Paper Review: The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep[1]

Catherine Zeng

Preface

The goal of reading this paper was learn about what patterns formed from distributed representations in neural circuits tell us about the circuit mechanism and latent encoded variables.

They looked at the manifold structure in the mammalian head direction circuit and noted that the states form a 1D ring that exhibits isometry and is invariant across waking and rapid eye movement sleep by running an experiment that allows the animal to do at least five dimensions of variations (i.e. location, linear speed, angular speed and orientation) and evaluating the manifolds during sleep states (not constrained by low-dimensional inputs from the world).

Links

- Dataset used: `http://crcns.org/data-sets/thalamus/th-1`
- SPUD code: `https://fietelabmit.files.wordpress.com/2019/08/spud_code.zip`

The mammalian head-direction (HD) system

The HD system uses external and internal cues to estimate the direction that the animal is heading with respect to the external world.

It has been long hypothesized that the HD circuit forms a stable, one-dimensional ring-shaped manifold in the high-dimensional population activity state space. In this paper, they show that this structure can be directional visualized in the population response manifold of the HD system.

Spline parameterization for unsupervised decoding (SPUD)

Background and definitions

Conventional dimensionality reduction methods assume the manifold is topologically trivial, but can fail to parametrize latent variables when the manifold is nontrivial. Persistent homology - they start by blurring the point cloud of data at different resolutions. Then they examine the emergence of connected groups of data points. A Spline is a special function defined piecewise by polynomials.

Steps for decoding internal state encoded by manifold:

1. Consider binned spiking data as points in a high-dimensional state space
2. Determine the topology of the point cloud using persistent homology (for measuring topological features of shapes and of functions).
3. Estimate intrinsic manifold dimension.
4. Fit the manifold with a spline of matching topology and dimension.
5. Parameterize the spline by smoothly changing variable of matching dimension and topology.
6. Given a population state, decode that state by projecting it to the nearest point on the spline.

Experiment

Applied method from anterodorsal thalamic nucleus (ADn) of mice. Observed that network states during waking exploration form a low-dimensional, nonlinear, manifold in the form of a convoluted ring. There was no evidence of a toroidal or more complex topological structure.

They fitted a nonlinear spline with the same topology to the manifold.

The LVE closely matches the directly measured head angle. **The match is a direct validation of the hypothesis that the topology of neural representations should match the topology of the represented variables.** Do the neurons encode additional variables?

Key predictions

1. The high-dimensional network response occupies a low-dimensional continuum of states with a dimension and topology matching the encoded variable (or variables).

2. There is isometry of encoded state intervals so that equal velocity inputs produce equal changes in the encoded state, regardless of the starting state.

3. States are autonomously generated and stabilized, and capable of self-sustained activation when sensory inputs are removed.

4. Fit the manifold with a spline of matching topology and dimension.The manifold is an attractor, whereby states initialized away from the manifold rapidly flow back.

5. Manifold states are energetically equal, with no net flow along the manifold.

6. A velocity input, encoding the time-derivative of the represented variable, drives the circuit in a special direction in the high-dimensional state space,

Their results support properties 1 and 2, which are not enough to establishing continuous attractor dynamics so they tested the other predictions further.

Manifold states are equivalent

Examined fluxes of states on and off the manifold. The flux through a region is the average over all trajectories that flow into and out of that region. For a continuous attractor not driven by directional input, they should expect zero net fluxes along the manifold (prediction 5). States not on the manifold should exhibit large net fluxes (property 4).

They discovered that net fluxes were larger at off-manifold states (with an increase in both radial and tangential components). The off-manifold fluxes were directed preferentially towards the ring showing that population states are attracted toward the manifold as predicted by attractor models.

Diffusive dynamics along the manifold during REM

Here, they gathered estimations of influence of noise on the circuit (noise is an important consideration for integrator, memory and representational circuits). In in high dimensions, independent noise is almost entirely directed off-manifold and does not move the system much along the manifold, which is consistent with prediction 5.

Evidence of input aligned to the manifold

Compare observed diffusion in neural circuits with theoretical predictions. The goal is to resolve the nature of noise driving diffusivity during REM.

These suggest that the network receives an input that is aligned to the manifold and is of the right amplitude for moving the state around the ring in response to waking head movements, therefore supporting prediction 6.

Higher-dimensional manifold and coherent dynamics in nREM sleep

Hippocampal circuits replay waking activity patterns during nREM sleep, and replays might be important for memory consolidation. nREM sleep is also described as disrupting the ability of the brain to maintain integrated representations.

They examined manifold structure and dynamics in the ADn during nREM sleep and found that the manifold is higher-dimensional and forms a conical surface.

The nREM manifold encodes at least two latent variables, which they decoded using SPUD. The LVE along the tangential direction represents an angular variable. The radial LVE encodes population firing rate. This result support that the angular structure of the manifold still represents the HD variable.

Conclusion

They showed an image of a 1D ring of activity states in the vertebrate HD circuit. This visualization parallels previous results for a topographically ordered HD ring in the invertebrate nervous system. This structure is hard to obtained from individual neural responses and pairwise correlations and shows a attractor that holds the circuits response amplitude fixed across waking and REM, with this constraint lifted in nREM.

Their innovation is to parameterize the manifold with splines of matching topology and to use this parameterization to decode latent variables, characterize manifold dynamics in waking and sleep, and test theoretical models. These illustrate the process of unsupervised decoding from topologically nontrivial population manifolds, and the manifold approach is useful for discovering unknown variables encoded in high-level brain areas and examining how structured states and dynamics emerge in neural circuits.

References

[1] R. Chaudhuri, B. Gercek, B. Pandey, A. Peyrache, and I. Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, pages 1–9, 08 2019. doi: 10.1038/s41593-019-0460-x.