A comparison of neural representation geometry induced by neuronal spiking and local field potentials in the human brain.

Hristos Courellis (hristos@caltech.edu)
Biological Engineering, California Institute of Technology
1200 East California Blvd, Pasadena, CA 91125, USA

Daniel Deng (hdeng3@caltech.edu)
Division of Biology and Biological Engineering, California Institute of Technology
1200 East California Blvd, Pasadena, CA 91125, USA

Juri Minxha, PhD (jminxha@caltech.edu)
Division of Biology and Biological Engineering, California Institute of Technology
1200 East California Blvd, Pasadena, CA 91125, USA

Adam Mamelak, MD (adam.mamelak@cshs.edu)
Neurosurgery Department, Cedars-Sinai Medical Center
8700 Beverly Blvd, Los Angeles, CA 90048, USA

Stefano Fusi, PhD (sf2237@columbia.edu)
Department of Neuroscience, Columbia University
3227 Broadway, New York, NY 10027, USA

Ueli Rutishauser, PhD (urut@caltech.edu)
Division of Biology and Biological Engineering, California Institute of Technology
1200 East California Blvd, Pasadena, CA 91125, USA

Abstract: Studying the geometry of neural representations formed by populations of neurons has provided powerful insight into cognition. Microelectrodes provide access to both single-neuron activity and local field potentials (LFPs). Little is known about whether the geometry of representations can be studied at the level of the LFP. Simultaneous comparison of representational geometries induced in each of these signals during a cognitive task can reveal properties of brain computation at multiple spatio-temporal scales. Here, we study neural representational geometries present in single-unit activity and in local field potentials in the brains of neurosurgical patients who are performing an inferential reasoning task. We demonstrate that the content and format of neural representations can be readily quantified for LFP, and that spike- and LFP-based geometries can be strongly correlated, albeit in a region-dependent manner.
Keywords: neural representation geometry, single neuron, local field potential

Introduction

A major goal of neuroscience is to understand how the relational structure of entangled, high-dimensional sensory inputs and internal states can be re-formatted to achieve complex cognition and behavior (Chung & Abbott, 2021; Marr, 1982). A growing body of literature concerns the study of representational geometry in neural state spaces, which are constructed using either the activity of single neurons (Bernardi et al., 2020; Boyle et al., 2024; Courellis et al., 2023; Fusi et al., 2016; Nogueira et al., 2023; Rigotti et al., 2013), or coarse-grained signals, such as EEG or fMRI (Bhandari et al., 2024; Ito et al., 2022; Kikumoto et al., 2023; Sheahan et al., 2021). However, to date, no studies have examined the relationship between representational geometries induced by neuronal activity (spikes) and simultaneously recorded local field potentials (LFP) in the human brain (Kreiman et al., 2006).

The LFP represents the coherent inputs and network states of a brain region (Buzsáki et al., 2012), whereas spikes are a reflection of the outputs of a region. Comparing the neural representations formed by these two types of signals can thus provide insights into the neural computation performed through the lens of transformations in representational geometry. Here, we did so by performing side-by-side geometric analyses of LFP and single-neuron activity recorded simultaneously (Courellis et al., 2023) from the brains of neurosurgical patients performing an inference-related cognitive task.

Methods

Recording Procedure Neurosurgical patients were implanted with Behnke-Fried electrodes (Fried et al., 1999) that allowed for recording of single-unit activity and LFP from several brain structures (Fig. 1d) including the ventral temporal cortex (VTC), hippocampus (HPC), amygdala (AMY), pre-supplementary motor area (preSMA), dorsal anterior cingulate cortex (dACC), and ventromedial prefrontal cortex (vmPFC). Unit activity from these regions was isolated using standard spike sorting techniques (Rutishauser et al., 2006). LFPs were simultaneously recorded from the same high-impedance electrodes on which the spikes were detected, and were processed using standard techniques (Minxha et al., 2020). Electrode localization was also conducted with peri-op MRI and CT as previously reported (Minxha et al., 2020).

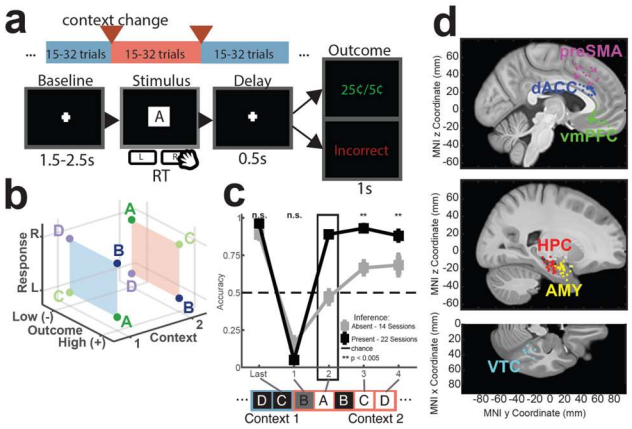

Figure 1. (a) Task structure. **(b)** 8 unique task conditions lead to 35 balanced dichotomies. Letters are unique stimuli.Note: the response to stimuli A/C is L in context 1 and R in context 2, v.v. for B/D. **(c)** Inference performance for sessions where patients performed above chance during non-inference trials (n=36 sessions, 13 patients). **(d)** Implant locations for all patients.

Task Patients performed a serial-reversal learning task where four different images (stimuli) were paired with either a left or right button press (Fig. 1a). There were two different latent contexts (never overtly signaled), each defined by its own unique stimulus-response (SR) map, and all SR associations were inverted between the two contexts. This task structure allowed subjects to perform inference: following the first incorrect trial after switching to a new context, a patient could immediately infer the new correct response for every image that follows by inverting responses. The ability of patients to perform inference was operationalized as their accuracy on "inference trials" – the first instance of encountering a stimulus after a switch (Fig. 1c, black square). A full description of this task is provided in (Courellis et al., 2023).

Quantifying Representation Geometry We used two metrics to quantify features of neural representation geometry: decoding accuracy and cross-condition generalization performance (CCGP). Extensive discussions of these metrics and their uses are available in prior work (Bernardi et al., 2020). In brief, these metrics operate over balanced dichotomies of task conditions (Fig. 1b), which are formed by splitting the 8 unique conditions into two equal groups of 4 conditions (e.g. Fig. 1b, 4 points in context 1 vs 4 points in context 2 is the context dichotomy). Each metric is computed independently for all possible 35 balanced dichotomies. Decoding accuracy across all dichotomies is greater on average in representations that encode more task

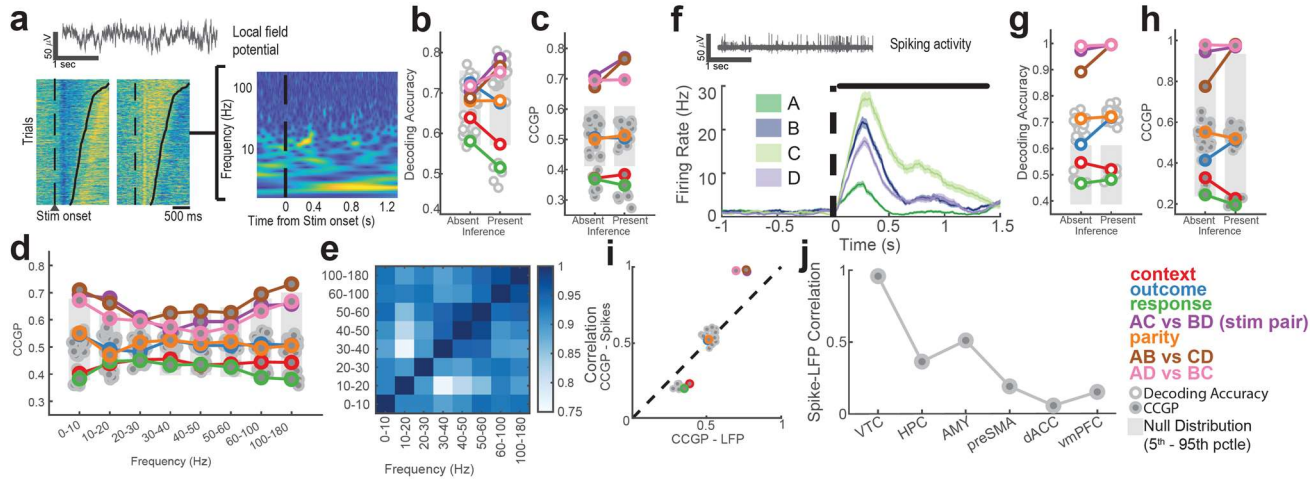

Figure 2. (a) Trial-by-trial LFP power was estimated with wavelet coefficients. Differences in task variable decodability **(b)** and CCGP **(c)** when comparing inference absent (left) and present (right) sessions when using LFP power. Geometric analysis performed on band-limited LFP wavelet power shows significant encoding of task variables at all frequencies **(d)** and that the geometry across frequencies is highly conserved **(e)**. Simultaneously recorded single-unit activity **(f)** was also used to perform task variable decoding **(g)** and CCGP **(h)** analysis. **(i)** Task variable CCGP is highly correlated between spiking activity and LFP. **(j)** Spike-LFP correlation decreases anteriorly along rostro-caudal axis of brain.

variables and interactions thereof. CCGP is an index of abstraction, or disentangling, of one variable with respect to others that are simultaneously encoded.

Note that the above metrics can be computed using any set of features that contain neural data recorded during the task. For the LFP, the features were wavelet coefficients estimated using a continuous wavelet transform for log-spaced frequencies from 0 to 180 Hz from 0.2 to 1.2s after stimulus onset (Fig. 2a). For the single-unit activity, spikes were counted during the same time window (Fig. 2f). Features were estimated trial-by-trial. Decoding accuracies and CCGPs were cross-validated on individual trials.

Results

Representational geometry of LFP in VTC is highly structured. Data recorded over 36 sessions (n = 13 patients) yielded LFP from 2960 channels (160 VTC, 616 HPC, 544 AMY, 528 preSMA, 568 dACC, 544 vmPFC). Sessions were classified as "Inference Present" (n=22) or "Inference Absent" (n=14) based on performance on the first inference trial (Fig. 1c). Analyses shown in Fig. 2d-j were performed using the "Inference Present" sessions. We begin by considering LFP geometry in the VTC, a region which is known to exhibit spike-LFP correlation in the nonhuman primate (Kreiman et al., 2006). Analysis of VTC LFPs revealed a highly structured representation wherein stimulus-identity related dichotomies (brown, purple, pink) were decodable (Fig. 2b) and in an abstract format (Fig. 2c), regardless of the inference status of the patient (Fig. 2b,c, left vs right). These findings are consistent with VTC being a sensory area containing task-invariant visual stimulus codes (Courellis et al., 2023).

Band-limited geometric analyses revealed that low and high frequency LFPs significantly contribute to the representational geometry (Fig. 2d), and that the geometry is highly correlated across frequencies (Fig. 2e, all $\rho > 0.75$, linear correlation between CCGP over all dichotomies). Thus, representational geometry in LFP does not arise from high-gamma alone, which is correlated with neuronal spiking (Buzsáki et al., 2012).

Spike- and LFP-based Representational Geometries are correlated in an area-dependent manner. The same neural recordings yielded 2694 well isolated single-units (261 VTC, 494 HPC, 889 AMY, 269 preSMA, 310 dACC, 463 vmPFC). Geometric analyses on VTC spikes (Fig. 2g,h) revealed a representational geometry that was highly correlated with the one found in the LFP through both variable decodability ($\rho = 0.81$) and CCGP (Fig. 2i, $\rho = 0.95$). This relationship, however, was not conserved across areas, with Spike-LFP CCGP correlations decreasing in more anterior regions (Fig. 2j).

Discussion

This analysis is a major step towards developing an understanding of how the relational structure of cognitive task representations is simultaneously encoded and transformed between neuronal spikes and LFP. We find that the representational format in VTC is strongly aligned between spikes and LFP. However, this relationship breaks down in other regions despite encoding of the individual task variables by both the neurons (Courellis et al., 2023) and the LFP. For example, stimulus-identity related dichotomies are decodable in both the LFP (not shown) and spikes in vmPFC, but their geometries are not correlated at the level of balanced dichotomies, warranting further investigation.

Many features of LFP representational geometry remain to be explored, including the learning- and performance-dependence of the LFP geometry in the context of the cognitive task, the details of how spike and LFP-based geometries differ in a region-specific manner, and the resultant implications for local computations occurring in each of these regions.

Acknowledgments

We thank the members of the Adolphs and Rutishauser labs for discussion, C. M. Reed, J. M. Chung, and the staff of the Cedars-Sinai Epilepsy Monitoring Unit for their support with recording, and we thank the patients and their families for their boundless generosity.

Funding: NIH U01NS117839, Simons Foundation Collaboration on the Global Brain (542941), and the Caltech NIMH Conte Center P50MH094258).

References

Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., & Salzman, C. D. (2020). The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell*, *183*(4), 954-967.e21. https://doi.org/10.1016/j.cell.2020.09.031

Bhandari, A., Keglovits, H., Chicklis, E., & Badre, D. (2024). *Task structure tailors the geometry of neural representations in human lateral prefrontal cortex* (p. 2024.03.06.583429). bioRxiv. https://doi.org/10.1101/2024.03.06.583429

Boyle, L. M., Posani, L., Irfan, S., Siegelbaum, S. A., & Fusi, S. (2024). Tuned geometries of hippocampal representations meet the computational demands of social memory. *Neuron*, *0*(0). https://doi.org/10.1016/j.neuron.2024.01.021

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*(6), 407–420. https://doi.org/10.1038/nrn3241

Chung, S., & Abbott, L. F. (2021). Neural population geometry: An approach for understanding biological and artificial neural networks. *Current Opinion in Neurobiology*, *70*, 137–144. https://doi.org/10.1016/j.conb.2021.10.010

Courellis, H. S., Mixha, J., Cardenas, A. R., Kimmel, D., Reed, C. M., Valiante, T. A., Salzman, C. D., Mamelak, A. N., Adolphs, R., Fusi, S., & Rutishauser, U. (2023). *Abstract representations emerge in human hippocampal neurons during inference behavior* (p. 2023.11.10.566490). bioRxiv. https://doi.org/10.1101/2023.11.10.566490

Fried, I., Wilson, C. L., Maidment, N. T., Engel, J., Behnke, E., Fields, T. A., MacDonald, K. A., Morrow, J. W., & Ackerson, L. (1999). Cerebral microdialysis combined with single-neuron and electroencephalographic recording in neurosurgical patients. Technical note. *Journal of Neurosurgery*, *91*(4), 697–705. https://doi.org/10.3171/jns.1999.91.4.0697

Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, *37*, 66–74. https://doi.org/10.1016/j.conb.2016.01.010

Ito, T., Klinger, T., Schultz, D. H., Murray, J. D., Cole, M. W., & Rigotti, M. (2022). *Compositional generalization through abstract representations in human and artificial neural networks* (arXiv:2209.07431). arXiv. https://doi.org/10.48550/arXiv.2209.07431

Kikumoto, A., Bhandari, A., Shibata, K., & Badre, D. (2023). A Transient High-dimensional Geometry Affords Stable Conjunctive Subspaces for Efficient Action Selection. *bioRxiv*, 2023.06.09.544428. https://doi.org/10.1101/2023.06.09.544428

Kreiman, G., Hung, C. P., Kraskov, A., Quiroga, R. Q., Poggio, T., & DiCarlo, J. J. (2006). Object Selectivity of Local Field Potentials and Spikes in the Macaque Inferior Temporal Cortex. *Neuron*, *49*(3), 433–445. https://doi.org/10.1016/j.neuron.2005.12.019

Marr, D. (1982). *Vision*. W. H. Freeman and Company.

Minxha, J., Adolphs, R., Fusi, S., Mamelak, A. N., & Rutishauser, U. (2020). Flexible recruitment of memory-based choice representations by human medial-frontal cortex. *Science (New York, N.Y.)*, *368*(6498), eaba3313. https://doi.org/10.1126/science.aba3313

Nogueira, R., Rodgers, C. C., Bruno, R. M., & Fusi, S. (2023). The geometry of cortical representations of touch in rodents. *Nature Neuroscience*, *26*(2), Article 2. https://doi.org/10.1038/s41593-022-01237-9

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, *497*(7451), Article 7451. https://doi.org/10.1038/nature12160

Rutishauser, U., Schuman, E. M., & Mamelak, A. N. (2006). Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. *Journal of Neuroscience Methods*, *154*(1–2), 204–224. https://doi.org/10.1016/j.jneumeth.2005.12.033

Sheahan, H., Luyckx, F., Nelli, S., Teupe, C., & Summerfield, C. (2021). Neural state space alignment for magnitude generalization in humans and recurrent networks. *Neuron*, *109*(7), 1214-1226.e8. https://doi.org/10.1016/j.neuron.2021.02.004