

Towards human Super EEG

Lucy L. W. Owen¹ and Jeremy R. Manning^{1*}

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

*To whom correspondence should be addressed; E-mail: jeremy.r.manning@dartmouth.edu.

Human *Super EEG*¹ entails measuring ongoing activity from every cell in a living human brain at millisecond-scale temporal resolutions. Although direct cell-by-cell Super EEG recordings are impossible using existing methods, here we present a technique for *inferring* neural activity at arbitrarily high spatial resolutions using human intracranial electrophysiological recordings. Our approach, based on Gaussian process regression, relies on two assumptions. First, we assume that some of the correlational structure of people’s brain activity is similar across individuals. Second, we resolve ambiguities in the data by assuming that neural activity from nearby sources will tend to be similar, all else being equal. One can then ask, for an arbitrary individual’s brain: given what we know about the correlational structure of other people’s brains, and given the recordings we made from electrodes implanted in this person’s brain, how would those recordings most likely have looked at *other* locations throughout this person’s brain?

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

Introduction

Current human brain recording techniques are fraught with compromise. Commonly used approaches include functional magnetic resonance imaging (fMRI), scalp electroencephalography (EEG), and magnetoencephalography (MEG). For each of these techniques, neuroscientists and

¹The term “Super EEG” was coined by Robert J. Sawyer in his popular science fiction novel *The Terminal Experiment* (1)

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

An increasingly popular approach is to improve the effective spatial resolution of MEG or scalp EEG data by using a geometric approach called *beamforming* to solve the biomagnetic or bioelectrical inverse problem (2). This approach entails using detailed brain conductance models (often informed by high spatial resolution anatomical MRI images) along with the known sensor placements (localized precisely in 3D space) to reconstruct brain signals originating from theoretical point sources deep in the brain (and far from the sensors). Traditional beamforming approaches must overcome two obstacles. First, the inverse problem beamforming seeks to solve has infinitely many solutions. Researchers have made traction towards constraining the solution space by assuming that signal-generating sources are localized on a regularly spaced grid spanning the brain and that individual sources are small relative to their distances to the sensors (3–5). The second, and in some ways much more serious, obstacle is that the magnetic fields produced by external (noise) sources are substantially stronger than those produced by the neuronal changes being sought (i.e. at deep structures, as measured by sensors at the scalp). This means that obtaining adequate signal quality often requires averaging the measured responses over tens to hundreds of responses or trials (e.g. see review by (5)).

Another approach to obtaining high spatial and temporal resolution neural data has been

to collect fMRI and EEG data simultaneously. Simultaneous fMRI-EEG has the potential to balance the high spatial resolution of fMRI with the high temporal resolution of scalp EEG, thereby, in theory, providing the best of both worlds. In practice, however, the signal quality of both recordings suffers substantially when the two techniques are applied simultaneously (e.g. see review by (6)). In addition, the experimental designs that are ideally suited to each technique individually are somewhat at odds. For example, fMRI experiments typically lock stimulus presentation events to the regularly spaced image acquisition time (TR), which maximizes the number of post-stimulus samples. By contrast, EEG experiments typically employ jittered stimulus presentation times to maximize the experimentalist’s ability to distinguish electrical brain activity from external noise sources such as from 60 Hz alternating current power sources.

The current “gold standard” for precisely localizing signals and sampling at high temporal resolution is to take (ECoG or iEEG) recordings from implanted electrodes (but from a limited set of locations in any given brain). This begs the following question: what can we infer about the activity exhibited by the rest of a person’s brain, given what we learn from the limited intracranial recordings we have from their brain and additional recordings taken from *other* people’s brains? Here we develop an approach, based on Gaussian process regression (7), that uses data from multiple people to estimate activity at arbitrary locations in each person’s brain (i.e., independent of their electrode placements). We test this *Super EEG* approach using a large dataset of intracranial recordings collected as neurosurgical patients studied and recalled random word lists (8–12). We show that the Super EEG algorithm recovers signals well from electrodes that were held out of the training dataset. We also examine the factors that influence how accurately activity may be estimated (recovered), which may have important implications for electrode design and for electrode placement in neurosurgical applications.

Approach

The Super EEG algorithm for inferring activity patterns throughout the brain using ECoG data in a multi-subject dataset is outlined and summarized in Figure 1. We describe (in this section) and evaluate (in *Results*) our approach using a large previously collected dataset comprising multi-session intracranial recordings taken from 6876 electrodes implanted in the brains of 88 epilepsy patients (8–12). We first applied a fourth order Butterworth notch filter to remove 60 Hz ($\pm .5$ Hz) line noise. We then excluded any electrodes that showed putative epileptiform activity. Specifically, we excluded from further analysis any electrode that exhibited an average kurtosis of 10 or greater across all of that patient’s recording sessions. We also excluded any patients with fewer than 2 electrodes that passed this criteria, as the Super EEG algorithm requires measuring correlations between 2 or more electrodes from each patient. Altogether this yielded clean recordings from 4149 electrodes implanted throughout the brains of 67 patients (Fig. 1A). Each individual patient contributes electrodes from a limited set of brain locations, which we localized in a common space (MNI152); an example patient’s 54 electrodes that passed the predefined kurtosis test are highlighted in red and blue.

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 (often) 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 :

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

where the width variable λ is a parameter of the algorithm (which may in principle be set according to location-specific tissue conductance profiles) that governs the level of spatial smoothing. In choosing λ for the analyses presented here, we sought to maximize spatial resolution (which implies a small value of λ) while also maximizing the algorithm's ability to generalize to any location throughout the brain, including those without dense electrode coverage (which implies a large value of λ). Using our prior work as a guide (13), we set $\lambda = 20$, although this could in theory be optimized, e.g. using cross validation.

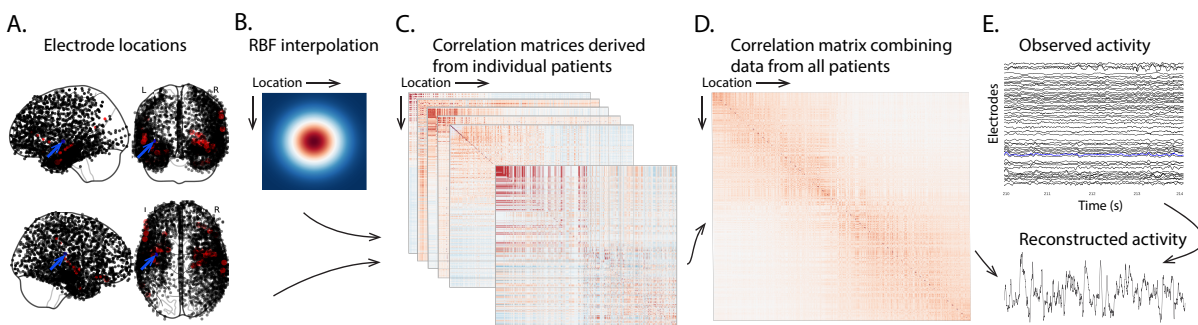
A second mechanism whereby a given region x can contribute to the recording at η is through anatomical connections between structures near x and η . We use spatial correlations in the data to estimate these anatomical connections. Let \bar{R} be the set of locations at which we wish to estimate local field potentials, and let R_s be set of locations at which we observe local field potentials from patient s (excluding the electrodes that did not pass the kurtosis test described above). In the analyses below we define $\bar{R} = \cup_{s=1}^S R_s$. We can calculate the expected inter-electrode correlation matrix for patient s , where $C_{s,k}(i, j)$ is the correlation between the time series of voltages for electrodes i and j from subject s during session k , using:

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

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

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

Next, we use Equation 1 to construct a number of to-be-estimated locations by number of patient electrode locations weight matrix, W . Specifically, W approximates how informative the



recordings at each location in R_s are in reconstructing activity at each location in \bar{R} , where the contributions fall off with an RBF according to the distances between the corresponding locations:

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

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

$$\hat{C}_s(x, y) = r \left(\frac{\sum_{i=1}^{|R_s|} \sum_{j=1}^i W(x, i) \cdot W(y, j) \cdot z(\bar{C}_s(i, j))}{\sum_{i=1}^{|R_s|} \sum_{j=1}^i W(x, i) \cdot W(y, j)} \right). \quad (6)$$

Intuitively, we construct an estimated correlation matrix from each individual patient's data (Fig. 1C) using Equation 6, and then we average these estimates across S patients to obtain the expected correlation matrix, \hat{K} (Fig. 1D):

$$\hat{K} = r \left(\frac{1}{S} \sum z(\hat{C}_s) \right). \quad (7)$$

Now we can use the following intuition: given (i) the observed responses from a limited set of locations in R_s (Y_s) and (ii) how each location's responses relate to all other responses (\hat{K}), we can estimate the LFP data from patient s , for any arbitrary location in \bar{R} (Fig. 1E).

Let α be the set of indices of patient s 's electrode locations in \bar{R} , and let β be the set of indices of all other locations in \bar{R} . In other words, β reflects the locations in \bar{R} where we did not observe a recording for patient s (these are the recording locations we will want to fill in using Super EEG). We can sub-divide \hat{K} as follows:

$$\hat{K}_{\beta, \alpha} = \hat{K}(\beta, \alpha), \text{ and} \quad (8)$$

$$\hat{K}_{\alpha, \alpha} = \hat{K}(\alpha, \alpha). \quad (9)$$

Here $\hat{K}_{\beta,\alpha}$ stores the correlations between the “unknown” activity at the locations in β and the observed activity at the locations in α , and $\hat{K}_{\alpha,\alpha}$ stores the correlations between the observed recordings (at the locations in α).

Let $Y_{s,k,\alpha}$ be the number-of-timepoints (T) by length(α) matrix of (observed) voltages from the electrodes in α during session k from patient s . Then we can estimate the voltage from patient s ’s k^{th} session at the locations in β using (7):

$$Y_{s,k,\beta} = ((\hat{K}_{\beta,\alpha} \cdot \hat{K}_{\alpha,\alpha}^{-1}) \cdot Y_{s,k,\alpha}^T)^T. \quad (10)$$

This equation is the foundation of the Super EEG algorithm. Whereas we observe the recordings only at the locations in α , Equation 10 allows us to estimate the recordings at all locations in β , which we can define *a priori* to include any locations we wish, throughout the brain. This yields estimates of the time-varying voltages at *every* location in \bar{R} .

We designed our approach to be agnostic to electrode impedances, as electrodes that do not exist do not have impedances. Therefore our algorithm recovers voltages in standard deviation (z -scored) units rather than attempting to recover absolute voltages. (This property reflects the fact that $\hat{K}_{\beta,\alpha}$ and $\hat{K}_{\alpha,\alpha}$ are correlation matrices rather than covariance matrices.) Also, note that Equation 10 requires computing a T by T matrix, which can become computationally intractable when T is very large (e.g. for the patient highlighted in Fig. 2, $T = 20458799$). However, we may approximate $Y_{s,k,\beta}$ in a piecewise manner by filling in $Y_{s,k,\beta}$ in blocks of size b samples (using the corresponding samples from $Y_{s,k,\alpha}$). In our computations we set $b = 25000$.

The Super EEG algorithm described above and in Figure 1 allows us to estimate (up to a constant scaling factor) LFPs for each patient at all arbitrarily chosen locations in the set \bar{R} , *even if we did not record that patient’s brain at all of those locations*.

Results

To test the accuracy with which the Super EEG algorithm reconstructs activity throughout the brain, we held out each electrode from the full dataset in turn and treated it as unobserved. We then asked: how closely did each of the Super EEG-reconstructed LFPs match the observed data? We sought to evaluate both the overall reconstruction accuracy as well as how reconstruction accuracy varied as a function of implantation location.

We first examined raw LFP traces and their associated Super EEG-derived reconstructions. Figure 2A displays the LFP from the blue electrode in Figure 1A, and its associated reconstruction, during a 4 s time window during one of the patient’s 6 recording sessions. Figure 2B displays a 2D histogram of the observed versus reconstructed voltages for every sample across 14.2 total hours of recordings from that patient (correlation: $r = 0.95, p < 10^{-10}$). Although the Super EEG algorithm recovered the recordings from this electrode well, we sought to quantify the algorithm’s performance across the full dataset.

Holding out each electrode from each patient in turn, we computed the average correlation (across recording sessions) between the Super EEG-reconstructed voltage traces and the observed voltage traces from that electrode. For each reconstruction, we estimated the full-brain correlation matrix using every *other* patient’s data (i.e. every patient except the one who contributed the to-be-reconstructed electrode data). In our analyses, we then substituted the average correlation matrix computed after excluding patient s ’s data for \hat{K} in Equations 8 and 9. This step ensured that the data we were reconstructing could not also be used to estimate the between-location correlations that drove the reconstructions via Equation 10 (otherwise the analysis would be circular).

We obtained a single correlation coefficient for each electrode location in \bar{R} , reflecting how well the Super EEG algorithm was able to recover the recording at that location (Fig. 3A).

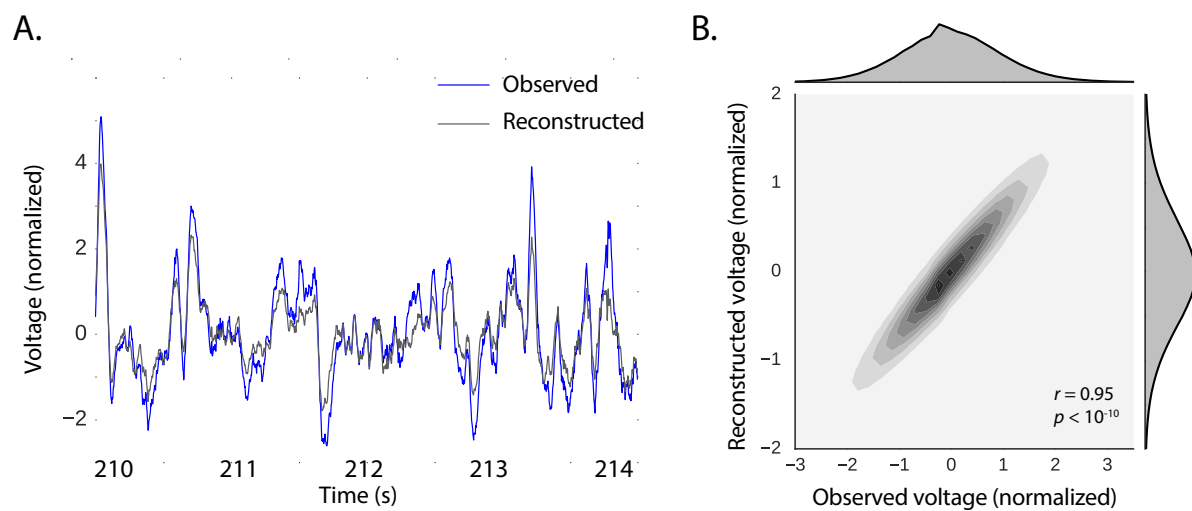


Figure 2: Observed and reconstructed LFP from a single electrode. A. Example LFP. A 4 s recording from the blue electrode in Figure 1A is displayed in blue, and the reconstructed LFP during the same time window is shown in gray. All voltages are plotted in standard deviation units. **B. Observed versus reconstructed voltages over 14.2 hours.** The 2D histogram reflects the relation between distributions of observed versus reconstructed voltages from patient 6, across the 14.2 hours of recorded data collected in 6 recording sessions. The correlation reported in the panel is between the observed and reconstructed voltages.

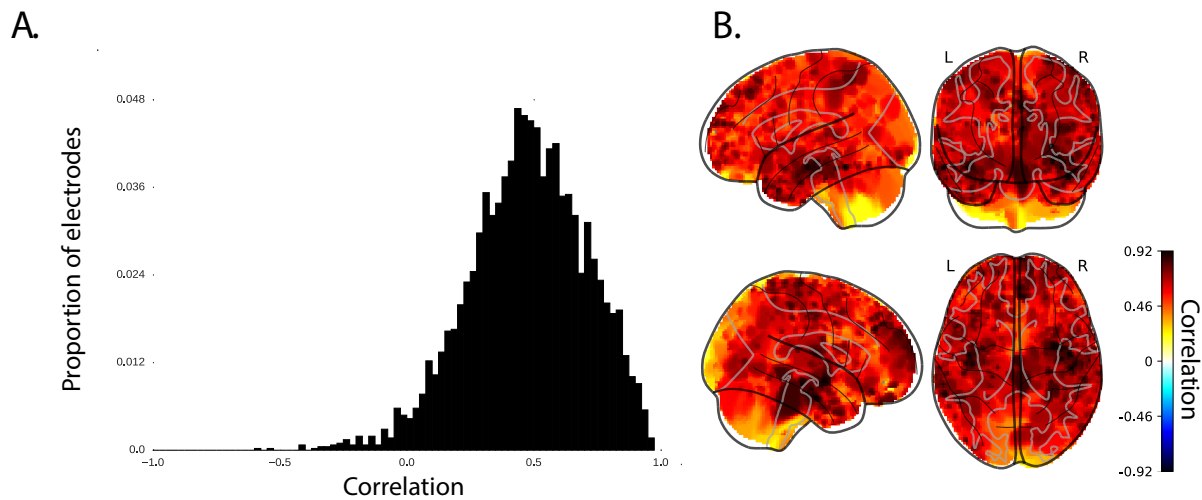


Figure 3: Reconstruction quality. **A. Distribution of correlation coefficients.** Across all electrodes from all patients, the panel displays the distribution of correlations between the observed and reconstructed LFP data. **B. Correlation map.** The glass brain map displays the average correlation between the observed and reconstructed LFP data by location. For this and all other glass brain maps that follow, we used Inverse Distance Weighting (14) to interpolate values at locations where no electrodes were implanted.

We compared this distribution of correlation coefficients to the distribution of across-session average correlation coefficients between the recordings from each electrode and its nearest neighbor from the same patient (paired t -test between z -transformed correlation coefficients: $t(4148) = 2.36, p = 0.018$). This is an especially conservative test, given that the Super EEG reconstructions exclude (from the correlation matrix estimates) all data from the patient whose data is being reconstructed, whereas the nearest neighbor correlations are between the recordings from the two nearest electrodes from the same patient. That the Super EEG-derived correlations were reliably stronger than these nearest-neighbor-derived correlations is exciting for two reasons. First, it implies that distant electrodes provide additional predictive power to the data reconstructions beyond the information contained in nearby electrodes. Second, it implies that the spatial correlations driving the Super EEG algorithm are, to some extent, shared across people.

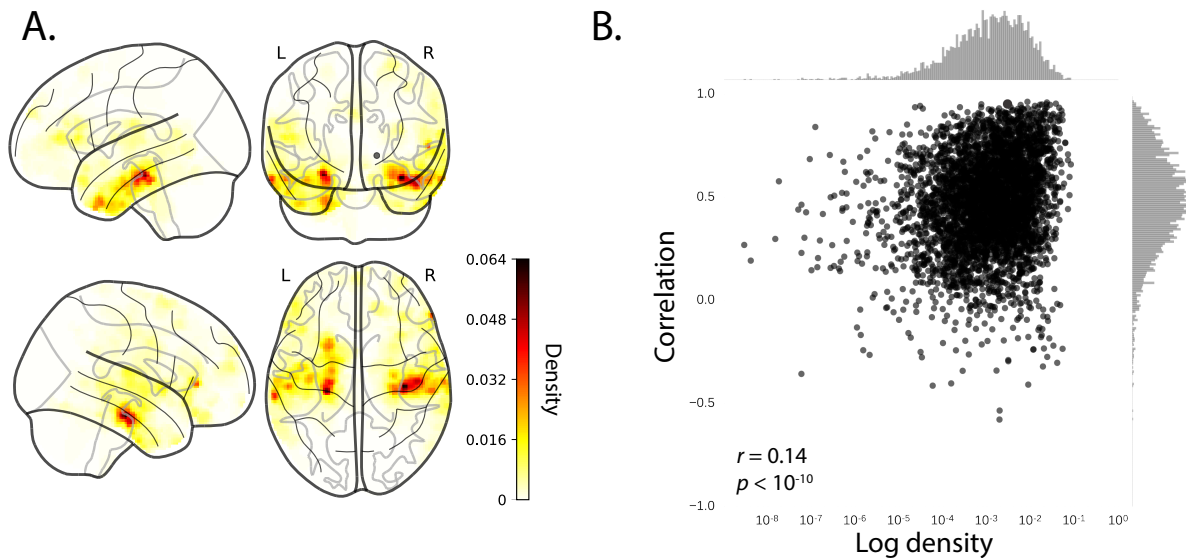


Figure 4: **Sampling density and reconstruction quality.** **A.** Sampling density by location. **B.** Reconstruction accuracy versus sampling density.

We also wondered whether reconstruction quality (measured as the correlation between the observed and reconstructed data) varied with the electrode locations (Fig. 3B). In general, reconstruction quality remained high throughout the brain. Qualitatively, reconstruction accuracy appeared especially high in the medial temporal lobe (bilaterally). Because the medial temporal lobe is a common epileptic focus (and therefore a common target for electrode implantation), we wondered whether reconstruction quality might relate to how densely each brain area was sampled in our test dataset (Fig. 4A). Indeed, we found a weak but statistically reliable correlation between reconstruction quality (z -transformed correlation) and electrode density (defined as $\exp\{-d\}$, where d is the average Euclidean distance between the electrode and its 10 nearest neighbors; $r = 0.14$, $p < 10^{-10}$; Fig. 4B).

In addition to exploring how reconstruction quality varies with location, we also wondered whether there might be effects of electrode placements on reconstruction quality. For example, are there particular implantation locations that yield especially high reconstruction accuracies at other locations throughout the brain? To gain insights into this question, for each electrode

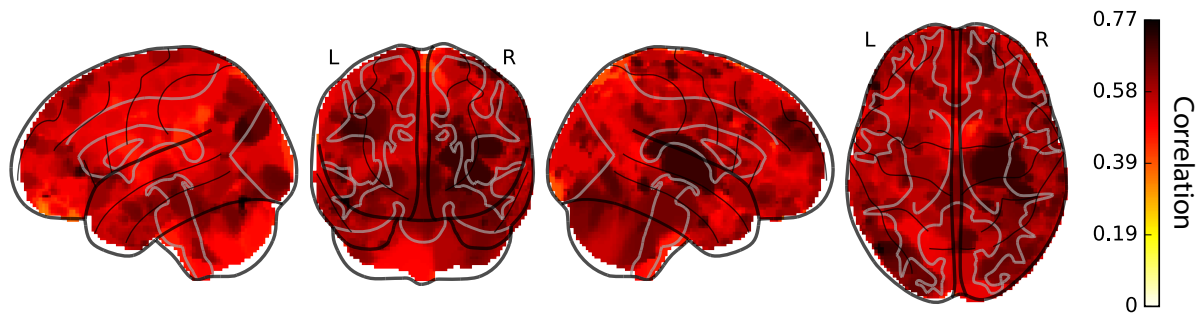


Figure 5: **Most informative electrode locations.** The glass brain map displays the average reconstruction correlations (across all electrodes) for patients with electrodes within a 10 MNI unit diameter sphere centered on each location.

location (across all patients), we computed the average reconstruction correlation (across all electrodes) for any patients who had electrodes within a 10 MNI unit diameter sphere centered on each electrode location. The resulting map highlights the locations of implanted electrodes from patients whose reconstructions were especially accurate (Fig. 5). The locations in dark red might therefore be good candidate implantation targets for neurosurgeons and neurologists who wish to use Super EEG to reconstruct full-brain electrophysiological signals.

The above findings, that one can infer brain activity throughout a person’s brain using recordings from a limited number of location from that person’s brain in conjunction with recordings from other people’s brains, have deep implications for the structure of brain data. The first implication is that the correlational structure of different people’s brain data is largely preserved across individuals. At one level this appears to contradict recent work showing that each individual has a unique resting state connectome (15); we return to this point in the *Discussion*. The second implication is that brain activity is highly redundant. For example, despite the fact that our brains comprise roughly 100 billion neurons (16), those neurons are highly interconnected and inter-dependent. This explains why techniques for embedding (seemingly) high dimensional brain recordings into much lower dimensional spaces [e.g. using techniques like Principle Components Analysis (PCA; (17)) and Independent Components

Analysis (ICA; (18, 19))] still retain much of the variance in the original data.

We applied PCA to each patient's data and found that, on average, 64.3% of the components explained at least 95% of the variance in the original data (Fig. 6). This suggests that the full set of recordings (across all electrodes) contain information about a substantially greater range of locations throughout the patients' brains. However, this measure of data redundancy also varied substantially across patients (minimum proportion: 0.08; maximum proportion: 0.80). We found that, patient-by-patient, the proportion of components needed to explain 95% of the variance in the original data was negatively correlated with the average reconstruction quality (defined as the z -transformed correlation between the observed and reconstructed data): $r = -0.62, p < 10^{-7}$. This indicates that the Super EEG algorithm performs best for patients whose brain data are the most redundant, irrespective of the number of electrodes implanted in their brains (correlation between reconstruction quality and number of electrodes: $r = 0.16, p = 0.19$).

Discussion

Super EEG infers full-brain activity patterns by leveraging correlations in those patterns of brain activity within and across people. Although the approach may, in principle, be used to infer brain activity *anywhere* in the brain, the inferences perform slightly better for regions with dense electrode sampling across patients. (Taken to the logical extreme, we could not hope to accurately recover activity patterns from brain areas where no recordings existed from any patient.) As more data are included in the inference procedure, this suggests that reconstruction accuracy should improve.

A fundamental assumption of the Super EEG algorithm is that the data covariance matrix is stable over time and across people. This is a useful simplification. However, a growing body of evidence from the fMRI community suggests that the data covariance matrix changes in

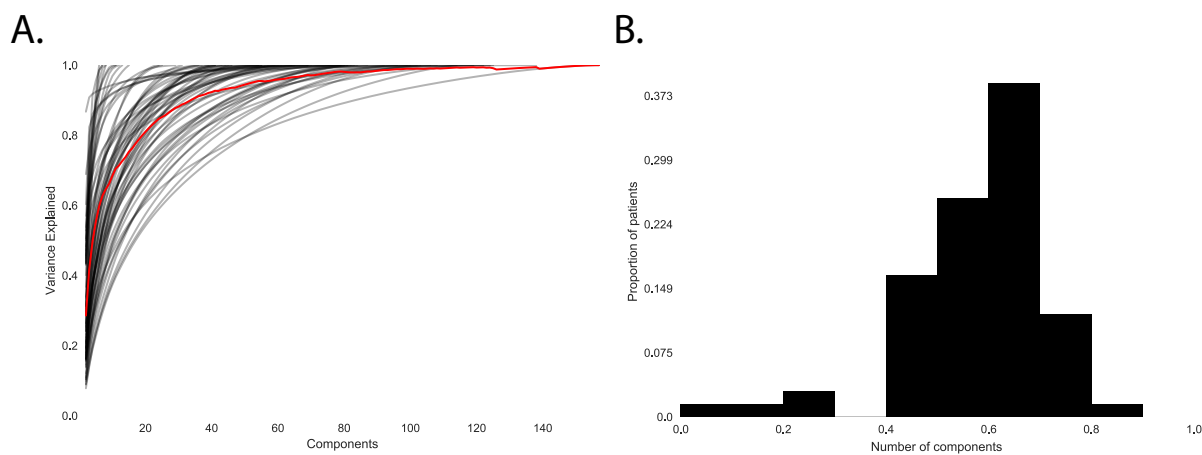


Figure 6: The redundancy of brain data. We used principal components analysis to project each patient's data onto a lower dimensional space. **A. Proportion of variance explained by each principal component.** For each patient, the panel displays the across-sessions average proportion of the variance explained as a function of the number of principal components. The across-patients average is shown in red. **B. Proportion of components that explains 95% of the variance.** The histogram displays the distribution across patients of the proportions of components (relative to the original number of electrodes) needed to explain at least 95% of the variance in the original data.

meaningful ways over time (for example, the data covariance matrix changes from moment-to-moment during story listening, serving as a unique “fingerprint” for each moment of the story; further, these task-driven timepoint-specific covariance fingerprints appear to be largely preserved across people (20, 21)). These findings indicate that the full-brain covariance matrix is not stable over time. Other recent work has shown that people’s resting state connectivity matrices may be used to uniquely identify individuals and predict fluid intelligence scores (15). This indicates that the full-brain covariance matrix is not stable across people. If the fundamental stability assumptions that Super EEG relies on are violated, how can the Super EEG algorithm still accurately recover LFP data? It is important to recognize that the fact that variability (over time or across people) is predictive (e.g. of cognitive states during story listening or fluid intelligence scores) does not necessarily mean that this variability is large in magnitude. Rather, we have long known that brain structure is tightly preserved across individuals (and over time, at least on the timescale of typical clinical and experimental recording sessions), and any functional changes must occur within the framework of the underlying structural anatomy. Nevertheless, one could imagine future improvements to the Super EEG approach that leverage resting state fMRI or structural data [e.g. diffusion tensor imaging (DTI)] to estimate Bayesian priors over the correlation matrices inferred, in the current framing, using only ECoG data. Further, relaxing the assumption that the covariance matrix is stable (over time and/or across people), and/or incorporating more detailed brain conductance models (e.g. informed by structural MRI scans) may improve the predictive performance of the approach.

One potential limitation of the Super EEG approach is that the above assumption of covariance stability across people may be violated even more if different patients are performing different cognitive tasks. (In the above analyses, all patients performed the same list learning task as their recordings were taken.) Additional work is needed to understand the extent to which the current findings generalize across cognitive tasks. If the covariance matrix is unsta-

ble across tasks, accurate reconstructions using the Super EEG approach would require building up new databases for estimating each task-specific covariance matrix. Or, using a more sophisticated approach, one could create a hierarchical model whereby each task-specific covariance matrix was modeled as a perturbation of a “global” task-unspecific covariance matrix (which could in turn be informed by fMRI or DTI data). Alternatively, if the covariance matrix is stable across tasks, this would suggest that recordings from multiple studies could be combined to improve the overall reconstruction accuracy.

A second potential limitation of the Super EEG approach is that it does not provide a natural means of estimating the precise timing of single-neuron action potentials. Prior work has shown that gamma band and broadband activity in the LFP may be used to estimate the firing rates of neurons that underly the population contributing to the LFP (22). Because Super EEG reconstructs LFPs throughout the brain, one could in principle use gamma or broadband power in the reconstructed signals to estimate the corresponding firing rates (though not the timings of individual action potentials).

Beyond providing a means of estimating ongoing activity throughout the brain using already implanted electrodes, our work also has implications for where to place the electrodes in the first place. Electrodes are typically implanted to maximize coverage of suspected epileptogenic tissue. However, our findings suggest that this approach could be further optimized. Specifically, one could leverage not only the non-invasive recordings taken during an initial monitoring period (as is currently done), but also recordings collected from other patients. We could then ask: given everything we know about the other patients and from the scalp recordings of this new patient, where should we place a fixed number of electrodes to maximize our ability to map seizure foci? As shown in Figure 5, recordings from different locations are differently informative in terms of reconstructing the spatiotemporal patterns throughout the brain. This property might be leveraged in decisions about where to surgically implant electrodes in future patients.

Concluding remarks

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

Acknowledgements

We are grateful for useful discussions with Luke J. Chang and Matthijs van der Meer. We are also grateful to Michael J. Kahana for generously sharing the ECoG dataset we analyzed in our paper, which was collected under NIMH grant MH55687 to MJK. Our work was also supported in part by NSF EPSCoR Award Number 1632738. The content is solely the responsibility of the authors and does not necessarily represent the official views of our supporting organizations.

References and Notes

1. R. J. Sawyer, *The Terminal Experiment* (HarperPrism, 1995).
2. J. Sarvas, *Phys. Med. Biol.* **32**, 11 (1987).
3. A. Z. Snyder, *Electroencephalography and Clinical Neurophysiology* **80**, 321 (1991).
4. S. Baillet, J. C. Mosher, R. M. Leahy, *IEEE Signal Processing Magazine* **18**, 14 (2001).
5. A. Hillebrand, K. D. Singh, I. E. Holliday, P. L. Furlong, G. R. Barnes, *Human Brain Mapping* **25**, 199 (2005).

6. R. J. Huster, S. Debener, T. Eichele, C. S. Herrmann, *The Journal of Neuroscience* **32**, 6053 (2012).
7. C. E. Rasmussen, *Gaussian processes for machine learning* (MIT Press, 2006).
8. P. B. Sederberg, M. J. Kahana, M. W. Howard, E. J. Donner, J. R. Madsen, *Journal of Neuroscience* **23**, 10809 (2003).
9. P. B. Sederberg, *et al.*, *Cerebral Cortex* **17**, 1190 (2007).
10. P. B. Sederberg, *et al.*, *Psychological Science* **18**, 927 (2007).
11. J. R. Manning, S. M. Polyn, G. Baltuch, B. Litt, M. J. Kahana, *Proceedings of the National Academy of Sciences, USA* **108**, 12893 (2011).
12. J. R. Manning, M. R. Sperling, A. Sharan, E. A. Rosenberg, M. J. Kahana, *The Journal of Neuroscience* **32**, 8871 (2012).
13. J. R. Manning, R. Ranganath, K. A. Norman, D. M. Blei, *PLoS One* **9**, e94914 (2014).
14. D. Shepard, *Proceedings of the ACM National Conference* pp. 517–524 (1968).
15. E. S. Finn, *et al.*, *Nature Neuroscience* **18**, 1664 (2015).
16. G. Buzsáki, *Rhythms of the Brain* (Oxford University Press, New York, 2006).
17. K. Pearson, *Philosophical Magazine* **2**, 559 (1901).
18. C. Jutten, J. Herault, *Signal Processing* **41**, 1 (1991).
19. P. Comon, C. Jutten, J. Herault, *Signal Processing* **24**, 11 (1991).
20. E. Simony, C. J. Honey, J. Chen, U. Hasson, *Nature Communications* (2016).

21. J. R. Manning, *et al.*, *bioRxiv* p. 106690 (2017).
22. J. R. Manning, J. Jacobs, I. Fried, M. J. Kahana, *The Journal of Neuroscience* **29**, 13613 (2009).