Exploring the Geometry of Latent Structures in Neural Manifolds

Yichen Wang[1,2], Chi-Ning Chou[2], 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

Motivation

Neural representation is a collection of neuron activations in the brain (e.g. spiking rates) or ANNs (e.g. activation function outputs) [1]. 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.

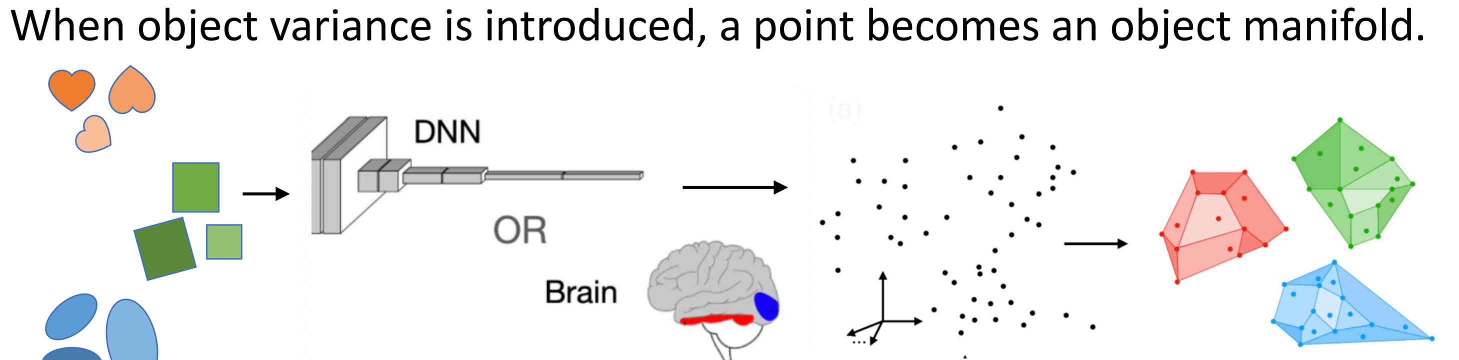

Previously: Our group's work successfully connects geometry of object manifolds (e.g. dimension, radius, capacity) to computation underlying classification [2,3].

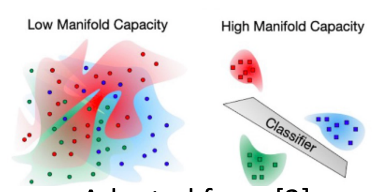

Adapted from [2]

Idea: Biological data has latent structures, which naturally brings up:

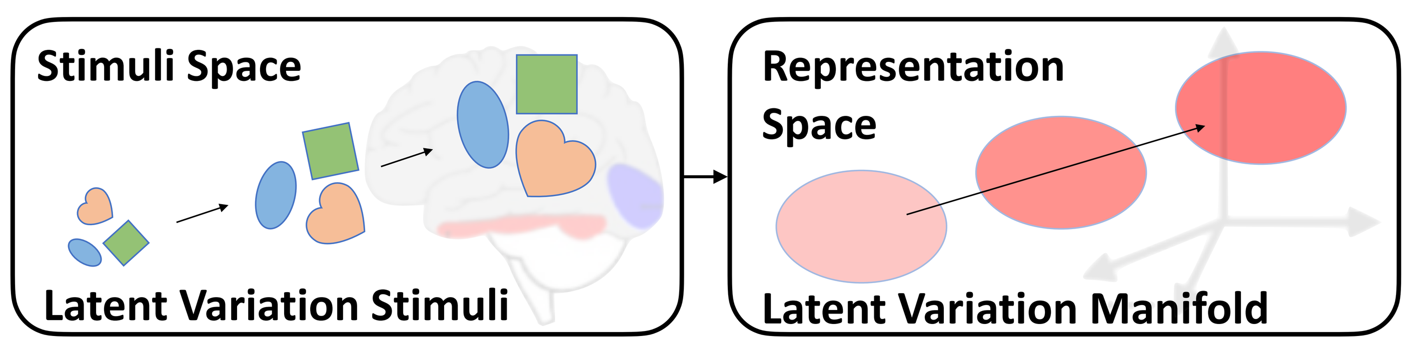

Question: How does the latent variation manifold geometry connect to computation underlying regression?

Intuition: Geometry (dimension, longitudinal radius, etc.) reflects regression performance - a theoretical work developing by Abdul Canatar et al [4].

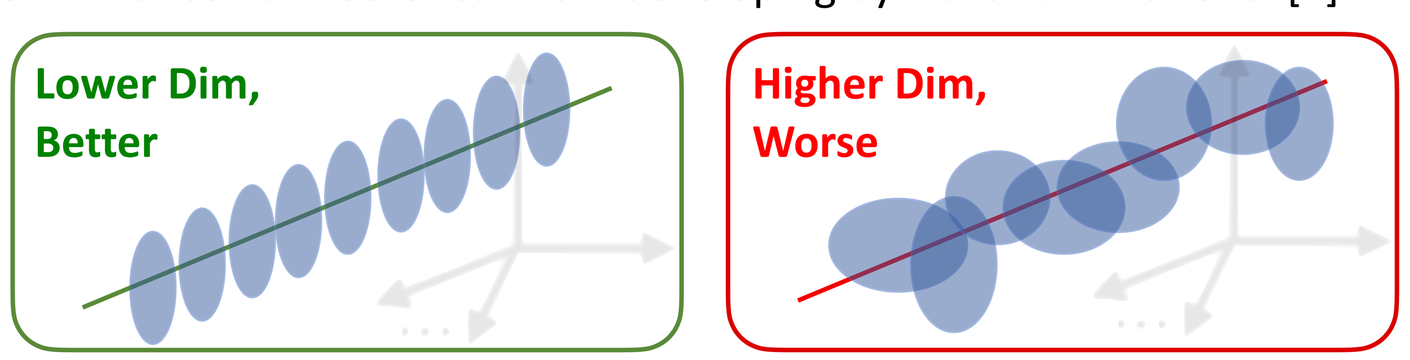

Latent Variation Manifolds: How Are They Constructed?

We are interested in three hyper parameters in a neural manifold with latent structure: Sampled Neurons (N), Manifold Size (M), Number of Manifold Centers (P).

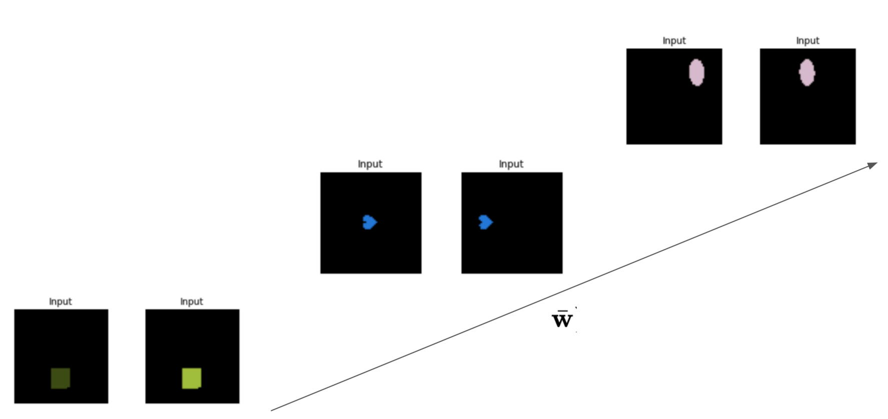

Example Manifold

Varying Latent: Y-Pos

Manifold Size: 2

Manifold Centers: 3

Noise Latent(s): X-Pos, Color

Random Projection: N neurons are sampled by random projections from a higher dimensional representations in the original neural manifold. (Right)

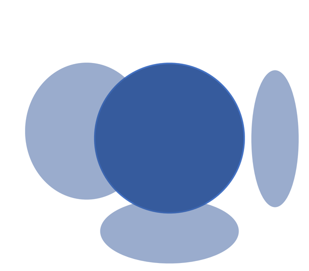

A Computational Objective: Regression Performance

Q: How do we measure the computational performance for regression?

Method: Find measures that capture different facets (ML and biological) of regression performance while connect to geometric theories that our group is developing [4].

Regression Decoding Precision (ε^*)

How well could we linearly decode the latent variable from neurons.

Define ε^* as the decoding precision, where ε^* is the minimum radius of the epsilon tube such that the following optimization problem has a solution.

Regression Storage Capacity (α)

How densely could neurons store a continuous latent variable.

Define α as the number of manifold centers per sampled neuron from the neural representation with a fixed ε^*.
$\alpha = P / N$

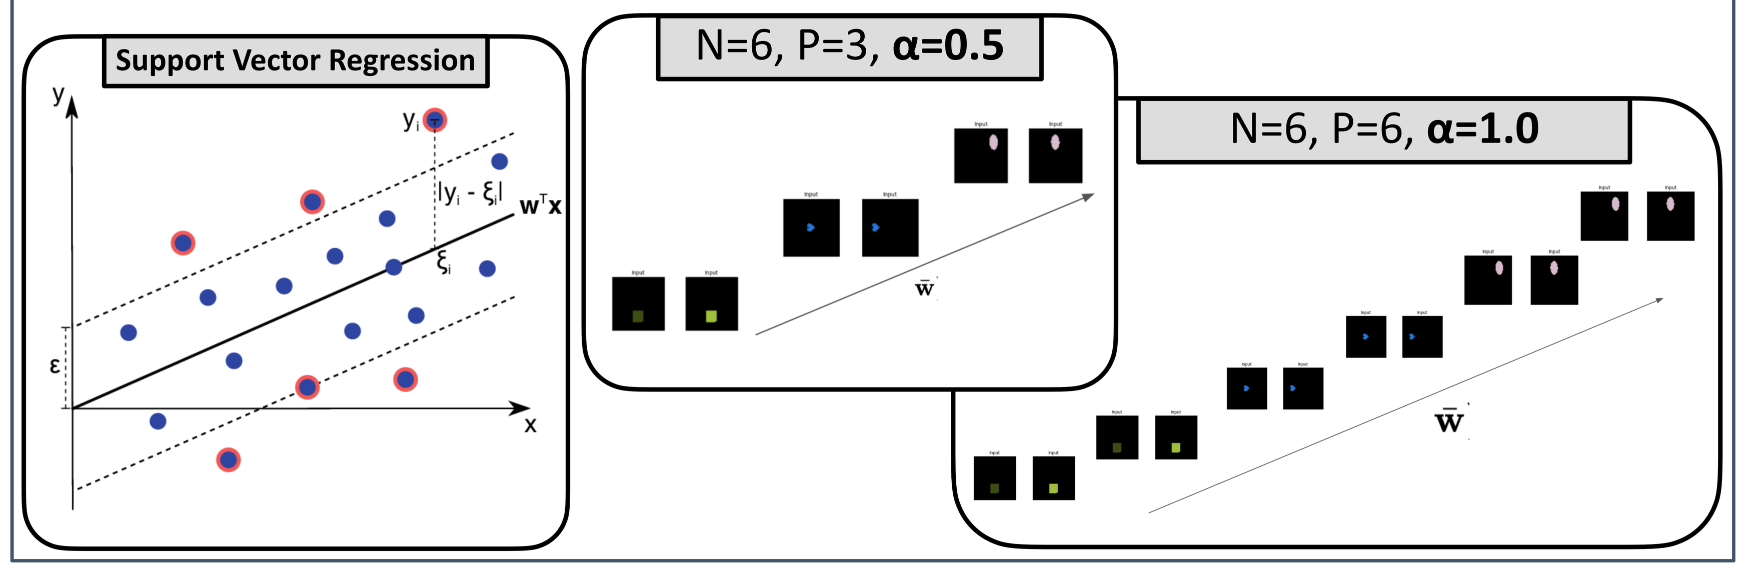

Datasets and Networks

DSprites Vision Dataset: 3 different 2D shapes. In each image, the object is associated with 5 parameters: color, size, position (x, y), and rotation. Representation extracted from ResNet-50 trained on ImageNet.

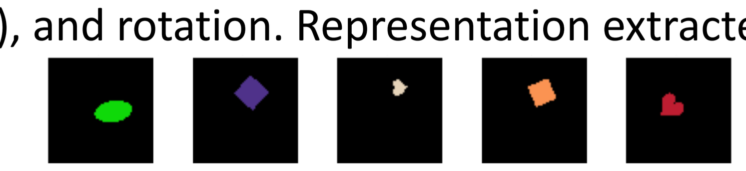

Inverse Autoencoder Dataset: 6 dimensional latent vectors are encoded into 1024 dimensional representations, which is then decoded back to the original latent vectors. Neural manifold variance originates from non-varying latent dimensions: they are randomly sampled.

Preliminary Investigations: Decoding and Storage

Q: Are the two proposed measures of regression capturing our geometric intuition?

Setup: For each fixed storage load (α) , we compute the evolution of decoding precision (ε^*) while inflating sampled neurons (N) and manifold centers (P).

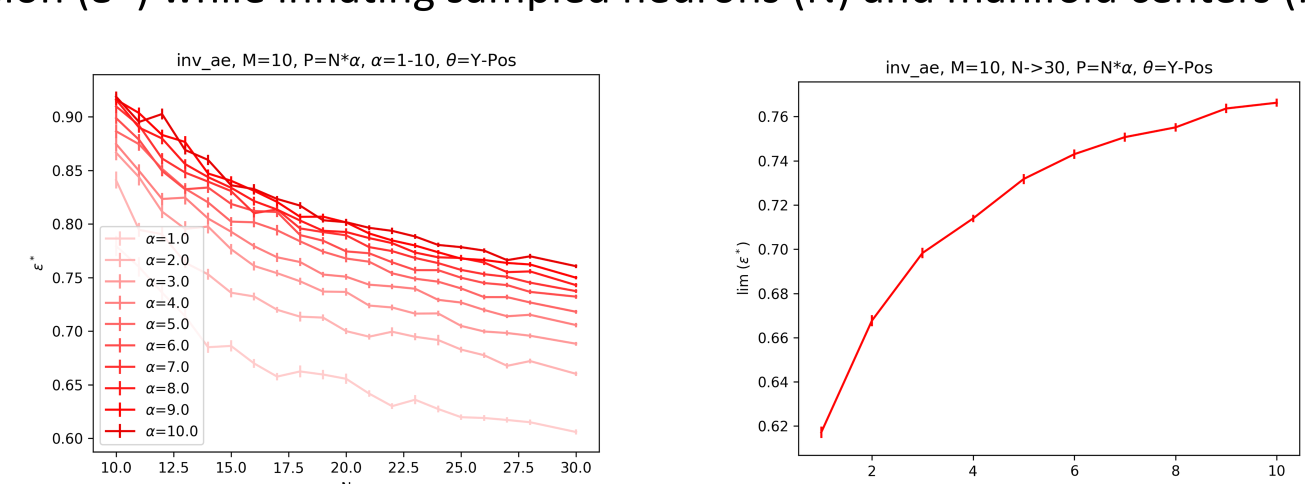

Observation: For a fixed load (P/N ratio), decoding precision is self-averaging in the thermodynamic limit, indicating a good observable.

Experiments: Connection Between Geometry and Regression

Q: How does manifold dimension affect regression performance?

Setup: We alter the number of noise latents to form manifolds with different latent dimensions.

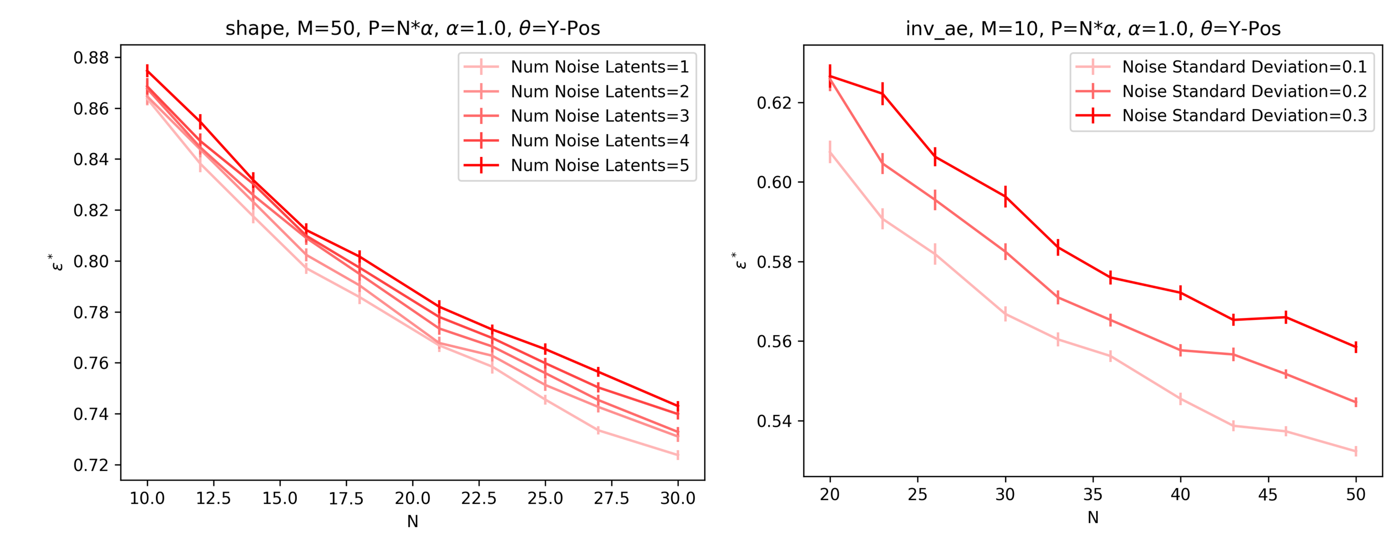

Observation: Higher manifold dimension degrades decoding precision given a fixed load (P/N ratio).

Potential Applications & Future Directions

Neuroscience of Auditory Deficit

- Hearing is naturally a process involved with recognizing continuous variables: frequency and amplitude.
- Patients with auditory deficit is hard at differentiating within and across these continuous variables.
- Regression capacity sheds light on the computational source of auditory deficit.

Geometry and Reasoning

- Relation between objects is encoded by linear combination of latent variables.
- Interpolation between stimuli yields a neural relation manifold.
- Geometry of these relation manifolds gives insights to the computational capacity of analogy, disentanglement, generalization and other tasks involving reasoning/abstraction.

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, and thank Abdul Canatar, Juspreet Singh Sandhu for helpful discussion, and thank the Simons Foundation for this internship opportunity.

Main References

[1] Kriegeskorte, Nikolaus, and Xue-Xin Wei. "Neural tuning and representational geometry." Nature Reviews Neuroscience22.11 (2021): 703-718.

[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.

[4] Abdul Canatar, Chi-Ning Chou, and SueYeon Chung, in preparation, 2023.