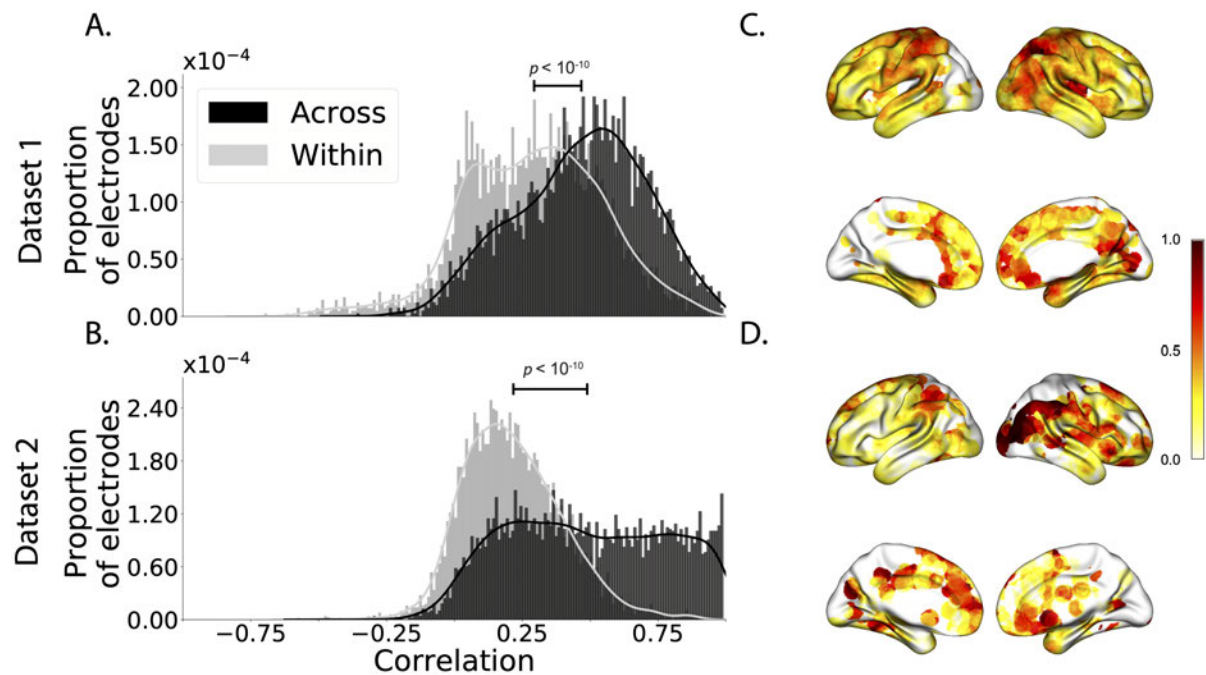


Figure 2: **Reconstruction accuracy across all electrodes in two ECoG datasets.** **A. Distributions of correlations between observed versus reconstructed activity by electrode, for Dataset 1.** The across-patient distribution (black) reflects reconstruction accuracy (correlation) using a correlation model learned from all but one patient’s data, and then applied to that held-out patient’s data. The within-patient distribution (gray) reflects performance using a correlation model learned from the same patient who contributed the to-be-reconstructed electrode. **B. Distributions of correlations for Dataset 2.** This panel is in the same format as Panel A, but reflects results obtained from Dataset 2. The histograms aggregate data across both Dataset 2 experiments; for results broken down by experiment see Figures S2 and S3. **C.–D. Reconstruction accuracy by location.** The colors denote the average across-session correlations, using the across-patient correlation model, between the observed and reconstructed activity at the given electrode location projected to the cortical surface (Combrisson et al. 2019). Panel C displays the map for Dataset 1 and Panel D displays the map for Dataset 2.

238 rived from the SuperEEG algorithm are, to some extent, similar across tasks. We tested this
239 finding more directly using Dataset 2. In Dataset 2, the recordings were limited to times when
240 each patient was participating in one of two experiments. Experiment 1 is a random-word
241 list free recall task; Experiment 2 is a categorized list free recall task (24 patients participated
242 in both). We wondered whether a correlation model learned from data from one experiment
243 might yield good predictions of data from the other experiment. Further, we wondered about
244 the extent to which it might be beneficial or harmful to combine data across tasks.

245 To test the task-specificity of the SuperEEG-derived correlation models, we restricted the
246 dataset to the 24 patients that participated in both experiments and repeated the above within-
247 and across-patient cross validation procedures separately for Experiment 1 and Experiment 2
248 data from Dataset 2. We then compared the reconstruction accuracies for held-out electrodes,
249 for models trained within versus across the two experiments, or combining across both exper-
250 iments (Fig. S1). In every case we found that across-patient models trained using data from
251 all other patients out-performed within-patient models trained on data only from the subject
252 contributing the given electrode ($t_s(23) > 6.50, p_s < 10^{-5}$). All reconstruction accuracies also re-
253 liably exceeded chance performance ($t_s(23) > 8.00, p_s < 10^{-8}$). Average reconstruction accuracy
254 was highest for the across-patient models limited to data from the same experiment (mean accu-
255 racy: $r = 0.68$); next-highest for the models that combined data across both experiments (mean
256 accuracy: $r = 0.61$); and lowest for models trained across tasks (mean accuracy: $r = 0.47$). This
257 pattern of results also held for each of the Dataset 2 experiments individually (Fig. S2). Taken
258 together, these results indicate that there are reliable commonalities in the spatial correlations
259 of full-brain activity across tasks, but that there are also reliable differences in these spatial
260 correlations across tasks. Whereas reconstruction accuracy benefits from incorporating data
261 from other patients, reconstruction accuracy is highest when constrained to within-task data,
262 or data that includes a variety of tasks (e.g., Dataset 1, or combining across the two Dataset 2

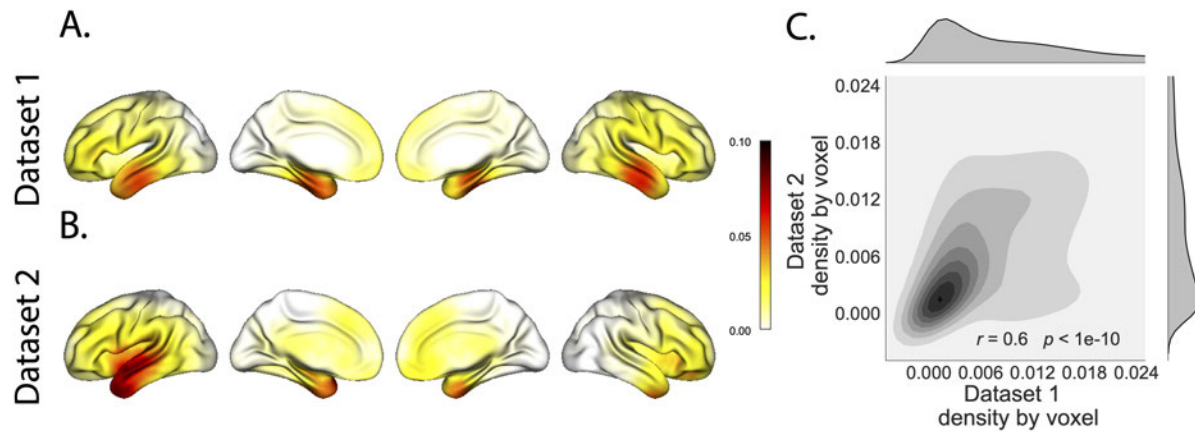


Figure 3: **Electrode sampling density by location.** **A. Electrode sampling density by voxel location in Dataset 1.** Each voxel is colored by the proportion of total electrodes in the dataset that are located within 20 MNI units of the given voxel. **B. Electrode sampling density by voxel location in Dataset 2.** This panel displays the sampling density map for Dataset 2, in the same format as Panel A. **C. Correspondence in sampling density by voxel location across Datasets 1 and 2.** The two-dimensional histogram displays the per-voxel sampling densities in the two Datasets, and the one-dimensional histograms display the proportions of voxels in each dataset with the given density value. The correlation reported in the panel is across voxels in the 4 mm³ MNI152 brain.

263 experiments).

264 Although both datasets we examined provide good full-brain coverage (when considering
265 data from every patient), electrodes were not sampled uniformly throughout the brain. For
266 example, in our patient population, electrodes are more likely to be implanted in regions like
267 the medial temporal lobe (MTL), and are rarely implanted in occipital cortex (Fig. 3A, B).
268 Separately for each dataset, for each voxel in the 4 mm³ voxel MNI152 brain, we computed
269 the proportion of electrodes in the dataset that were contained within a 20 MNI unit radius
270 sphere centered on that voxel. We defined the *density* at that location as this proportion.
271 Across Datasets 1 and 2, the electrode placement densities were similar (correlation by voxel:
272 $r = 0.6, p < 10^{-10}$). We wondered whether regions with good coverage might be associated with
273 better reconstruction accuracy. For example, Figures 2C and D indicate that some electrodes
274 in the MTL (which tends to be relatively densely sampled) have relatively high reconstruction

275 accuracy, and occipital electrodes (which tends to be relatively sparsely sampled) tend to have
276 relatively low reconstruction accuracy. To test whether this held more generally across the
277 entire brain, for each dataset we computed the electrode placement density for each electrode
278 from each patient (using the proportion of *other* patients' electrodes within 20 MNI units of the
279 given electrode). We then correlated these density values with the across-patient reconstruction
280 accuracies for each electrode. We found no reliable correlation between reconstruction accuracy
281 and density for Dataset 1 ($r = 0.05, p = 0.70$) and a reliable negative correlation for Dataset
282 2 ($r = -0.21, p = 0.05$). This suggests that the reconstruction accuracies we observed are *not*
283 driven solely by sampling density, but rather may also reflect higher order properties of neural
284 dynamics such as functional correlations between distant voxels (Betz et al. 2017).

285 Prior work in humans and animals has shown that the spatial profile of the local field
286 potential differs by frequency band (e.g., with respect to volume conductance properties and
287 contribution to the local field potential; Buzsaki et al. 2012; Crone et al. 2011; Fries et al. 2007).
288 For example, lower frequency components of the local field potential tend to have higher
289 power and extend further in space than high-frequency components (e.g., Manning et al.
290 2009; Miller et al. 2007). We wondered whether the reconstructions we observed might be
291 differently weighting or considering the contributions of activity at different frequency bands.
292 We therefore examined a range of frequency bands (δ : 2–4 Hz; θ : 4–8 Hz; α : 8–12 Hz; β : 12–30
293 Hz; γ_L : 30–60 Hz; and γ_H : 60–100 Hz), along with a measure of broadband (BB) power. We
294 used second-order Butterworth bandpass filters to compute the activity patterns within each
295 narrow frequency band. We defined broadband power as the mean height of a linear robust
296 regression fit in log-log space to the order 4 Morelet wavelet-computed power spectrum at
297 50 log-spaced frequencies from from 2–100 Hz (Manning et al. 2009). We then repeated our
298 within-subject and across-subject cross-validated reconstruction accuracy tests (analogous to
299 Fig. 2) separately for each frequency band (Fig. 4). (We also carried out a similar analysis on the

300 Hilbert transform-computed spectral power within each narrow band; see Fig. S4.) Across both
301 datasets, we found that our approach is best at reconstructing patterns of broadband activity
302 (right-most bars in Figs. 4A and D), a correlate of population firing rate (Manning et al. 2009).
303 We also achieved good reconstruction accuracy within each narrow frequency band (Figs. 4 and
304 S4). Activity at lower frequencies (δ , θ , α , and β) tended to be reconstructed better than high-
305 frequency patterns (γ_L and γ_H), with reconstruction accuracy peaking in the θ band. Overall,
306 these results indicate that our approach is able to accurately recover information within the
307 2–100 Hz range.

308 A basic assumption of our approach (and of most prior ECoG work) is that electrode record-
309 ings are most informative about the neural activity near the recording surface of the electrode.
310 But if we consider that activity patterns throughout the brain are meaningfully correlated, are
311 there particular implantation locations that, if recorded from a given patient's brain, yield es-
312 pecially high reconstruction accuracies throughout the rest of their brain? For example, one
313 might hypothesize that brain structures that are heavily interconnected with many other struc-
314 tures could be more informative about full-brain activity patterns than comparatively isolated
315 structures. To test this hypothesis, we computed the average reconstruction accuracy across all
316 of each patient's electrodes (using our across-patients cross validation test; black histograms
317 in Fig. 2A and B). We first labeled each patient's electrodes, in each dataset, with the average
318 reconstruction accuracy for that patient. In other words, we assigned every electrode from
319 each patient the same value, reflecting how well the activity patterns for that patient were
320 reconstructed. Next, for each voxel in the 4 mm³ MNI brain, we computed the average value
321 across any electrode (from any patient) that came within 20 MNI units of that voxel's center.
322 This yielded an *information score* for each voxel, reflecting the (weighted) average reconstruc-
323 tion accuracy across any patients with electrodes near each voxel, where the averages were
324 weighted to reflect patients who had more electrodes implanted near that location. We created

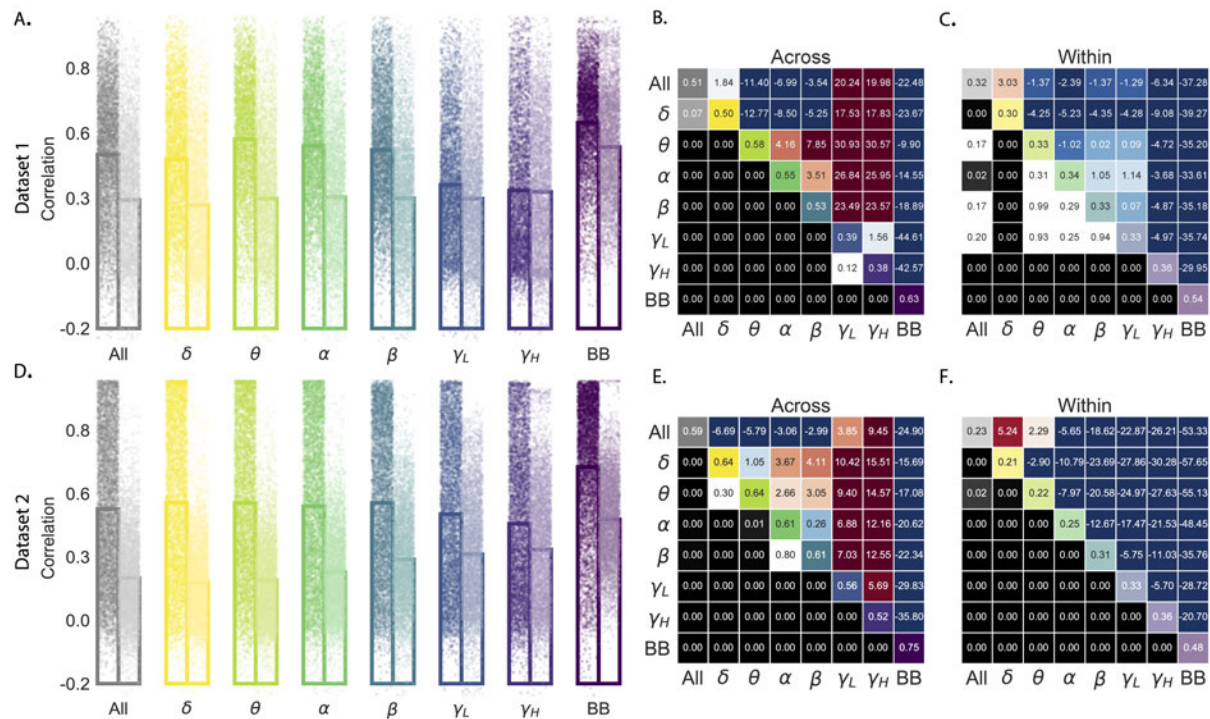


Figure 4: **Reconstruction accuracy across all electrodes in two ECoG datasets for each frequency band.** **A.** Distributions of correlations between observed versus reconstructed activity by electrode for each frequency band in Dataset 1. Each color denotes a different frequency band. Within each color group, the darker dots and bar on the left display the distribution (and mean) across-patient reconstruction accuracies (analogous to the black histograms in Fig. 2). The lighter dots and bar on the right display the distribution (and mean) within-patient reconstruction accuracies (analogous to the gray histograms in Fig. 2). Each dot indicates the reconstruction accuracy for one electrode in the dataset. To facilitate visual comparison with the frequency-specific results, the leftmost bars (gray) re-plot the histograms in Figure 2A. **B.** Statistical summary of across-patient reconstruction accuracy by electrode for each frequency band in Dataset 1. In the upper triangles of each map, warmer colors (positive t -values) indicate that the reconstruction accuracy for the frequency band in the given row was greater (via a two-tailed paired-sample t -test) than for the frequency band in the given column. Cooler colors (negative t -values) indicate that reconstruction accuracy for the frequency band in the given row was lower than for the frequency band in the given column. The lower triangles of each map denote the corresponding p -values for the t -tests. The diagonal entries display the average reconstruction accuracy within each frequency band. **C.** Statistical summary of within-patient reconstruction accuracy by electrode for each frequency band in Dataset 1. This panel displays the within-patient statistical summary, in the same format as Panel B. **D.** Distributions of correlations between observed versus reconstructed activity by electrode, for each frequency band in Dataset 2. This panel displays reconstruction accuracy distributions for each frequency band for Dataset 2. **E.–F.** Statistical summaries of across-patient and within-patient reconstruction accuracy by electrode for each frequency band in Dataset 2. These panels are in the same as Panels B and C, but display results from Dataset 2.

325 a single map of these information scores for each dataset, highlighting regions that are espe-
326 cially informative about activity in *other* brain areas (Figs. 5A and B). Despite task and patient
327 differences across the two datasets, we nonetheless found that the information score maps from
328 both datasets were correlated (voxelwise correlation between information scores across the two
329 datasets: $r = 0.18, p < 10^{-10}$). Our finding that there were some commonalities between the
330 two datasets' information score maps lends support to the notion that different brain areas
331 are (reliably) differently informative about full-brain activity patterns. We also examined the
332 intersection between the top 10% most informative voxels across the two datasets (gray areas
333 in Fig. 5C, networks shown in Fig. 6A, top row). Supporting the notion that structures that
334 are highly interconnected with the rest of the brain are most informative about full-brain activ-
335 ity patterns, the intersecting set of voxels with the highest information scores included major
336 portions of the dorsal attention network (e.g., inferior parietal lobule, precuneus, inferior tem-
337 poral gyrus, thalamus, and striatum) as well as some portions of the default mode network
338 (e.g., angular gyrus) that are highly interconnected with a large proportion of the brain's gray
339 matter (e.g., Tomasi and Volkow 2011).

340 We also wondered whether the map of information scores might vary as a function of the
341 spectral components of the activity patterns under consideration. We computed analogous
342 maps of information scores for each individual frequency band. Across Datasets 1 and 2
343 (with the exception of α -band activity), we observed reliable positive correlations between
344 the voxelwise maps of information scores (δ : $r = 0.09, p < 10^{-57}$; θ : $r = 0.24, p < 10^{-60}$; α :
345 $r = -0.03, p < 0.001$; β : $r = 0.02, p = 0.0011$; γ_L : $r = 0.1, p < 10^{-67}$; γ_H : $r = 0.03, p < 10^{-7}$;
346 broadband: $r = 0.21, p < 10^{-297}$).

347 To gain additional insight into which regions were most informative about full-brain ac-
348 tivity patterns at different frequency bands, we next computed (for each frequency band) the
349 intersection of the top 10% highest information scores across the maps for Datasets 1 and 2

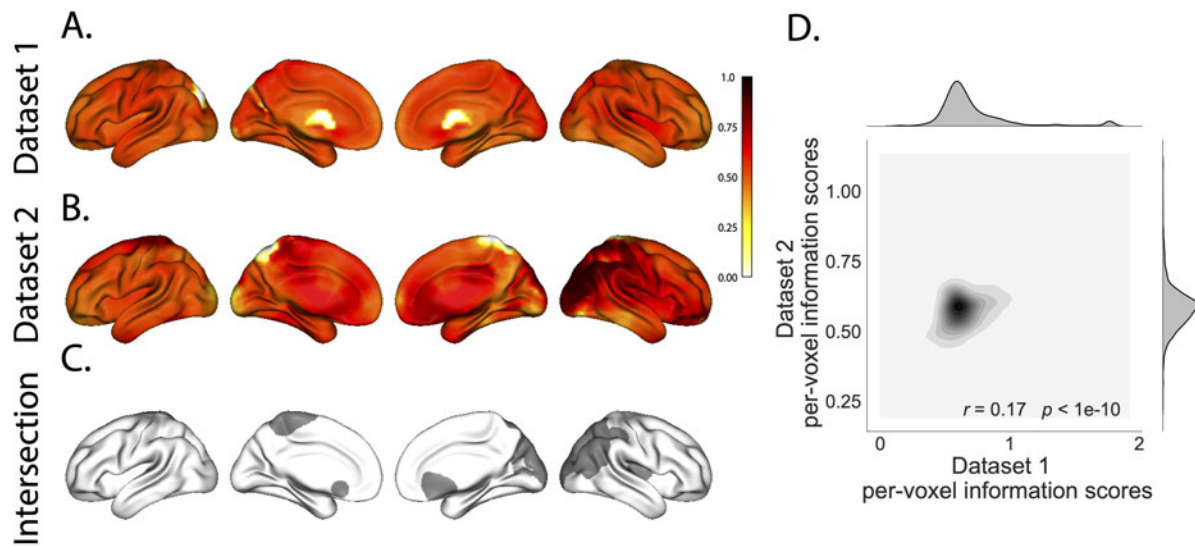


Figure 5: **Most informative recording locations.** **A. Dataset 1 information scores by voxel.** The voxel colors reflect the weighted average reconstruction accuracy across all electrodes from any patients with at least one electrode within 20 MNI units of the given voxel. **B. Dataset 2 information scores by voxel.** This panel is in the same format as Panel A. **C. Intersection.** Gray areas indicate the intersections between the top 10% most informative voxels in each map and projected onto the cortical surface (Combrisson et al. 2019). **D. Correspondence in information scores by voxel across Datasets 1 and 2.** The correlation reported in the Panel is between the per-voxel information scores across Datasets 1 and 2.

350 (analogous to our approach in Fig. 5C). This yielded a single map of the (reliably) most infor-
351 mative locations, for each frequency band we examined. We then carried out *post hoc* analyses
352 on each of these maps to characterize the underlying structural and functional properties of
353 each set of regions we identified as being particularly informative about one or more types of
354 neural pattern (Figs. 6 and S5).

355 A growing body of neuroscientific research is concerned with characterizing the *parcella-*
356 *tions* of anatomical and functional brain networks (for review see Arslan et al. 2018; Zalesky
357 et al. 2010). The dominant approaches entail obtaining a full-brain connectivity matrix using
358 either diffusion tensor imaging to identify the brain's network of white matter connections,
359 or functional connectivity (typically applied to resting state data) to correlate the patterns of
360 activity exhibited by different brain structures. One can then apply graph theoretic approaches
361 to assign each brain structure (typically a single fMRI voxel) to one or more networks (for review
362 see Bullmore and Sporns 2009). The result is a set of distinct (or partially overlapping) brain
363 "networks" that may be further examined to elucidate their potential functional role. We over-
364 laid a well-cited seven-network parcellation map identified by Yeo et al. (2011) onto the maps of
365 brain locations that were most informative about each type of neural pattern. For each of these
366 information maps, we computed the proportion of voxels in the most informative brain regions
367 that belonged to each of the seven networks identified by Yeo et al. (2011); Figure 6D. We found
368 that the regions we identified as being most informative about different neural patterns varied
369 markedly with respect to which functional networks they belonged to (Fig. 6A, B).

370 The variability we observed in the frequency-specific information score maps is consistent
371 with the notion that there is no "universal" brain region that reflects all types of activity pat-
372 terns throughout the rest of the brain. Rather, each region's activity patterns appear to be
373 characterized by different spectral profiles, and the ability to infer full-brain activity patterns
374 at a particular frequency band depends on the structural and functional connectome specific

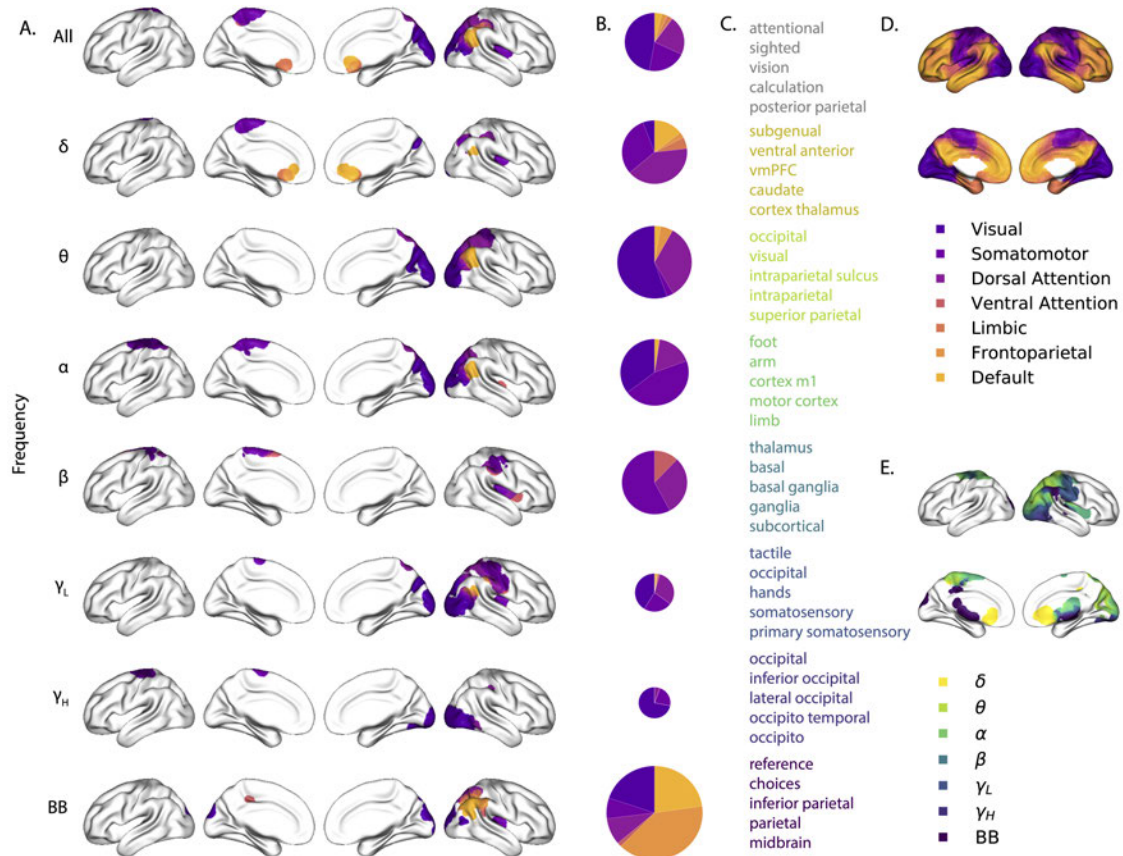


Figure 6: Most informative recording locations by frequency band. A. Intersections between information score maps by frequency band. The regions indicated in each row depict the intersection between the top 10% most informative locations across Datasets 1 and 2. **B. Network memberships of the most informative brain regions.** The pie charts display the proportions of voxels in each region that belong to the seven networks identified by Yeo et al. (2011). The relative sizes of the charts for each frequency band reflect the average across-subject reconstruction accuracies (Figs. 4A, D). The voxels in Panel A are colored according to the same network memberships. **C. Neurosynth terms associated with the most informative brain regions, by frequency band.** The lists in each row display the top five neurosynth terms (Rubin et al. 2017) decoded for each region. **D. Network parcellation map and legend.** The parcellation defined by Yeo et al. (2011) is displayed on the inflated brain maps. The colors and network labels serve as a legend for Panels A and B. **E. Combined map of the most informative brain regions.** The map displays the union of the most informative maps in Panel A, colored by frequency band. The labels also serve as a legend for Panel C.

375 to that frequency band (Fig. 6E). We wondered how the maps we found might fit in with prior
376 work. To this end, in addition to examining the anatomical profiles of each map, we used Neu-
377 rosynth (Rubin et al. 2017) to identify (using meta analyses of the neuroimaging literature) the
378 top five most common terms associated with each frequency-specific map (Fig. 6C). We found
379 that δ patterns across the brain were best predicted by regions of ventromedial prefrontal cortex,
380 striatum, and thalamus (yellow). These regions are also implicated in modulating δ oscillations
381 during sleep, and are heavily interconnected with cortex (e.g., Amzica and Steriade 1998). The
382 brain areas most informative about full-brain θ patterns were occipital and parietal regions
383 associated with visual processing and visual attention (light green). Prior work has implicated
384 θ oscillations in these areas in periodic sampling of visual attention (e.g., Busch and VanRullen
385 2010). We found that full-brain α patterns were best predicted by motor areas (dark green),
386 which also exhibit α band changes during voluntary movements (e.g., Jurkiewicz et al. 2006).
387 Striatum and thalamus (teal) were most informative about full-brain β patterns. Prior work has
388 implicated striatal β activity in sensory and motor processing (Feingold et al. 2015) and thalamic
389 β activity has been implicated in modulating widespread β patterns across neocortex (Sherman
390 et al. 2016). Somatosensory areas (dark blue) were most informative about full-brain γ_L pat-
391 terns. Prior work has implicated somatosensory γ_L in somatosensory processing and motor
392 planning (Ihara et al. 2003). Occipital cortex (purple) was most informative about full-brain
393 γ_H patterns. Occipital γ_H has also been linked with visual processing and reading (Wu et al.
394 2011) and the transmission of visual representations from low-order to higher-order visual
395 areas (Matsumoto et al. 2013). Full-brain broadband patterns were best predicted by inferior
396 parietal cortex precuneus (maroon). Functional neuroimaging BOLD responses (Simony et al.
397 2016) and broadband ECoG patterns (Honey et al. 2012) in these default-mode hubs have been
398 implicated in processing context-dependent representations that unfold over long timescales.

399 Taken together, the frequency-specific information maps suggest a potential new interpreta-

400 tion of many of the above previously reported findings. Prior work has largely treated region-
401 specific narrowband and broadband activity as an indicator that activity at those frequency
402 ranges reflects that the given region is representing or supporting a particular function. Our
403 work suggests an alternative interpretation that when we observe a particular neural pattern in
404 a particular brain region, it may instead reflect how that region is transmitting information to
405 the rest of the brain via signalling at the given frequency range.

406 Discussion

407 Are our brain's networks static or dynamic? And to what extent are the network properties of
408 our brains stable across people and tasks? One body of work suggests that our brain's *functional*
409 networks are dynamic (e.g., Manning et al. 2018; Owen et al. 2019), person-specific (e.g., Finn
410 et al. 2015), and task-specific (e.g., Turk-Browne 2013). In contrast, although the gross anatomical
411 structure of our brains changes meaningfully over the course of years as our brains develop, on
412 the timescales of typical neuroimaging experiments (i.e., hours to days) our anatomical networks
413 are largely stable (e.g., Casey et al. 2000). Further, many aspects of brain anatomy, including
414 white matter structure, are largely preserved across people (e.g., Jahanshad et al. 2013; Mori
415 et al. 2008; Talairach and Tournoux 1988). There are several possible means of reconciling this
416 apparent inconsistency between dynamic person- and task-specific functional networks versus
417 stable anatomical networks. For example, relatively small magnitude anatomical differences
418 across people may be reflected in reliable functional connectivity differences. Along these
419 lines, one recent study found that diffusion tensor imaging (DTI) structural data is similar
420 across people, but may be used to predict person-specific resting state functional connectivity
421 data (Becker et al. 2018). Similarly, other work indicates that task-specific functional connectivity
422 may be predicted by resting state functional connectivity data (Cole et al. 2016; Tavor et al. 2016).
423 Another (potentially complementary) possibility is that our functional networks are constrained

424 by anatomy, but nevertheless exhibit (potentially rapid) task-dependent changes (e.g., Sporns
425 and Betzel 2016).

426 Here we have taken a model-based approach to studying whether high spatiotemporal
427 resolution activity patterns throughout the human brain may be explained by a static connec-
428 tome model that is shared across people and tasks. Specifically, we trained a model to take
429 in recordings from a subset of brain locations, and then predicted activity patterns during the
430 same interval, but at *other* locations that were held out from the model. Our model, based on
431 Gaussian process regression, was built on three general hypotheses about the nature of the
432 correlational structure of neural activity (each of which we tested). First, we hypothesized that
433 functional correlations are stable over time and across tasks. We found that, although aspects of
434 the patients' functional correlations were stable across tasks, we achieved better reconstruction
435 accuracy when we trained the model on within-task data. This suggests that our general ap-
436 proach could be extended to better model across-task changes, e.g., following Cole et al. (2016);
437 Tavor et al. (2016); and others. Second, we hypothesized that some of the correlational structure
438 of people's brain activity is similar across individuals. Consistent with this hypothesis, our
439 model explained each patient's data best when trained using data from *other* patients— even
440 when compared models trained within-patient. Third, we resolved ambiguities in the data by
441 hypothesizing that neural activity from nearby sources tends to be similar, all else being equal.
442 This hypothesis was supported through our finding that all of the models we trained that
443 incorporated this spatial smoothness assumption predicted held-out data well above chance.

444 One potential limitation of our approach is that it does not provide a natural means of
445 estimating the precise timing of single-neuron action potentials. Prior work has shown that
446 gamma band and broadband activity in the LFP may be used to estimate the firing rates of
447 neurons that underly the population contributing to the LFP (Crone et al. 2011; Jacobs et al.
448 2010; Manning et al. 2009; Miller et al. 2008). Because SuperEEG reconstructs LFPs throughout

449 the brain, one could in principle use broadband power in the reconstructed signals to estimate
450 the corresponding firing rates (though not the timings of individual action potentials). We
451 found that we were able to reconstruct full-brain patterns of broadband power well (Fig. 4).

452 A second potential limitation of our approach is that it relies on ECoG data from epilepsy
453 patients. Recent work comparing functional correlations in epilepsy patients (measured using
454 ECoG) and healthy individuals (measured using fMRI) suggests that there are gross similarities
455 between these populations (e.g., Kucyi et al. 2018; Reddy et al. 2018). Nevertheless, because all of
456 the patients we examined have drug-resistant epilepsy, it remains uncertain how generally the
457 findings reported here might apply more broadly to the population at large (e.g., non-clinical
458 populations).

459 Beyond providing a means of estimating ongoing activity throughout the brain using
460 already-implanted electrodes, our work also has implications how to optimize electrode place-
461 ments in neurosurgical evaluations. Electrodes are typically implanted to maximize coverage
462 of suspected epileptogenic tissue. However, our findings suggest that this approach might be
463 improved upon. Specifically, one could leverage not only the non-invasive recordings taken
464 during an initial monitoring period (as is currently done routinely), but also recordings col-
465 lected from *other* patients. We could then ask: given what we learn from other patients' data
466 (and potentially from the scalp EEG recordings of this new patient), where should we place a
467 fixed number of electrodes to maximize our ability to map seizure foci? As shown in Figures 5,
468 6, and S5, recordings from different regions vary with respect to how informative they are about
469 different narrowband and broadband full-brain activity patterns.

470 By providing a means of reconstructing full-brain activity patterns, the SuperEEG approach
471 maps ECoG recordings from different patients into a common neural space, despite that different
472 patients' electrodes were implanted in different locations. This feature of our approach enables
473 across-patient ECoG studies, analogous to across-subject fMRI studies (e.g., Haxby et al. 2001,

474 2011; Norman et al. 2006). Whereas the focus of this manuscript is to specifically evaluate
475 which aspects of neural activity patterns SuperEEG recovers well (or poorly), in parallel work
476 we are training across-patient classifiers by leveraging the common neural spaces obtained by
477 applying SuperEEG to multi-patient ECoG data. For example, we have shown that SuperEEG-
478 derived activity patterns may be used to accurately predict psychiatric conditions such as
479 depression (Scangos et al. 2020). Analogous approaches could in principle be used to develop
480 improved brain-computer interfaces and/or to carry out other analyses that would benefit from
481 high spatiotemporal resolution full-brain data in individuals, projected into a common ECoG
482 space across people.

483 **Concluding remarks**

484 Over the past several decades, neuroscientists have begun to leverage the strikingly profound
485 mathematical structure underlying the brain's complexity to infer how our brains carry out
486 computations to support our thoughts, actions, and physiological processes. Whereas tradi-
487 tional beamforming techniques rely on geometric source-localization of signals measured at the
488 scalp, here we propose an alternative approach that leverages the rich correlational structure
489 of two large datasets of human intracranial recordings. In doing so, we are one step closer to
490 observing, and perhaps someday understanding, the full spatiotemporal structure of human
491 neural activity.

492 **Code availability**

493 We have published an open-source toolbox implementing the SuperEEG algorithm. It may be
494 downloaded [here](#). Additionally, we have provided code for all analyses and figures reported in
495 the current manuscript, available [here](#).

496 **Data availability**

497 The datasets analyzed in this study were generously shared by Michael J. Kahana. A portion of
498 Dataset 1 may be downloaded [here](#). Dataset 2 may be downloaded [here](#).

499 **Acknowledgements**

500 We are grateful for useful discussions with Luke J. Chang, Uri Hasson, Josh Jacobs, Michael J.
501 Kahana, and Matthijs van der Meer. We are also grateful to Michael J. Kahana for generously
502 sharing the ECoG data we analyzed in our paper, which was collected under NIMH grant
503 MH55687 and DARPA RAM Cooperative Agreement N66001-14-2-4-032, both to M.J.K. Our
504 work was also supported in part by NSF EPSCoR Award Number 1632738 and by a sub-award
505 of DARPA RAM Cooperative Agreement N66001-14-2-4-032 to J.R.M. The content is solely
506 the responsibility of the authors and does not necessarily represent the official views of our
507 supporting organizations.

508 **Author Contributions**

509 J.R.M conceived and initiated the project. L.L.W.O., T.A.M, and A.C.H. performed the anal-
510 yses using software packages that all authors contributed to. J.R.M. and L.L.W.O. wrote the
511 manuscript with input from all other authors.

512 **References**

513 Adachi Y, Osada T, Sporns O, Watanabe T, Matsui T, Miyamoto K, Miyashita Y 2012. Functional
514 connectivity between anatomically unconnected areas is shaped by collective network-level
515 effects in the macaque cortex. *Cerebral Cortex* 22:1586–1592.

- 516 Amzica F, Steriade M 1998. Electrophysiological correlates of sleep delta waves. *Electroen-*
517 *cephalography and Clinical Neurophysiology* 107:69–83.
- 518 Arslan S, Ktena S I, Makropoulos A, Robinson E C, Rueckert D, Parisot S 2018. Human brain
519 mapping: a systematic comparison of parcellation methods for the human cerebral cortex.
520 *NeuroImage* 170:5–30.
- 521 Baillet S, Mosher J C, Leahy R M 2001. Electromagnetic brain mapping. *IEEE Signal Processing*
522 *Magazine* 18:14–30.
- 523 Becker C O, Pequito S, Pappas G J, Miller M B, adn D S Bassett S T Grafton, Preciado V M 2018.
524 Spectral mapping of brain functional connectivity from diffusion imaging. *Scientific Reports*
525 8: <https://www.doi.org/10.1038/s41598-017-18769-x>.
- 526 Betzel R F, Medaglia J D, Kahn A E, Soffer J, Schonhaut D R, Bassett D S 2017. Inter-regional
527 ECoG correlations predicted by communication dynamics, geometry, and correlated gene
528 expression. *arXiv* 1706.06088.
- 529 Bullmore E., Sporns O. 2009. Complex brain networks: graph theoretical analysis of structural
530 and functional systems. *Nature Reviews Neuroscience* 10:186–198.
- 531 Busch N A, VanRullen R 2010. Spontaneous EEG oscillations reveal periodic sampling of visual
532 attention. *Proceedings of the National Academy of Science USA* 107:16048–16053.
- 533 Buzsaki G., Anastassiou C., Koch C. 2012. The origin of extracellular fields and currents - eeg,
534 ecog, lfp and spikes. *Nature Reviews Neuroscience* 13:407–419.
- 535 Casey B J, Giedd J N, Thomas K M 2000. Structural and functional brain development and its
536 relation to cognitive development. *Biological Psychology* 54:241–257.

- 537 Cole M W, Ito T, Bassett D S, Schultz D H 2016. Activity flow over resting-state networks shapes
538 cognitive task activations. *Nature Neuroscience* 19:1718–1726.
- 539 Combrisson E, Vallat R, O'Reilly C, Jas M, Pascarella A, l Saive A, Thiery T, Meunier D, Altukhov
540 D, Lajnef T, Ruby P, Guillot A, Jerbi K 2019. Visbrain: a multi-purpose GPU-accelerated open-
541 source suite for multimodal brain data visualization. *Frontiers in Neuroinformatics* 13:1–14.
- 542 Crone N. E., Korzeniewska A., Franaszczuk P.J. 2011. Cortical gamma responses: searching
543 high and low. *International Journal of Psychophysiology* 79:9–15.
- 544 Ezzyat Y, Kragel J E, Burke J F, Levy D F, Lyalenko A, Wanda P, O'Sullivan L, Hurley K B,
545 Busygin S, Pedisich I, Sperling M R, Worrell G A, Kucewicz M T, Davis K A, Lucas T H, Inman
546 C S, Lega B C, Jobst B C, Sheth S A, Zaghoul K, Jutras M J, Stein J M, Das S R, Gorniak
547 R, Rizzuto D S, Kahana M J 2017. Direct brain stimulation modulates encoding states and
548 memory performance in humans. *Current Biology* 27:1–8.
- 549 Ezzyat Y, Wanda P A, Levy D F, Kadel A, Aka A, Pedisich I, Sperling M R, Sharan A D, Lega
550 B C, Burks A, Gross R E, Inman C S, Jobst B C, Gorenstein M A, Davis K A, Worrell G A,
551 Kucewicz M T, Stein J M, Gorniak R, Das S R, Rizzuto D S, Kahana M J 2018. Closed-loop
552 stimulation of temporal cortex rescues functional networks and improves memory. *Nature*
553 *Communications* 9: <https://www.doi.org/10.1038/s41467-017-02753-0>.
- 554 Feingold J, Gibson D J, DePasquale B, Graybiel A M 2015. Bursts of beta oscillation differentiate
555 postperformance activity in the striatum and motor cortex of monkeys performing movement
556 tasks. *Proceedings of the National Academy of Science USA* 112:13687–13692.
- 557 Finn E S, Shen X, Scheinost D, Rosenberg M D, Huang J, Chun M M, Papademetris X, Constable
558 R T 2015. Functional connectome fingerprinting: identifying individuals using patterns of
559 brain connectivity. *Nature Neuroscience* 18:1664 – 1671.

- 560 Fries P., Nikolić D., Singer W. 2007. The gamma cycle. *Trends in Neurosciences* 30:309–316.
- 561 Goñi J, van den Heuvel M P, Avena-Koenigsberger A, de Mendizabal N V, Betzel R F, Griffa
562 A, Hagmann P, Corominas-Murtra B, Thiran J-P, Sporns O 2014. Resting-brain functional
563 connectivity predicted by analytic measures of network communication. *Proceedings of the*
564 *National Academy of Sciences* 111:833–838.
- 565 Grabner G., Janke A. L., Budge M. M., Smith D., Pruessner J., Collins D. L. 2006. Symmetric
566 atlas and model based segmentation: an application to the hippocampus in older adults.
567 *Medical Image Computing and Computer-Assisted Intervention* 9:58–66.
- 568 Haxby J. V., Gobbini M. I., Furey M. L., Ishai A., Schouten J. L., Pietrini P. 2001. Distributed
569 and overlapping representations of faces and objects in ventral temporal cortex. *Science*
570 293:2425–2430.
- 571 Haxby J V, Guntupalli J S, Connolly A C, Halchenko Y O, Conroy B R, Gobbini M I, Hanke
572 M, Ramadge P J 2011. A common, high-dimensional model of the representational space in
573 human ventral temporal cortex. *Neuron* 72:404–416.
- 574 Hillebrand A, Singh K D, Holliday I E, Furlong P L, Barnes G R 2005. A new approach to
575 neuroimaging with magnetoencephalography. *Human Brain Mapping* 25:199–211.
- 576 Honey C J, Sporns O, Cammoun L, Gigandet X, Thiran J P, Meuli R, Hagmann P 2009. Predicting
577 human resting-state functional connectivity from structural connectivity. *Proceedings of the*
578 *National Academy of Science USA* 106:2035–2040.
- 579 Honey C. J., Thesen T., Donner T. H., Silbert L. J., Carlson C. E., Devinsky O., Doyle J. C., Rubin
580 N., Heeger D. J., Hasson U. 2012. Slow cortical dynamics and the accumulation of information
581 over long timescales. *Neuron* 76:423–434.

- 582 Horak P C, Meisenhelter S, Song Y, Testorf M E, Kahana M J, Viles W D, Bujarski K A, Connolly
583 A C, Robbins A A, Sperling M R, Sharan A D, Worrell G A, Miller L R, Gross R E, Davis K A,
584 Roberts D W, Lega B, Sheth S A, Zaghloul K A, Stein J M, Das S R, Rizzuto D S, Jobst B C
585 2017. Interictal epileptiform discharges impair word recall in multiple brain areas. *Epilepsia*
586 58:373–380.
- 587 Huster R J, Debener S, Eichele T, Herrmann C S 2012. Methods for simultaneous EEG-fMRI: an
588 introductory review. *The Journal of Neuroscience* 32:6053–6060.
- 589 Ihara A, Hirata M, Yanagihara K, Ninomiya H, Imai K, Ishii R, Osaki Y, Sakihara K, Izumi H,
590 Imaoka H, Kato A, Yoshimine T 2003. Neuromagnetic gamma-band activity in the primary
591 and secondary somatosensory areas. *NeuroReport* 14:273–277.
- 592 Jacobs J., Manning J.R., Kahana M.J. 2010. Response to Miller: “Broadband” vs. “high gamma”
593 electrocorticographic signals. *Journal of Neuroscience* 30.
- 594 Jahanshad N, Kochunov P V, Sprooten E, Mandl R C, Nichols T E, Almasry L, Blangero J, Brouwer
595 R M, Curran J E, de Zubicaray G I, Duggirala R, Fox P T, Hong L E, Landman B A, Martin
596 N G, McMahon K L, Medland S E, Mitchell B D, Olvera R L, Peterson C P, Starr J M, Sussmann
597 J E, Toga A W, Wardlaw J M, Wright M J, Pol H E Hulshoff, Pastin M E, McIntosh A M,
598 Deary I J, Thompson P M, Glahn D C 2013. Multi-site genetic analysis of diffusion images and
599 voxelwise heritability analysis: A pilot project of the enigma-dti working group. *NeuroImage*
600 81:455–469.
- 601 Jurkiewicz M T, Gaetz W C, Bostan A C, Cheyne D 2006. Post-movement beta rebound is
602 generated in motor cortex: Evidence from neuromagnetic recordings. *NeuroImage* 32:1281–
603 1289.
- 604 Kragel J E, Ezzyat Y, Sperling M R, Gorniak R, Worrell G A, Berry B M, Inman C, Lin J-J, Davis

- 605 K A, Das S R, Stein J M, Jobst B C, Zaghoul K A, Sheth S A, Rizzuto D S, Kahana M J 2017.
606 Similar patterns of neural activity predict memory formation during encoding and retrieval.
607 *NeuroImage* 155:70–71.
- 608 Kucewicz M T, Berry B M, Kremen V, Brinkmann B H, Sperling M R, Jobst B C, Gross R E, Lega
609 B, Sheth S A, Stein J M, Das S R, Gorniak R, Stead S M, Rizzuto D S, Kahana M J, Worrell
610 G A 2017. Dissecting gamma frequency activity during human memory processing. *Brain*
611 140:1337–1350.
- 612 Kucewicz M T, Berry B M, Miller L R, Khadjevand F, Ezzyat Y, Stein J M, Kremen V, Brinkmann
613 B H, Wanda P, Sperling M R, Gorniak R, Davis K A, Jobst B C, Gross R E, Lega B, Gompel
614 J Van, Stead S M, Rizzuto D S, Kahana M J, Worrell G A 2018. Evidence for verbal memory
615 enhancement with electrical brain stimulation in the lateral temporal cortex. *Brain* 141:971–
616 978.
- 617 Kucyi A, Schrouff J, Bickel S, Foster B L, Shine J M, Parvizi J 2018. Intracranial electrophysiology
618 reveals reproducible intrinsic functional connectivity with human brain networks. *Journal of*
619 *Neuroscience* 38:4230–4242.
- 620 Lin J-J, Rugg M D, Das S, Stein J, Rizzuto D S, Kahana M J, Lega B C 2017. Theta band power
621 increases in the posterior hippocampus predict successful episodic memory encoding in
622 humans. *Hippocampus* 27:1040–1053.
- 623 Manning J. R., Jacobs J., Fried I., Kahana M. J. 2009. Broadband shifts in LFP power spectra
624 are correlated with single-neuron spiking in humans. *The Journal of Neuroscience* 29:13613–
625 13620.
- 626 Manning J. R., Polyn S. M., Baltuch G., Litt B., Kahana M. J. 2011. Oscillatory patterns in tem-

- 627 poral lobe reveal context reinstatement during memory search. *Proceedings of the National*
628 *Academy of Sciences, USA* 108:12893–12897.
- 629 Manning J R, Ranganath R, Norman K A, Blei D M 2014. Topographic factor analysis: a Bayesian
630 model for inferring brain networks from neural data. *PLoS One* 9:e94914.
- 631 Manning J. R., Sperling M. R., Sharan A., Rosenberg E. A., Kahana M. J. 2012. Spontaneously
632 reactivated patterns in frontal and temporal lobe predict semantic clustering during memory
633 search. *The Journal of Neuroscience* 32:8871–8878.
- 634 Manning J R, Zhu X, Willke T L, Ranganath R, Stachenfeld K, Hasson U, Blei D M, Norman
635 K A 2018. A probabilistic approach to discovering dynamic full-brain functional connectivity
636 patterns. *NeuroImage* 180:243–252.
- 637 Matsumoto Joseph Y, Stead Matt, Kucewicz Michal T, Matsumoto Andrew J, Peters Pierce A,
638 Brinkmann Benjamin H, Danstrom Jane C, Goerss Stephan J, Marsh W Richard, Meyer Fred B
639 et al. 2013. Network oscillations modulate interictal epileptiform spike rate during human
640 memory. *Brain* 136.
- 641 Miller K. J., Leuthardt E. C., Schalk G., Rao R. P. N., Anderson N. R., Moran D. W., Miller J. W.,
642 Ojemann J. G. 2007. Spectral changes in cortical surface potentials during motor movement.
643 *Journal of Neuroscience* 27:2424–2432.
- 644 Miller K. J., Shenoy P., den Nijs M., Sorensen L. B., Rao R. P. N., Ojemann J. G. 2008. Beyond
645 the gamma band: the role of high-frequency features in movement classification. *IEEE*
646 *Transactions on Biomedical Engineering* 55:1634 – 1637.
- 647 Mori S, Oishi K, Jiang H, Jiang L, Li X, Akhter K, Hua K, Faria A V, Mahmood A, Woods R,
648 Toga A W, Pike G B, Neto P R, Evans A, Zhang J, Huang H, Miller M I, van Zijl P, Mazziotta J

- 649 2008. Stereotaxic white matter atlas based on diffusion tensor imaging in an icbm template.
650 NeuroImage 40:570–582.
- 651 Norman K. A., Newman E., Detre G., Polyn S. M. 2006. How inhibitory oscillations can train
652 neural networks and punish competitors. Neural Computation 18:1577–1610.
- 653 Owen L L W, Chang T H, Manning J R 2019. High-level cognition during story listen-
654 ing is reflected in high-order dynamic correlations in neural activity patterns. bioRxiv
655 doi.org/10.1101/763821.
- 656 Rasmussen C E 2006. Gaussian processes for machine learning. MIT Press.
- 657 Reddy P G, Betzel R F, Khambhati A N, Shah P, Kini L, Litt B, Lucas T H, Davis K A, Bassett
658 D S 2018. Genetic and neuroanatomical support for functional brain network dynamics in
659 epilepsy authors genetic and neuroanatomical support for functional brain network dynamics
660 in epilepsy. arXiv 1809.03934.
- 661 Rubin T N, Kyoejo O, Gorgolewski K J, Jones M N, Poldrack R A, Yarkoni T 2017. Decoding
662 brain activity using a large-scale probabilistic functional-anatomical atlas of human cognition.
663 PLoS Computational Biology 13:e1005649.
- 664 Sarvas J 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse
665 problem. Phys. Med. Biol. 32:11–22.
- 666 Sawyer Robert J 1995. The Terminal Experiment. HarperPrism.
- 667 Scangos K W, Khambhati A N, Daly P M, Owen L L W, Manning J R, Ambrose J B, Austin E,
668 Dawes H E, Krystal A D, Chang E F 2020. Biomarkers of depression symptoms defined by
669 direct intracranial neurophysiology. bioRxiv doi.org/10.1101/2020.02.14.943118.

- 670 Sederberg P. B., Kahana M. J., Howard M. W., Donner E. J., Madsen J. R. 2003. Theta and
671 gamma oscillations during encoding predict subsequent recall. *Journal of Neuroscience*
672 23:10809–10814.
- 673 Sederberg P. B., Schulze-Bonhage A., Madsen J. R., Bromfield E. B., Litt B., Brandt A., Kahana
674 M. J. 2007a. Gamma oscillations distinguish true from false memories. *Psychological Science*
675 18:927–932.
- 676 Sederberg P. B., Schulze-Bonhage A., Madsen J. R., Bromfield E. B., McCarthy D. C., Brandt A.,
677 Tully M. S., Kahana M. J. 2007b. Hippocampal and neocortical gamma oscillations predict
678 memory formation in humans. *Cerebral Cortex* 17:1190–1196.
- 679 Sejnowski T J, Churchland P S, Movshon J A 2014. Putting big data to good use in neuroscience.
680 *Nature Neuroscience* 17:1440–1441.
- 681 Sherman M A, Lee S, Law R, Haegens S, Thorn C A, Hämäläinen M S, Moore C I, Jones S R
682 2016. Neural mechanisms of transient neocortical beta rhythms: Converging evidence from
683 humans, computational modeling, monkeys, and mice. *Proceedings of the National Academy*
684 *of Science USA* 113:E4885–E4894.
- 685 Simony E, Honey C J, Chen J, Hasson U 2016. Uncovering stimulus-locked network dynamics
686 during narrative comprehension. *Nature Communications* 7:1–13.
- 687 Snyder A Z 1991. Dipole source localization in the study of EP generators: a critique. *Electroen-*
688 *cephalography and Clinical Neurophysiology* 80:321–325.
- 689 Solomon E A, Gross R, Lega B, Sperling M R, Worrell G, Sheth S A, Zaghoul K A, Jobst B C,
690 Stein J M, Das S, Gorniak R, Inman C, Seger S, Kragel J E, Rizzuto D S, Kahana M J 2018. Mtl
691 functional connectivity predicts stimulation-induced theta power. *Nature Communications*
692 In press.

- 693 Sporns O, Betzel R F 2016. Modular brain networks. *Annual Review of Psychology* 67:613–640.
- 694 Talairach J., Tournoux P. 1988. Co-planar stereotaxic atlas of the human brain. Verlag, Stuttgart.
- 695 Tavor I, Jones O Parker, Mars R B, Smith S M, Behrens T E, Jbabdi S 2016. Task-free MRI predicts
696 individual differences in brain activity during task performance. *Science* 352:216–220.
- 697 Tomasi D, Volkow N D 2011. Association between functional connectivity hubs and brain
698 networks. *Cerebral Cortex* 21:2003–2013.
- 699 Turk-Browne N B 2013. Functional interactions as big data in the human brain. *Science* 342:580–
700 584.
- 701 Weidemann C T, Kragel J E, Lega B C, Worrell G A, Sperling M R, Sharan A D, Jobst B C,
702 Khadjevand F, Davis K A, Wanda P A, Kadel A, Rizzuto D S, Kahana M J 2019. Neural activity
703 reveals interactions between episodic and semantic memory systems during retrieval. *Journal*
704 *of Experimental Psychology: General* 148:1–12.
- 705 Wu H C, Nagasawa T, Brown E C, Juhasz C, Rothermel R, Hoechstetter K, Shah A, Mittal S,
706 Fuerst D, Sood S, Asano E 2011. Gamma-oscillations modulated by picture naming and word
707 reading: Intracranial recording in epileptic patients. *Clinical Neurophysiology* 122:1929–1942.
- 708 Yeo B T T, Krienen F M, Sepulcre J, Sabuncu M R, Lashkari D, Hollinshead M, Roffman J L,
709 Smoller J W, Zollei L, Polimieni J R, Fischl B, Liu H, Buckner R L 2011. The organization of
710 the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neuro-*
711 *physiology* 106:1125–1165.
- 712 Zalesky A, Fornito A, Harding I H, Cocchi L, Yücel M, Pantelis C, Bullmore E T 2010. Whole-
713 brain anatomical networks: does the choice of nodes matter? *NeuroImage* 50:970–983.