Basis profile curve identification to understand electrical stimulation effects in human brain networks

Kai J. Miller^{*1,2}, Klaus-Robert Müller³⁻⁶, Dora Hermes²

1 Department of Neurological Surgery, Mayo Clinic, Rochester, MN, USA

2 Department of Biomedical Engineering & Physiology, Mayo Clinic, Rochester, MN, USA

3 Google Research, Brain Team, Berlin, Germany

4 Machine Learning Group, Department of Computer Science, Berlin Institute of Technology, Berlin, Germany

5 Department of Artificial Intelligence, Korea University, Seoul, Republic of Korea

6 Max Planck Institute for Informatics, Saarbrücken, Germany

* miller.kai@mayo.edu

Abstract

Brain networks can be explored by delivering brief pulses of electrical current in one area while measuring voltage responses in other areas. We propose a convergent paradigm to study brain dynamics, focusing on a single brain site to observe the average effect of stimulating each of many other brain sites. Viewed in this manner, visually-apparent motifs in the temporal response shape emerge from adjacent stimulation sites. This work constructs and illustrates a data-driven approach to determine characteristic spatiotemporal structure in these response shapes, summarized by a set of unique “basis profile curves” (BPCs). Each BPC may be mapped back to underlying anatomy in a natural way, quantifying projection strength from each stimulation site using simple metrics. Our technique is demonstrated for an array of implanted brain surface electrodes in a human patient. This framework enables straightforward interpretation of single-pulse brain stimulation data, and can be applied generically to explore the diverse milieu of interactions that comprise the connectome.

Author summary

We present a new machine learning framework to probe how brain regions interact using single-pulse electrical stimulation. Unlike previous studies, this approach does not assume a form for how one brain area will respond to stimulation in another area, but rather discovers the shape of the response in time from the data. We call the set of characteristic discovered response shapes “basis profile curves” (BPCs), and show how these can be mapped back onto the brain quantitatively. An illustrative example is included from one of our human patients to characterize inputs to the parahippocampal gyrus. A code package is downloadable from <https://purl.stanford.edu/rc201dv0636> so the reader may explore the technique for own data, or study sample data provided to reproduce the illustrative case presented in the manuscript.

Introduction

Brain networks have been explored electrophysiologically with a variety of techniques, spanning a variety of spatial scales, such as electroencephalography (EEG),

magnetoencephalography (MEG), intracranial EEG (iEEG), and microelectrode local field potentials (LFPs). Efforts to infer connectivity between brain regions may search for correlated signals in response to supervised perturbation by a behavioral task, or in an unsupervised state (“resting” awake, or sleeping). Alternately, it has been shown that interactions between brain regions may be studied by applying or inducing pulses of electrical stimulation to a particular site, while measuring the electrophysiological response elsewhere [1–3]. In recent years, a sub-field of neuroscience has matured around systematic stimulation and measurement through implanted (iEEG) arrays of brain surface (electrocorticography, ECoG) or deeply-penetrating (stereoelectroencephalography, SEEG) electrodes, typically called “cortico-cortical evoked potentials” (CCEPs) or, for the special case of short pulses separated by several seconds, “single-pulse electrical stimulation” (SPES) [4–6]. The more general term, “CCEP” will be used to refer to both in this text.

For an array of N total electrodes, there are a potential set of order N^2 CCEP interactions that may be explored (for bipolar stimulation, the actual number will depend on which sets of adjacent stimulation pairs are chosen to deliver electrical pulses through). There are different network paradigms with which to address CCEP data, illustrated in figure 1. In the “all-to-all” case where one wishes to examine the full set of interactions (incorporating the temporal property of each response), the limited number of stimulation events possible to record in the clinical environment (where these measurements are made) does not allow for a well-defined exploration of the network. Therefore, scientists have imposed a type of constraint, or a hybridization of several constraints. One such approach is to reduce the problem by beginning the exploration with a pre-defined interaction based on location, and then study the temporal dynamics within that paired framework (figure 1D) [4].

A “divergent” paradigm is commonly adopted, where the effect of stimulating a chosen site on all of the other sites is used to infer motifs of connectivity. However, the underlying cortical laminar architecture of each recipient (measured from) site is very different, and therefore the voltage timecourse of each CCEP cannot be interpreted in a common “physiological language” to distinguish different types of interactions [7]. Nonetheless, many studies have found it useful to assume a canonical form (temporal structure) and then parametrize within this assumed form (e.g. voltage at fixed delay time from stimulus). The most common of these is a negative deflection between about 10 and 100 ms (or “N1”) and a later second negative deflection (“N2”) to characterize which brain regions are connected [4]. That approach forces each response into a fixed interpretation where the form being fit may not actually be present.

A different, “convergent”, paradigm is to focus on measurement from a single site, and examine the effect of stimulating in the remainder of the array [8]. Although initial studies from this paradigm have retained the N1/N2 framework, the convergent approach is, in principle, more tractable than other paradigms. When one measures from a single site and stimulates many, there is only one “physiological language” that evoked responses must be interpreted in, because the underlying anatomy is unchanged. In this context, different shaped temporal responses must imply a fundamentally different nature of interaction between brain areas.

With the convergent paradigm as our initial framework, we made the visual observation that inputs to primary motor cortex had very different shaped CCEPs from one another, but with a compelling spatial clustering [9]. This finding was again observed, but for responses in the parahippocampal gyrus (PHG) of a different patient, and are shown in figure 2.

While anatomical clustering of stimulation sites that produce similar voltage responses can be observed anecdotally by visual inspection, there is no generally-available quantitative tool for organizing these physiological

Fig 1. Cortico-cortical evoked potential analysis paradigms. **A:** *Convergent* - Evoked responses at one chosen site (gray circle) are compared with the effect of stimulating all other sites (yellow circles with lightning bolt). For N electrodes, this characterizes N interactions. **B:** *Divergent* - The temporal response of all sites are examined and compared in response to stimulation of a chosen site (N interactions). **C:** *All-to-all* - All N^2 interactions between sites are characterized. **D:** *Hypothesis preselected* - Two sites are chosen based upon a pre-defined anatomical or functional hypothesis, and a 1-way or 2-way interaction between them is characterized. **E:** In the convergent paradigm, all measured responses from a brain surface electrode are associated with the same underlying laminar architecture, so each response shape measured implies a distinct type of input. **F:** In the divergent paradigm, different shaped responses may be measured from different sites, in response to stimulation at a single site. This creates ambiguity because different shaped responses cannot distinguish between the same type of output arriving at cortical sites with different underlying laminar architecture and different types of inputs to sites with similar laminar architecture.

measurements¹ [7]. This work aims to close this gap and develop a novel tool to uncover and cluster these temporal motifs in CCEPs and enable systematic exploration of connectivity. We name the resulting canonical voltage response shapes “basis profile curves” (BPC) and will contribute by a framework to identify them.

Within the convergent CCEP paradigm, there are several criteria to constrain the BPC framework. First, there should be no assumption of the form of BPCs – they should emerge from the data naturally. Second, they should be able to be mapped into the original data, and onto the brain anatomy in an intuitive way. Third, each stimulation response trial should be able to be parameterized by a single BPC (rather than a superposition of BPCs). Fourth, there should be no orthogonality constraint in BPC shapes in case features (such as the N1 or N2) are shared between BPCs.

Technically, this framework amounts to a hierarchical clustering problem, where the subgroup of single-stimulation events from each stimulation site are known, but how the stimulation site subgroups cluster into the larger group of characteristic stimulation response shapes is to be discovered. In this manuscript, we describe a generic solution to this problem, and showcase its potential to obtain novel insights for a representative dataset from implanted ECoG electrodes in a human patient (figure 2).

¹Our exploration of the literature could not uncover a technique perhaps originating from a different application setting that could be translated for the present brain data.

Fig 2. Single-pulse cortico-cortical evoked potentials. **A:** An array of brain surface (ECoG) electrodes were surgically placed on the left hemisphere of a brain tumor patient. **B:** The voltage (red trace) was measured at a parahippocampal gyrus (PHG) electrode site. **C:** Biphasic stimulation pulses were delivered between adjacent electrodes throughout the array. **D:** Responses from each stimulation pulse are aligned into a matrix $V_k(t)$. **E:** Averaged subgroup responses $G_n(t)$ (from the PHG measurement site) are shown at the site of each stimulation pair that produced them.

Materials and methods

Clinical measurement of cortico-cortico evoked potentials

The patient shown for the illustrative example of this technique was a patient with a left temporal-occipital-parietal tumor (discovered to be a dysembryoplastic neuroepithelial tumor) who underwent placement of an electrocorticographic (ECoG) electrode array to localize their seizures and map brain function (Fig. 2). This array consisted of a 6x8 frontal-temporal-parietal grid, a 2x8 anterior temporal grid, and 3 1x4 sub-temporal strips of platinum electrodes (Ad-Tech, Racine, WI). The circular electrode contacts had 4 mm diameter (2.3 mm exposed), 1 cm inter-electrode distance, and were embedded in silastic. These arrays were surgically placed on the sub-dural brain surface during staged surgical treatment for functional mapping and seizure localization prior to tumor resection, with consented research conducted under Mayo Clinic IRB 15-006530, which also authorizes sharing of the data.

Voltage data $V(t_0)$ were recorded at 2048Hz on a Natus Quantum amplifier. Electrode pairs were stimulated ~ 10 -12 times with a single biphasic pulse of 200 microseconds duration and 6mA amplitude every 3-7 seconds using a Nicolet Cortical Stimulator (figure 2A-C). Electrodes were localized on the CT and coregistered to an MRI using the “CTMR” package [10], available in the “ctmr” folder of the ECoG library [11] or on github [12].

Data structure (Fig. 2)

Data are first structured in a stimulation-evoked voltage matrix $\mathbf{V}$: The time-by-1 matrix $V(t_0)$ for the whole experiment, from the chosen electrode, was sorted into the matrix $V_k(t)$, where t denotes the time from the k^{th} electrical stimulation, τ_k : $(\tau_k + t_1) \leq t \leq (\tau_k + t_2)$. The dimensions of $\mathbf{V}$ are $T \times K$, with K total stimulation events (trials) by T total timepoints (over the interval $t_1 \leq t \leq t_2$). For this illustration, t_1 was set to 0.050 sec and t_2 to 2.000 sec.

Within this matrix $V_k(t)$, similar events were given a common subgroup label n : These are naturally-grouped stimulation events, corresponding to the physical location of stimulation pair. The number of repeats is not the same in each stimulation-pair subgroup (typically 10-12). There are N total of these stimulation-pair subgroups. Although we assume that stimulations within each subgroup are independent for the purpose of these analyses, this is an approximation. An example of potential non-independence may be seen in the non-zero offset of individual trials from the yellow site pair in figure 2C, presumably due to direct charging of the cortical lamina due to proximity of the stimulation-pair sites to the recording site.

Trials from the same pair of stimulation sites, n , can be combined to obtain the subgroup-averaged evoked voltage change matrix $\mathbf{G}$: For stimulation subgroup n , the average voltage temporal profile is $G_n(t) = \langle V_k(t) \rangle_{k \in n}$. The dimensions of $\mathbf{G}$ are $N \times T$. With this type of brain stimulation data, $G_n(t)$ are commonly given the name “Cortico-cortical evoked potentials” (CCEPs).

Fig 3. Technique for identification of Basis Profile Curves (BPCs). An illustration of the series of steps to extract BPCs from the voltage matrix V and subgroup assignments $k \in n$. **A:** Within stimulation-pair subgroup self-projections. **B:** Between-subgroup cross-projections. **C:** An illustration of sets of cross- and self-projections for stimulation-pair subgroup 11, $S_{11,m}$. **D:** The significance of each set $S_{n,m}$ is determined initially by t-value vs. zero. Negative t-values are set to zero. The matrix of these values is then scaled to 1, and labeled Ξ . **E:** Non-negative matrix factorization (NNMF) is performed to identify structure. **F:** The inner dimension of NNMF is iteratively reduced. **G:** BPCs are identified from the groups clustered in the rows of H .

Single-trial cross-projections & Significance matrix (Fig. 3A-D)

In order to understand shared structure between stimulation trials, we first obtain a matrix of unit-normalized single trials: $\tilde{V}_k(t) = V_k(t)/|V_k(t)|$. Each $\tilde{V}_k(t)$ is then projected into all other trials, $\mathbf{P} = \tilde{\mathbf{V}}^T \mathbf{V}$:

$$P(k, l) = \sum_t \tilde{V}_k(t) V_l(t)$$

Note that $P(k, l) \neq P(l, k)$. The full matrix $\mathbf{P}$ is subsequently sorted into an array of sets $S_{n,m}$ that characterize each cross-subgroup interaction:

$$S_{n,m} = P\{k \in n, l \in m\}$$

with $k = l$ omitted.

We then construct a matrix of t-values whose elements $\Xi(n, m)$ are the t-values of the corresponding distribution in $S_{n,m}$ (each set's mean divided by its standard error [13]). This matrix is first subjected to a non-negativity constraint (needed for subsequent factorization): if $\Xi(n, m) < 0$, then $\Xi(n, m) \mapsto 0$. Ξ is then scaled to its maximum $\Xi \mapsto \Xi / \max(\Xi)$, making $0 \leq \Xi(n, m) \leq 1$. This significance matrix plays a role analogous to a cross-correlation matrix to help understand preserved structure of individual trials within, and between different, stimulation-pair subgroups.

Note that the diagonal elements of Ξ can be very small. Interestingly, in some cases, off-diagonal elements are larger than corresponding diagonal element from the same row (Figure 3D). This occurs when there is reliable structure in the response, but, first, within-group variation is larger than the cross-group variation, and, second, there is structure in the variation itself. A simple example of this is when two subgroups produce nearly identical response shapes, but there is adaptation in the response to repeated stimulation for each electrode-pair subgroup - a finding illustrated in some superior temporal gyrus sites from our example (visually apparent in subgroups 11&23 of Fig. 3A-C, and the green in Fig. 4C).

Non-Negative Matrix Factorization (NNMF) for clustering (Fig. 3E)

When clustering stimulation sites that produce similar measured responses, a non-negative projection weight constraint must be applied. This is done because the problem is not a source-localization - the same cluster cannot have both positive and negative contributions to the recording site. Physiologically, this follows because laminar anatomy not invertible, and when there is positive-negative flip in voltage in one electrode, even if shape is similar, that points to a different biology we want separately segregated into a different cluster (as illustrated in Fig 1E).

The process of non-negative matrix factorization is applied to the matrix Ξ , performing a decomposition [14]:

$$\Xi \sim \mathbf{W} \mathbf{H}$$

Where Ξ has dimensions $N \times M$, $\mathbf{W}$ has dimensions $N \times Q$, and $\mathbf{H}$ has dimensions $Q \times M$. The goal of NNMF in this context is to minimize $\eta = \|\Xi - \mathbf{W}\mathbf{H}\|^2$, with the non-negativity constraint $W_{nq}, H_{qm} \geq 0$.

NNMF multiplicative update rules

Begin with randomly generated $\mathbf{W}$ and $\mathbf{H}$, with elements between 0 and 1. The elements of $\mathbf{W}$ and $\mathbf{H}$ are then iteratively updated to better approximate Ξ until

convergence criteria are met. Using a multiplicative (rather than additive) update rule preserves non-negativity in matrix weights [15].

1. First, elements of $\mathbf{H}$ are individually scaled according to :

$$H_{qm} \mapsto H_{qm} \frac{(\mathbf{W}^T \boldsymbol{\Xi})_{qm}}{(\mathbf{W}^T \mathbf{W} \mathbf{H})_{qm}}$$

2. Second, the rows of $\mathbf{H}$ are normed, pushing the normalization factor onto the columns of $\mathbf{W}$:

$$H(q, M) \mapsto H(q, M) / |H(q, M)|, \text{ and}$$

$$W(N, q) \mapsto W(N, q) \cdot |H(q, M)|$$

3. Third, elements of $\mathbf{W}$ are individually scaled according to :

$$W_{nq} \mapsto W_{nq} \frac{(\boldsymbol{\Xi} \mathbf{H}^T)_{nq}}{(\mathbf{W} \mathbf{H} \mathbf{H}^T)_{nq}}$$

4. Fourth, calculate error $\eta = |\boldsymbol{\Xi} - \mathbf{W} \mathbf{H}|^2$ and assess for convergence, exiting the update loop when the ratio of change in error between subsequent steps to the error is below a set threshold: $\Delta\eta/\eta < 10^5$.

One could alternately choose NMF with sparseness constraints built into the construction of the factorization algorithm [16], though we defer this to future study. As the process begins with randomly generated initial $\mathbf{W}$ and $\mathbf{H}$, we re-run the NMF algorithm a number of times to identify a reliable minimal error η between $\boldsymbol{\Xi}$ and $\mathbf{W} \mathbf{H}$. Separately, one might perform an algorithmic minimization with convergence rather than brute-force repetition, though we found this impractical.

Dimensionality reduction (Fig. 3F)

The output of NMF can be highly degenerate, and this degeneracy can be quantified by the off-diagonal elements of the matrix $\mathbf{H} \mathbf{H}^T$. We define the sum of the upper-half off-diagonal elements as ζ (Figure 2F). Then, we iteratively reduce the number of inner components Q by 1 and re-perform the NMF until $\zeta < 1$. This iterative reduction performs our clustering, where the non-zero elements in each 1-by- N row of H define how stimulation-pair subgroups are clustered together. The goal of pruning is to minimize number of response shape motifs needed to explain the significant data. We might instead have chosen to constrain the magnitude of any one of the off-diagonal elements of $\mathbf{H} \mathbf{H}^T$ (rather than the sum of all), which would limit the magnitude of any pairwise comparison between BPCs, but not have a global constraint in overall shared structure.

While we have chosen to start with the beginning number of potential clusters (internal dimension Q in NMF) to be the same as the number of stimulation-pairs for this example, it is more expedient by starting with a lower number of clusters (~ 10 appears appropriate). This is sensible, since the number of basic motifs in laminar organization (at each brain site) is constrained by a limited number of cell types. One could alternately scan through many different inner dimensions, and select the best in terms of explained variance (see supp Fig. 1).

Clustering of subgroups (Fig. 3G)

At the conclusion of this process, we perform a winner-take-all operation on the columns of $\mathbf{H}$ such that one stimulation-pair subgroup can only belong to one component. Then, all but the largest elements of each column are set to zero: For each row q of $\mathbf{H}$, a set of stimulation-pair subgroups is assigned to each cluster - If $H_{qn} > \frac{1}{2\sqrt{N}}$, then $n \in q$. This threshold is set because if all subgroups contributed equally to a cluster, then the element weight of each would be $1/\sqrt{N}$. A tolerance factor of $1/2$ was added to allow for some variation.

This “winner-take-all” approach implies a process constructed for canonical responses rather than superposition of contributing motifs. For the present work, this assumption is appropriate, though it may not hold in situations where separate physiologic motifs contribute independently, and each may superimpose in the measured response.

Furthermore, in the present example, there is not a significant penalty for inclusion of a group since the significance of this inclusion is reflected by the scoring of single events, as described below in the subsection “Projecting basis profile curves back into data”.

Identification of basis profile curves using Linear Kernel PCA (Fig. 3G)

Parsing the set of all single-stimulation responses from groups that are clustered together by the NNMF process, we can identify characteristic “Basis Profile Curve” (BPC) shapes, $B_q(t)$, in the following manner:

The subset of single trial clustered responses are first concatenated as $V_{k \in n \in q}(t) \equiv \mathbf{V}_{(q)}$, to reflect single trials belonging to stimulation-pair subgroups ($k \in n$) that in turn belong to cluster q ($n \in q$), with K_q such trials. We would like to identify a representative basis curve, $B_q(t)$, that represents the “principal direction” of $V_{k \in q}(t)$. However, the practical fact that the number of timepoints, T , generally far exceeds the number of trials, K , in these data ($T \gg K_q$) prohibits a standard principal component decomposition (PCA, [17]), which would require $K_q > T^2$ to characterize the T -by- T matrix of interdependencies between timepoints.

We address this issue by inverting the decomposition using the “Linear Kernel PCA” technique [18–20]. This method allows for the interchange of an eigenvalue decomposition of the matrix $\mathbf{V}_{(q)}\mathbf{V}_{(q)}^T$ (T^2 elements) with $\mathbf{V}_{(q)}^T\mathbf{V}_{(q)}$ (K_q^2 elements). Following this approach, we obtain a matrix $\mathbf{F}$, whose columns are the eigenvectors of $\mathbf{V}_{(q)}^T\mathbf{V}_{(q)}$, with associated eigenvalues contained in the diagonal matrix $\boldsymbol{\xi}^2$, satisfying $(\mathbf{V}_{(q)}^T\mathbf{V}_{(q)})\mathbf{F} = \mathbf{F}\boldsymbol{\xi}^2$. We can then solve for the eigenvectors of $\mathbf{V}_{(q)}\mathbf{V}_{(q)}^T$, contained in the columns of $\mathbf{X}$:

$$\mathbf{X}\boldsymbol{\xi} = \mathbf{V}_{(q)}\mathbf{F}^T$$

We keep the first column of $\mathbf{X}$ as our basis curve $B_q(t)$.

Note that, if we were to take the simple average of all candidate stimulation-pair subgroups instead of the 1st principal component for this canonical shape, it could significantly dilute the BPC form with noise. As illustrated in cluster 3 of our example, several insignificant stimulation-pair subgroups are clustered along with a strong and significant response subgroup. As we have constructed it, uncorrelated noise doesn’t contribute significantly to variance employed by the kernel PCA approach.

Fig 4. Projection of Basis Profile Curves (BPCs). **A:** The contribution of each BPC B_q to a single trial from its cluster can be quantified according to a scalar multiplier $\alpha_k^{(q)}$, and residual noise ε_k (trial 238 illustrated). **B:** The 3 BPCs for our example case. **C:** The spatial representation of BPCs, color-coded, with diameter and color intensity indicating magnitude (group-averaged signal-to-noise ratio). Each BPC distribution is individually scaled to maximum. Gray indicates sites discarded by thresholding (figure 3G).

Projecting basis profile curves back into data

Utilizing the formalism from functional data analysis [21,22], we can represent each individual trial as a projection of a basis profile curve $B_q(t)$, scaled by a scalar $\alpha_k^{(q)}$, with residual error $\varepsilon_k(t)$:

$$V_k(t) = \alpha_k^{(q)} B_q(t) + \varepsilon_k(t)$$

We expect that $E(\varepsilon) = 0$ and $E(\varepsilon_k^2) \sim E(\varepsilon_l^2)$, for all k and l . This allows us to estimate the projection of $B_q(t)$ into each individual trial:

$$\alpha_k^{(q)} = \sum_t B_q(t) V_k(t) \quad \text{and} \quad \sum_t B_q(t) \varepsilon_k(t) = 0$$

Since $E(\varepsilon) = 0$, and $\sum_t B_q(t) \alpha_k^{(q)} B_q(t) = \alpha_k^{(q)} \sum_t B_q(t) B_q(t) = \alpha_k^{(q)}$. Having $\alpha_k^{(q)}$ thus determined, we can quantify the residual noise after regressing out the shape of $B_q(t)$:

$$\varepsilon_k(t) = V_k(t) - \alpha_k^{(q)} B_q(t)$$

With this description $V_k(t) = \alpha_k^{(q)} B_q(t) + \varepsilon_k(t)$, several useful quantities for each trial V_k can be described (omitting q for notational simplicity): a “projection weight” α_k ; a scalar “noise” summary term $\sqrt{\varepsilon_k^T \varepsilon_k}$; a “signal-to-noise” $\alpha_k / \sqrt{\varepsilon_k^T \varepsilon_k}$; the “explained variance” by application of the BPC is $1 - \frac{\varepsilon_k^T \varepsilon_k}{V_k^T V_k}$. For stimulation-pair subgroup n , we can estimate the presence of residual structure after application of B_q , by the remaining pair-wise correlation in noise terms: $\langle \varepsilon_k^T \varepsilon_l \rangle_{k,l \in n; k \neq l}$.

Code and data availability

Code written in MATLAB to reproduce all of the steps and illustrations contained in this manuscript is freely available along with the sample dataset at <https://purl.stanford.edu/rc201dv0636>. The file “kjm_bpcmethod_readme.pdf” describes the code and dataset, along with instructions for how to perform the analyses.

Results

Cortico-cortical and subcortical inputs converge in each brain region, and disentangling of this convergence will shed light on the networks that interconnect the brain. In order to explore this idea, we stimulated pairs of intracranial electrodes implanted across many different brain areas and measured voltage responses in a single site. We expect that biologically different types of inputs will produce characteristic shapes in the voltage timecourse, and that these might be clustered into distinct groups based upon where was stimulated.

Our framework, aiming to better understand brain connectivity, is grounded in a convergent paradigm, examining a set of temporal voltage responses to stimulation, all measured from the same site (figures 1 and 2). Each response event is labeled by the site of stimulation. Then, a novel algorithm is applied within this framework to identify canonical temporal response motifs, which we call “basis profile curves” (BPCs). Each BPC clusters subgroups of stimulation-pairs together into a larger group that induce a similar response profile, and are likely engaging same microcircuitry in their connectivity from the stimulated brain site to the measured brain site.

Multiple different stimulation evoked voltage response motifs are measured from one brain site

The BPC approach allowed us to extract a concise set of Basis Profile Curves that describes the multitude of responses observed a single site. As illustrated in figure 3, the algorithmic approach begins by obtaining a set of all projection magnitudes (correlations in response timecourse) between pairs of single trials within their own stimulation-pair subgroups and across different stimulation-pair subgroups. A matrix characterizing the significance of each subgroup-subgroup projection magnitudes is generated from the t-values of these sets, before setting negative values to 0 and scaling to 1. Then non-negative matrix factorization (NNMF) is repeatedly performed to decompose this matrix into a pair of other matrices, one of which characterizes correlated features within the matrix, and the other of which characterizes the weight of each feature (and is normalized). NNMF is performed many times, iteratively reducing the inner dimension of factorization until a cross-correlation threshold between features is surpassed. The iterative reduction clusters stimulation-pair subgroups by their relative weights within the weight matrix (Supplemental figure 1 shows that the technique identifies the optimal number of clusters from the data for portion of variance explained in our example case). For each cluster, linear kernel PCA is applied to the concatenated larger group of single responses from all included stimulation-pair subgroups used to extract the temporal shape of the BPC. In this way, a unique BPC is associated with each cluster.

The algorithm, by construction, produces BPCs that have a set of desired properties (figure 4): Each BPC has a characteristic shape in time of “canonical” responses with simple visual formulation where connectivity between areas is paired with a temporal BPC motif that can be a window into the nature of the interaction. Although the literature has predominantly supported a canonical form for the timecourse of CCEPs

(e.g. polyphasic with 2 characteristic “N1/N2” negative deflections), this was not the case in our measurements. We found instead that the measured cortico-cortical evoked potentials could be described by three basis profile curves, which are unique in shape (Figure 4B), only one of which (B_3) is consistent with the reported N1/N2 form.

Stimulation of adjacent anatomical sites produces similar voltage responses that are clustered together

The initial step in examining these stimulation data from the *convergent paradigm* is to plot the evoked response at the measurement site onto the brain surface at the site of stimulation (figure 2). Visual inspection of these plots suggests that the response shapes cluster along anatomical boundaries.

Once BPCs have been identified, single stimulation events can be characterized with the architecture from functional data analysis [21–23], which enables quantification of BPC projection weight (signal) and residual noise. These quantifications may be projected back to brain anatomy in an intuitive way to visualize what sites are meaningfully connected to the measured-from site. In our example case (figure 4), the back projections reveal a relatively sparse network interacting with the PHG originating from the superior temporal gyrus (STG, green in figure 4), the posterior portion of the inferior temporal gyrus (orange), and the fusiform gyrus (yellow). However, rather than identifying this clustering by visual inspection, the clustering is quantitative with summary weight and clear statistical description. The result is very similar in appearance whether quantified by projection weight, signal-to-noise, or explained variance (supplemental figure 2).

If a different measurement site is selected in the example patient data, at the temporal pole (TP), a very similar clustered region from the STG emerges (Supplemental figure 3). Interestingly, the shape of the associated BPC is very different. This is precisely the dilemma envisioned by the *divergent paradigm* shown in figure 1F. The interpretation of the difference in these shapes is ambiguous. One cannot tease out whether the difference in BPC shape implies a different kind of connectivity, or simply a reflection of the different microcircuitry of the TP and STG in response to a similar type of input.

Discussion

This work begins with the general question of how electrical stimulation paired with voltage measurement can be used to understand spatial and temporal structure in brain networks. Our framework begins with a *convergent approach*, where one measurement site is selected, and responses to sets of repeated stimulations in many other sites are quantified (figure 1). Each electrical stimulation pulse is brief ($<1\text{ms}$) and the time between pulses is long ($>3\text{s}$), allowing for transient voltage changes to return to baseline. A relatively long time between pulses means that mono-synaptic and polysynaptic effects can contribute to the temporal structure of the responses, and be captured during data analyses. This experimental paradigm where enough time has passed between consecutive stimulations for transient effects to die out is typically called “single pulse electrical stimulation” [24], a subset of “cortico-cortical evoked potential” (CCEP) measurement. Our initial observations of these types of data within the convergent paradigm did not suggest a universal form to evoked voltage changes from the same brain site [9], so the framework we have developed does not assume what shape the measured responses should have. We constructed a data-driven technique to extract motifs from all responses at a single seed site that we call “basis profile curves”

(BPCs), beginning with a natural set of subgroups defined by repeated stimulations at the same brain site.

The convergent approach to connectivity

A common approach to explore network structure in CCEP research is to examine every possible interaction within an array of electrodes. In this *all-to-all* approach, the large number (N^2) of all possible interactions between the N intracranial electrodes are studied (figure 1), and each response profile has hundreds-to-thousands of timepoints. The several seconds between each electrical stimulation and the large set of electrode pairs to stimulate between severely limits the number of repeated stimulations that may be performed in the clinical setting. From a practical perspective, this means that if one does not approach the study of CCEPs with clear constraints, then the problem becomes too high-dimensional to handle, given the limited amount of data available. Many existing studies have attempted to address this by presuming a specific temporal structure in the evoked response, such as the voltage at a pre-specified timepoint post-stimulation [25]. However, as illustrated in figures 2&S3 and reported by Kundu and colleagues [26], there is no single canonical CCEP response shape or feature, even when measuring from a single electrode.

Our approach to this high-dimensional data dilemma is to first constrain our study to measurements from a single electrode at a time, leveraging the convergent paradigm. The BPC framework then enables one to naively extract a family of response shapes specific to that site. Although there is no prior assumption about the forms the BPCs should have, they are constructed with the constraint that they should be reproducible within subgroups of repeated stimulations at the same brain site.

Aside from computational convenience, the convergent paradigm is also useful because we can reduce the plethora of potential interactions to a smaller, more tractable, set that may be linked to physiological interpretation (i.e. a few different appearing motifs in interaction between the stimulated and measured sites). For example, inputs to superficial or deeper layers should produce different response motifs in the electric potential measured at the brain surface, and be isolated as distinct BPCs (Fig. 1).

Although a single site is selected for the convergent approach, one may iteratively uncover the larger connectivity space: Projections between brain regions can be identified, beginning with a known seed site and then using the proposed algorithm to find a strongly projecting site within a BPC cluster which, in turn, becomes the new seed. This would allow one to trace projections in reverse to explore and model a network of previously unknown interactions.

Properties of the BPC algorithm

The BPC algorithmic framework is constructed to satisfy a pre-defined set of desired properties, each of which is associated with distinct computations.

Simple assignment of each subgroup to a single BPC: We wanted to meaningfully decompose brain responses to stimulation with a process that can be naturally mapped back on to the underlying anatomy. This means that stimulated electrode pairs that evoke the same response motif should be grouped together in a way that can be plainly viewed on a brain rendering. In our BPC responses, this means that positive voltage deflections of a particular shape will not be clustered with negative deflections of a similar shape (and vice-versa). Ensuring this is very important when interpreting the signals physiologically. For example, sign flips could reflect inputs at superficial vs deep lamina or at different classes of synapses, and so should be clustered independently [27]. Mathematically, this decomposition can be formulated as a clustering problem with preservation of sub-group structure and inclusion of a non-negativity constraint.

Commonly applied techniques such as independent/principal component analyses (ICA/PCA) may not be productive in this setting (even with non-negativity constraints [28]). They assume linear superposition of motifs with arbitrary sign and weight rather than the identification of distinct, unique motifs. In addition, ICA/PCA have different loss functions than the BPC framework, emphasizing different decomposition targets and are unable to reveal the subgroup structure of stimulation-pair sites necessary for our analysis. Specifically, they are specialized to generate components with minimum independency, rather than to cluster motifs by similarity (which is our goal). Canonical correlation analysis (CCA) or variants thereof [29,30], which might allow for some labeling subgroup structure, do not easily allow one to incorporate necessary constraints and are of limited help for our purpose.

Allowance of limited overlap in BPC shape: Although two different responses may reflect different types of inputs, a limited sub-interval of time may have transient similarity. Therefore, it can be useful to allow for some limited shared structure in the timecourses of different BPCs. In our example, this is illustrated in the negative deflections seen in the initial ~ 500 ms of B_2 and B_3 , Figure 4B. We implement this allowance of limited overlap by using a winner-take-all approach rather than enforcing orthogonality in the rows of $\mathbf{H}$ (e.g. $\mathbf{H}\mathbf{H}^T = \mathbf{I}$). The amount of overlap allowed can be adapted by setting the maximum value of individual off-diagonal elements of $\mathbf{H}\mathbf{H}^T$, or their sum (ζ).

Disregarding of meaningless subgroups: Another physiologically meaningful constraint for the BPC framework is that existing sub-group structure, where responses to stimulation come from the same pair of electrodes should either be reliable on a trial-by-trial basis or not contribute to the clustering. In the BPC approach, a correlation-significance matrix quantifying similarity in single-trial pairwise correlations organized between sub-groups (stimulation-pair sites) can be obtained. Unlike many common decomposition techniques which construct covariance or correlation matrices (ICA/PCA/CCA), the significance matrix generated by this process can have very small diagonal elements, and, in some cases, off-diagonal elements are larger than corresponding diagonal elements from the same row (Figure 3D)². These small diagonal elements, coupled with the thresholding of the elements of factor matrix H excludes stimulation-pair subgroups that do not produce reliable responses in measurement from the identification of BPCs. Once BPCs have been extracted, more explicit quantification of significance for each stimulation-pair subgroup may be tested for by testing the magnitudes $\alpha_{k,k \in n}$ (for subgroup n) versus zero.

Simple metric to describe single trials: At the beginning of our data exploration, each single trial k is described by its timecourse $V_k(t)$, and the stimulation-pair, $k \in n$, that produced it (its subgroup). At the conclusion of the BPC extraction, each trial is assigned to a single BPC (q) with a scalar projection weight ($\alpha_k^{(q)}$), and a residual noise timecourse ($\varepsilon_k(t)$). The parameterization of single-trial stimulation responses takes the form $V_k(t) = \alpha_k^{(q)} B_q(t) + \varepsilon_k(t)$, which is a formalism borrowed from the field of functional data analysis [23]. One might use the tools developed in that discipline for a larger exploration of these data when generalizing across different patients, tasks, stimulation paradigms, and recording settings. This formalism allows for straightforward characterization of signal and uncorrelated noise in each stimulation single-trial, and comparisons of these across stimulation-pair subgroups and also across BPCs. As noted in the method subsection “Projecting basis profile curves back into data”, one may quantify a scalar values of signal, noise, and variance explained by the form of the BPC for each single-trial response. These quantities may be useful to plot versus one another (supplemental figure 2, and allows one to easily scale distributions independently or globally (as seen in supplemental figure 3D&E). The advantage of this

²See description in methods subsection “Single-trial cross-projections & Significance matrix”

formalism moves well beyond that of notational convenience. Namely, the magnitudes of response to stimulation in different brain areas can be directly compared versus one another to quantify how much the local microcircuitry is being influenced, even though the shape of the response timeseries may be very different. This quantification is typically difficult for timeseries analysis in these types of data, where template-projection or other similar approaches are otherwise necessary [31].

Response motifs and underlying physiology

At the macroscale that is measured by ECoG or SEEG electrode contact size, an entire population of about half of a million neurons is being averaged over [11]. As such, we expect that only a few BPCs will explain variance from a measurement site. Each region is defined by a unique laminar architecture (defined as a unique Brodmann area, [32]), so we expect that the pro-dromic (reflecting inputs) or the anti-dromic (reflecting outputs) will be constrained to a few unique motifs, and reflected by relatively few BPCs. Therefore our BPC technique allows for visual inspection of temporal motifs that are well-defined in measurement space (see figure 4) but raise important questions about the neurophysiology they reflect. Potentially typical shapes may reveal connection to microcircuitry in intuitive way, where surface positive deflections may be tied to deep positive ion influx or superficial ion efflux from specific synapse types in the large pyramidal cells beneath [27]. Alternately, they may reveal projections from the stimulation site directly to different classes of cells within the laminar architecture of cortex (e.g. interneurons vs pyramidal neurons). Some BPC morphologies might reveal different motifs in connectivity at the macroscale. For example, one might speculate about different types of connectivity within the temporal lobe that may be inferred from examination of figures 4 and S3. Direct connections from intracortical axonal projects within gray matter (via lateral projections) might be differentiated from those relayed subcortically through white matter tracts. Note that indirect projections relayed through a third cortical site or a set of subcortical nuclei might each be revealed in characteristic BPC shapes, and this will be explored in future studies.

Potential future applications

We would like to emphasize that initial explorations with our BPC technique should be framed in a well-controlled setting with connection to simple functional studies from primary motor, auditory, or visual regions to look for commonality in BPC shapes. In this manner, existing understanding from other types of measurement can serve to validate BPC interpretation, and examine whether they are conserved at homologous brain sites across different patients.

Subsequent studies might vary experimental conditions with a task, medication, or property of the stimulation (current magnitude and temporal profile, or timing between pulses), and quantify signal, noise, and residual structure in the responses. For example, one could examine the effect of simply changing the amplitude of stimulation, where it has been demonstrated that different stimulation magnitudes can elicit very different morphologies [26] and examine distribution of BPCs as a function of that.

Inspecting recordings from a single site, our novel framework could be applied to behavioral, rather than electrical stimulation. For example, one could study higher order visual areas, like the fusiform gyrus, and examine responses to presented images of different semantic types (pictures of faces, houses, tools, etc., etc.) [33,34]. In such a study, one would replace stimulation pair groups with semantic stimulus groups (pictures of faces, houses, tools, etc.) to see how different semantic groups cluster together in production of the fusiform electrophysiological response.

This “winner-take-all” approach (figure 3G) implies a process constructed for canonical responses rather than superposition of contributing motifs. In other words, this means that each stimulation response is assigned to a single BPC, rather than as a combination of multiple BPCs. For the present work, this assumption is appropriate, though it may not hold in situations where separate physiologies contribute independently and may superimpose in the measured response. There may be later approaches where initial voltage changes (1st order projections) and later (2nd & higher order projections) are assessed independently and allowed to superimpose, particularly when attempting to detect/discover potential higher-order responses. In future work, this could be performed by picking multiple sets of t_1 and t_2 , and re-performing analyses for isolated time windows.

One result of this functional data analysis formalism is that candidate structures not correlated to initial subgroup labeling could be tested for by looking at correlations of $\langle \varepsilon_k^T \varepsilon_l \rangle_{k \neq l}$, where k and l would be of a candidate different type of subgroup (for example, the third stimulation pulse from each set of stimulations). In this case, after a full exploration of that residual structure, the error term ε for a given trial would split out into a second order set of BPCs, C_p , with coefficients $\beta_k^{(p)}$, where $\varepsilon \rightarrow \beta^{(p)} C_p + \varepsilon'$, and $V_k(t) = \alpha_k^{(q)} B_q(t) + \beta_k^{(p)} C_p(t) + \varepsilon'_k(t)$.

Conclusion

We have detailed a new data-driven framework for uncovering motifs in epoch-based timeseries data belonging to labeled subgroups. These motifs are called “basis profile curves” (BPCs), and they determine characteristic spatiotemporal structure. Each timeseries epoch is assigned to a unique BPC, with motif projection strength and residual noise simply parameterized. This framework is applied to understand the effect of electrical stimulation in the human brain using arrays of implanted electrodes, where the labeled subgroups are repeated stimulations at the same site. We introduce a set of paradigms for interpreting these measurements, of which the *convergent* one allows application of the BPC framework. In our illustrative example of measurements from the surface of the parahippocampal gyrus, we find that identified BPCs clearly uncover several connected regions and allow them to be viewed and interpreted intuitively.

Acknowledgments

We are grateful for the participation of the patient in the study, and for the assistance of the staff in Saint Marys Hospital at Mayo Clinic, Rochester, MN. Greg Worrell, Tom Richner, and Pradeep Udayavar Shenoy generously provided helpful discussion and commentary during the drafting of the manuscript. KJM was supported by the Van Wagenen Fellowship, the Brain Research Foundation with a Fay/Frank Seed Grant, and the Brain & Behavior Research Foundation with a NARSAD Young Investigator Grant. This work was also supported by NIH-NCATS CTSA KL2 TR002379 (KJM). DH was supported by the NIH-NIMH CRCNS R01MH122258-01. Manuscript contents are solely the responsibility of the authors and do not necessarily represent the official views of the NIH. KRM was supported in part by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea Government (No. 2017-0-00451, Development of BCI based Brain and Cognitive Computing Technology for Recognizing User’s Intentions using Deep Learning) and (No. 2019-0-00079, Artificial Intelligence Graduate School Program, Korea University), and by the German Ministry for Education and Research (BMBF) under Grants 01IS14013A-E, 01GQ1115, 01GQ0850, 01IS18025A, 031L0207D and 01IS18037A; the German Research Foundation (DFG) under Grant Math+, EXC 2046/1, Project ID 390685689.

References

1. Adrian ED. The spread of activity in the cerebral cortex. *The Journal of physiology*. 1936;88(2):127.
2. Purpura DP, Pool J, Frumin M, Housepian E, et al. Observations on evoked dendritic potentials of human cortex. *Electroencephalography and clinical neurophysiology*. 1957;9(3):453–459.
3. Ilmoniemi RJ, Virtanen J, Ruohonen J, Karhu J, Aronen HJ, Näätänen R, et al. Neuronal responses to magnetic stimulation reveal cortical reactivity and connectivity. *Neuroreport*. 1997;8(16):3537–3540.
4. Matsumoto R, Nair DR, LaPresto E, Najm I, Bingaman W, Shibasaki H, et al. Functional connectivity in the human language system: a cortico-cortical evoked potential study. *Brain*. 2004;127(10):2316–2330.
5. Valentin A, Anderson M, Alarcon G, Seoane JG, Selway R, Binnie C, et al. Responses to single pulse electrical stimulation identify epileptogenesis in the human brain in vivo. *Brain*. 2002;125(8):1709–1718.
6. Enatsu R, Gonzalez-Martinez J, Bulacio J, Kubota Y, Mosher J, Burgess RC, et al. Connections of the limbic network: a corticocortical evoked potentials study. *Cortex*. 2015;62:20–33.
7. Crowther LJ, Brunner P, Kapeller C, Guger C, Kamada K, Bunch ME, et al. A quantitative method for evaluating cortical responses to electrical stimulation. *Journal of neuroscience methods*. 2019;311:67–75.
8. Keller CJ, Honey CJ, Entz L, Bickel S, Groppe DM, Toth E, et al. Corticocortical evoked potentials reveal projectors and integrators in human brain networks. *Journal of Neuroscience*. 2014;34(27):9152–9163.
9. Miller KJ, Huiskamp G, van Blooijis D, Hermes D, Gebbink TA, Ferrier CH, et al. An observation of anatomical clustering in inputs to primary motor cortex in cortico-cortical brain surface evoked potentials. In: 2019 7th International Winter Conference on Brain-Computer Interface (BCI). IEEE; 2019. p. 1–2.
10. Hermes D, Miller KJ, Noordmans HJ, Vansteensel MJ, Ramsey NF. Automated electrocorticographic electrode localization on individually rendered brain surfaces. *Journal of neuroscience methods*. 2010;185(2):293–298.
11. Miller KJ. A library of human electrocorticographic data and analyses. *Nature Human Behaviour*. 2019;3(11):1225–1235.
12. Hermes D. The CTMR Package; 2010. Available from: https://github.com/dorahermes/Paper_Hermes_2010_JNeuroMeth.
13. Gosset WS. The probable error of a mean. *Biometrika*. 1908; p. 1–25.
14. Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. *Nature*. 1999;401(6755):788–791.
15. Lee DD, Seung HS. Algorithms for non-negative matrix factorization. In: *Advances in neural information processing systems*; 2001. p. 556–562.
16. Hoyer PO. Non-negative matrix factorization with sparseness constraints. *Journal of machine learning research*. 2004;5(Nov):1457–1469.

17. Bishop CM. Pattern recognition and machine learning. springer; 2006.
18. Schölkopf B, Mika S, Smola A, Rätsch G, Müller KR. Kernel PCA pattern reconstruction via approximate pre-images. In: International Conference on Artificial Neural Networks. Springer; 1998. p. 147–152.
19. Schölkopf B, Smola A, Müller KR. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*. 1998;10(5):1299–1319.
20. Schölkopf B, Mika S, Burges CJ, Knirsch P, Müller KR, Rätsch G, et al. Input space versus feature space in kernel-based methods. *IEEE transactions on neural networks*. 1999;10(5):1000–1017.
21. Chiou JM, et al. Dynamical functional prediction and classification, with application to traffic flow prediction. *The Annals of Applied Statistics*. 2012;6(4):1588–1614.
22. Li PL, Chiou JM. Identifying cluster number for subspace projected functional data clustering. *Computational Statistics & Data Analysis*. 2011;55(6):2090–2103.
23. Ramsay JO, Silverman BW. Applied functional data analysis: methods and case studies. Springer; 2007.
24. Mouthaan BE, van't Klooster MA, Keizer D, Hebbink GJ, Leijten FS, Ferrier CH, et al. Single Pulse Electrical Stimulation to identify epileptogenic cortex: Clinical information obtained from early evoked responses. *Clinical neurophysiology*. 2016;127(2):1088–1098.
25. Keller CJ, Honey CJ, Mégevand P, Entz L, Ulbert I, Mehta AD. Mapping human brain networks with cortico-cortical evoked potentials. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2014;369(1653):20130528.
26. Kundu B, Davis TS, Philip B, Smith EH, Arain A, Peters A, et al. A systematic exploration of parameters affecting evoked intracranial potentials in patients with epilepsy. *Brain Stimulation*. 2020;.
27. Mitzdorf U. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiological reviews*. 1985;65(1):37–100.
28. Sigg CD, Buhmann JM. Expectation-maximization for sparse and non-negative PCA. In: Proceedings of the 25th international conference on Machine learning; 2008. p. 960–967.
29. Bießmann F, Meinecke FC, Gretton A, Rauch A, Rainer G, Logothetis NK, et al. Temporal kernel CCA and its application in multimodal neuronal data analysis. *Machine Learning*. 2010;79(1-2):5–27.
30. Dähne S, Biessmann F, Samek W, Haufe S, Goltz D, Gundlach C, et al. Multivariate machine learning methods for fusing multimodal functional neuroimaging data. *Proceedings of the IEEE*. 2015;103(9):1507–1530.
31. Miller KJ, Schalk G, Hermes D, Ojemann JG, Rao RP. Spontaneous decoding of the timing and content of human object perception from cortical surface recordings reveals complementary information in the event-related potential and broadband spectral change. *PLoS computational biology*. 2016;12(1):e1004660.

32. Brodmann K. Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Barth; 1909.
33. Miller KJ, Hermes D, Pestilli F, Wig GS, Ojemann JG. Face percept formation in human ventral temporal cortex. *Journal of Neurophysiology*. 2017;118(5):2614–2627.
34. Miller KJ, Hermes D, Witthoft N, Rao RP, Ojemann JG. The physiology of perception in human temporal lobe is specialized for contextual novelty. *Journal of neurophysiology*. 2015;114(1):256–263.

Supporting information

S1 Fig.

Supplemental Fig 1. Explained variance as a function of number of inner dimensions, Q . We can quantify the total variance explained by the fitting of resulting BPCs, $B_q(t)$ across all included stimulation-pair trials: $\langle V_k^T V_k - \varepsilon_k^T \varepsilon_k \rangle_{k \in n \in q}$. When we normalize this by the total variance, $\langle V_k^T V_k \rangle_{k \in \text{all } n}$, we have the portion of total variance explained (blue circles). However, the goal of the decomposition is to identify motifs, explained by BPCs, from the full set of stimulations. For example, if one were to perform stimulations at a site that have no effect on the measurement site, it should not undermine our confidence in the decomposition. Therefore, a more appropriate normalization is to instead divide the explained variance by the variance of the trials included in the clustering, $\langle V_k^T V_k \rangle_{k \in n \in q}$. For our example case, this immediately validates the dimensionality $Q = 3$, selected by the algorithm (Fig. 3).

S2 Fig.

Supplemental Fig 2. Alternate metrics for projection weights of BPC curves. As in figure 4, but using alternate scoring metrics for each stimulation-pair subgroup, shown in inset squares overlying each cortical rendering.

S3 Fig.

Supplemental Fig 3. Illustration from a different site. Stimulation responses from a site in the temporal pole are shown. **A:** Responses from each stimulation pulse are aligned into a matrix $V_k(t)$. **B:** Averaged responses $G_n(t)$ are shown at the site of each stimulation pair that produced them. **C:** BPCs produced by the algorithm. **D:** Weights (group-averaged signal-to-noise ratio) associated with each BPC (color-coded), and non-included sites (gray). **E:** Spatial representation of BPCs, color-coded, with diameter and color intensity indicating magnitude. All values are scaled to the global maximum

across all BPCs. **F:** As in (E), but with each BPC distribution individually scaled to its own maximum. Note the similarity in the spatial distribution of the stimulation-site cluster labeled in green to figure 3, but the completely different shape of the BPC (likely reflecting the different laminar architecture of the two recipient measurement sites).