Neural Manifold Capacity Captures Representation Geometry, Correlations, and Task-Efficiency Across Species and Behaviors

Chi-Ning Chou (cchou@flatironinstitute.org)
Center for Computational Neuroscience, Flatiron Institute, New York, USA

Luke Arend (luke.arend@nyu.edu)
Center for Neural Science, New York University, New York, NY, USA

Albert J. Wakhloo (awakhloo@flatironinstitute.org)
Center for Theoretical Neuroscience, Columbia University, New York, NY, USA

Royoung Kim (roespigas@gmail.com)
Center for Neuroscience Imaging Research, Institute for Basic Science, Suwon, Republic of Korea

Will Slatton (wslatton@flatironinstitute.org)
Center for Neural Science, New York University, New York, NY, USA

SueYeon Chung (schung@flatironinstitute.org)
Center for Computational Neuroscience, Flatiron Institute, 160 5th Avenue

Abstract

Relating the coordinated activity of neurons to cognitive functions is a fundamental challenge in neuroscience. While experimental evidence indicates these neuronal populations act as core computational units, quantifying these population codes in relation to their functional roles has remained elusive. Prior approaches based on representational geometries, functional alignment or dimensionality reduction have faced limitations in robustly linking neural population structure to computation across scales and modalities. Here, we fill this gap by introducing effective Geometric measures from Correlated Manifold Capacity theory (GCMC), a framework that employs analytical methods from statistical physics, to connect the geometry of neural population activities to readout performance, thereby quantifying coding efficiency. Applying this to diverse neural recordings across organisms and tasks, we demonstrate multi-scale analyses previously inaccessible. These include tracking changes in coding efficiency and geometry across brain regions, revealing task-relevant manifold dynamics over time, and characterizing representational changes during learning. The geometric measures serve as interpretable descriptors relating the structure of coordinated neural population activity to embedded computations. Our framework provides a general and principled approach for mapping neural population codes to their functional roles, enabling data-driven insights into the neural underpinnings of perception and behavior.

Keywords: neural representations; population geometry; efficient coding

Introduction

Neurons collectively represent task-relevant information in the brain (Yuste, 2015; Saxena & Cunningham, 2019). In neuroscience we informally call the collection of neural response vectors to some given task condition or input stimulus a "neural manifold" (see Fig. 1a for some examples) (Chung & Abbott, 2021). The neural manifold principle postulates that the geometrical, statistical, and structural properties of these manifolds are highly relevant to the study of the functional roles of population representations (Chung & Abbott, 2021; Kriegeskorte & Kievit, 2013; Kriegeskorte & Wei, 2021).

Storage capacity (Gardner, 1988; Brunel, Nadal, & Toulouse, 1992) measures the amount of linearly decodable information per neuron. Despite being theoretically used to quantify neural population coding in a wide range of models (Clopath, Nadal, & Brunel, 2012; Rubin, Abbott, & Sompolinsky, 2017), it has been challenging to estimate the capacity in high-dimensional and heterogeneous data. The Manifold Capacity Theory (MCT) (Chung, Lee, & Sompolinsky, 2018) adopts the formalism of statistical physics and battle the curse of dimensionality via deriving the low-dimensional effective geometry underlying the data. By connecting capacity to effective geometry through an analytical formula, the MCT further suggests the definition of computationally relevant geometric terms such as effective dimension and effective radius of neural manifolds. These effective geometric measures have led to applications in analyzing neural representations across biological datasets (Yao et al., 2023; Paraouty et al., 2023; Froudarakis et al., 2020) and artificial neural networks (Cohen, Chung, Lee, & Sompolinsky, 2020; Dapello et al., 2021; Kuoch et al., 2023). However, the effective manifold geometric measures from MCT are ignorant to neural correlations, which play crucial roles in neural information processing (Averbeck, Latham, & Pouget, 2006; Clopath et al., 2012). Consequently, they have been limited to datasets with low correlations and sometimes yield inaccurate approximations to capacity (Wakhloo, Sussman, & Chung, 2023).

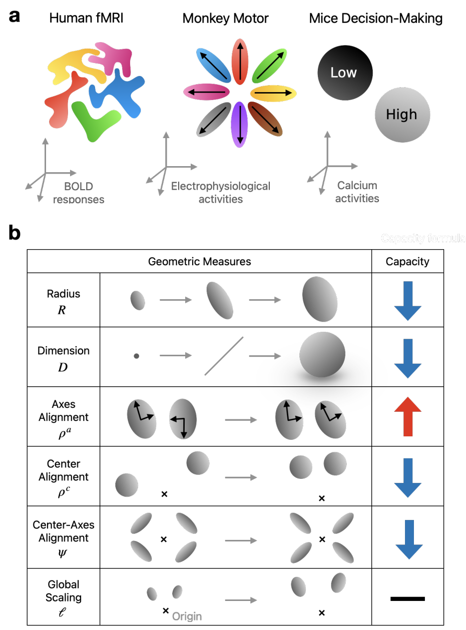

Figure 1: **Manifold geometry as intermediate descriptors for multi-scale neural data analysis.** **a**, Examples of using manifolds as analysis units. These manifolds lie in the neural state space with each coordinate being the neural activity of a recording unit. Left to right: human fMRI recordings on THINGS dataset (Hebart et al., 2023), where a manifold corresponds to a stimuli category and a region of interest (ROI); a monkey delayed center-out reaching task, where a manifold corresponds to a target; an auditory decision making task in mouse posterior parietal cortex (Plitt & Giocomo, 2021), where a manifold is associated to the decision outcome. **b**, Effective manifold geometric measures serve as bridges between the neural activity space and the behavior/perceptual space. We define effective geometric measures on top of the anchor geometry. The table shows the qualitative relationship between each measure and the capacity.

Results

In this work, we introduce effective Geometric manifold measures from Correlated Manifold Capacity theory (GCMC) (Fig. 1b). Our contributions are three-fold: (1) GCMC incorporates the complex correlation structure via new effective geometric measures, which explicate how manifold geometry influences downstream computational efficiency. (2) GCMC connects noise correlations to manifold geometry, hence it unifies the concept of correlations at different system scales. (3) GCMC enables multi-scale data analysis such as quantifying the spatial progression of encoding efficiency across brain regions (Fig. 2a), revealing task-relevant temporal dynamics (Fig. 2b), and characterizing the variances and invariances in learning (Fig. 2c). We demonstrate the power and the applicability of GCMC in a wide spectrum of datasets (Freeman, Ziemba, Heeger, Simoncelli, & Movshon, 2013; Majaj, Hong, Solomon, & DiCarlo, 2015; Hebart et al., 2023; Perich et al., 2018; Kiani, Cueva, Reppas, & Newsome, 2014; Kiani et al., 2015; Plitt & Giocomo, 2021; Najafi et al., 2020) with various task modalities (e.g., vision, perceptual decision, motor, spatial memory), various model organisms (e.g., mice, monkeys, humans, artificial neural networks), and various recording methods (e.g., electrophysiology, calcium imaging, fMRI) (Fig. 1a). Finally, the effective manifold geometric measures can be conceptualized as order parameters for phases associated with computational efficiency, aiding in the generation of data-driven hypotheses and latent embedding. In summary, GCMC opens up opportunities to explore new neuroscience questions at the neural population level.

References

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature reviews neuroscience*, *7*(5), 358–366.

Brunel, N., Nadal, J.-P., & Toulouse, G. (1992). Information capacity of a perceptron. *Journal of Physics A: Mathematical and General*, *25*(19), 5017.

Chung, S., & Abbott, L. (2021). Neural population geometry: An approach for understanding biological and artificial neural networks. *Current opinion in neurobiology*, *70*, 137–144.

Chung, S., Lee, D. D., & Sompolinsky, H. (2018). Classification and geometry of general perceptual manifolds. *Physical Review X*.

Clopath, C., Nadal, J.-P., & Brunel, N. (2012). Storage of correlated patterns in standard and bistable purkinje cell models. *PLoS computational biology*, *8*(4), e1002448.

Cohen, U., Chung, S., Lee, D. D., & Sompolinsky, H. (2020). Separability and geometry of object manifolds in deep neural networks. *Nature communications*.

Dapello, J., Feather, J., Le, H., Marques, T., Cox, D., McDermott, J., ... Chung, S. (2021). Neural population geometry reveals the role of stochasticity in robust perception. *Advances in Neural Information Processing Systems*, *34*, 15595–15607.

Freeman, J., Ziemba, C. M., Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (2013). A functional and perceptual signature of the second visual area in primates. *Nature neuroscience*, *16*(7), 974–981.

Froudarakis, E., Cohen, U., Diamantaki, M., Walker, E. Y., Reimer, J., Berens, P., ... Tolias, A. S. (2020). Object manifold geometry across the mouse cortical visual hierarchy. *BioRxiv*.

Gardner, E. (1988). The space of interactions in neural network models. *Journal of physics A: Mathematical and general*, *21*(1), 257.

Hebart, M. N., Contier, O., Teichmann, L., Rockter, A. H., Zheng, C. Y., Kidder, A., ... Baker, C. I. (2023). Things-data, a multimodal collection of large-scale datasets for investigating object representations in human brain and behavior. *Elife*, *12*, e82580.

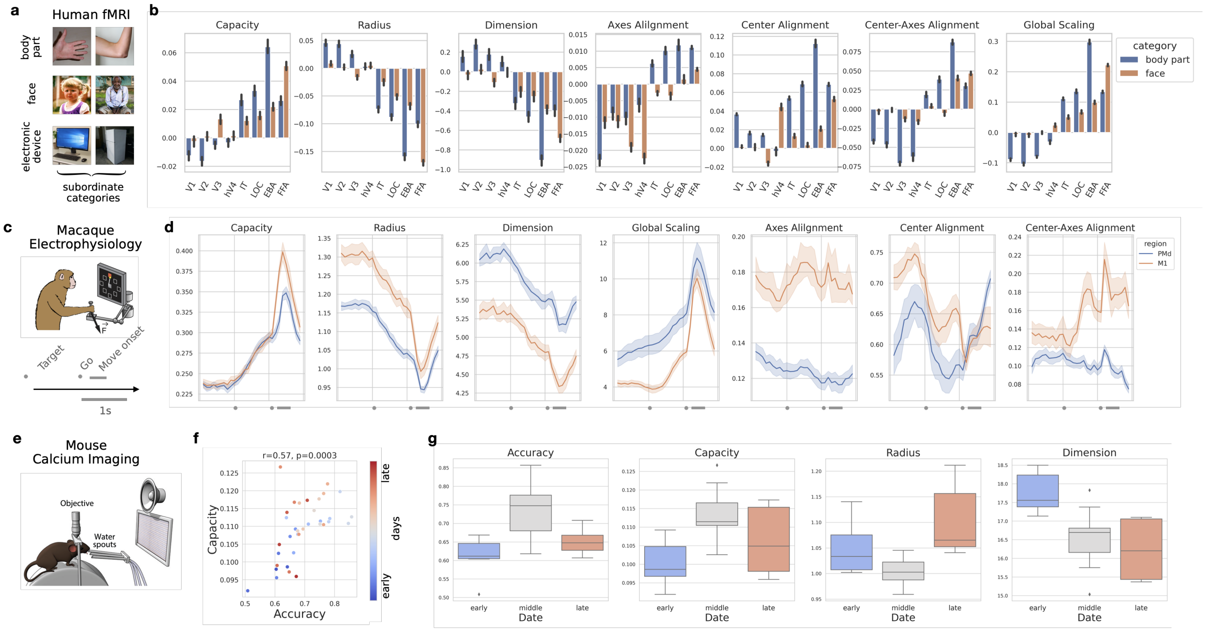

Figure 2: **Manifold geometry as intermediate descriptors for multi-scale neural data analysis.** **a**, A human fMRI dataset from Hebart et al. (Hebart et al., 2023). We consider three image categories, each having 5 subordinate categories with 12 images. A manifold is defined as the preprocessed BOLD responses from the same region of interest (ROI) and corresponds to stimuli from the same subordinate category (Methods). **b**, Effective manifold geometric measures of an example human subject. GCMC analysis was conducted on all three categories and all ROIs. For each ROI, we subtract the resulting values of manifold geometric measures by that corresponding to the electronic device category. **c**, Top: A monkey electrophysiological dataset on a delayed center-out reaching task with eight targets. Neural recordings are from premotor and primary motor cortices from Perich et al. (Perich et al., 2018). Bottom: The trial structure. Figure adapted from ref. (Perich et al., 2018), with permission. **d**, The in-trial dynamics of the effective manifold geometry of an example monkey. **e**, A mice decision-making dataset. Calcium imaging recordings from posterior parietal cortex by Najafi et al. (Najafi et al., 2020). Behavioral setup: mice were trained to report whether the multisensory event rate is low or high. Figure adapted from ref. (Najafi et al., 2020), with permission. **f**, Scatterplot of capacity versus behavioral performance of 35 training sessions of an example mouse. r, Pearson correlation coefficient. **g**, Barplot of accuracy, capacity, and effective geometric measures in three learning stages.

Kiani, R., Cueva, C. J., Reppas, J. B., & Newsome, W. T. (2014). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Current Biology*, *24*(13), 1542–1547.

Kiani, R., Cueva, C. J., Reppas, J. B., Peixoto, D., Ryu, S. I., & Newsome, W. T. (2015). Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. *Neuron*, *85*(6), 1359–1373.

Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: integrating cognition, computation, and the brain. *Trends in cognitive sciences*, *17*(8), 401–412.

Kriegeskorte, N., & Wei, X.-X. (2021). Neural tuning and representational geometry. *Nature Reviews Neuroscience*, *22*(11), 703–718.

Kuoch, M., Chou, C.-N., Parthasarathy, N., Dapello, J., DiCarlo, J. J., Sompolinsky, H., & Chung, S. (2023). Probing biological and artificial neural networks with task-dependent neural manifolds. In *Conference on parsimony and learning (proceedings track)*.

Majaj, N. J., Hong, H., Solomon, E. A., & DiCarlo, J. J. (2015). Simple learned weighted sums of inferior temporal neuronal firing rates accurately predict human core object recognition performance. *Journal of Neuroscience*, *35*(39), 13402–13418.

Najafi, F., Elsayed, G. F., Cao, R., Pnevmatikakis, E., Latham, P. E., Cunningham, J. P., & Churchland, A. K. (2020). Excitatory and inhibitory subnetworks are equally selective during decision-making and emerge simultaneously during learning. *Neuron*, *105*(1), 165–179.

Paraouty, N., Yao, J. D., Varnet, L., Chou, C.-N., Chung, S., & Sanes, D. H. (2023). Sensory cortex plasticity supports auditory social learning. *Nature communications*, *14*(1), 5828.

Perich, M. G., Gallego, J. A., & Miller, L. E. (2018). A neural population mechanism for rapid learning. *Neuron*.

Plitt, M. H., & Giocomo, L. M. (2021). Experience-dependent contextual codes in the hippocampus. *Nature neuroscience*, *24*(5), 705–714.

Rubin, R., Abbott, L., & Sompolinsky, H. (2017). Balanced excitation and inhibition are required for high-capacity, noise-robust neuronal selectivity. *Proceedings of the National*

Academy of Sciences, *114*(44), E9366–E9375.

Saxena, S., & Cunningham, J. P. (2019). Towards the neural population doctrine. *Current opinion in neurobiology*, *55*, 103–111.

Wakhloo, A. J., Sussman, T. J., & Chung, S. (2023). Linear classification of neural manifolds with correlated variability. *Physical Review Letters*.

Yao, J. D., Zemlianova, K. O., Hocker, D. L., Savin, C., Constantinople, C. M., Chung, S., & Sanes, D. H. (2023). Transformation of acoustic information to sensory decision variables in the parietal cortex. *Proceedings of the National Academy of Sciences*, *120*(2), e2212120120.

Yuste, R. (2015). From the neuron doctrine to neural networks. *Nature reviews neuroscience*, *16*(8), 487–497.