



Towards the neural population doctrine

Shreya Saxena and John P Cunningham

Across neuroscience, large-scale data recording and population-level analysis methods have experienced explosive growth. While the underlying hardware and computational techniques have been well reviewed, we focus here on the novel science that these technologies have enabled. We detail four areas of the field where the joint analysis of neural populations has significantly furthered our understanding of computation in the brain: correlated variability, decoding, neural dynamics, and artificial neural networks. Together, these findings suggest an exciting trend towards a new era where neural populations are understood to be the essential unit of computation in many brain regions, a classic idea that has been given new life.

Address

Department of Statistics, Grossman Center for the Statistics of Mind Zuckerman Mind, Brain Behavior Institute, Center for Theoretical Neuroscience, Columbia University, New York City, United States

Corresponding author: Cunningham, John P (jpc2181@columbia.edu)

Current Opinion in Neurobiology 2019, **55**:103–111

This review comes from a themed issue on **Machine learning, big data, and neuroscience**

Edited by **Maneesh Sahani** and **Jonathan Pillow**

<https://doi.org/10.1016/j.conb.2019.02.002>

0959-4388/Published by Elsevier Ltd.

Introduction

The neuron has been heralded as the structural and functional unit of neural computation for more than a hundred years [1–3], and this doctrine has driven a vast array of our field's most important discoveries [4]. Increasingly prevalent, however, is the idea that *populations* of neurons may in fact be the essential unit of computation in many brain regions [4–6]. Certainly the idea of information being encoded in an ensemble of neurons is not a new one. In fact, the conflict between single neurons and populations of neurons as the perceptual, behavioral, and cognitive unit of computation dates back to the beginning of the 20th century [4,6], with the first concrete theories of networks of neurons introduced in the 1940s [7,8]. The ideas, while being novel, were not testable due to the technological shortcomings of both recording techniques and computational resources. However, we currently find

ourselves in the ideal era for scientific discovery, given the astounding progresses in both these enabling technologies.

Electrophysiology recordings have been the hallmark of neuronal recordings over the last 80 years — extracellular recordings of one or multiple electrodes, each capturing up to a few neurons. More recently, multi-electrode arrays and imaging techniques (optical, and more recently voltage) have been used to efficiently capture the simultaneous activity of hundreds and thousands of neurons, with this number steadily growing using tools such as the Neuropixel [9]. In tandem, the increase in computational resources has led to the development of efficient and scalable statistical and machine learning methods; see the methodological reviews [10,11].

As our ability to simultaneously record from large populations of neurons is growing exponentially [9,12], the analysis of the covariation of populations of neurons has provided us with scientific insights in many domains. Here, we highlight several recent findings in four domains of neuroscience where the joint analysis of a population of neurons has been central to scientific discovery that would not be possible using single neurons alone. Firstly, trial-to-trial ‘noise’ correlations have been shown to influence the information carrying capacity of a neural circuit. Secondly, decoding of behavior using correlated populations of neurons can yield levels of accuracy beyond what would be anticipated from single neurons alone. Thirdly, the dynamic analysis of stimulus-driven population recordings over time can be projected into a lower dimensional subspace to reveal computational strategies employed by different brain regions. Lastly, artificial neural networks (ANNs) can aid in simulations that reproduce population structure, as well as directly modeling neuronal activity. We focus on the analysis of a population of N neurons in ‘state space’, where each neuron's activity at any time point is represented as a dot in either the N dimensional observation space or in a lower dimensional subspace.

We pinpoint one or two recent studies in each domain that stand out (indicated using * and **). Unlike previous reviews on population-level neuroscience [10,11], we focus here not on the data analysis methodologies, but rather the notable scientific findings that have resulted. These scientific findings are, first, reshaping the way the field thinks about computation, and, second, fundamentally population-based. Taken together, these two features point to a future where the central scientific theme is not the neuron doctrine, but the *neural population*

doctrine. We conclude with topics that we think future studies may need to address.

Correlated trial-to-trial variability in populations of neurons is a key indicator for behavioral states

As we gain the ability to simultaneously record the activity of more and more neurons, we must ask how much information we can hope to achieve by doing so, that is what is the information gained per added neuron? Generally speaking, ‘signal’ correlations, or tuning curves, are useful in terms of decoding as well as understanding the dynamics of neural population over time. However, in the pursuit of this specific question, the systematic study of covariation in the activity of pairs of neurons during repeated presentations of the same stimulus (‘noise’ correlation) has also been a well-studied and particularly tractable analysis tool (Figure 1a).

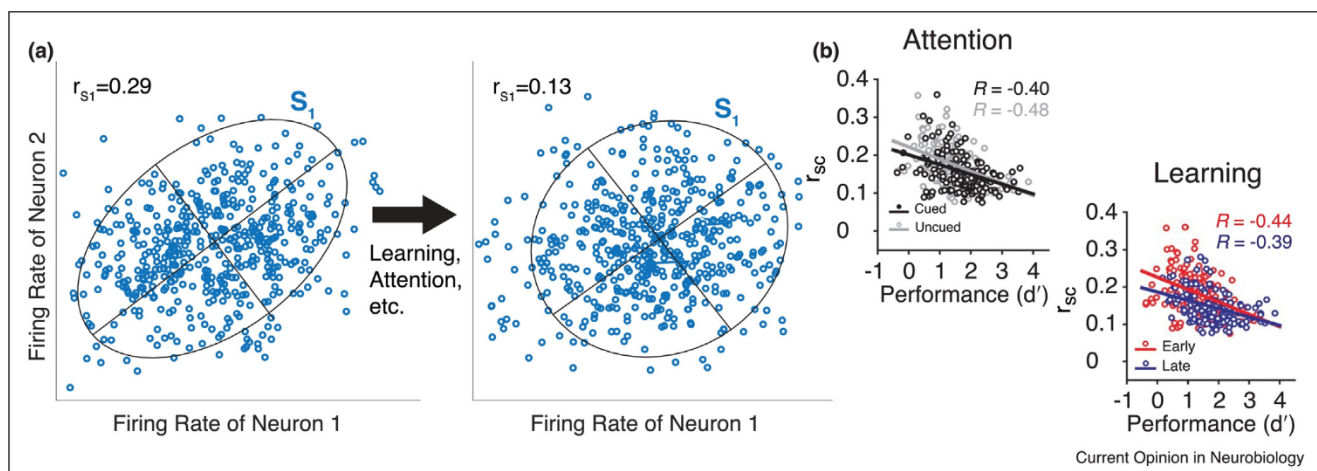
In fact, a tremendous amount of research has led to the fundamental belief that correlations are linked to the information carrying capacity of the brain, which itself has widely been postulated to be directly related to the level of performance of an individual on a given task [13,14]. One metric for the information carrying capacity is the linear Fisher information, which is a measure of the smallest change in the stimulus that can be distinguished with an optimal linear decoder of the neuronal population response. It has been shown that under certain assumptions on the distribution of the neuronal populations, and assuming homogenous tuning curves, the presence of correlations causes the linear Fisher information of the neurons to saturate as $N \rightarrow \infty$ (i.e. a growing number of neurons N) [13,15,16]. In fact, pairwise correlations in

various brain areas are affected by attention, learning, arousal, and adaptation (see [17] for a review).

Recently, in an important study [18**], Ni *et al.* showed that the correlated variability in V4 neurons during attention and learning — processes that have inherently different timescales — robustly decreases (Figure 1b). They also found that the correlated variability is strongly correlated with performance, across learning and attention conditions. In this study, in addition to mean pairwise correlation, the authors demonstrated the use of a truly population-based metric, that is percentage of variance explained in the first principal component (PC) of simultaneously recorded neurons, to further quantify the decrease in correlated variability. Moreover, they showed that a ‘choice’ decoder built on neural activity in the first PC performs as well as one built on the full dataset, suggesting that the relationship of neural variability to behavior lies in a relatively small subspace of the state space. These findings reinforce the viewpoint that the relevant unit of computation may be a subspace of the neural activity. As a next step, it would be interesting to examine whether the decrease in correlated variability between learning and attention happens in the same dimensions, that is do the neurons decrease their variability *in the same way* during both processes. However, this requires the simultaneous recording of a population of neurons over long time periods.

Experimental findings in this field highlight the importance of advancing the theory on correlated variability, especially in the small N regime, in order to better understand the mechanisms that lead to a decrease in correlations and its link to behavior. Going forward, we highlight the need for theoretical contributions based on

Figure 1



(a) Schematic showing the decrease in correlated variability of neurons as represented in ‘state space’: the firing rate of neuron 1 on the x-axis, and the firing rate of neuron 2 on the y-axis, during presentation of the same stimulus S_1 . Each circle represents a trial. (b) The average Pearson’s correlation coefficient robustly decreases with an increase in performance during attention as well as learning. Figure adapted from [18**].

less restrictive assumptions, as well as bringing together theoretical insights and experimental findings to fully address these questions. One important step in this direction was performed by Moreno-Bote *et al.* [14], who showed that in the regime of inhomogeneous tuning curves, if the ‘noise’ correlations are in the direction of the ‘signal’ correlations, that is tuning curves, this leads to a saturation in the linear Fisher information. These noise correlations in the direction of signal correlations are called ‘differential’ correlations. See [17] for an excellent review.

We also highlight the need to utilize analyses that truly benefit from the recordings of a population of neurons. Examining population metrics of correlated variability in neural subspaces of simultaneously recorded neurons (such as the principal component space, as above) is an important next step that will move towards bridging theory and experimental predictions.

Decoding accuracy is more than the sum of its parts

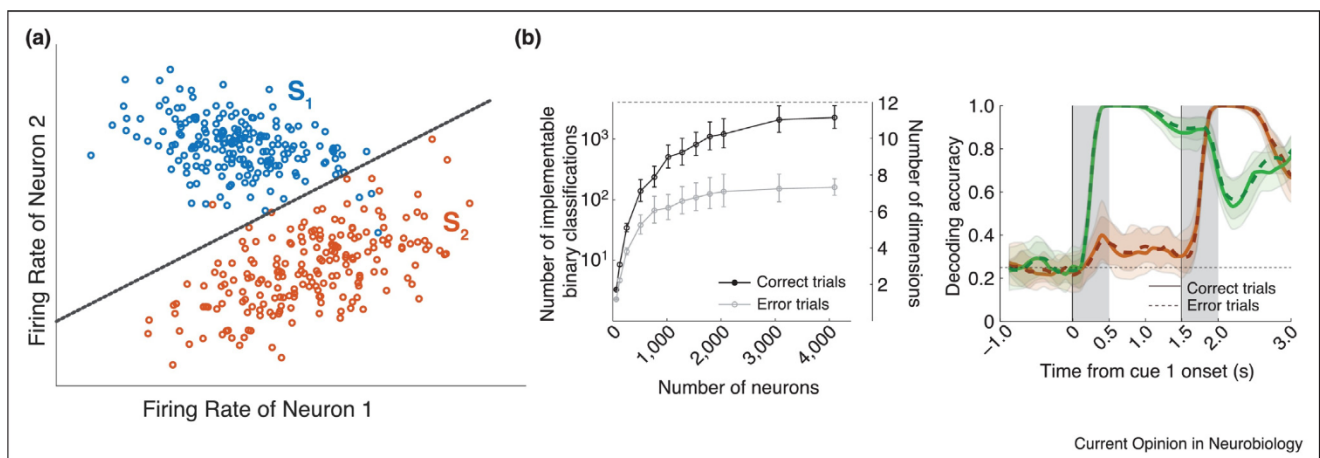
Here we focus on the case where population analysis is virtually essential: when neurons have so-called ‘mixed selectivity’ [5,19**,20]. As an increasing number of experiments involving multiple categories of stimuli are being performed, neurons in several brain areas have been found to have mixed selectivity: neurons significantly modulate their activity in response to more than one stimulus category/type [21,5,19**]. For example, in Figure 2a, we see two neurons encoding for two different stimuli; decoding stimulus identity from either of the two neurons alone would be suboptimal. However, if both

neurons are recorded, the two stimuli are linearly separable.

It is well known that higher dimensional data can be more easily linearly separated for classification purposes (the ‘kernel trick’ in machine learning). In [19**], Rigotti *et al.* argue that increasing the dimensionality of data (by, for example, recording more neurons) only helps decoding if neurons display ‘nonlinear’ mixed selectivity, that is an additional neuron’s response cannot be explained by a linear superposition of the existing responses to the individual parameters. The authors show that individual neurons in the prefrontal cortex do in fact encode for multiple stimuli and that the collection of neurons display nonlinear mixed selectivity. Moreover, they show that the recorded neural representations have high dimensionality, and that this dimensionality can predict behavioral performance, as shown in Figure 2b. A review of this phenomenon is also provided in [5]. The authors detail the observation that the relevant unit of computation may in fact be individual neurons in lower-order sensory regions, where we may not need to consider population analysis in order to decode activity. However, as we move towards higher-order regions and increasingly complex tasks, we need population analyses if in the presence of mixed selectivity [5,22,6].

Neurons in multiple brain regions have been shown to have mixed selectivity. In [23], the authors show that neurons in the amygdala have representations pertaining to both stimulus and context in a flexible decision-making task, which may be the mechanism for the parietal prefrontal cortex neurons to access the same information.

Figure 2



(a) Schematic showing that it is possible to linearly decode when recording from two neurons displaying mixed selectivity, although it would not be possible to decode the activity well using either of the two neurons. Here, the firing rate activity is represented in state space during presentation of two different stimuli S₁ and S₂. (b) The number of dimensions in prefrontal cortex activity are higher for correct trials than error trials (left), although the decoding accuracy is the same for both (right). Thus, the dimensionality of neural data can directly predict behavior. Figure adapted from [19**].

In the inferior temporal cortex, prefrontal cortex, and the cortical-hippocampus loop, several studies have found that the neurons act with mixed selectivity [24,6,25]. In [21], the authors conclude that neurons in the posterior parietal cortex in rats have misaligned subspaces for movement and decision-making, elucidating that neurons are essentially multitaskers. Accurately decoding this activity necessitates a population-level analysis.

Analysis of neural activity over time reveals computational strategies

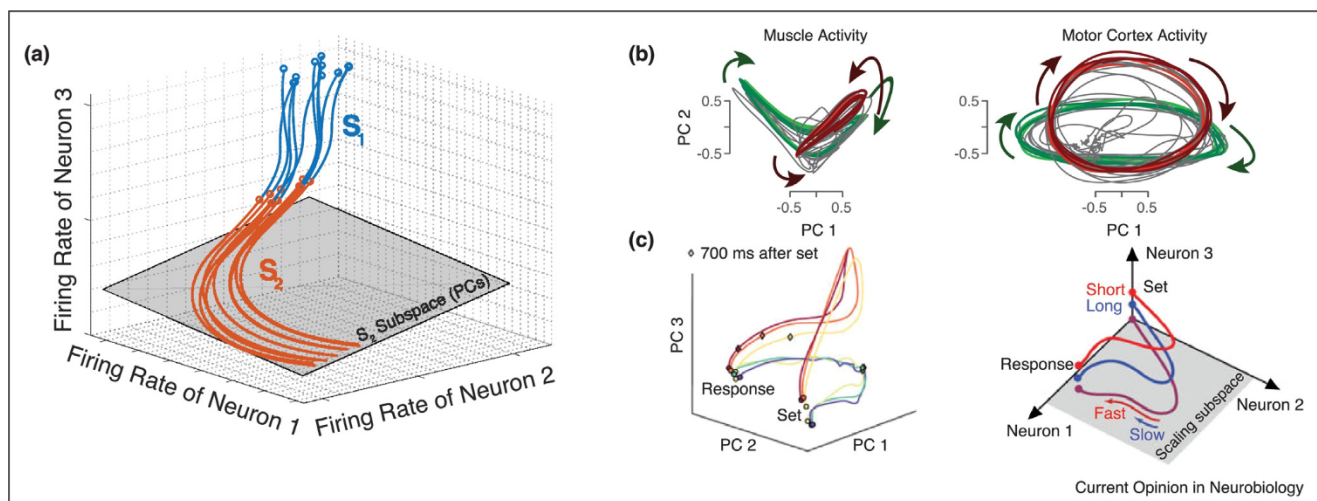
Considering neural population activity over time as states in a dynamical system has a long history, for example as in [26], where authors examine the potential mechanisms of memory and error correction using neuron-like components. This dynamical systems perspective is now prevalent in neuroscience, with the motor regions being the most natural testing ground, since we have access to the time-varying behavior as a direct output, and in fact, significant work has shown the value of this perspective [27–30], and new work continues to appear [31]. The analysis of neural activity over time, particularly during time-dependent behavior such as movement generation, can be studied using techniques developed to study dynamic (time-varying) activity. It is now common practice while analyzing motor regions to visualize the activity of the population of neurons corresponding to a trajectory in state space [32,27], as shown in Figure 3a. Here, the firing rate activity over time is shown in a lower dimensional space using principal components analysis (PCA).

The covariation of neurons allows us to locate subspaces occupied by the neural activity during specific tasks. An

interesting question under examination is whether the neural activity occupies orthogonal or overlapping subspaces during different tasks. Studies show computations of both kind. Specifically, the primate oculomotor system has pre-saccadic and saccadic activity in overlapping subspaces [33,34]. On the other hand, [35] finds that preparatory activity in the primate dorsal premotor cortex and primary motor cortex lives in the null space of the subspace in which movement activity resides. However, this finding in itself can be explained by either independence or orthogonality of the underlying subspaces, and in [36], it was revealed that the preparatory and movement subspaces are in fact orthogonal to each other. The orthogonality of the subspaces during different tasks was also explored in [37], where it was found that in a ‘pull’ task, the activity was almost orthogonal to a ‘walk’ task in mice. What determines whether neural activity lies in overlapping or orthogonal subspaces, or does the neural activity in fact lie in a continuum of more or less overlapping subspaces, depending on how similar the resulting tasks are? Moreover, does the union of the dimensions of the neural manifolds grow with the number of tasks considered, or does it reach a plateau at a fraction of the total number of neurons? Answering these questions is essential to making any fundamental claims about the role of subspaces in population activity.

Recently, specific properties of population dynamics that test ideas beyond subspace analysis have furthered our understanding of computational strategies in motor regions of the brain. For example, in [38*], a metric known as ‘trajectory tangling’ was used to characterize how, and potentially why, the motor cortex population response is

Figure 3



(a) Schematic showing 3 neurons’ activity over time during two different conditions, S_1 and S_2 , for example, preparatory period and movement period. The PC space for S_2 is also shown. (b) Tangling is lowered in the primary motor cortex activity (right) as compared to the electromyogram activity (left). Figure adapted from [38*]. (c) The trajectories decompose smoothly into speed scaling dimensions and non-speed scaling dimensions. Figure adapted from [39].

structured very differently from the muscle population response, as shown in Figure 3b. The degree to which tangling remains low reflects the degree to which population trajectories could plausibly have been produced by a noise-robust dynamical system that is fully observed in the recorded population. Results indicate that this is true of motor cortex, where activity has been proposed to reflect such dynamics, but not in the muscle population, where the relevant driving dynamics presumably lie upstream. Adopting trajectories with particularly low tangling may be a computational strategy employed by the brain to improve noise robustness. In [39], separate speed-varying and speed-invariant dynamics in the subspace are found in the medial frontal cortex (MFC) during a task that required the animal to internally keep track of time. In this setting, the authors probed the ‘temporal scaling’ of the firing rate activity, which is defined as the self-similarity of the neural activity when artificially stretched in time across different speeds. This idea is represented in Figure 3c, and indeed, they found that temporal scaling is present in the MFC. They also found that temporal scaling is not found in the thalamus, which projects to the MFC; through a series of maneuvers, they posit that temporal scaling may originate either in the MFC, or in circuits projecting to the MFC. These findings identify a possible mechanism for controlling processes at a continuum of different speeds.

These two studies have fundamental implications on the potential mechanisms of processing through different brain areas down to the kinematic output. In a recent study [40], the authors find that while different motor regions display superficially similar single-neuron responses during visually guided reaching, their different population dynamics indicate that they are likely performing quite different computations. These studies go beyond the relatively simplistic idea of neural activity residing in separate subspaces in order to be uniquely decoded by downstream neurons to produce separate outputs, and examine specific properties that the dynamical activity may possess. Moreover, the control of such activity is an interesting and open question, with recent studies like [41] addressing biologically plausible mechanisms to produce these activity.

Dynamic activity is extremely important in non-motor areas as well. For example, in [42], it was shown that motion and color detection can be explained by a single dynamical process in the prefrontal cortex. In [43], new information is seen to dynamically change the activity in the mouse parietal posterior cortex in order to implement evidence accumulation.

Nodes in Artificial Neural Networks emulate activity in certain regions of the brain

Although we have been modeling neural activity using artificial nodes for a long time [26], modeling with

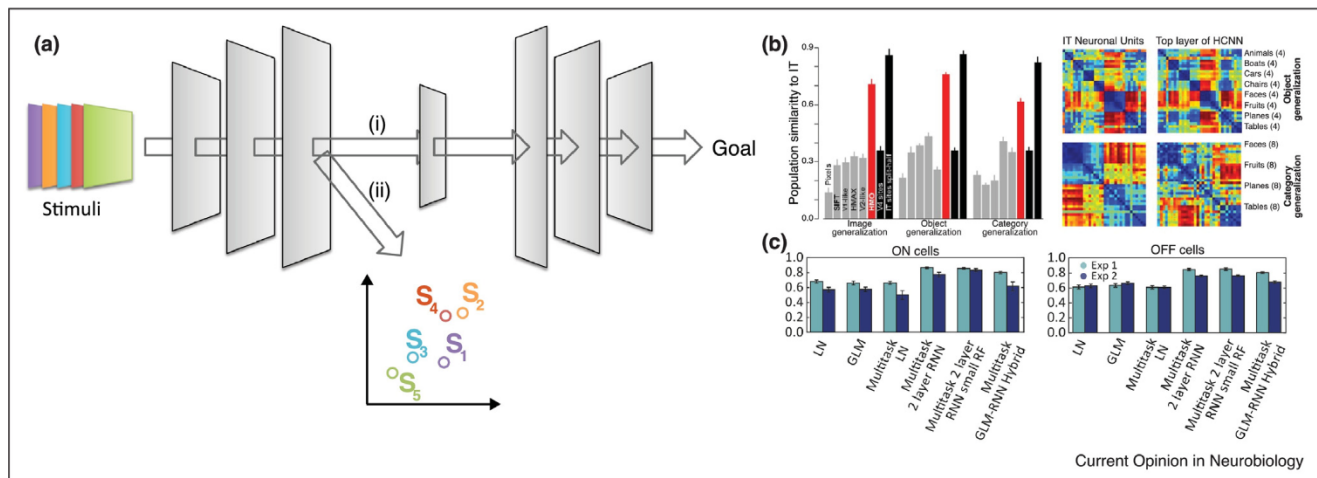
networks of neurons has been limited to carefully designed studies and an intricate hand-tuning of parameters (see e.g. [44]). The massive advances in learning deep neural networks has made using artificial neural networks (ANNs) with a large number of nodes and layers, with a large variety of structures, more approachable. However, the potential of ANNs to accurately describe neuronal computations in the brain remains a subject of active debate. Indeed, the biological plausibility of learning and computation in ANNs is an active field of study [45–48]. In [48], the authors introduce a novel method to assess whether the artificial networks are predictive of primate performance: they compare the set of errors made by the ANNs to those made by the human or monkey doing the same task. They find that although the ANNs are predictive of primate patterns of object-level confusion, they are not predictive of the performance on an object discrimination task in individual images. This finding sets a quantitative benchmark for comparisons between primates and neural networks.

There are currently two main approaches to utilizing artificial neural networks in modeling networks of neurons: (a) utilizing a goal-driven neural network, and using the learned network’s layers to compare and predict the population activity in different brain regions, and (b) modeling the activity of individual neurons as emanating from a feedforward or recurrent neural network architecture. These two approaches are characterized in Figure 4a.

(i) Goal-driven networks. Modeling a goal-driven ANN to describe neural activity attempts to combine the big-picture view of the different regions or layers of the brain, and the representation of the activity of individual nodes. Artificial models with a feedforward and recurrent architecture have been used to successfully capture dynamical activity [49,50]. In the context of motor dynamics, Sussillo *et al.* trained a recurrent neural network to produce recorded muscle activity, and showed that the activity of the artificial nodes had the same structure as the recorded dorsal premotor cortex and primary motor cortex activity [51]. In [41], the authors showed that it is possible to optimally tune a network of excitatory and inhibitory neurons to qualitatively model neural activity in the preparatory phase of a movement, that will lead to a stable and reliable movement trajectory, providing a potential computational basis to the observed results in [32]. Other motor studies in varied tasks have also succeeded in relating the motor cortex activity to artificial nodes in an ANN [38,36]. ANNs have also been used to investigate different tasks including working memory [52,53] and perceptual decision-making [42].

In the visual cortex, Yamins *et al.* used hierarchical convolutional neural networks (CNNs) for modeling the processing of information in the brain, with each layer

Figure 4



(a) Artificial Neural Network (ANN) with (i) goal-driven training, and (ii) data-driven modeling. **(b)** Left Panel: The model (red bar) outperforms other known models (light grey) in population similarity to IT neurons. In black is data from V4 and other IT neurons; Right Panel: The representation dissimilarity matrix of the top layer in a 4 layer CNN (CNN) as learned using goal-driven training (right) is very similar to that of electrophysiology recordings in the IT (left). Figure adapted from [54^{*}]. **(c)** Data-driven modeling of retinal cells quantifies the performance of different models, in particular highlighting the performance gains using recurrent neural networks (RNNs). Figure adapted from [57^{*}].

of the neural network being analogous to a brain region [54^{*}]. They showed that directly modeling an experimental task using a 4 layer CNN can lead to a neural network in which the layers have a similar structure to the neuronal responses in V4 and IT (see Figure 4b). They found correlation ($\sim 50\%$) between the top and penultimate layers of the model for the IT and V4 populations in terms of explained variance. Recently, the authors have extended their results to train a 6 layer CNN to successfully predict category-orthogonal object properties along the ventral stream [55].

(ii) Data-driven modeling. Directly modeling the activity of individual neurons using ANNs is a growing line of research, with some recent successes in prediction of single neuron activity. Typically, a (partially) shared model is learned for all neuronal activity. It is then possible to compute a receptive field per neuron, or latent activity per population of neurons, which may be of scientific interest, as well as use these models for predicting and decoding activity.

Retinal ganglion cells have been a particularly attractive target for population modeling, beginning (at least) a decade ago [56] and continuing with more recent ANN-based approaches [57^{*},58]. In [57^{*}], the activity of retinal ganglion cells is modeled as arising from a recurrent architecture. As shown in Figure 4c, a 2 layer recurrent neural network consistently captures around 80% of the explained variance across experiments and cell types, which outperforms other known models for retinal activity. In [58], the

authors model retinal ganglion cells in an ANN with a convolutional architecture. In [59], the authors build on previous methods by introducing a more efficient model while performing end-to-end training, and applying this to neuronal recordings from the primary visual cortex.

It is as yet unclear which of the two approaches, (i) or (ii), are better suited to model the activity of neurons as well as provide insight about the computations in separate regions of the brain. In [60], the authors address this question in the context of V1 neurons, and find that both these modeling approaches perform similarly, while outperforming other known models. Moreover, they find that (ii) provides a simpler learned architecture.

While modeling the activity of the recorded neural activity, one can simultaneously learn a dynamical model that can produce this neural data in order to gain insights on computational strategies employed by neural populations, as discussed in the previous section. Classic versions of this have included linear dynamical systems [61,62]. Gao *et al.* [30] introduced ANNs as the nonlinear map from latent dynamics to observable neural population activity, a theme which has been elaborated both by including ANNs in the dynamics model, and otherwise [63–66].

Looking ahead: when can we trust the results of population-level analyses?

Throughout this review we have highlighted exciting findings that have resulted from the joint analysis of neural populations, works that exemplify the broad and

rapidly growing trend in the field towards the *neural population doctrine*. However, this exciting progress has ignited a serious and increasingly contentious debate about whether these analyses are actually producing novel findings about the brain, or if they are simply recapitulating ‘old knowledge dressed up in new clothes’ [67]. Indeed, this concern has led to high-profile issues in other fields confronting similarly explosive data (and data analysis) growth (e.g. [68]). Without question, this controversy is legitimate, in so much as the perils of high-dimensional data are well known. To rationalize and quantify this debate, then, one of the most essential data analytical tasks going forward will be quantitative and computational results that shed light on this debate. A first key step in this direction is in [69**], where a nontrivial null hypothesis leads to a ‘null model’ from which surrogate datasets can be drawn. These datasets are then passed through a given population analysis, to quantify a null distribution about the extent to which some population structure is expected under the null hypothesis; existing results in both mixed selectivity (see Section ‘Decoding accuracy is more than the sum of its parts’ of this review) and population dynamics (see Section ‘Analysis of neural activity over time reveals computational strategies’) are tested under this framework. See [70–72] for other studies testing and extending these concepts.

Conclusions

As recording techniques and computing capabilities continue to improve, experimental and computational studies continue to demonstrate that neuronal *populations* may in fact be the relevant unit of computation in many brain regions. Throughout this review, we have pointed out various studies that support this scientific trend, in domains spanning correlated variability, decoding, dynamical activity, and artificial neural networks.

Going forward, there are still major concerns in these domains that need to be addressed. Firstly, when examining the dimensionality of neural recordings, the dimensionality of the stimulus or the resulting behavior may need to be large in order to have a better understanding of the underlying neural dynamics. Secondly, the state space approach, though very useful, may obscure ways to develop causal interventions in order to systematically verify and advance scientific findings. We need to develop new methods or experiments which bring us closer to a mechanistic understanding of the underlying phenomena.

(i) Pushing the envelope on the dimensionality of the behavior. Analysis of stereotyped behavior in a laboratory setting has long been considered the primary experimental setting for reliable and reproducible results. One of the next big challenges in neuroscience is to move past the confines of tasks with a small number of degrees

of freedom (and the limitations perhaps imposed on the richness of the data recorded) and examine tasks with ‘free’ behaviors. In fact, the recorded behavior may be high-dimensional enough that one may need to use population analysis to characterize the behavior in the first place, then relate neural activity to this complex behavior [73–75].

ii) Discovering causal relationships to behavior. Projecting the neural activity in a rotated, stretched, lower dimensional space, although enables visualization and allows us to relate the activity to behavior, may preclude a clear understanding of the circuit-level activity that lead to these behaviors. Methods such as perturbation analyses can be used to detangle this relationship. An example of this method is a novel Brain Computer Interface paradigm that perturbs the neural association of tasks ‘within’ a learned neural manifold and compares this to a perturbation ‘outside’ the manifold [76]. Learning ‘outside manifold’ perturbations is harder, which may reflect the connections between the relevant neurons. See [77] for a good perspective on why causal inference in the presence of multi-scale dynamics is difficult.

Conflict of interest statement

Nothing declared.

Acknowledgements

This work was supported by the Swiss National Science Foundation (Research Award P2SKP2_178197), NIH R01NS100066, Simons Foundation 542963, NSF NeuroNex DBI-1707398, The Gatsby Charitable Foundation, the Sloan Foundation, and the McKnight Foundation.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Ramón, Cajal S: *Histology of the Nervous System of Man and Vertebrates*. USA: Oxford University Press; 1995. vol 1.
 2. Sherrington CS: **Observations on the scratch-reflex in the spinal dog**. *J Physiol* 1906, **34**:1-50.
 3. Barlow HB: **Summation and inhibition in the frog’s retina**. *J Physiol* 1953, **119**:69-88.
 4. Yuste R: **From the neuron doctrine to neural networks**. *Nat Rev Neurosci* 2015, **16**:487.
 5. Fusi S, Miller EK, Rigotti M: **Why neurons mix: high dimensionality for higher cognition**. *Curr Opin Neurobiol* 2016, **37**:66-74.
 6. Eichenbaum H: **Barlow versus Hebb: when is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition?** *Neurosci Lett* 2017.
 7. McCulloch WS, Pitts W: **A logical calculus of the ideas immanent in nervous activity**. *Bull Math Biophys* 1943, **5**:115-133.
 8. Hebb DO: *The Organization of Behavior: A Neuropsychological Theory*. Psychology Press; 2005.
 9. Steinmetz NA, Koch C, Harris KD, Carandini M: **Challenges and opportunities for large-scale electrophysiology with neuropixels probes**. *Curr Opin Neurobiol* 2018, **50**:92-100.

10. Paninski L, Cunningham JP: **Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience.** *Curr Opin Neurobiol* 2018, **50**:232.
11. Cunningham JP, Yu BM: **Dimensionality reduction for large-scale neural recordings.** *Nat Neurosci* 2014, **17**:1500-1509.
12. Stevenson IH, Kording KP: **How advances in neural recording affect data analysis.** *Nat Neurosci* 2011, **14**:139.
13. Zohary E, Shadlen MN, Newsome WT: **Correlated neuronal discharge rate and its implications for psychophysical performance.** *Nature* 1994, **370**:140.
14. Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A: **Information-limiting correlations.** *Nat Neurosci* 2014, **17**:1410.
15. Abbott LF, Dayan P: **The effect of correlated variability on the accuracy of a population code.** *Neural Comput* 1999, **11**:91-101.
16. Averbeck BB, Latham PE, Pouget A: **Neural correlations, population coding and computation.** *Nat Rev Neurosci* 2006, **7**:358.
17. Kohn A, Coen-Cagli R, Kanitscheider I, Pouget A: **Correlations and neuronal population information.** *Annu Rev Neurosci* 2016, **39**.
18. Ni A, Ruff D, Alberts J, Symmonds J, Cohen M: **Learning and attention reveal a general relationship between population activity and behavior.** *Science* 2018, **359**:463-465.
- Ni *et al.* examine the effects of learning and attention on V4 neurons, and find that these two processes both cause the correlated variability of neurons to decrease. They use a population-based metric (percentage of variance explained in first principal component) to further quantify the decrease in correlated variability.
19. Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S: **The importance of mixed selectivity in complex cognitive tasks.** *Nature* 2013, **497**:585.
- Rigotti *et al.* detail the specific conditions under which decoding accuracy improves as we record more neurons — termed nonlinear mixed selectivity of the neurons. They show that neurons in the prefrontal cortex of monkeys do in fact display nonlinear mixed selectivity and higher dimensionality, especially during better performance.
20. Rust NC: **Population-based representations: from implicit to explicit.** *Cognit Neurosci* 2014, **337**.
21. Raposo D, Kaufman MT, Churchland AK: **A category-free neural population supports evolving demands during decision-making.** *Nat Neurosci* 2014, **17**:1784-1792.
22. DiCarlo JJ, Zoccolan D, Rust NC: **How does the brain solve visual object recognition?** *Neuron* 2012, **73**:415-434.
23. Saez A, Rigotti M, Ostojic S, Fusi S, Salzman C: **Abstract context representations in primate amygdala and prefrontal cortex.** *Neuron* 2015, **87**:869-881.
24. McKenzie S, Keene CS, Farvik A, Bladon J, Place R, Komorowski R, Eichenbaum H: **Representation of memories in the cortical-hippocampal system: results from the application of population similarity analyses.** *Neurobiol Learn Mem* 2016, **134**:178-191.
25. Machens CK, Romo R, Brody CD: **Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex.** *J Neurosci* 2010, **30**:350-360.
26. Hopfield JJ, Tank DW: **Computing with neural circuits: a model.** *Science* 1986, **233**:625-633.
27. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV: **Neural population dynamics during reaching.** *Nature* 2012, **487**:51.
28. Saxena S, Sarma SV, Dahleh M: **Performance limitations in sensorimotor control: tradeoffs between neural computing and accuracy in tracking fast movements.** *bioRxiv* 2018:464230.
29. Churchland MM, Cunningham JP: **A dynamical basis set for generating reaches.** *Cold Spring Harbor Symposia on Quantitative Biology.* Cold Spring Harbor Laboratory Press; 2014:67-80. vol 79.
30. Gao Y, Archer EW, Paninski L, Cunningham JP: **Linear dynamical neural population models through nonlinear embeddings.** *Advances in Neural Information Processing Systems.* 2016:163-171.
31. Sauerbrei B, Guo J-Z, Mischiati M, Guo W, Kabra M, Verma N, Branson K, Hantman A: **Motor cortex is an input-driven dynamical system controlling dexterous movement.** *bioRxiv* 2018:266320.
32. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV: **Cortical preparatory activity: representation of movement or first cog in a dynamical machine?** *Neuron* 2010, **68**:387-400.
33. Hanes DP, Schall JD: **Neural control of voluntary movement initiation.** *Science* 1996, **274**:427-430.
34. Schall JD, Thompson KG: **Neural selection and control of visually guided eye movements.** *Annu Rev Neurosci* 1999, **22**:241-259.
35. Kaufman MT, Churchland MM, Ryu SI, Shenoy KV: **Cortical activity in the null space: permitting preparation without movement.** *Nat Neurosci* 2014, **17**:440-448.
36. Elsayed GF, Lara AH, Kaufman MT, Churchland MM, Cunningham JP: **Reorganization between preparatory and movement population responses in motor cortex.** *Nat Commun* 2016, **7**:13239.
37. Miri A, Warriner CL, Seely JS, Elsayed GF, Cunningham JP, Churchland MM, Jessell TM: **Behaviorally selective engagement of short-latency effector pathways by motor cortex.** *Neuron* 2017, **95**:683-696.
38. Russo AA, Bittner SR, Perkins SM, Seely JS, London BM, Lara AH, Miri A, Marshall NJ, Kohn A, Jessell TM *et al.*: **Motor cortex embeds muscle-like commands in an untangled population response.** *Neuron* 2018, **97**:953-966.
- Russo *et al.* demonstrate the use of a population-based metric of dynamical data known as ‘tangling’, which is seen to be much lower in the activity of the primary motor cortex as compared to EMG activity and activity in other brain structures. This may reflect the lower propensity of motor cortex neurons to be driven by input dynamics, thus conferring noise robustness.
39. Wang J, Narain D, Hosseini EA, Jazayeri M: **Flexible timing by temporal scaling of cortical responses.** *Nat Neurosci* 2018, **21**:102.
- Wang *et al.* do a series of analyses to identify the similarity in state space of firing rate activity in the medial frontal cortex (MFC) during processes at different timescales. They found that temporal scaling is found in the MFC, and through a series of inactivation and computational studies, they posit that temporal scaling may originate either in the MFC, or in circuits projecting to the MFC.
40. Lara A, Cunningham J, Churchland M: **Different population dynamics in the supplementary motor area and motor cortex during reaching.** *Nat Commun* 2018, **9**:2754.
41. Hennequin G, Vogels TP, Gerstner W: **Optimal control of transient dynamics in balanced networks supports generation of complex movements.** *Neuron* 2014, **82**:1394-1406.
42. Mante V, Sussillo D, Shenoy KV, Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex.** *Nature* 2013, **503**:78-84.
43. Morcos AS, Harvey CD: **History-dependent variability in population dynamics during evidence accumulation in cortex.** *Nat Neurosci* 2016, **19**:1672.
44. Abbott L: **Decoding neuronal firing and modelling neural networks.** *Q Rev Biophys* 1994, **27**:291-331.
45. Kietzmann TC, McClure P, Kriegeskorte N: **Deep neural networks in computational neuroscience.** *bioRxiv* 2018:133504.
46. Micconi T: **Biologically plausible learning in recurrent neural networks reproduces neural dynamics observed during cognitive tasks.** *Elife* 2017, **6**.
47. Song HF, Yang GR, Wang X-J: **Reward-based training of recurrent neural networks for cognitive and value-based tasks.** *Elife* 2017, **6**.

48. Rajalingham R, Issa EB, Bashivan P, Kar K, Schmidt K, DiCarlo JJ: **Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks.** *J Neurosci* 2018, **38**:7255-7269.
49. Buonomano DV, Maass W: **State-dependent computations: spatiotemporal processing in cortical networks.** *Nat Rev Neurosci* 2009, **10**:113.
50. Sussillo D, Abbott LF: **Generating coherent patterns of activity from chaotic neural networks.** *Neuron* 2009, **63**:544-557.
51. Sussillo D, Churchland MM, Kaufman MT, Shenoy KV: **A neural network that finds a naturalistic solution for the production of muscle activity.** *Nat Neurosci* 2015, **18**:1025-1033.
52. Barak O, Sussillo D, Romo R, Tsodyks M, Abbott L: **From fixed points to chaos: three models of delayed discrimination.** *Prog Neurobiol* 2013, **103**:214-222.
53. Rajan K, Harvey CD, Tank DW: **Recurrent network models of sequence generation and memory.** *Neuron* 2016, **90**:128-142.
54. Yamins DL, Hong H, Cadieu CF, Solomon EA, Seibert D, DiCarlo JJ: **Performance-optimized hierarchical models predict neural responses in higher visual cortex.** *Proc Natl Acad Sci* 2014, **111**:8619-8624.
- Yamins *et al.* use hierarchical convolutional neural networks to build an end-to-end neural network modeling the neural pathways required to implement an object recognition task, also carried out by primates. They find that the layers of the network are able to successfully predict neural activity in the IT and the V4.
55. Hong H, Yamins DL, Majaj NJ, DiCarlo JJ: **Explicit information for category-orthogonal object properties increases along the ventral stream.** *Nat Neurosci* 2016, **19**:613.
56. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky E, Simoncelli EP: **Spatio-temporal correlations and visual signalling in a complete neuronal population.** *Nature* 2008, **454**:995.
57. Batty E, Merel J, Brackbill N, Heitman A, Sher A, Litke A, Chichilnisky E, Paninski L: **Multilayer recurrent network models of primate retinal ganglion cell responses.** 2016.
- Using a series of models with increasing complexity, the authors succeed in capturing a very high amount of explained variance of the activity of individual neurons. They demonstrate the benefit of recurrent neural networks with spatial, temporal and static nonlinearities.
58. McIntosh L, Maheswaranathan N, Nayebi A, Ganguli S, Baccus S: **Deep learning models of the retinal response to natural scenes.** *Advances in Neural Information Processing Systems*. 2016:1369-1377.
59. Kliindt D, Ecker AS, Euler T, Bethge M: **Neural system identification for large populations separating “what” and “where”.** *Advances in Neural Information Processing Systems*. 2017:3509-3519.
60. Cadena SA, Denfield GH, Walker EY, Gatys LA, Tolias AS, Bethge M, Ecker AS: **Deep convolutional models improve predictions of macaque v1 responses to natural images.** *bioRxiv* 2017:201764.
61. Smith AC, Brown EN: **Estimating a state-space model from point process observations.** *Neural Comput* 2003, **15**:965-991.
62. Macke JH, Buesing L, Cunningham JP, Byron MY, Shenoy KV, Sahani M: **Empirical models of spiking in neural populations.** *Advances in Neural Information Processing Systems* 2011:1350-1358.
63. Wu A, Roy NG, Keeley S, Pillow JW: **Gaussian process based nonlinear latent structure discovery in multivariate spike train data.** *Advances in Neural Information Processing Systems*. 2017:3496-3505.
64. Zhao Y, Park IM: **Variational latent Gaussian process for recovering single-trial dynamics from population spike trains.** *Neural Comput* 2017, **29**:1293-1316.
65. Hernandez D, Moretti AK, Wei Z, Saxena S, Cunningham J, Paninski L: **A novel variational family for hidden nonlinear Markov models.** 2018. arXiv preprint arXiv:1811.02459.
66. Pandarinath C, O’Shea DJ, Collins J, Jozefowicz R, Stavisky SD, Kao JC, Trautmann EM, Kaufman MT, Ryu SI, Hochberg LR *et al.*: **Inferring single-trial neural population dynamics using sequential auto-encoders.** *Nat Methods* 2018.
67. Pillow JW, Aoi MC: **Is population activity more than the sum of its parts?** *Nat Neurosci* 2017, **20**:1196.
68. Eklund A, Nichols TE, Knutsson H: **Cluster failure: why FMRI inferences for spatial extent have inflated false-positive rates.** *Proc Natl Acad Sci* 2016, **113**:7900-7905.
69. Elsayed GF, Cunningham JP: **Structure in neural population recordings: an expected byproduct of simpler phenomena?** *Nat Neurosci* 2017, **20**:1310.
- Elsayed and Cunningham develop a statistical control framework for results that depend on neural populations, using both Fisher randomization and maximum entropy-based methods. They test two prominent systems results against the null hypothesis that those results could be caused by tuning, temporal structure, or neuronal correlations.
70. Harrison MT, Amarasingham A, Truccolo W: **Spatiotemporal conditional inference and hypothesis tests for neural ensemble spiking precision.** *Neural Comput* 2015, **27**:104-150.
71. Loaiza-Ganem G, Gao Y, Cunningham JP: **Maximum entropy flow networks.** *International Conference on Learning Representations (ICLR)* 2017.
72. Savin C, Tkačik G: **Maximum entropy models as a tool for building precise neural controls.** *Curr Opin Neurobiol* 2017, **46**:120-126.
73. Nguyen JP, Shipley FB, Linder AN, Plummer GS, Liu M, Setru SU, Shaevitz JW, Leifer AM: **Whole-brain calcium imaging with cellular resolution in freely behaving *Caenorhabditis elegans*.** *Proc Natl Acad Sci* 2016, **113**:E1074-E1081.
74. Kato S, Kaplan HS, Schrödel T, Skora S, Lindsay TH, Yemini E, Lockery S, Zimmer M: **Global brain dynamics embed the motor command sequence of *Caenorhabditis elegans*.** *Cell* 2015, **163**:656-669.
75. Wiltchko AB, Johnson MJ, Iurilli G, Peterson RE, Katon JM, Pashkovski SL, Abreira VE, Adams RP, Datta SR: **Mapping sub-second structure in mouse behavior.** *Neuron* 2015, **88**:1121-1135.
76. Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Byron MY, Batista AP: **Neural constraints on learning.** *Nature* 2014, **512**:423.
77. Jazayeri M, Afraz A: **Navigating the neural space in search of the neural code.** *Neuron* 2017, **93**:1003-1014.