Studying How Neural Networks Learn Computationally Hierarchical Representations Using Manifold Capacity

Yichen Wang[1,2], Chi-Ning Chou[2], Artem Kirsanov[2,3], SueYeon Chung[2,3]

Affiliation 1: University of California, Los Angeles, Los Angeles, CA 90024
Affiliation 2: Center for Computational Neuroscience, Flatiron Institute, New York, NY 10010
Affiliation 3: Center for Neural Science, NYU, New York, NY 10003

Background & Motivation

Neural representation is a collection of neuron activations in the brain (e.g. spiking rates) or ANNs (e.g. activation function outputs). During each computation, a stimulus would be mapped to a coordinate specified by the neuron activations in a high-dimensional neural representation space. When object variance is introduced, a point becomes an object manifold.

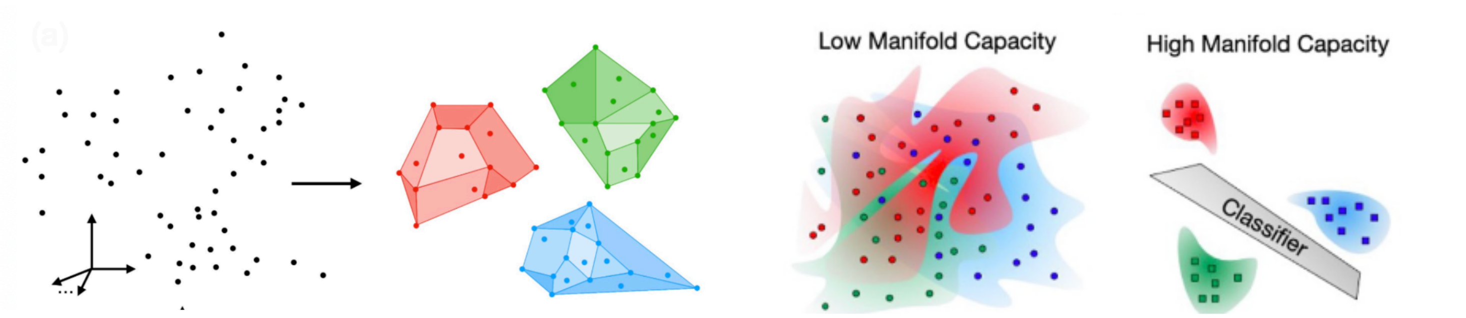

Adapted from [2]

Manifold Capacity: A quantitative measure for the linear decoding and packing efficiency of neural manifolds that are linked to mesoscopic geometric observables.

Motivation: Humans & Neural Networks break the curse of dimensionality, but how?

Current Hypothesis: The learned data must be structured hierarchically in some way.

Neuroscience Observations
Modalities: Visual, Auditory, Motor, Reasoning
Areas: PFC, IT, Hippocampus, Basal Ganglia

Manifold Capacity

Machine Learning Observations
Compositionality of Features
LLM Grammars
CNN Filters
Diffusion Process

Our Answer: The learned hierarchical structure is computationally efficient for tasks.

Synthetic Hierarchical Data: How Are They Constructed?

Setup: Gaussian mixture model with three hyperparameters: Manifold radius (R), Entry-manifold spread (R1), Subordinate-manifold spread (R2).

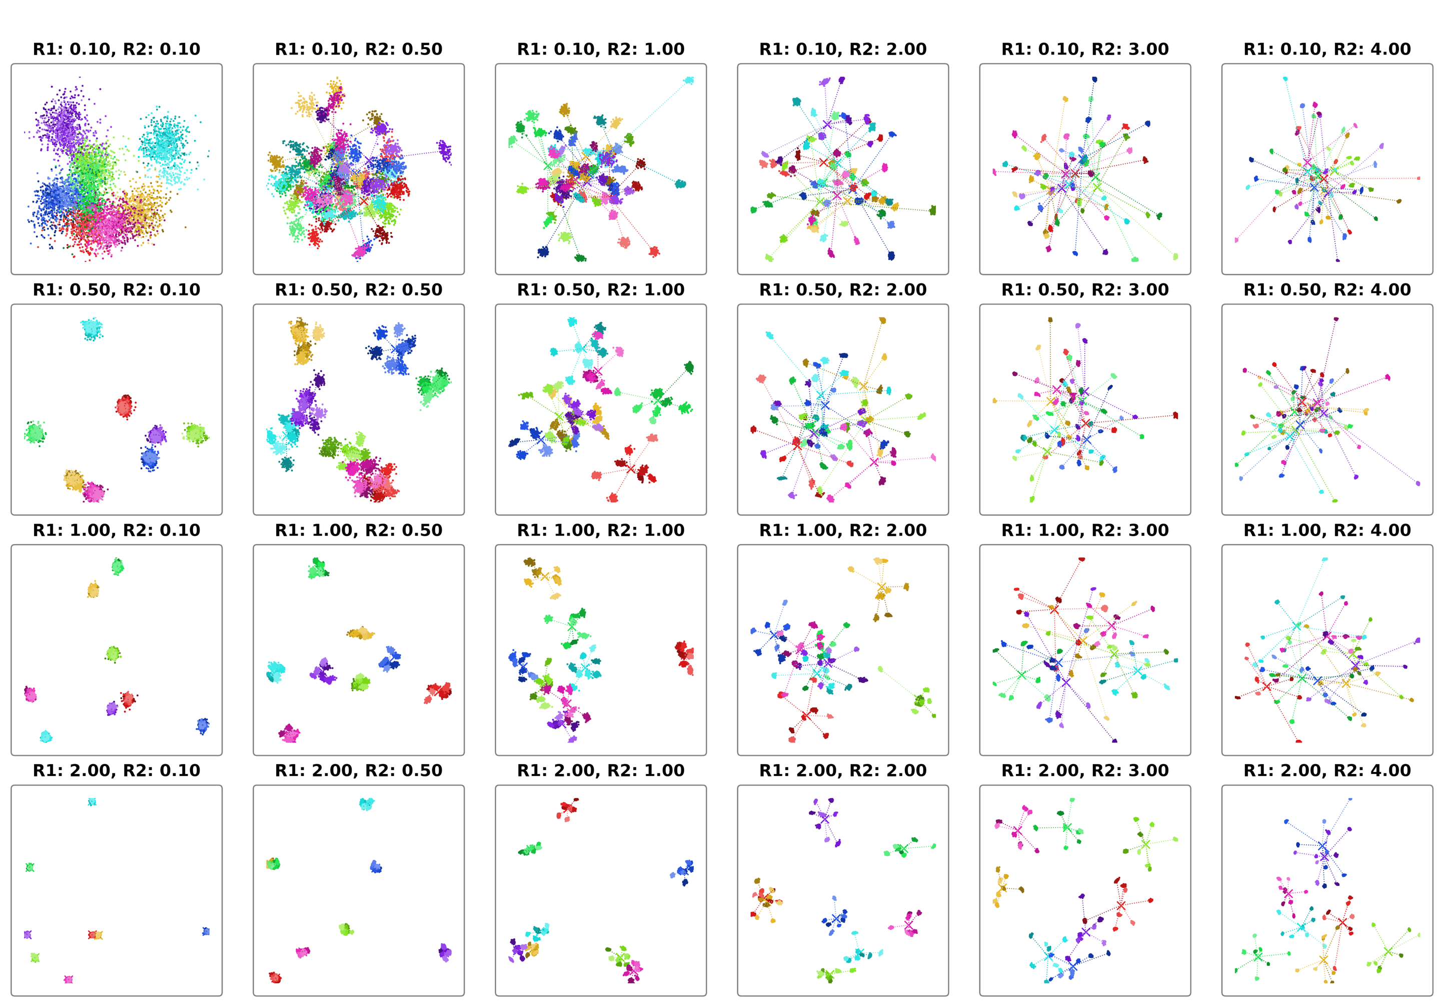

The Capacity Profile Delineates Computationally Hierarchical Structures

Q: How to quantify and delineate different computationally hierarchical structures?

Method: Evaluate manifold capacity in three different strategies, capturing computational efficiency at different hierarchical levels.

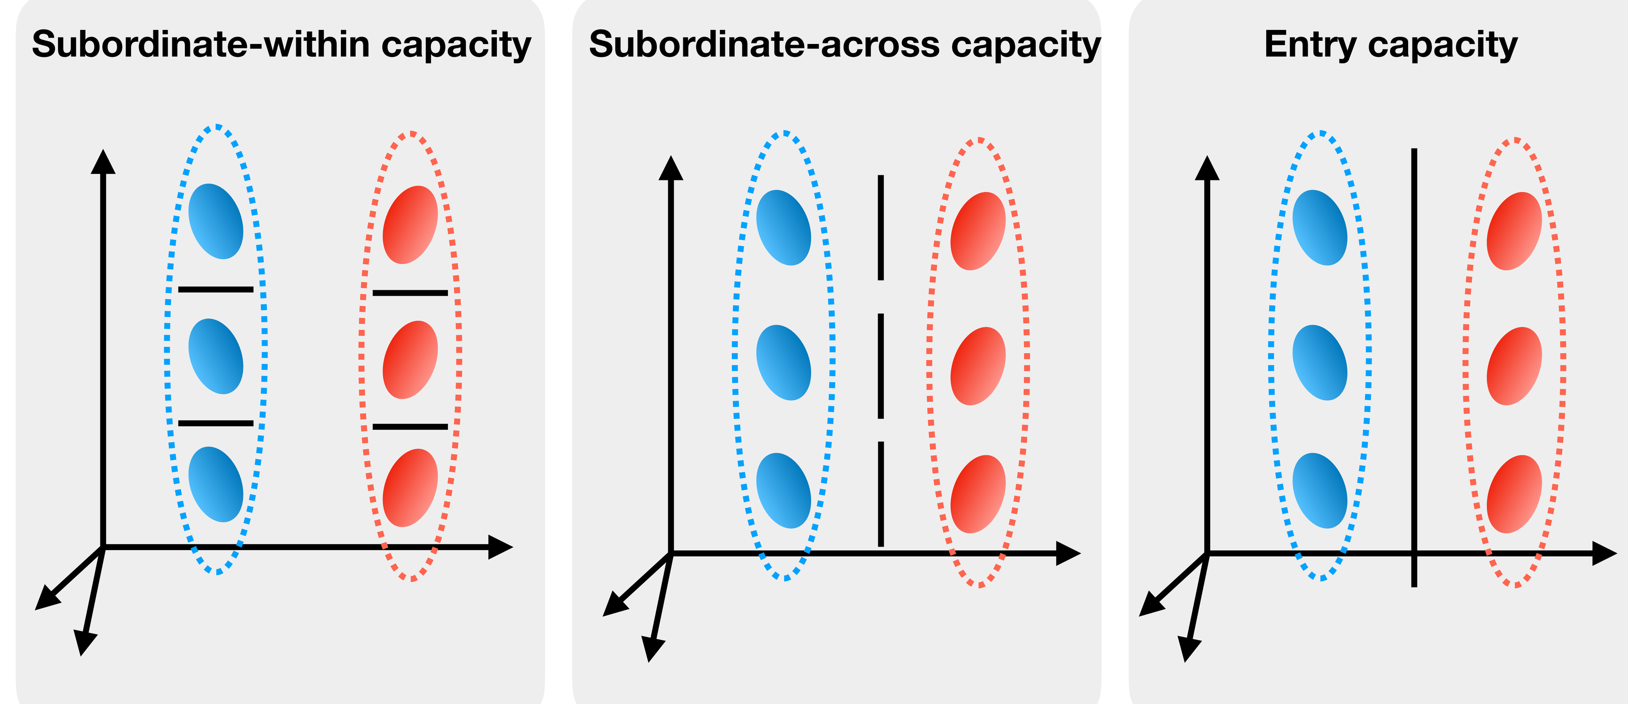

Result: We can delineate four distinct hierarchical organization species in the capacity profile space, which is linked to ground truth synthetical hierarchical data space.

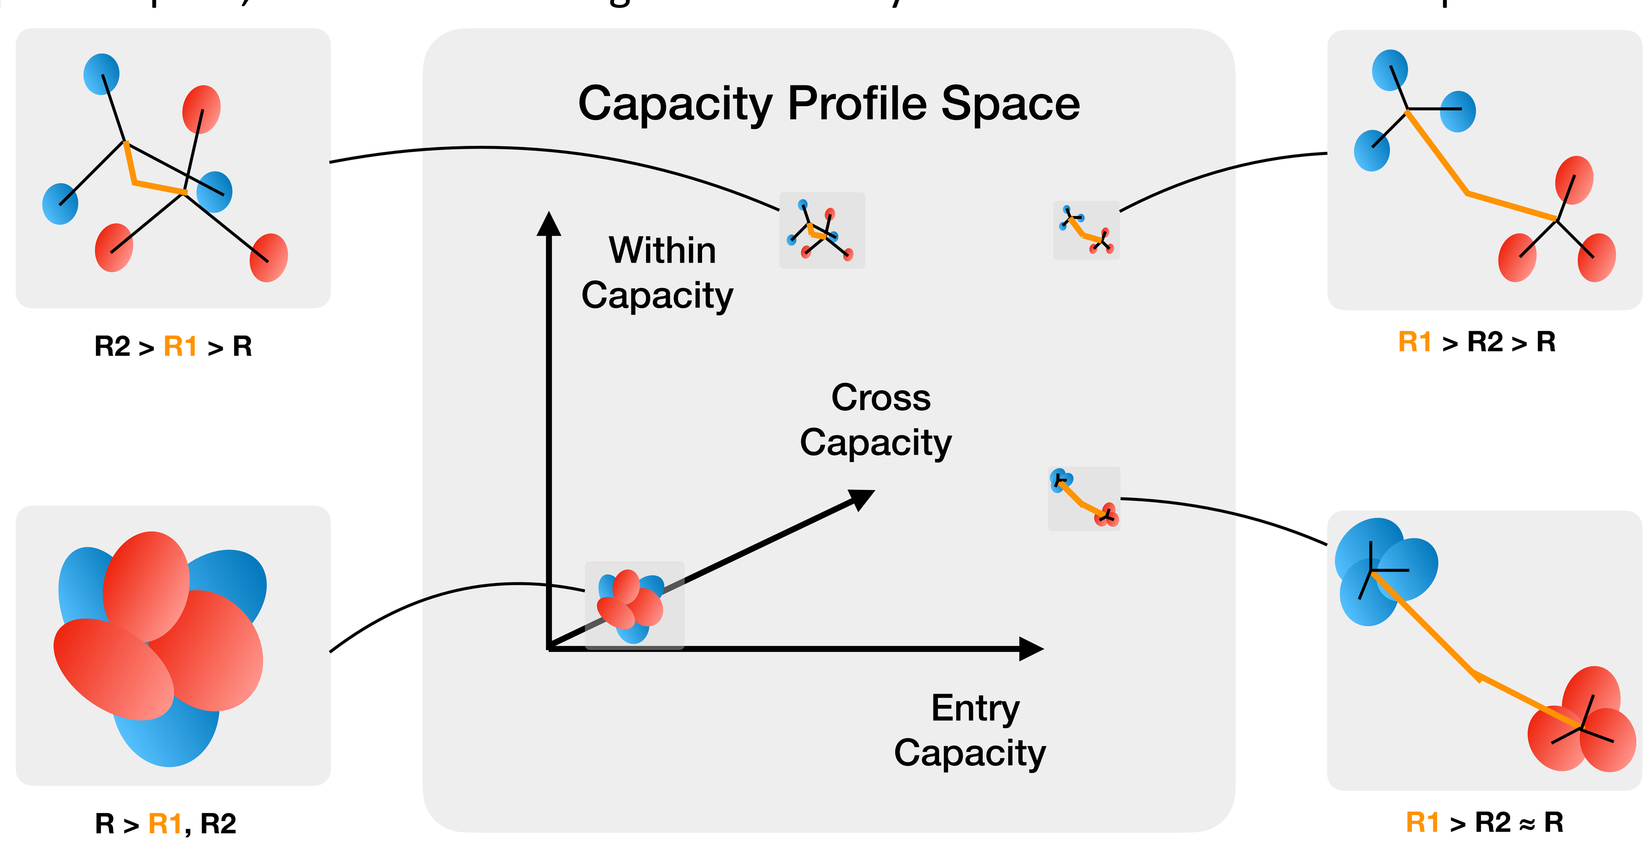

Multi-tasking Induce Distinct Hierarchical Learning Strategies

Q2: What induce neural networks to learn computationally hierarchical structures?

Hierarchical strategy 1

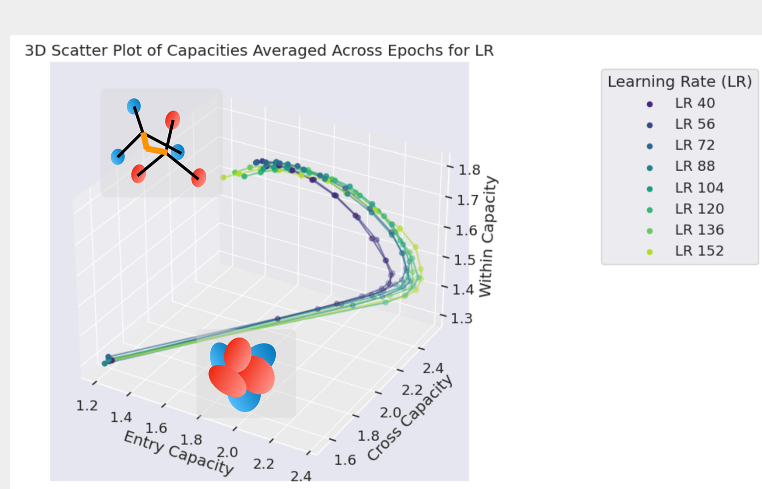

Single-tasking: Trained to categorize only at the subordinate level.

Hierarchical strategy 2

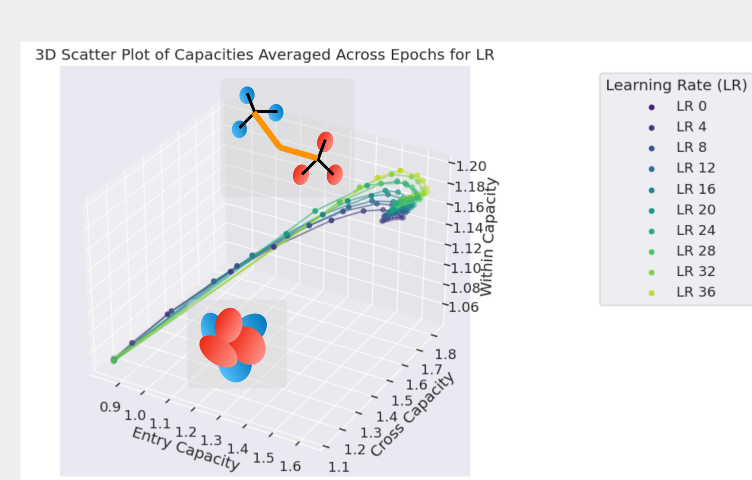

Multi-tasking: Trained to categorize at both entry and subordinate level.

Data Organization Induce Distinct Hierarchical Learning Strategies

Q2: What induce neural networks to learn computationally hierarchical structures?

Non-hierarchical strategy

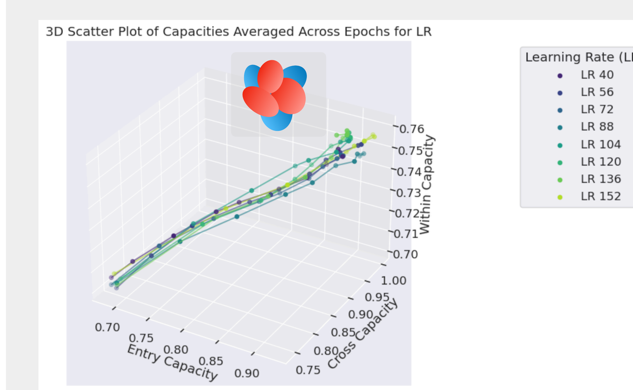

Non-hierarchical training data. Trained to categorize only at the subordinate level.

Hierarchical strategy

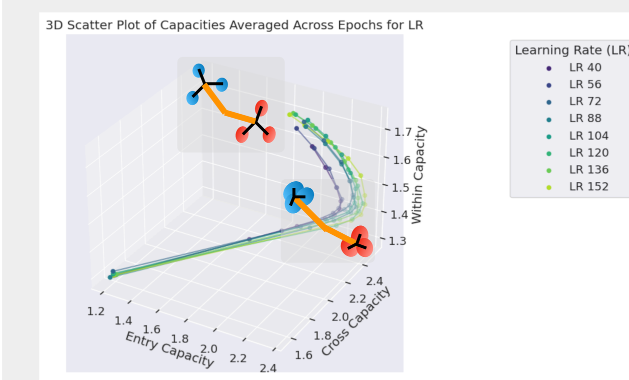

Hierarchical training data. Trained to categorize only at the subordinate level.

Computational Consequence: Implicit Learning

Q3: What are the benefits of computationally hierarchical features? When we learn to categorize objects, do we implicitly learn to categorize higher level features?

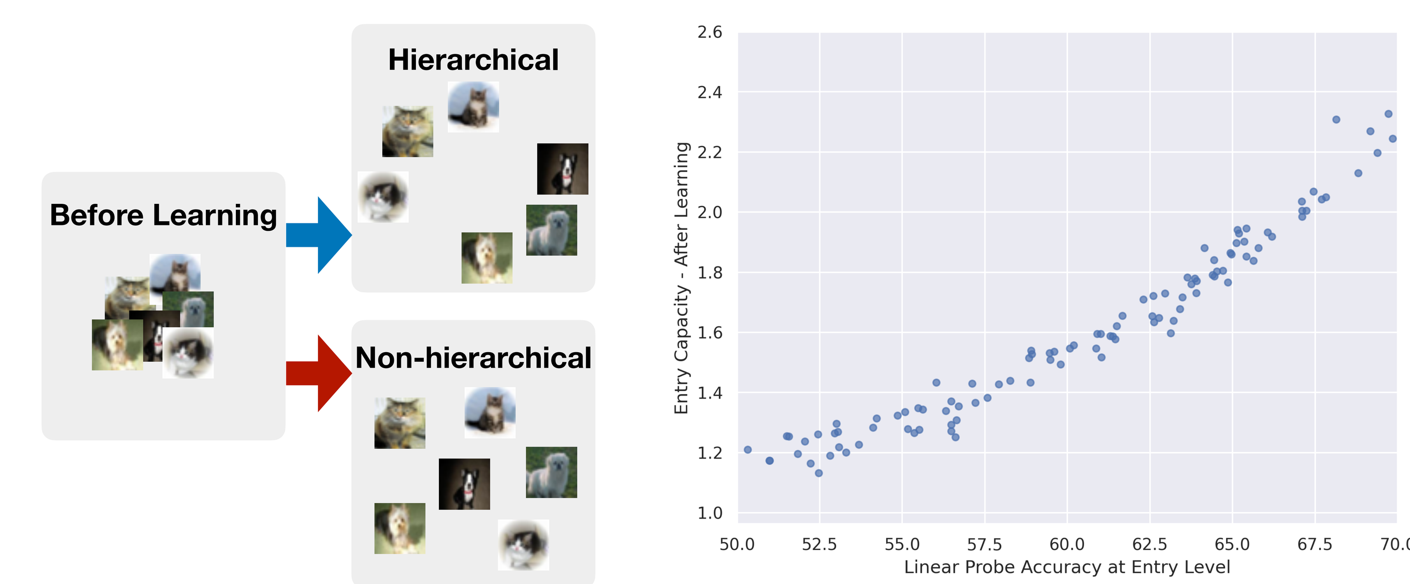

Result: Entry level capacity tracks generalization performance to entry-level classification tasks.

Acknowledgements

Y.W. would like to thank his supervisor Prof. SueYeon Chung and mentor Dr. Chi-Ning Chou for project planning, discussion, suggestion, and personal support throughout the internship process, thank Artem Kirsanov for collaboration, thank Will Yang and Hyunmo Kong for helpful discussion, and thank the Simons Foundation for this internship opportunity.

Main References

[1] Chou, Chi-Ning, Arend, Luke, Wakhloo, Albert J., Kim,Royoung, Slatton, Will, and Chung, SueYeon. "Neural Manifold Capacity Captures Representation Geometry, Correlations, and Task-Efficiency Across Species and Behaviors", bioRxiv 2024.02.26.582157.

[2] Chung, SueYeon, and L. F. Abbott. "Neural population geometry: An approach for understanding biological and artificial neural networks." Current opinion in neurobiology 70 (2021): 137-144.

[3] Chung, SueYeon, Daniel D. Lee, and Haim Sompolinsky. "Classification and geometry of general perceptual manifolds." Physical Review X 8.3 (2018): 031003.