



ELSEVIER



The language of the brain: real-world neural population codes

J Andrew Pruszynski^{1,2,3} and Joel Zylberberg^{4,5,6}

Addresses

¹Department of Physiology and Pharmacology, Western University, London, ON, Canada²Department of Psychology, Western University, London, ON, Canada³Robarts Research Institute, London, ON, Canada⁴Center for Vision Research, York University, Toronto, ON, Canada⁵Department of Physics and Astronomy, York University, Toronto, ON, Canada⁶Canadian Institute for Advanced Research, Toronto, ON, CanadaCorresponding author: Zylberberg, Joel (joelzy@yorku.ca)**Current Opinion in Neurobiology** 2019, **58**:30–36This review comes from a themed issue on **Computational Neuroscience**Edited by **Máté Lengyel** and **Brent Doiron**<https://doi.org/10.1016/j.conb.2019.06.005>

0959-4388/© 2019 Elsevier Ltd. All rights reserved.

Introduction

External stimuli trigger electrical activity in the nervous system's sensory receptors such as photoreceptors in the eyes, hair cells in the cochlea, and mechanoreceptors in the skin. This electrical activity causes action potentials in projection neurons that bring information about the sensory stimulus to the central nervous system, triggering further action potentials as the signal ascends the sensory processing hierarchy. A fundamental problem in neuroscience is determining the *language* of this signalling: to identify how external stimuli are represented by the patterns of action potentials they evoke.

How do we tackle the challenge of understanding neural population codes within the sensory systems? This effort consists of two main approaches. First, experiments are conducted where sensory stimuli are presented to an animal or person while the corresponding patterns of neural activity are recorded with electrical or, optical methods (e.g., [Figure 1](#)). Second, models are constructed that attempt to simulate the activity of neuronal populations in response to some set of external stimuli: among those models are the traditional *tuning curve plus noise* models [1–6], and the rapidly-advancing *deep learning* [7,8] models from the field of artificial intelligence. Within these models, researchers can perturb different aspects of the neuronal population responses to stimuli, and relate

those perturbations to the system's ability to represent and/or process sensory information.

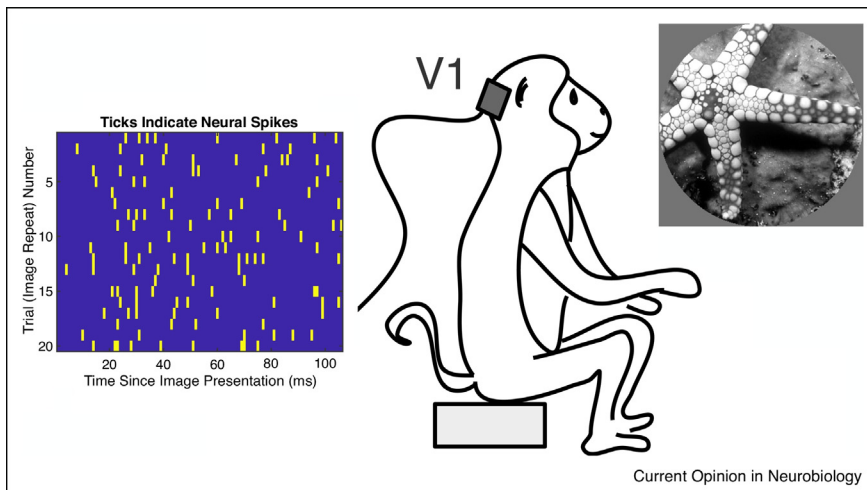
In this review, we summarize some key insights from this recent body of work and attempt to identify key challenges and opportunities for future studies of neural population coding. We focus mainly – though not exclusively – on studies of the visual and somatosensory systems, but note at the outset that many of the insights likely apply to other modalities.

Stimulus-dependence of inter-neural correlations matters

Over repeat presentations of the same stimulus, sensory neurons usually show highly variable responses [10] ([Figure 1](#)). Because the brain must infer and act on stimuli on single trials (i.e., without averaging responses over trials), this variability has the potential to corrupt the brain's sensory representation. This issue leads many researchers to ask how the brain forms reliable sensory representations out of seemingly-unreliable components. Older research showed that, when the variability (“noise”) is correlated between neurons, it can – depending on the specifics of which neurons are positively or negatively correlated – have much less impact on the neural code than would noise that was independent between neurons [1–5,11–13]. The key intuition is that the geometrical relationships between the neurons' tuning (stimulus-induced changes in mean firing rate) and the noise correlation structure determines whether the correlations mitigate or exacerbate the impact of noise on the population code ([Figure 2](#)).

Noise correlations are defined separately for each specific stimulus and neuron pair, and for a given neuron pair can differ substantially between stimuli [14^{**}, 15^{**}, 17^{*}, 18]. Nevertheless, older studies tended to simply average noise correlations over stimuli, and to study computational models in which a single (stimulus-independent) noise correlation coefficient was assigned to each cell pair [1–3] (although see Ref. [6] for an exception). Recent work shows that stimulus dependence of noise correlations can dramatically alter their impact on population coding: two groups recently showed that the stimulus dependence of noise correlations between direction selective retinal ganglion cells can double the amount of stimulus information encoded in the population responses, compared with populations of uncorrelated neurons [14^{**}, 15^{**}]. When stimulus dependence of correlations was artificially removed from the ganglion cell populations, nearly all of this population coding benefit was lost [14^{**}]. This effect arises because the stimulus-dependent correlations between ganglion cells

Figure 1




In vivo visual neuroscience experiment. A monkey was shown images while the spiking activity of roughly 100 neurons were simultaneously recorded in primary visual cortex (V1; data from Ref. [9]). Activity of one neuron is shown for 20 repeated presentations of the seastar image; ticks indicate times at which the neuron emitted spikes. Inter-trial variability in the response is substantial, posing the question of how the brain forms reliable sensory representations with such variable neuronal units.

“shape” the noise so that it is more separable from the underlying stimulus-driven signals, similar to the cartoon in the left panel of Figure 2. Another recent study [16**] asked which patterns of correlation between neurons cause the *worst* corruption of the population code by neural variability, as quantified via Fisher information. They found that Fisher information increases without bound as the neural

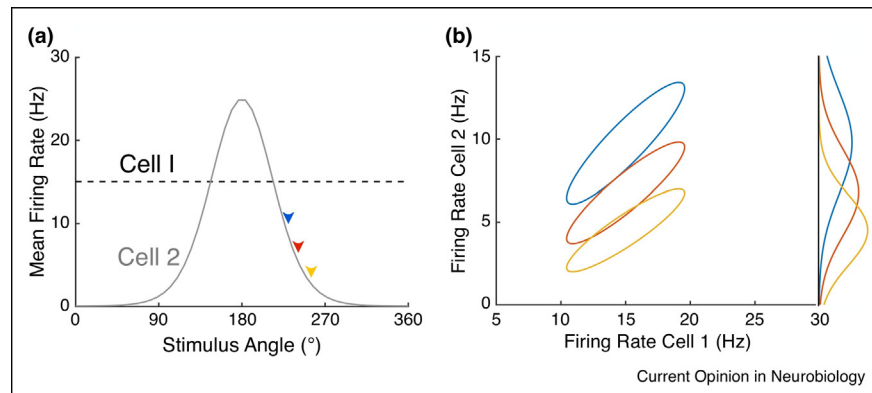
population size increases, with any type of correlation between neurons, except for one important exception. That exception, dubbed *differential correlations* arises when the neural variability exactly mimics the stimulus-induced changes in mean neural firing rates (Figure 2, right panel). In that case, there is no way for a decoder to separate the “signal” from the “noise”, and the Fisher information

Figure 2



Stimulus-dependence of correlations matters. Cartoon depicts the *neural response space*: each axis in this space is the firing rate of a single neuron on a single stimulus presentation trial. Within that space, one can plot the distribution of neural firing rates in response to a single stimulus; these *stimulus-conditioned response distributions* are shown as ellipses. Their centers indicate the mean firing rates of all neurons in response to that stimulus. The ellipse boundaries are the distributions’ 1-standard-deviation probability contours, and their orientations depend on the correlations between neurons. Four such distributions are shown, under two different conditions (left, and right). The line, or surface, connecting the mean firing rates for all possible stimuli is known as the *signal manifold*: here it is shown as a single curve for simplicity. In the left panel, the ellipses have their long axes orthogonal to the green curve: this minimizes the overlap between neighboring ellipses, maximizing the ability to differentiate between similar stimuli based on the neural responses. Because the curve has different orientations at different positions, maintaining this orthogonality requires that the ellipses have different orientations: the correlations between neurons must depend on the stimulus. Direction-selective ganglion cells in retina have stimulus-dependent correlations with this property [14**, 15**]. In the right panel, the same signal curve is shown, but now the ellipses have their long axes parallel to the signal curve. This maximizes the overlap between neighboring ellipses, minimizing the ability to differentiate between similar stimuli based on the neural responses. In this worst-case scenario, the correlations (known as *differential correlations*) also depend on the stimulus [16**].

Figure 3



Untuned neurons matter. (a) The tuning curves of two neurons are shown: these depict the mean firing rate of each neuron in response to stimuli moving in different directions (defined by their direction angle). Cell 1 is untuned: it has the same firing rate for all stimuli. Cell 2 is tuned. (b) Stimulus-conditioned response distributions, indicated by the 1-standard-deviation probability contours, for this two-cell population, in response to three different stimulus directions (indicated by arrows in panel A). These were obtained by assuming Poisson variability (variance in each neuron's firing rate equal to the mean), and a positive correlation between the two neurons. In the two-dimensional response space, the distributions have minimal overlap, meaning that the different stimulus directions can be differentiated based on the neural responses. If one ignores the untuned neuron, and looks only at the firing rate of the tuned one, these two-dimensional distributions are instead replaced by the one-dimensional projections shown on the right vertical axis. (Stimulus-conditioned distributions of Cell 2's firing rate). Those distributions overlap significantly, indicating that it is hard to tell, based on Cell 2's firing rate alone, which stimulus direction is present. A similar geometrical configuration arises in higher dimensions (i.e., with more than 2 neurons).

saturates as the population size increases. Differential correlations are stimulus-dependent.

Thus, several recent studies highlight the need for our understanding of the neural code to account for the stimulus-dependence of correlations between neurons.

Neurons with stimulus-independent firing can still contribute to population codes

Older work in sensory neuroscience focused exclusively on neurons whose firing rates vary as the stimulus changes (i.e., ones with some *tuning*), tacitly assuming that any neurons without tuning did not contribute to the sensory neural code. Recent papers challenged this assumption, showing that even neurons with zero stimulus-dependence to their firing rates can contribute to the population code: populations consisting only of the tuned neurons contain less stimulus information than do mixed populations containing both the tuned and the untuned cells [19, 20^{••}, 21[•]]. Crucially, this result arises even in pure rate codes, and does not rely on spike timing information.

The key idea is that, because noise is correlated over the neural population, untuned neurons' activity levels carry information about trial-specific firing rate deviations ("noise") in tuned cells. Accordingly, observing untuned cells enables those deviations to be estimated and subtracted from tuned cells' firing rates, which is equivalent to making them less variable. This effect is shown geometrically in the cartoons in Figure 3.

Relatedly, a recent machine-learning study investigated units within artificial neural networks (ANNs) trained for object recognition tasks, whose activations had varying levels of selectivity for the object category [22]. Ablating the least category-selective units within the ANN profoundly negatively impacted the network's ability to correctly identify images. Within the ANNs, there is no "noise" per se (as in the biological networks). However, within each object category, there are large variations in luminance, pose, etc., that act as effective "noise" corrupting the object categorization task. The importance of these category-unselective units thus likely has similar origins to the role of untuned neurons in the biology: the category-unselective units may track the confounding variables in a manner that better allows downstream units to account for their effects on the category-selective units' activations.

These works emphasize the need to consider how populations of neurons work together to form sensory codes; neurons that would on their own not convey any information (e.g. untuned ones, or category-unselective ones) can serve to make the overall population code more effective. Every neuron probably matters, and none should be ignored.

Information propagation vs information content put different constraints on population codes

Canonical population coding studies used information theory [23] to quantify how well the neural population responses in a given neural population – e.g. in retina – could be

decoded to identify the stimulus presented to the animal; these include experimental studies [14^{••}, 15^{••}, 17[•],24,25], and theoretical ones [1–6,26,27]. A shared key assumption is that high information content corresponds directly to high performance of the sensory neural population code.

One recent study asked instead how well information in a peripheral neural population could propagate through potentially noisy and nonlinear downstream circuitry [28^{••}], as in a multi-layer feedforward neural network. Critically, that work found that neural codes optimized for information content at the periphery, and those optimized for the ability of information to propagate, could be very different: information robustness and information content put different constraints on the population code. This means that the “raw” information content at the periphery may not be the best metric by which to evaluate neural population codes, and that instead future studies should perhaps ask how much of the information from the periphery makes it through the multiple layers of processing that exist downstream.

Population codes can multiplex spike timing and firing rate information in potentially context-dependent ways

When it comes to understanding how neuronal populations represent the outside world, one has to ultimately consider what aspect of the spike trains are relevant for transmitting information along the neural processing pathway. The overwhelming focus in this respect has been on the rate code, where stimulus information is encoded by how intensely a neuron fires over relatively long timescales [29]. The focus on rate coding is sensible given its simplicity and the pervasiveness of correlations between sensory variables and neuronal firing rates. However, rate codes ignore extra information that can potentially be embedded in the temporal structure of action potentials. Such temporal codes, which can take various forms (see Ref. [30] for details), can carry significant information about sensory stimuli at various levels of the neuroaxis and in various modalities, including touch [31–35] and vision [36–41].

Determining which parts of the nervous system use rate or temporal codes when encoding a given stimulus property is an important and active area of research in all sensory modalities. In general, definitive evidence for any particular code remains elusive. Even in the classical example of motoneurons, where rate codes have been heavily assumed, recent work shows that action potentials can influence behavior on the millisecond timescale [42^{••}]. Indeed, rate and temporal codes have often been presented as competing concepts, the goal being to directly establish that a favored neural code is better at encoding a particular stimulus parameter in some brain area. This intense debate has recently given way to an appreciation that rate and temporal codes can readily

co-exist, permitting the nervous system to multiplex complementary information along the same neural processing pathway [43–46].

For example, neurons in the cerebral cortex nominally unresponsive to auditory stimulation - in the sense that they do not change their firing rate after stimulus onset - can robustly encode task-relevant stimuli and predict animal behavior in their inter-spike intervals, strong evidence of temporal coding in the auditory processing pathway [47]. However, selectively removing only these unresponsive neurons or responsive neurons (those that carry stimulus and decision information in their firing rates) from the overall population decreases performance less than removing the same number of neurons drawn randomly from the two populations, suggesting that there exists an interaction between rate and temporal codes that has a synergistic effect on animal behavior. Multiplexing may also allow different downstream targets to read-out specific aspects of the same sensory input. For example, recent work in touch has shown that orientation acuity during hand control is nearly an order of magnitude better and faster than orientation acuity measured in classical psychophysical paradigms [48^{••}]. These behavioral findings must arise from the same peripheral neurons, which can signal edge orientation both in terms of rate and temporal codes [32]. One simple idea, as yet untested, is that different downstream targets read out rate codes, and temporal codes to support perception and hand control, respectively.

At the end of the day, the nervous system does not know about rate codes and temporal codes in any explicit way. Information processing and transmission happens via synapses and network dynamics that demonstrate various integration time constants [49]. Where there are longer time constants (at NMDA receptors, for example) neural elements tend to integrate over many spiking events and thus the resulting signals appear more rate-like. And where there are shorter time constants (at AMPA receptors, for example) neural elements will tend to integrate across temporally coincident spikes and thus the resulting signals will better maintain spike timing information. Moreover, neuromodulatory effects can alter the timescale of integration of neurons and/or networks to instantiate contextual shifts in how stimulus information is coded suggesting not only that multiple neural codes can be multiplexed but that this process may be controlled in a task-dependent manner [50,51]: under different contexts, a given circuit could extract either the rate-coded, or the timing-coded information present in its afferents.

Deep learning adds a powerful new addition to our research toolbox

In the past decade, machine learning algorithms, called deep artificial neural networks (ANNs), that loosely mimic the mammalian visual cortical hierarchy, have

transformed the field of computer vision [7]. More recently, neuroscientists have turned to these models to help understand sensory processing in the brain. Some of that work has used ANNs as theoretical abstractions to describe multi-stage information processing, akin to that which takes place in the brain. We have discussed a few of these examples above, including work that used ANNs to study how noise affects information propagation [28**] through multi-layered circuits, or how untuned neurons might contribute to object categorization [22].

Another promising approach has been to train the ANNs to perform tasks that an animal might encounter, like identifying the objects in visual scenes [8,52,53], or features in tactile stimulus patterns [54*], and then comparing the representations of sensory inputs in these trained ANNs to those observed in neural recordings. This line of work shows that, as ANNs improve at object recognition tasks, the way they represent the inputs becomes more similar to those of the primate sensory cortex [52]. Recent results suggest that these task-optimized ANNs form more accurate models of monkey V1 than do data-driven models that are directly optimized as visual cortical models [55*] by fitting the models to neural data. In other recent work the data-driven models outperform visual task-optimized ones in mimicking the responses of V1 neurons to natural image stimuli [56*]. Comparisons between studies are somewhat difficult owing to differences in the stimuli presented to the animals, differences in task conditions (e.g., passive fixation vs. active behavior vs. anaesthetized animals), and differences in the optimization procedures used to train the ANN models. A key open question is thus whether and when task-optimized ANNs form the best models of sensory neural circuits, and whether and when incorporating neural data directly in training the ANNs yields better models.

Despite this challenge, ANN models have already demonstrated their potential to help identify the visual features to which neurons respond: most simply, one takes an ANN model that is trained to respond to natural images very similarly to how a V1 neuron does, and “inverts” that model to identify the visual feature(s) that drive high predicted firing rates [56*,59**]. This approach could be especially useful for deeper parts of the visual hierarchy, like areas V2 and V4, in which neurons respond to highly complex image features, and in highly nonlinear ways [57,58,60**].

Alternatively, one can train multi-layer ANNs to reproduce the responses of neurons several synapses away from the sensory receptors (e.g., retinal ganglion cells), and then use the response properties of units in the intermediate layers of the ANN to understand the processing performed by neurons that are “between” the receptors and the the observed neurons (e.g., bipolar cells) [61*].

This approach has the potential to help understand the function of neurons, like retinal bipolar cells, that can be difficult to access in neurophysiology experiments.

Future outlook: translational implications and key challenges

A strong motivation behind our research efforts is the fact that a thorough understanding of the neural code will enable novel treatments for those with sensorimotor deficits. For example, with sufficient information about the visual neural code images could be translated into the appropriate patterns of visual cortical neural activity, which could then be “written into” the visual cortex via high-density electrode arrays or optical methods, thereby restoring sight to blind people [62]. Similarly, with sufficient information about the somatosensory neural code, a robot’s joint configuration or the surface properties of the object it is touching could be written into the somatosensory cortex, thereby improving the performance of brain-machine interfaces used by tetraplegics [63]. The recent advances described in this review could facilitate the development of these next-generation prosthetics, and machine learning is likely to be especially important for the creation of “stimulus to brain” translation algorithms [56*]. Along with providing assistive technologies for those with sensory impairments, these prosthetics are likely to revolutionize our understanding of the neural code, because they will enable experiments in which people report their subjective experience, in response to different patterns of neural activity stimulated within their sensory neurons.

Achieving these translational outcomes is likely to require us to go beyond correlative descriptions of sensory information coding, and to instead obtain causal ones: instead of describing the neurons that spike when a given stimulus is present (and assuming that their firing represents the stimulus), we may need to determine how neurons’ spiking *causes* an internal representation of the stimulus to be formed [64]. Moreover, while much of this essay describes neural population codes as static mappings between represented stimuli and neural firing rates, or inter-spike intervals, recent work highlights the possibility that the *dynamics* of the population responses could be a key – and much less thoroughly studied – feature of neural population codes [65,66]. Addressing these issues, of causality, and dynamics, is likely to yield substantial advances in our understanding of the language of the brain, and our ability to use that knowledge to help those with sensorimotor deficits.

Acknowledgements

JAP acknowledges funding from the Canadian Institutes of Health Research (Foundation Grant # 353197) and the Canada Research Chairs Program. JZ acknowledges funding from the A.P. Sloan Foundation, Google, the Canadian Institute for Advanced Research, and the Canada Research Chairs Program.

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. Averbeck BB, Latham PE, Pouget A: **Neural correlations, population coding and computation.** *Nat Rev Neurosci* 2006, **7**:358.
2. Abbott LF, Dayan P: **The effect of correlated variability on the accuracy of a population code.** *Neural Comput* 1999, **11**:91-101.
3. Ecker AS, Berens P, Tolias AS, Bethge M: **The effect of noise correlations in populations of diversely tuned neurons.** *J Neurosci* 2011, **31**:14272-14283.
4. Sompolinsky H, Yoon H, Kang K, Shamir M: **Population coding in neuronal systems with correlated noise.** *Phys Rev E* 2001, **64**:051904.
5. Shamir M: **Emerging principles of population coding: in search for the neural code.** *Curr Opin Neurobiol* 2014, **25**:140-148.
6. Josic K, Shea-Brown E, Doiron B, de la Rocha J: **Stimulus-dependent correlations and population codes.** *Neural Comput* 2009, **10**:2774-2804.
7. LeCun Y, Bengio Y, Hinton G: **Deep learning.** *Nature* 2015, **521**:436-444.
8. Yamins D, DiCarlo JJ: **Using goal-driven deep learning models to understand sensory cortex.** *Nat Neurosci* 2016, **19**:356-365.
9. Coen-Cagli R, Kohn A, Schwartz O: **Flexible gating of contextual influences in natural vision.** *Nat Neurosci* 2015, **18**:1648-1655.
10. Churchland MM *et al.*: **Stimulus onset quenches neural variability: a widespread cortical phenomenon.** *Nat Neurosci* 2010, **13**:369-378.
11. Hu Y, Zylberberg J, Shea-Brown E: **The sign rule and beyond: boundary effects, flexibility, and noise correlations in neural population codes.** *PLoS Comp Biol* 2014, **10**:e1003469.
12. da Silveira RA, Berry MJ: **High-fidelity coding with correlated neurons.** *PLoS Comp Biol* 2014, **10**:e1003970.
13. Kohn A, Coen-Cagli R, Kanitscheider I, Pouget A: **Correlations and neuronal population information.** *Ann Rev Neurosci* 2016, **39**:237-256.
14. Zylberberg J, Cafaro J, Turner MH, Shea-Brown E, Rieke F: **Direction-selective circuits shape noise to ensure a precise population code.** *Neuron* 2016, **89**:369-383.
One of the first papers to show a very strong beneficial effect for noise correlations in a neural population code. Combining theory and experiment, the authors show that this effect arises because the correlations “shape” the noise so as to be orthogonal to the signals that convey the visual information.
15. Franke F, Fiscella M, Sevelev M, Roska B, Hierlemann A, da Silveira RA: **Structures of neural correlation and how they favor coding.** *Neuron* 2016, **89**:409-422.
Shows results that are similar to the Zylberberg *et al.* 2016 paper listed above, but in rabbit retina.
16. Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A: **Information-limiting correlations.** *Nat Neurosci* 2014, **17**:1410-1417.
Shows that there is a unique pattern of noise correlation that causes information to saturate in the limit of large population sizes. Based on the data processing inequality – which shows that information must saturate in this limit – the authors argue that these “information-limiting” correlations should be common in neural circuits.
17. Lin IC, Okun M, Carandini M, Harris KD: **The nature of shared cortical variability.** *Neuron* 2015, **87**:644-656.
Shows that much of the “noise” in the responses of neurons in cat visual cortex is shared between neurons, and that much of the noise acts multiplicatively on the neurons’ firing rates. This is similar to the noise structures seen in the Franke *et al.* 2016 and Zylberberg *et al.* 2016 papers listed above. That the same effect is seen in mouse retina, rabbit retina, and cat cortex, suggests that it may be somewhat universal.
18. Kohn A, Smith MA: **Stimulus dependence of neuronal correlation in primary visual cortex of the macaque.** *J Neurosci* 2005, **25**:3661-3673.
19. Safaai H, von Heimendahl M, Sorando JM, Diamond ME, Maravall M: **Coordinated population activity underlying texture discrimination in rat barrel cortex.** *J Neurosci* 2013, **33**:5843-5855.
20. Leavitt ML, Pieper F, Sachs AJ, Martinez-Trujillo JC: **Correlated variability modifies working memory fidelity in primate prefrontal neuronal ensembles.** *Proc Natl Acad Sci USA* 2017, **114**:E2494-E2503.
Shows that, in a working memory tasks, neurons whose firing rates do not depend on task variables can still contribute to the population code. Also provides some geometric intuition for this effect.
21. Zylberberg J: **Untuned But Not Irrelevant: A Role For Untuned Neurons In Sensory Information Coding.** *bioRxiv* 134379. 2017.
Shows a result similar to Leavitt *et al.* 2017, listed above, but in visual cortex of mouse as opposed to prefrontal cortex of monkey. Also provides theoretical calculations showing that this effect generalizes to large neural populations.
22. Morcos AS, Barrett DG, Rabinowitz NC, Botvinick M: (2018). On the importance of single directions for generalization. *arXiv*: 1803.06959 [stat.ML].
23. Borst A, Theunissen FE: **Information theory and neural coding.** *Nat Neurosci* 1999, **2**:947.
24. Strong SP, Koberle R, van Steveninck RRDR, Bialek W: **Entropy and information in neural spike trains.** *Phys Rev Lett* 1998, **80**:197.
25. Bialek W, Rieke F, Van Steveninck RDR, Warland D: **Reading a neural code.** *Science* 1991, **252**:1854-1857.
26. Cayco-Gajic NA, Zylberberg J, Shea-Brown E: **Triplet correlations among similarly tuned cells impact population coding.** *Frontiers Comp Neurosci* 2015, **9**:57.
27. Zylberberg J, Shea-Brown E: **Input nonlinearities can shape beyond-pairwise correlations and improve information transmission by neural populations.** *Phys Rev E* 2015, **92**:062707.
28. Zylberberg J, Pouget A, Latham PE, Shea-Brown E: **Robust information propagation through noisy neural circuits.** *PLoS Comput Biol* 2017, **13**:e1005497.
Shows that very different neural codes arise if one optimizes for information content at the periphery vs. optimizing the amount of information that can successfully propagate through noisy and/or nonlinear downstream circuits. This suggests a re-examination of the assumption that high information at the periphery suffices to define highly effective population codes.
29. Adrian ED, Zotterman Y: **The impulses produced by sensory nerve-endings: Part II. The response of a Single End-Organ.** *J Physiol* 1926, **23**:151-171.
30. Panzeri S, Brunel N, Logothetis NK, Kayser C: **Sensory neural codes using multiplexed temporal scales.** *Trends Neurosci* 2010, **33**:111-120.
31. Panzeri S, Petersen RS, Schultz SR, Lebedev M, Diamond ME: **The role of spike timing in the coding of stimulus location in rat somatosensory cortex.** *Neuron* 2001, **29**:769-777.
32. Pruszyński JA, Johansson RS: **Edge-orientation processing in first-order tactile neurons.** *Nat Neurosci* 2014, **17**:1404-1409.
33. Johansson RS, Birznieks I: **First spikes in ensembles of human tactile afferents code complex spatial fingertip events.** *Nat Neurosci* 2004, **7**:170-177.
34. Mackevicius EL, Best MD, Saal HP, Bensmaia SJ: **Millisecond precision spike timing shapes tactile perception.** *J Neurosci* 2012, **32**:15309-15317.
35. Arabzadeh E, Panzeri S, Diamond ME: **Deciphering the spike train of a sensory neuron: counts and temporal patterns in the rat whisker pathway.** *J Neurosci* 2006, **26**:9216-9226.
36. Butts DA, Weng C, Jin J, Yeh CI, Lesica NA, Alonso JM, Stanley GB: **Temporal precision in the neural code and the timescales of natural vision.** *Nature* 2007, **449**:92-95.

37. Gollisch T, Meister M: **Rapid neural coding in the retina with relative spike latencies.** *Science* 2008, **319**:1108-1111.
38. Meister M, Berry MJ: **The neural code of the retina.** *Neuron* 1999, **22**:350-435.
39. Reinagel P, Reid RC: **Temporal coding of visual information in the thalamus.** *J Neurosci* 2000, **20**:5392-5400.
40. Mainen ZF, Sejnowski TJ: **Reliability of spike timing in neocortical neurons.** *Science* 1995, **268**:1503-1506.
41. Bair W, Koch C: **Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey.** *Neural Comput* 1996, **8**:1185-1202.
42. Srivastava KH, Holmes CM, Vellema M, Pack AR, Elemans CP, Nemenman I, Sober SJ: **Motor control by precisely timed spike patterns.** *PNAS* 2017, **114**:1171-1176.
- Shows that the precise timing of action potentials reaching the muscles affects can dramatically alter their contraction patterns. This means that, contrary to common belief, the "output" signals reaching the muscles from the nervous system are not strictly rate-coded.
43. Harvey MA, Saal HP, Dammann JF, Bensmaia SJ: **Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex.** *PLoS Biol* 2013, **11**:e1001558.
44. Zuo Y, Safaai H, Notaro G, Mazzoni A, Panzeri S, Diamond ME: **Complementary contributions of spike timing and spike rate to perceptual decisions in rat S1 and S2 cortex.** *Curr Biol* 2015, **25**:357-436.
45. Weber AI, Saal HP, Lieber JD, Cheng J, Manfredi LR, Dammann JF, Bensmaia SJ: **Spatial and temporal codes mediate the tactile perception of natural textures.** *Proc Natl Acad Sci USA* 2013, **110**:17107-17112.
46. Pirsichel F, Kretzberg J: **Multiplexed population coding of stimulus properties by leech mechanosensory cells.** *J Neurosci* 2016, **36**:3636-3647.
47. Insanally MN, Carcea I, Field RE, Rodgers CC, DePasquale B, Rajan K, DeWeese MR, Albanna BF, Froemke RC: **Nominally non-responsive frontal and sensory cortical cells encode task-relevant variables via ensemble consensus-building.** *bioRxiv* 347617. 2018.
- Presents a novel inter-spike interval (ISI) based decoder to estimate task variables from neural spike trains. Analyzing data from rat prefrontal cortex and auditory cortex, the authors show that even neurons whose firing rates do not depend on task variables can have task information encoded in their ISI patterns.
48. Pruszyński JA, Flanagan JR, Johansson RS: **Fast and accurate edge orientation processing during object manipulation.** *eLife* 2018, **7**:e31200.
- In the context of hand control, people are much better at processing edge orientation information than when they are doing a perceptual task. This work shows that task demands strongly modulate the way sensory information is interpreted by the brain. The authors offer an account related to multiplexed neural codes at the level of the tactile periphery.
49. König P, Engel AK, Singer W: **Integrator or coincidence detector? The role of the cortical neuron revisited.** *Trends Neurosci* 1996, **19**:130-137.
50. Nadim F, Bucher D: **Neuromodulation of neurons and synapse.** *Curr Opin Neurobiol* 2014, **29**:48-56.
51. Zylberberg J, Strowbridge BW: **Mechanisms of persistent activity in cortical circuits: possible neural substrates for working memory.** *Ann Rev Neurosci* 2017, **40**:603-627.
52. 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 USA* 2014, **111**:8619-8624.
53. Khaligh-Razavi SM, Kriegeskorte N: **Deep supervised, but not unsupervised, models may explain IT cortical representation.** *PLoS Comp Biol* 2014, **10**:e1003915.
54. Zhao CW, Daley MJ, Pruszyński JA: **Neural network models of the tactile system develop first-order units with spatially complex receptive fields.** *PLoS One* 2018, **13**:e0199196.
- Trained ANNs under a set of realistic biological constraints to discriminate between tactile input patterns and found that the responses of units in the trained networks resembled those of first-order tactile neurons. This suggests that ANNs could drive progress in tactile neuroscience, much as they have in visual neuroscience.
55. Cadena SA, Denfield GH, Walker EY, Gatys LA, Tolia AS, Bethge M, Ecker AS: **Deep convolutional models improve predictions of macaque V1 responses to natural images.** *bioRxiv* 201764. 2017.
- Constructed models to predict V1 neuron firing rates from natural images, using LASSO regression against the activations of units in the intermediate layers of a deep object recognition network called VGG. Showed that this model, constructed from a task-trained ANN, outperformed data-driven models where ANNs were directly trained to take in images and return neural firing rates.
56. Kindel WF, Christensen ED, Zylberberg J: **Using deep learning to probe the neural code for images in primary visual cortex.** *J Vision* 2019, **19**:29.
- Presents a counterpoint to the Cadena et al. result described above: Kindel et al. instead find that data-driven models outperform ones built from task-trained neural networks, in predicting V1 firing rates. This paper also presents a method to use deep learning to identify the stimuli that drive V1 neurons to spike. This method shows the potential to identify neurons that respond to non-localized image features—like diffuse textures—that have not previously been observed in V1.
57. Kalfas I, Vinken K, Vogels R: **Representations of regular and irregular shapes by deep Convolutional Neural Networks, monkey inferotemporal neurons and human judgments.** *PLoS Comp Biol* 2018, **14**:e1006557.
58. Pasupathy A, Connor CE: **Shape representation in area V4: position-specific tuning for boundary conformation.** *J Neurophys* 2001, **86**:2505-2519.
59. Walker EY et al.: **Inception in visual cortex: in vivo-silico loops reveal most exciting images.** *bioRxiv* 506956. 2018.
- Walker et al. trained ANNs to take in images and predict the firing rates of mouse V1 neurons. They then synthesized from the neural networks images that were optimized to drive high firing rates in each neuron. Walker et al. then presented these tailored images to the animals, and found that they drove firing rates similar to those predicted by the ANNs.
60. Bashivan P, Kar K, DiCarlo JJ: **Neural population control via deep image synthesis.** *Science* 2019, **364**:eaav9436.
- Bashivan et al. trained ANNs to take in images and predict the firing rates of monkey V4 neurons. They then synthesized from the neural networks images that were optimized to drive high firing rates in each neuron. Bashivan et al. presented these tailored images to the monkeys, and found that they drove firing rates similar to those predicted by the ANNs.
61. 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* (p 2016:1369-1377.
- One of the first studies to apply deep learning directly to neural data. Trained deep neural networks to predict retinal ganglion cell firing rates in response to natural images. Found that units in intermediate layers of the deep network have response properties that resemble those of bipolar cells in retina.
62. Bosking WH, Beauchamp MS, Yoshor D: **Electrical stimulation of visual cortex: relevance for the development of visual cortical prosthetics.** *Ann Rev Vis Sci* 2017, **3**:141-166.
63. Bensmaia SJ, Miller LE: **Restoring sensorimotor function through intracortical interfaces: progress and looming challenges.** *Nat Rev Neuro* 2014, **15**:313-325.
64. Panzeri S, Harvey CD, Piasini E, Latham PE, Fellin T: **Cracking the neural code for sensory perception by combining statistics, intervention, and behavior.** *Neuron* 2017, **93**:491-507.
65. Pandarinath C et al.: **Inferring single-trial neural population dynamics using sequential auto-encoders.** *Nat Methods* 2018, **15**:805-815.
66. Zylberberg J, Hyde RA, Strowbridge BW: **Dynamics of robust pattern separability in the hippocampal dentate gyrus.** *Hippocampus* 2016, **26**:623-632.