

Bridging large-scale neuronal recordings and large-scale network models using dimensionality reduction

Ryan C. Williamson^{1,2,3}, Brent Doiron^{1,4}, Matthew A. Smith^{1,5,6}, Byron M. Yu^{1,7,8}

1: Center for the Neural Basis of Cognition, Pittsburgh, PA, USA.

2: Department of Machine Learning, Carnegie Mellon University, Pittsburgh, PA, USA.

3: School of Medicine, University of Pittsburgh, Pittsburgh, PA, USA.

4: Department of Mathematics, University of Pittsburgh, Pittsburgh, PA, USA.

5: Department of Ophthalmology, University of Pittsburgh, Pittsburgh, PA, USA.

6: Department of Bioengineering, University of Pittsburgh, Pittsburgh, PA, USA.

7: Department of Electrical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA.

8: Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA.

Highlights

- The interplay of neuronal recordings and network models is becoming ever stronger.
- Population activity structure provides common ground for incisive comparisons.
- Dimensionality reduction can be used to identify population activity structure.
- This approach is used to study working memory, decision making, motor control, etc.

Abstract

A long-standing goal in neuroscience has been to bring together neuronal recordings and neural network modeling to understand brain function. Neuronal recordings can inform the development of network models, and network models can in turn provide predictions for subsequent experiments. Traditionally, neuronal recordings and network models have been related using single-neuron and pairwise spike train statistics. We review here recent studies that have begun to relate neuronal recordings and network models based on the multi-dimensional structure of neuronal population activity, as identified using dimensionality reduction. This approach has been used to study working memory, decision making, motor control, and more. Dimensionality reduction has provided common ground for incisive comparisons and tight interplay between neuronal recordings and network models.

Introduction

For decades, the fields of experimental neuroscience and neural network modeling proceeded largely in parallel. Whereas experimental neuroscience focused on understanding how the activities of individual neurons relate to sensory stimuli and behavior, the modeling community sought to understand theoretically how neural networks can give rise to brain function. In recent years, developments in neuronal recording technology have enabled the simultaneous recording of hundreds of neurons or more [1]. Concurrently, increases in computational power have enabled the simulation of large neural networks [2]. Together, these developments should enable experimental data to more stringently constrain network model design and network models to better predict neuronal activity for subsequent experiments [3, 4].

A key question is how to relate large-scale neuronal recordings with large-scale network models. Network models typically do not attempt to replicate the precise anatomical connectivity of the biological network from which the neurons are recorded, since the underlying anatomical connectivity is usually unknown (although technological developments are making this possible [5]). In such settings, there is not a one-to-one correspondence of each recorded neuron with a model neuron. To date, comparisons between recordings and models have primarily relied on aggregate spike train statistics based on single neurons (e.g., distribution of firing rates [6], distribution of tuning preferences [7], and Fano factor [8]) and pairs of neurons (e.g., spike time [9] and spike count correlations [10, 11]), as well as single-neuron activity time courses [12, 13]. To go beyond single-neuron and pairwise statistics, recent studies have examined the multi-dimensional structure of neuronal population activity to uncover important insights into mechanisms underlying neuronal computation (e.g., [14–24]). This has motivated the inquiry of whether network models reproduce such population activity structure, in addition to single-neuron and pairwise statistics, raising the bar on what constitutes an agreement between a network model and neuronal recordings [3].

Population activity structure can be characterized using dimensionality reduction [25–27], which provides a concise summary (i.e., a low-dimensional representation) of how a population of neurons covaries and how their activities unfold over time. Several dimensionality reduction methods have been applied to neuronal population activity, including principal component analysis (e.g., [14, 15, 20, 28, 29]), demixed principal component analysis [30], factor analysis [16, 19, 31], Gaussian-process factor analysis [32], latent factor analysis via dynamical systems [33], tensor component analysis [34], and more (see [25] for a review). The




Fig 1. Relating biological and model networks using population analyses: Because a model network typically does not attempt to replicate the precise anatomical connectivity of a biological network, there is not a one-to-one correspondence of each biological neuron with a model neuron. Dimensionality reduction can be used to obtain a concise summary of the population activity from each network. This provides common ground for incisive comparisons between biological and model networks. Discrepancies in the population activity structure between biological and model networks can then help to refine model networks.

low-dimensional representation describes a neuronal process being carried out by the larger circuit from which the neurons were recorded [32, 35]. The same dimensionality reduction method can be applied to the recorded activity and to the network model activity, resulting in population activity structures that can be directly compared (Fig. 1). This benefit is also true of related methods for comparing neuronal recordings and network models involving neuronal decoding, population response similarity, and predicting the activity of one neuron from a population of other neurons [3].

Dimensionality reduction has been adopted by recent studies to relate neuronal recordings and network models to study working memory, decision making, motor control, and more. Although many studies have separately employed large-scale neuronal recordings, large-scale network models, and dimensionality reduction, this review focuses on studies that incorporate all three components. Below we describe these studies, organized by the aspect of population activity structure used to relate neuronal recordings and network models: population activity time courses, functionally-defined neuronal subspaces, and population-wide neuronal variability. These were chosen first because they represent the key ways in which dimensionality reduction has been used in the literature to relate population recordings and network models. More importantly, these three categories represent fundamental aspects of population activity structure – how it unfolds over time, how different types of information can be encoded in different subspaces, and how it varies from trial to trial.

Population activity time courses

Dynamical structures, such as point attractors, line attractors, and limit cycles, arising from network models have long been hypothesized to underlie the computational ability of biological

networks of neurons [36–38]. Such dynamical structures have been implicated in decision making [39,40], memory [41–43], oculomotor integration [44,45], motor control [46], olfaction [47], and more. A fundamental question in systems neuroscience is whether these dynamical structures are actually used by the brain. Although single-neuron and pairwise metrics can be informative [42,45], analyzing the activity of a population of neurons together has enabled deeper connections. In particular, the time course of the activity of a population of neurons can be summarized by low-dimensional neuronal trajectories [25], as identified by dimensionality reduction. These neuronal trajectories can provide a signature of a particular dynamical structure. For example, a point attractor shows convergent trajectories. The neuronal trajectories extracted from the recorded activity can then be compared with those extracted from the network model activity. Such a comparison does not require a one-to-one correspondence between each recorded neuron and a model neuron, but instead relies on a summary of the population activity time courses.

This approach was recently used to study how the brain flexibly controls the timing of behavior [48,49]. By applying dimensionality reduction to neuronal activity recorded from medial frontal cortex, Wang et al. found that population activity time courses for different time intervals followed a stereotypical path, but traversed that path at different speeds (Fig. 2A, top). To understand how a network of neurons can accomplish this, the authors trained a recurrent network model with 200 neurons to produce only the appropriate stimulus-behavior relationships. Wang et al. then applied dimensionality reduction to the activity from the network model. They surprisingly observed that the neuronal trajectories of the network model also followed a stereotypical path, even though the network model was not trained to reproduce the recorded activity (Fig. 2A, bottom). This population-level correspondence enabled by dimensionality reduction laid the foundation for them to then dissect the network model to understand the core neuronal mechanisms [50]. They found that the input to the network drove the network activity from one fixed point to another, where the transition speed was determined by the depth of the energy basin created by the input (Fig. 2A, bottom).

Other studies have also used this approach to understand how the time course of neuronal activity relates to computations underlying motor control [12,51–53], decision making [17,54–56], and working memory [13,57]. In each of these studies, a network model was constructed without referencing the recorded activity. Dimensionality reduction was applied to extract neuronal trajectories to obtain a correspondence between the neuronal recordings and network models. To study the neuronal mechanisms underlying the observed time courses, the network models were then dissected to reveal dynamical structures, such as fixed points or point attractors [17,54,55], line attractors [17], and oscillatory modes [12,51–53]. Whether or not these dynamical structures are indeed at play in real neuronal networks is still an open question. Nevertheless, these studies are beginning to demonstrate that it is at least fruitful to interpret neuronal activity in terms of these dynamical structures, a process facilitated by dimensionality reduction.

Functionally-defined neuronal subspaces

Recent studies have investigated how distinct types of information encoded by the *same* neuronal population can be parsed by downstream brain circuits [58–60]. An enticing proposal is that different types of information are encoded in different subspaces within the population activity space, where the subspaces are identified using dimensionality reduction. For example,




Fig 2. Examples of comparing neuronal recordings and network models using dimensionality reduction. A. (Top) Population activity time courses from medial frontal cortex during a time production task. Each trajectory represents a time course of neuronal activity during a different produced time interval. Circles represent the start and end of the time production interval and diamonds represent a fixed time interval after the start circle. Diamonds appear closer to the start of the trajectory on long-interval trials (blue) than short-interval trials (red), indicating that neuronal activity traverses the path at different speeds during the two intervals. (Bottom) Circles represent fixed points in the model network's dynamics and diamonds represent a fixed time interval after the start of the time production task. A similar difference in traversal speed is observed in the model network as was observed in the neural recordings. Adapted with permission from [48]. B. (Top) Delay period activity from prefrontal cortex during a delayed saccade task. Each trajectory represents a different stimulus condition. The trajectories for different stimuli remain well-separated in a stimulus subspace throughout the delay period. (Bottom) Network model activity demonstrating similar subspace stability. Adapted with permission from [20]. C. (Top) Dimensionality of population-wide neuronal variability in primary visual cortex increases with the number of neurons recorded. (Bottom) A similar dimensionality trend is observed for a spiking network model with clustered excitatory connections (blue), but not for a model with unstructured connectivity (red). Adapted with permission from [31].

Kaufman et al. [18] asked how it is possible for neurons in the motor cortex to be active during motor preparation, yet not generate an arm movement. They found that motor cortical activity during motor preparation resided outside of the activity subspace most related to muscle

contractions. This allows the motor cortex to prepare arm movements without driving
downstream circuits, a characteristic which can be implemented by a linear readout mechanism.
This concept of functionally-defined neuronal subspaces has also been used in other studies of
motor control [23, 61–63], decision making [30, 64], short-term memory [30, 65], learning [19],
and visual processing [24].

To understand how a neuronal circuit can implement and exploit such functionally-defined
neuronal subspaces, one can construct a network model to see whether it reproduces the
empirical observations. If so, one can then dissect the network to study the underlying
mechanisms. Mante et al. [17] applied dimensionality reduction to recordings in prefrontal
cortex to find that motion and color of the visual stimulus were encoded in distinct subspaces.
They then trained a recurrent network model with 100 neurons to produce only the appropriate
stimulus-behavior relationships. When they applied the same dimensionality reduction method
to the network model activity, they surprisingly found that the motion and color of the visual
stimulus were also encoded in distinct subspaces, even though the network model was not
trained to reproduce the recorded activity. This population-level correspondence between the
network model and recordings was enabled by dimensionality reduction and went beyond
comparisons based on individual neurons or pairs of neurons. Mante et al. then dissected the
network model to uncover how the two types of information encoded in distinct subspaces can
be selectively used to form a decision.

Dimensionality reduction has also revealed that, in some cases, standard network models do
not reproduce the functionally-defined subspaces identified from neuronal recordings. For
example, Murray et al. [20] applied dimensionality reduction to recordings in prefrontal cortex
during a working memory task to find that, even though firing rates of individual neurons
changed over time, there was a subspace in which the activity stably encoded the memorized
target location (Fig. 2B, top). They then applied the same analyses to activity from several
prominent network models and found that none of them reproduced both the time-varying
activity of individual neurons *and* the subspace in which the memory was stably encoded. This
provided the impetus to develop a new network model that did reproduce these features of the
recorded activity (Fig. 2B, bottom) (see also [66]). As another example, Elsayed et al. [67]
found that standard network models do not reproduce the empirical observation described
above that neuronal activity during movement preparation and movement execution lie in
orthogonal subspaces. Such insights obtained using dimensionality reduction can guide the
development of more sophisticated network models.

Population-wide neuronal variability

The previous sections focus largely on neuronal activity that is averaged across trials and on
firing rate-based network models. This naturally obscures the trial-to-trial variability that is a
fundamental feature of neuronal responses across the cortex [68], both at the level of single
neuron responses [69] as well as variability shared by the population [11, 70]. Theoretical and
experimental studies have focused on how the structure of that variability places limits on
information coding [71–74], and in turn influences our behavior. At the same time, a growing
body of work has demonstrated that variability can be thought of not only as noise to be
removed, but also a signature of ongoing decision processes and cognitive variables
(e.g., [75–77]). To move beyond single-neuron and pairwise measurements of neuronal
variability, recent studies have begun to consider population-wide measures of neuronal

variability [78–82], as enabled by dimensionality reduction. Such measures allow one to i) assess whether the large number of single-neuron and pairwise variability measurements can be succinctly summarized by a small number of variables (e.g., the entire population increasing and decreasing its activity together can be described by a single scalar variable), and ii) relate the population activity on individual experimental trials to behavior [22, 32–34, 83–85].

In parallel with the growing interest in neuronal variability, there have been attempts to create network models that exhibit variability matching recorded neurons. In particular, a class of models has used the balance between excitation and inhibition as a way to generate variability as an emergent property of network structure, rather than via an external variable source [71, 86, 87]. In these models, the particular structure of the network has a large impact on the population-wide variability that emerges. Using the lens of factor analysis, Williamson et al. [31] found that the dimensionality of spontaneous activity fluctuations in V1 neurons increases with the number of recorded neurons (Fig. 2C, top). This was more consistent with activity generated by networks with clustered excitatory connections [8] than networks with unstructured connectivity [86] (Fig. 2C, bottom). The combination of population-wide measures of variability (in this case, dimensionality) and the ability to manipulate model network structures facilitated an understanding of how features of variability observed in biological networks relate to network structure.

The approach of using dimensionality reduction to compare the population-wide variability of neuronal recordings and network models has also been applied to study spontaneous versus evoked activity [88, 89], the activity of different classes of neurons [90], and the activity during different behavioral conditions, such as attention [81, 82]. Dimensionality reduction has also been used to analyze population activity from balanced network models to help identify the crucial network architecture and synaptic timescales required to produce the low-dimensional shared variability that is widely reported in neuronal recordings [82, 87]. Together these studies demonstrate the power of combining dimensionality reduction and network models to understand the mechanisms and effects of neuronal variability.

Conclusion

Dimensionality reduction has enabled incisive comparisons between biological and model networks in terms of population activity time courses, functionally-defined neuronal subspaces, and population-wide neuronal variability. Such comparisons result in either i) a correspondence between the neuronal recordings and the network model, in which case the model can be dissected to understand underlying network mechanisms, or ii) discrepancies between the neuronal recordings and standard network models, leading to the development of improved models. This approach (cf. Fig. 1) has already provided insight into the neuronal mechanisms underlying brain functions such as working memory, decision making, and motor control, and is likely to become even more important as the scale of neuronal recordings and network models grows.

A key consideration in network modeling is what aspects of neuronal recordings the model should reproduce. We posit that the population activity structure (including population activity time courses, functionally-defined neuronal subspaces, and population-wide neuronal variability) will provide key signatures of how neurons work together to give rise to brain function. Thus, if a network model is to provide a systems-level account of brain function, we should require it to reproduce the population activity structure of neuronal recordings, in

addition to existing population metrics [3] and standard spike train statistics based on individual and pairs of neurons.

Most studies described here have used neuronal recordings to inform network models via dimensionality reduction. An important future direction is to use network models and dimensionality reduction to design new experiments and form predictions. For example, if one day we can experimentally perturb neuronal activity in specified directions in the population activity space [91], we can test whether driving the population activity in particular directions leads to particular decisions or movements predicted by the network model. The hope is to establish a virtuous cycle, where neuronal recordings and network models closely inform each other through the common ground provided by dimensionality reduction.

Acknowledgements

RCW was supported by a Richard King Mellon Foundation Presidential Fellowship in the Life Sciences. BD was supported by NIH CRCNS R01DC015139, ONR N00014-18-1-2002, and Simons Foundation 542967. BD and MAS were supported by a grant from the Hillman Foundation. MAS was supported by NIH R01EY022928 and P30EY008098, NSF NCS BCS 1734901, an unrestricted award from Research to Prevent Blindness, and the Eye and Ear Foundation of Pittsburgh. BMY was supported by NSF NCS BCS 1734916, NSF NCS BCS 1533672, NIH R01HD071686, NIH CRCNS R01NS105318, and Simons Foundation 364994 and 543065.

References

1. Stevenson, I.H. and Kording, K.P. (2011). How advances in neural recording affect data analysis. *Nature Neuroscience* *14*, 139–142.
2. Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J.M., Diesmann, M., Morrison, A., Goodman, P.H., Harris, F.C., et al. (2007). Simulation of networks of spiking neurons: a review of tools and strategies. *Journal of Computational Neuroscience* *23*, 349–398.
3. Yamins, D.L.K. and DiCarlo, J.J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience* *19*, 356–365.
4. Barak, O. (2017). Recurrent neural networks as versatile tools of neuroscience research. *Current Opinion in Neurobiology* *46*, 1–6.
5. Lee, W.C.A., Bonin, V., Reed, M., Graham, B.J., Hood, G., Glattfelder, K., and Reid, R.C. (2016). Anatomy and function of an excitatory network in the visual cortex. *Nature* *532*, 370 EP –.
6. Roxin, A., Brunel, N., Hansel, D., Mongillo, G., and van Vreeswijk, C. (2011). On the distribution of firing rates in networks of cortical neurons. *Journal of Neuroscience* *31*, 16217–16226.
7. Chariker, L., Shapley, R., and Young, L.S. (2016). Orientation selectivity from very sparse LGN inputs in a comprehensive model of macaque V1 cortex. *Journal of Neuroscience* *36*, 12368–12384.

8. Litwin-Kumar, A. and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience* *15*, 1498–1505. 240
241
9. Trousdale, J., Hu, Y., Shea-Brown, E., and Josić, K. (2012). Impact of network structure and cellular response on spike time correlations. *PLoS computational biology* *8*, e1002408. 242
243
10. Stringer, C., Pachitariu, M., Steinmetz, N.A., Okun, M., Bartho, P., Harris, K.D., Sahani, M., and Lesica, N.A. (2016). Inhibitory control of correlated intrinsic variability in cortical networks. *eLife* *5*, e19695. 244
245
246
11. Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G.K., and Josić, K. (2016). The mechanics of state-dependent neural correlations. *Nature neuroscience* *19*, 383–393. 247
248
12. Sussillo, D., Churchland, M.M., Kaufman, M.T., and Shenoy, K.V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience* *18*, 1025–1033. 249
250
251
13. Rajan, K., Harvey, C.D., and Tank, D.W. (2016). Recurrent network models of sequence generation and memory. *Neuron* *90*, 128–142. 252
253
14. Mazor, O. and Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* *48*, 661–673. 254
255
15. Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyujukian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. *Nature* *487*, 51–56. 256
257
258
16. Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* *484*, 62–68. 259
260
17. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* *503*, 78–84. 261
262
18. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience* *17*, 440–448. 263
264
265
19. Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. *Nature* *512*, 423–426. 266
267
20. Murray, J.D., Bernacchia, A., Roy, N.A., Constantinidis, C., Romo, R., and Wang, X.J. (2017). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proceedings of the National Academy of Sciences* *114*, 394–399. 268
269
270
271
21. Remington, E.D., Egger, S.W., Narain, D., Wang, J., and Jazayeri, M. (2018). A dynamical systems perspective on flexible motor timing. *Trends in Cognitive Sciences* *22*, 938–952. 272
273
274
22. Ruff, D.A., Ni, A.M., and Cohen, M.R. (2018). Cognition as a window into neuronal population space. *Annual Review of Neuroscience* *41*, 77–97. 275
276

23. Perich, M.G., Gallego, J.A., and Miller, L.E. (2018). A neural population mechanism for rapid learning. *Neuron* *100*, 964–976.e7. 277
278
24. Semedo, J.D., Zandvakili, A., Machens, C.K., Yu, B.M., and Kohn, A. (2018). Cortical areas interact through a communication subspace. *Neuron* , in press. 279
280
25. Cunningham, J.P. and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience* *17*, 1500–1509. 281
282
26. Gao, P. and Ganguli, S. (2015). On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology* *32*, 148–155. 283
284
27. Gallego, J.A., Perich, M.G., Miller, L.E., and Solla, S.A. (2017). Neural manifolds for the control of movement. *Neuron* *94*, 978 – 984. 285
286
28. Cowley, B.R., Smith, M.A., Kohn, A., and Yu, B.M. (2016). Stimulus-driven population activity patterns in macaque primary visual cortex. *PLoS Computational Biology* *12*, e1005185. 287
288
289
29. Gallego, J.A., Perich, M.G., Naufel, S.N., Ethier, C., Solla, S.A., and Miller, L.E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications* *9*, 4233. 290
291
292
30. Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C.E., Kepecs, A., Mainen, Z.F., Qi, X.L., Romo, R., Uchida, N., and Machens, C.K. (2016). Demixed principal component analysis of neural population data. *eLife* *5*, e10989. 293
294
295
31. Williamson, R.C., Cowley, B.R., Litwin-Kumar, A., Doiron, B., Kohn, A., Smith, M.A., and Yu, B.M. (2016). Scaling properties of dimensionality reduction for neural populations and network models. *PLoS Computational Biology* *12*, e1005141. 296
297
298
32. Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology* *102*, 614–635. 299
300
301
33. Pandarinath, C., O’Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods* *15*, 805–815. 302
303
304
305
34. Williams, A.H., Kim, T.H., Wang, F., Vyas, S., Ryu, S.I., Shenoy, K.V., Schnitzer, M., Kolda, T.G., and Ganguli, S. (2018). Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis. *Neuron* *98*, 1099–1115.e8. 306
307
308
309
35. Buonomano, D.V. and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience* *10*, 113 EP –. Review Article. 310
311
312
36. Wilson, H.R. and Cowan, J.D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik* *13*, 55–80. 313
314

37. Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences* *79*, 2554–2558. 315–316
38. Brody, C.D., Romo, R., and Kepecs, A. (2003). Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Current Opinion in Neurobiology* *13*, 204–211. 317–319
39. Machens, C.K., Romo, R., and Brody, C.D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* *307*, 1121–1124. 320–321
40. Wang, X.J. (2008). Decision making in recurrent neuronal circuits. *Neuron* *60*, 215–234. 322
41. Wang, X.J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences* *24*, 455–463. 323–324
42. Wimmer, K., Nykamp, D.Q., Constantinidis, C., and Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature Neuroscience* *17*, 431–439. 325–327
43. Chaudhuri, R. and Fiete, I. (2016). Computational principles of memory. *Nature Neuroscience* *19*, 394–403. 328–329
44. Seung, H.S. (1996). How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences* *93*, 13339–13344. 330–331
45. Miri, A., Daie, K., Arrenberg, A.B., Baier, H., Aksay, E., and Tank, D.W. (2011). Spatial gradients and multidimensional dynamics in a neural integrator circuit. *Nature Neuroscience* *14*, 1150–1159. 332–334
46. Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annual Review of Neuroscience* *36*, 337–359. 335–337
47. Rabinovich, M., Huerta, R., and Laurent, G. (2008). Transient dynamics for neural processing. *Science* *321*, 48–50. 338–339
48. Wang, J., Narain, D., Hosseini, E.A., and Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nature Neuroscience* *21*, 102–110. 340–341
49. Remington, E.D., Narain, D., Hosseini, E.A., and Jazayeri, M. (2018). Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* *98*, 1005–1019.e5. 342–344
50. Sussillo, D. and Barak, O. (2013). Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Computation* *25*, 626–649. 345–346
51. Hennequin, G., Vogels, T., and Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* *82*, 1394 – 1406. 347–349
52. Michaels, J.A., Dann, B., and Scherberger, H. (2016). Neural population dynamics during reaching are better explained by a dynamical system than representational tuning. *PLoS Computational Biology* *12*, e1005175. 350–352

53. Russo, A.A., Bittner, S.R., Perkins, S.M., Seely, J.S., London, B.M., Lara, A.H., Miri, A., Marshall, N.J., Kohn, A., Jessell, T.M., et al. (2018). Motor cortex embeds muscle-like commands in an untangled population response. *Neuron* *97*, 953–966.e8. 353–355
54. Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. (2015). Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neuron* *86*, 1067–1077. 356–358
55. Chaisangmongkon, W., Swaminathan, S.K., Freedman, D.J., and Wang, X.J. (2017). Computing by robust transience: how the fronto-parietal network performs sequential, category-based decisions. *Neuron* *93*, 1504–1517.e4. 359–361
56. Mastrogiuseppe, F. and Ostojic, S. (2018). Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron* *99*, 609–623.e29. 362–363
57. Barak, O., Sussillo, D., Romo, R., Tsodyks, M., and Abbott, L.F. (2013). From fixed points to chaos: Three models of delayed discrimination. *Progress in Neurobiology* *103*, 214–222. 364–366
58. Park, I.M., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience* *17*, 1395–1403. 367–369
59. Pagan, M. and Rust, N.C. (2014). Quantifying the signals contained in heterogeneous neural responses and determining their relationships with task performance. *Journal of Neurophysiology* *112*, 1584–1598. 370–372
60. Fusi, S., Miller, E.K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. *Current Opinion in Neurobiology* *37*, 66–74. 373–374
61. Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. *Nature* *532*, 459–464. 375–376
62. Miri, A., Warriner, C.L., Seely, J.S., Elsayed, G.F., Cunningham, J.P., Churchland, M.M., and Jessell, T.M. (2017). Behaviorally selective engagement of short-latency effector pathways by motor cortex. *Neuron* *95*, 683 – 696.e11. 377–379
63. Hennig, J.A., Golub, M.D., Lund, P.J., Sadtler, P.T., Oby, E.R., Quick, K.M., Ryu, S.I., Tyler-Kabara, E.C., Batista, A.P., Yu, B.M., et al. (2018). Constraints on neural redundancy. *eLife* *7*, e36774. 380–382
64. Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience* *17*, 1784–1792. 383–385
65. Daie, K., Goldman, M., and Aksay, E.F. (2015). Spatial patterns of persistent neural activity vary with the behavioral context of short-term memory. *Neuron* *85*, 847–860. 386–387
66. Druckmann, S. and Chklovskii, D. (2012). Neuronal circuits underlying persistent representations despite time varying activity. *Current Biology* *22*, 2095–2103. 388–389

67. Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nature Communications* *7*, 13239. 390-392
68. Renart, A. and Machens, C.K. (2014). Variability in neural activity and behavior. *Current Opinion in Neurobiology* *25*, 211–220. 393-394
69. Faisal, A.A., Selen, L.P., and Wolpert, D.M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience* *9*, 292–303. 395-396
70. Cohen, M.R. and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nature Neuroscience* *14*, 811–819. 397-398
71. Shadlen, M.N. and Newsome, W.T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience* *18*, 3870–3896. 399-401
72. Abbott, L.F. and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation* *11*, 91–101. 402-403
73. Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience* *7*, 358–366. 404-405
74. Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., and Pouget, A. (2014). Information-limiting correlations. *Nature Neuroscience* *17*, 1410–1417. 406-407
75. Cohen, M.R. and Maunsell, J.H.R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience* *12*, 1594–1600. 408-409
76. Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* *63*, 879–888. 410-411
77. Nienborg, H., R. Cohen, M., and Cumming, B.G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annual Review of Neuroscience* *35*, 463–483. 412-414
78. Ecker, A., Berens, P., Cotton, R.J., Subramanian, M., Denfield, G., Cadwell, C., Smirnakis, S., Bethge, M., and Tolias, A. (2014). State dependence of noise correlations in macaque primary visual cortex. *Neuron* *82*, 235–248. 415-417
79. Rabinowitz, N.C., Goris, R.L., Cohen, M., and Simoncelli, E.P. (2015). Attention stabilizes the shared gain of V4 populations. *eLife* *4*, e08998. 418-419
80. Lin, I.C., Okun, M., Carandini, M., and Harris, K.D. (2015). The nature of shared cortical variability. *Neuron* *87*, 644–656. 420-421
81. Kanashiro, T., Ocker, G.K., Cohen, M.R., and Doiron, B. (2017). Attentional modulation of neuronal variability in circuit models of cortex. *eLife* *6*, e23978. 422-423
82. Huang, C., Ruff, D., Pyle, R., Rosenbaum, R., Cohen, M., and Doiron, B.D. (2018). Circuit models of low dimensional shared variability in cortical networks. *Neuron*, in press. 424-425-426

83. Cohen, M.R. and Maunsell, J.H. (2010). A neuronal population measure of attention predicts behavioral performance on individual trials. *Journal of Neuroscience* *30*, 15241–15253. 427
428
429
84. Kiani, R., Cueva, C., Reppas, J., and Newsome, W. (2014). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Current Biology* *24*, 1542–1547. 430
431
432
85. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2015). Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *eLife* *4*, e04677. 433
434
435
86. Van Vreeswijk, C. and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* *274*, 1724–1726. 436
437
87. Rosenbaum, R., Smith, M.A., Kohn, A., Rubin, J.E., and Doiron, B. (2017). The spatial structure of correlated neuronal variability. *Nature Neuroscience* *20*, 107–114. 438
439
88. Mazzucato, L., Fontanini, A., and La Camera, G. (2016). Stimuli reduce the dimensionality of cortical activity. *Frontiers in Systems Neuroscience* *10*. 440
441
89. Hennequin, G., Ahmadian, Y., Rubin, D.B., Lengyel, M., and Miller, K.D. (2018). The dynamical regime of sensory cortex: stable dynamics around a single stimulus-tuned attractor account for patterns of noise variability. *Neuron* *98*, 846–860.e5. 442
443
444
90. Bittner, S.R., Williamson, R.C., Snyder, A.C., Litwin-Kumar, A., Doiron, B., Chase, S.M., Smith, M.A., and Yu, B.M. (2017). Population activity structure of excitatory and inhibitory neurons. *PLoS One* *12*, e0181773. 445
446
447
91. Jazayeri, M. and Afraz, A. (2017). Navigating the neural space in search of the neural code. *Neuron* *93*, 1003–1014. 448
449

Mante et al. [17]*: In a context-dependent decision-making task, the authors found that color and motion information were encoded in distinct subspaces of PFC population activity. A recurrent network model trained to perform the same task revealed a network-level mechanism of how the two types of information can be selectively used to form a decision.

Murray et al. [20]:** The authors analyzed population activity recorded in prefrontal cortex and identified a low-dimensional subspace in which stimulus information was reliably encoded, despite the fact that individual neurons showed substantial time-varying activity. They then found that standard models did not reproduce this empirical observation, and proceeded to develop a “stable subspace” model that did reproduce this observation.

Wang et al. [48]:** Recording from medial frontal cortex during a timing task, the authors found that population activity time courses followed a stereotypical path, but traversed the path at different speeds based on the duration of the timing interval. They found similar trends in a recurrent network model trained to perform the task and showed that speed of traversal was determined by the network inputs.

Sussillo et al. [12]*: The authors trained a recurrent network model to produce muscle activity patterns observed in an arm reaching task. The model activity surprisingly showed rotational dynamics that mimicked those observed empirically in M1 population recordings.

Chaisangmongkan et al. [55]*: The authors found that PFC and LIP neurons show mixed selectivity during a delayed match-to-category task, and that the neuronal trajectories extracted using dimensionality reduction are interpretable during each epoch of the task. They then constructed a recurrent network model to understand the network principles that govern the activity time courses during this task.

Williamson et al. [31]:** This study compared the population activity structure of V1 recordings and spiking network models while varying the number of neurons and trials analyzed. The scaling trends of the V1 recordings better resembled a model with clustered excitatory connections than one with unstructured connectivity.

Elsayed et al. [67]*: This study found that M1 population activity during movement preparation and movement execution resides in orthogonal subspaces. Standard network models did not reproduce this empirical observation.

Mazzucato et al. [88]*: Comparing gustatory cortex recordings and spiking network models, the authors examined how the dimensionality of population activity grows with population size during spontaneous and evoked activity. They then developed a theoretical upper bound on dimensionality based on the level of pairwise correlations.