

Automatic methods for cortex-wide layer identification of electrophysiological signals reveals a cortical motif for the expression of neuronal rhythms

André M Bastos (abastos@mit.edu)

Omar Costilla-Reyes (costilla@mit.edu)

Earl K Miller (ekmiller@mit.edu)

Department for Brain and Cognitive Sciences
Massachusetts Institute of Technology
Building 46-6257
46 Vassar St.
Cambridge, MA 02139

Abstract:

Cerebral cortex is composed of 6 anatomical layers. How these layers contribute to computations that give rise to cognition remains a challenge in neuroscience. Part of this challenge is to reliably identify laminar markers from in-vivo neurophysiological data. Classic methods for laminar identification are based on assumptions which are often violated and require expert users to identify the pattern, potentially introducing bias. We recorded local field potentials (LFP) with probes containing 16 or 32 electrodes that span all cortical layers in frontal, parietal, and visual cortex in monkeys. We describe two novel methods to identify layers in a fully automatic and quantitative way. The first method represents relative power across electrodes from as a 2-dimensional image, and maximizes image similarity across probes. The second method leverages ensemble machine learning to maximize classification accuracy of LFP data to a laminar label. Both methods detect consistent patterns, and the image similarity approach reveals a cortex-wide motif of laminar expression for delta/theta, alpha/beta and gamma rhythms. Delta/theta (1-4 Hz) and gamma (50-150 Hz) power peak in superficial layers 2/3, and alpha/beta (10-30 Hz) power peaks in deep layers 5/6.

Keywords: cerebral cortex; cortical layers; machine learning; gamma rhythms; alpha/beta rhythms

Introduction

Cerebral cortex is composed of a sheet of neurons arranged in 6 layers. Nissl staining and tract tracing studies have long established anatomical differences between layers. Superficial layers 2/3 contain excitatory pyramidal neurons providing the majority of the feedforward output (Felleman and van Essen, 1991). Deep layers 5/6 contain the excitatory pyramidal cells that send feedback output and that project to subcortical sources such as thalamus. Layer 4 is the largest target of feedforward inputs from the thalamus and cortical areas that are hierarchically lower. Feedback inputs

tend to avoid targeting layer 4 and project most prominently in layer 1, which itself has few neurons but contains the apical dendrites of pyramidal neurons in layers 2-6 (Felleman and van Essen, 1991). This basic anatomical wiring diagram has led to prominent hypotheses about what computational functions might map on to these layers. One popular class of models in cognitive neuroscience, predictive coding, has hypothesized that superficial vs. deep layers contain different computational units, with superficial layers containing prediction error units and deep layers containing prediction units (Rao and Ballard, 2001, Friston, 2010, Bastos et al., 2012). However, despite increasing interest in this hypothesis and others about the functions of cortical layers, there is currently a dearth of electrophysiological data on this topic, possibly due to the difficulty in identifying layers based on in-vivo electrophysiology.

Laminar identification based on Current Source Density (CSD) Mapping

Current Source Density is arguably the state of the art method for identifying cortical layers based on neurophysiology. This is based on the principle that the layer showing the strongest net excitation will depolarize first in response to an excitatory stimulus. This generates intracellular current flow, which is measurable with extracellular electrodes as a current sink. These currents flows are defined according to the following equation for channel i (equation 1), assuming electrodes are arranged in a linear fashion and that the inter-contact spacing is constant:

$$CSD(i) = \sigma \frac{V(i-1) + 2V(i) - V(i+1)}{s^2}$$



where s is the inter-contact spacing and σ is the tissue conductivity. Negative values of the CSD correspond to current sinks. This has been used in primary and higher-order visual cortex (Mitzdorf and Singer, 1981; Schroeder et al., 1999).

Laminar identification based on transition between Cerebral Spinal Fluid and Gray Matter

Laminar identification can also be performed by placing some electrodes on the probe just outside the cortex in the Cerebral Spinal Fluid (CSF) and have the remaining electrodes penetrate into the Gray Matter (GM). Three markers define the transition: first, the presence of a heartbeat artifact is frequently observed on contacts just outside the gray matter, as a result of small draining veins located at the pial surface. Second, the voltage fluctuations in the LFP increase in GM. Third, neuronal spiking is observable in the GM but not in the CSF. Relative to this transition, electrodes on the probe are then sorted into specific layers based on known anatomical distances between the pial surface and each layer, or this method is used to map electrodes into a more simplified distinction of superficial vs. deep (Johnston et al., 2019).

Limitations of the current state of the art

One limitation with CSD analysis for layer identification is the underlying assumption that layer 4 will be the first to depolarize in response to a well-defined input. This assumption holds true predominately for lower-order sensory areas, but may not generalize to the entire cortex, where stimuli that produce “strong excitation” are unknown. Another limitation is the low signal to noise (SNR) ratio. CSD analysis is inherently noisy because it is based on the second spatial derivative, which is a difference of differences. SNR variation across electrodes, which is typical in electrophysiology, will corrupt the measure. A limitation of the CSF/GM transition method is that it can only be used in the overlying superficial cortex. In higher-order primates, the cortex is folded, creating a complex pattern of gyri and sulci, and deeper structures cannot be mapped with this technique. Both techniques require the laminar probe to be perpendicular to the cortical sheet and therefore cannot be used in cortical sulci. Finally, both techniques are based on the judgments of an expert user, which potentially introduces bias.

Methods

Laminar Similarity Maximization (LSM) algorithm

To overcome the current limitations of laminar identification, we developed a laminar alignment

method that fulfilled the following three criteria: first, it could work in different cortical areas (V4 in visual/temporal cortex, area 7A in posterior parietal cortex, and areas in the prefrontal cortex, PFC) and was not dependent on assumptions from primary sensory cortex. Second, it was robust to changes in SNR ratio over electrodes on the probe. Third, it was fully automated and quantitative. To begin, we calculated each probe’s relative power over laminar depth (the electrode dimension) as a function of frequency from 1 to 250 Hz (equation 2):

$$\text{Relative Power}(c, f) = \frac{\text{Power}(c, f)}{\max[\text{Power}(1:n_{chan}, f)]}$$

This produced a 2-dimensional image, with electrodes on the y-axis and frequency on the x-axis. Each pixel has an intensity ranging between 0 and 1. Values of 1 indicate the electrode with highest power at that frequency. Hereafter we will refer to this relative power profile as the probe image. We tested whether each probe’s image was reproducible across all other probes in our sample (N=123) both within and across areas by calculating the Structural Similarity Index (SSIM) between images. SSIM is an image-processing algorithm that compares the consistency of spatial information between two images in terms of local luminance, contrast, and overall structure (Zhou et al., 2004). This method was preferable to other similarity metrics like the mean squared error because rather than compare the images pixel by pixel, the SSIM metric uses the overall structure (global patterns) of the image in addition to local information.

Our algorithm, termed the Laminar Similarity Maximization (LSM) seeks the best possible match between probe images to a reference image. It proceeds as follows. Step 1: Beginning with a randomly drawn probe as a starting point, we defined this first probe image as the reference image. Step 2: Randomly draw a second image from the sample and calculate SSIM relative to the first. Spatially shift the image profile of the second image relative to the first by all possible spatial lags in 100um steps that still preserves at least half the image for comparison. Step 3: Take the spatial lag that maximizes SSIM. Step 4: average the two images to form a new reference image. Step 5: Repeat steps 1-4 for the remaining probes.

To demonstrate that this alignment procedure did not depend on a pre-defined starting point, we repeated the same procedure, but using 1,000 random starting points, generating 1,000 profiles (one per starting point). We then tested whether every average profile “looked like” the others by repeating the LSM algorithm, only this time on each of the average profiles. The results were consistent: the average profile of each starting point had a high image similarity to all others.

We then repeated LSM on the averaged profiles to produce an “average of averages” as the final reference image. It is shown in Figure 1 (left panel). All probe images were then aligned to this reference image. This final step is deterministic and not influenced by which probe was taken as the starting point.

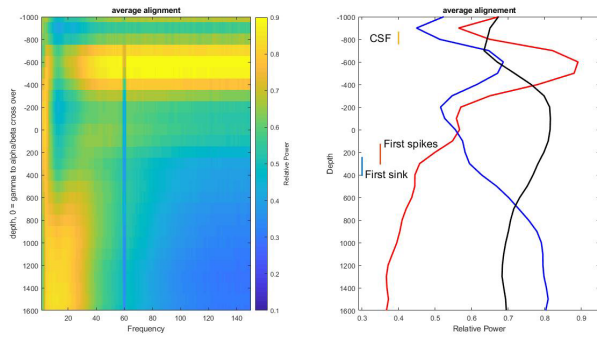


Figure 1. (left) Average aligned profile image representing relative power with color intensity. (right) Average relative power in three distinct frequency ranges: theta (1-4Hz), alpha/beta (10-30 Hz), and gamma (50-150 Hz). Depth zero corresponds to the cross over position between relative power at gamma and alpha/beta. The relationship between this point and the contact with transition between CSF and brain (+/- 1SEM), contact with first significant sink in CSD profile (+/- 1SEM), and the contact with first significant spiking (+/- 1SEM) are depicted.

Next, we assessed the statistical significance of each probe’s similarity to the reference image. We created a randomization procedure that would a) control for the multiple comparisons problem, since our procedure takes the maximum similarity at multiple spatial shifts, and b), control for the possibility that our procedure was maximizing the similarity between random images. To do this, we constructed a randomization distribution under the null hypothesis that relative power did not depend on frequency or layer. For each randomization, the reference image was randomized in both the laminar and frequency dimensions. We then repeated the LSM algorithm, taking the maximum image similarity between the random profile and the each probe image. This randomization was performed 1,000 times. Only 2% (3 out of 125) of probes had less similarity than this reference distribution at $p=0.05$ (Figure 2).

We next compared the distribution of image similarities from each area separately with the randomization distribution. The mean alignment from probes placed in each of the three areas was all above the mean of the randomization distribution (mean = 0.15 of the randomization distribution, V4: mean SSIM = 0.37, student’s t-test, $p < 2^{-16}$, PFC: mean SSIM = 0.42, student’s t-test, $p < 2^{-16}$, 7A: mean SSIM = 0.45,

student’s t-test, $p < 2^{-16}$). This demonstrates that the reference image can be found individually in all areas, and that each area contributes to the average profile.




Figure 2. Image similarity analysis for all probes separately colored by their area. The gray bars are the randomization histogram for noise.

Comparison of LSM algorithm to CSD and CSF/GM transition methods

We compared our method to the CSD and CSF/GM methods by asking where these other methods identified layer 4 and layer 1 (respectively) relative to the zero point of the LSM method. We defined as the cross-over point between relative power in the gamma (50-150 Hz) and alpha/beta (10-30 Hz) bands as depth zero. Note that previous studies have linked this cross-over point to layer 4 (Bastos et al., 2018). Relative to the gamma to alpha/beta cross over zero point, negative values indexing laminar positions more superficial to zero, and positive values indexing positions more deep to zero. The average location to the CSD sink was +325 μm (SEM = 84 μm), and the CSF/GM border was on -811 μm (SEM = 57 μm), these laminar markers are shown in Figure 1, right panel.

Machine learning method

In this method we used the raw LFP data to build an ensemble machine learning model that can differentiate between the deep and superficial layers in the brain based only on the univariate raw LFP signals. The full methodological details of the method are given in a companion submission to this paper (Costilla-Reyes et al., *submitted*). We present it here to highlight an alternative approach to automated laminar identification from neurophysiological data.

We pooled half the data from within PFC and V4 to train an extra trees algorithm (Hastie et al., 2009) to model each area's LFP from superficial and deep layers, separately. We then tested whether the model could correctly predict whether an LFP was recorded from superficial or deep layers based on recordings from sessions not used during training. We implemented a cross-validation approach with 5 folds to obtain reliable pattern recognition performance. We used the f-score metric, which is a harmonic average of the precision and recall, to judge the model's performance. The mean f-score of the method was 0.8 for V4 and 0.84 for PFC (Figure 3).

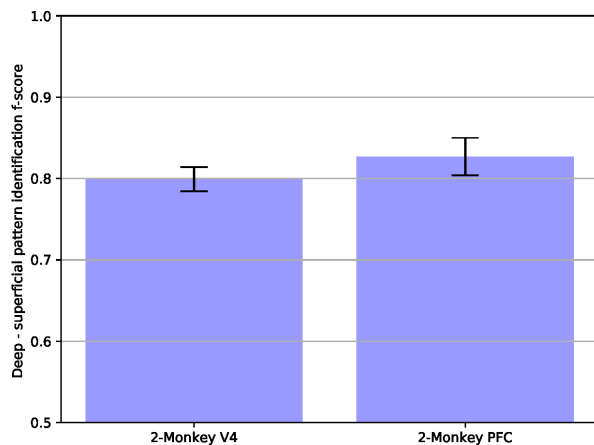


Figure 3. F-score for correctly detecting superficial channels and deep channels for area V4 and PFC, based on data from 2 monkeys and a total of 10 sessions. Mean +/- 1 SEM over cross-validation folds.

Conclusion

In summary, the methods presented in this paper have enabled a purely data-driven approach to laminar identification which overcomes many of the problems associated with traditional methods. The LSM algorithm approach shows that each probe image is similar to the reference image, which is based on the average of probes across all areas. The features producing this similarity were shown to be neuronal rhythms in the delta/theta, alpha/beta and gamma frequency ranges, which are consistently expressed in distinct layers. These similarities across areas in frequency content then enabled us to use a machine learning approach starting from the raw LFP to produce reasonable classification accuracy of electrodes to laminar subdivision for a binary classification approach superficial vs. deep. In future work we will investigate which specific LFP features are most important for this classification. We will also seek to understand which neuronal processes are producing these differences in the laminar expression of cortical rhythms. We also aim

to develop the technique to generalize across areas, and use a number of different univariate data features rather than the raw LFP. Finally, these methods for laminar identification will enable us to shine light on the role of layer-specific contribution for cognition.

Acknowledgments

This research was funded by NIMH (K99MH116100, A.M.B) and by NIMH 5R37MH087027 and MURI N00014-16-1-2832 (E.K.M, A.M.B, O.C.R).

References

- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.
- Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76, 695–711.
- Costilla-Reyes, O., Bastos, A.M., Miller, E.K., (submitted). Pattern recognition of deep and superficial layers of the macaque brain with large-scale local field potentials.
- Felleman, D.J., & van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J., & Tibshirani, R. (2009). *The elements of statistical learning* (Vol. 2) (No. 1). Springer.
- Johnston, K., Ma, L., Schaeffer, L., Everling, S. (2019). Alpha oscillations modulate preparatory activity in marmoset area 8Ad. *Journal of Neuroscience*, 39(10), 1855–1866.
- Mitzdorf, U., & Singer, W. (1979). Excitatory synaptic ensemble properties in the visual cortex of the macaque monkey: a current source density analysis of electrically evoked potentials. *Journal of Comparative Neurology*, 187, 71–84.
- Rao, R.P.N., & Ballard, D.H. (2001). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.
- Schroeder, C.E., Mehta, A.D., Givre, S.J. (1998). A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cerebral Cortex*, 8, 575–592.
- Zhou, W., A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. (2004). Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing*, 13(4), 600–612.