Towards the design principles of neural population codes
Elad Schneidman

The ability to record the joint activity of large groups of neurons would allow for direct study of information representation and computation at the level of whole circuits in the brain. The combinatorial space of potential population activity patterns and neural noise imply that it would be impossible to directly map the relations between stimuli and population responses. Understanding of large neural population codes therefore depends on identifying simplifying design principles. We review recent results showing that strongly correlated population codes can be explained using minimal models that rely on low order relations among cells. We discuss the implications for large populations, and how such models allow for mapping the semantic organization of the neural codebook and stimulus space, and decoding.

Address
Department of Neurobiology, Weizmann Institute of Science, Rehovot, Israel

Corresponding author: Schneidman, Elad
(elad.schneidman@weizmann.ac.il)

Redundancy reduction vs. robustness and learnability

Efficient coding has been a key prospective design principle of the neural code — suggesting that neuronal spiking patterns convey information in a way that maximizes their capacity, given constraints such as energetic cost, number of neurons, time, etc. [1]. Accordingly, groups of neurons would maximize the information they convey about their stimuli by minimizing the overlap in the information they carry [2]. This idea of *redundancy reduction* implies that the responses of cells would be decorrelated, and was used to predict or explain the nature of the receptive fields and response properties of cells in different parts of the visual system, the nature of eye movements, and more [3–7].

However, correlations among neurons are fundamental for robustness of the code to abundant neuronal noise. Moreover, redundancy is exactly what makes a code learnable from examples [8]. It is clear then that correlations must be an inherent part of the design of the neural code. Indeed, analysis of large neuronal populations has shown that large groups can be strongly correlated (see below) and the potential importance and benefits of these correlations have been extensively discussed [5,9–13].

The design of the code of large populations of neurons in terms of correlation structures, connectivity, and sparseness have been explored theoretically (see e.g. [14–16]), but the ability to record simultaneously from large populations of cells makes it possible to study population codes directly [17–24]. Importantly, neuronal variability or noise means that in studying these codes we need to use probability distributions over the vast space of combinatorial coding patterns. Thus, we will not be able to exhaustively map the code of large populations by direct sampling, and would need to find simplifying principles that govern the organization of the codebook and its functional role [25,26•,27,28].

Correlation structures of population codes

In many neural systems, the correlation between pairs of neurons is typically weak. This is the case for the correlation between cells averaged over stimuli (sometimes called signal correlations): denoting the spiking patterns of cells i and j with x_i and x_j, $P(x_i, x_j)$ is typically close to $P(x_i)P(x_j)$ [11,25,29]. Similarly, the pairwise correlations in responses to repeated presentation of the same stimulus, s, are weak: $P(x_i, x_j|s)$ is typically not far from $P(x_i|s)P(x_j|s)$ [5,30,31], and so averaging over stimuli, cell pairs have been reported to have low 'noise correlations' (a somewhat misleading name as the difference between these two distributions measures stimulus dependent coding correlations). Moreover, the information that pairs carry about their stimuli is typically weakly redundant: $I(x_i, x_j ; s)$ is usually a little lower than $I(x_i ; s) + I(x_j ; s)$ [30,32,33], where $I(x ; y)$ is the mutual information between x and y. All these seem consistent with the idea of redundancy reduction.

However, groups of a few tens of cells can be strongly correlated, even when their typical pairwise correlations are weak [25,34,35,36,37•,38•,39]. These group correlations — quantified by the difference between $P(x_1, x_2, \ldots, x_n)$ and $P(x_1)P(x_2) \cdots P(x_N)$ — mean that the 'vocabulary' of the neural population is more structured and less

diverse than it would have been if the cells were independent, given the coding properties of single cells. Similarly, the stimulus-dependent distribution of neural responses $P(x_1, x_2, \ldots, x_n|s)$ can be far from the conditionally independent one $P(x_1|s)P(x_2|s)\cdots P(x_N|s)$, for particular stimulus s, showing that while noise correlations may be weak on average, they can be strong for interesting stimuli and play a role in coding [40,41[•],42[•]].

Models of population activity reveal sparse high-order organization of the neural code

To understand and quantify the nature and origin of the correlated structure of population codes, several families of models have been commonly used. *Generalized linear models* (GLMs) extend the Linear-nonlinear Poisson rate models of single cells to predict the instantaneous rate of neuron i, given by $\lambda_i(t) = \exp\left(k_i \cdot s + h_i \cdot x_i + \sum_j l_{ij} \cdot x_j + \mu_i\right)$ where s is the stimulus, x_j is the spiking history of cell j at time t, μ_i is the log of the cell's baseline firing rate, k_i is a spatio-temporal filter describing the cell's response to the stimulus, l_{ij} is the incoming coupling of cell j to i, and h_i is a temporal filter on the cell's own history; the last three are fit to maximize the likelihood of the data under this model. These models surpass conditionally independent models in predicting the activity of single cells and low-order correlations between cells [43,44], give a causal interpretation of the temporal patterns of cells, and have been successfully used in stimulus decoding [45[•],46,47,48,49]. The maximum entropy (ME) framework offers the mathematically minimal probabilistic model of population activity that satisfies a set of desired constraints without making arbitrary assumptions [50]. For example, if we denote the activity of cell i in a small time bin by $x_i = 0$ for silence, and $x_i = 1$ for spiking, then the most random and unique model that obeys the firing rates and pairwise correlations between cells, is given by $P(x_1, x_2, \ldots, x_n) = (1/Z)\exp\left\{\sum_i \alpha_i x_i + (1/2)\sum_{i\neq j}\beta_{ij}x_i x_j\right\}$, where the set of α_i's and β_{ij}'s are chosen to obey the constraints, and Z is a normalization factor (or partition function). Such pairwise models capture with high accuracy the probability of individual spatial spiking patterns of tens of cells in different systems at the level of individual combinatorial patterns [25,34,35,39,51,52,53[•],54]. Restricted Boltzmann Machines offer a natural extension for capturing high-order relations between cells [55[•]]. Figure 1 presents an example of the accuracy of pairwise ME models in capturing population activity patterns at the resolution of individual population patterns.

The accuracy of these models reflect the low-dimensional nature of population codes. For example, the accurate pairwise ME model relies on $n(n+1)/2$ parameters, compared with the 2^n that would be needed for the distribution of binary activity patterns of n cells in the general case, and the n parameters needed for the independent model. Even these pairwise model may be further

reduced while retaining accuracy by taking local structure or discretized interaction values into account [52].

For larger populations, an accurate model of the vocabulary of activity patterns may require models that go beyond pairs, in particular for spatio-temporal patterns [36,39,53[•],54,55[•],56]. Adding a small number of high-order terms such as the distribution of synchrony values of the population or a selected set of high-order terms are sufficient for ~100 cells responding to natural stimuli [53[•],54,57]. These high-order dependencies may arise from the nature of the stimulus, the internal structure of larger populations, or the activity of unseen (hidden) neurons [53[•],55[•],58,59].

Extending this approach to the dynamics of population responses to stimuli, the stimulus-dependent maximum entropy models (SDME) of population encoding capture population responses significantly better than conditionally independent models (Figure 2A). Moreover, the importance of including correlations in population encoding models becomes more apparent with population size [40,41[•],42[•]] (see e.g. Figure 2b).
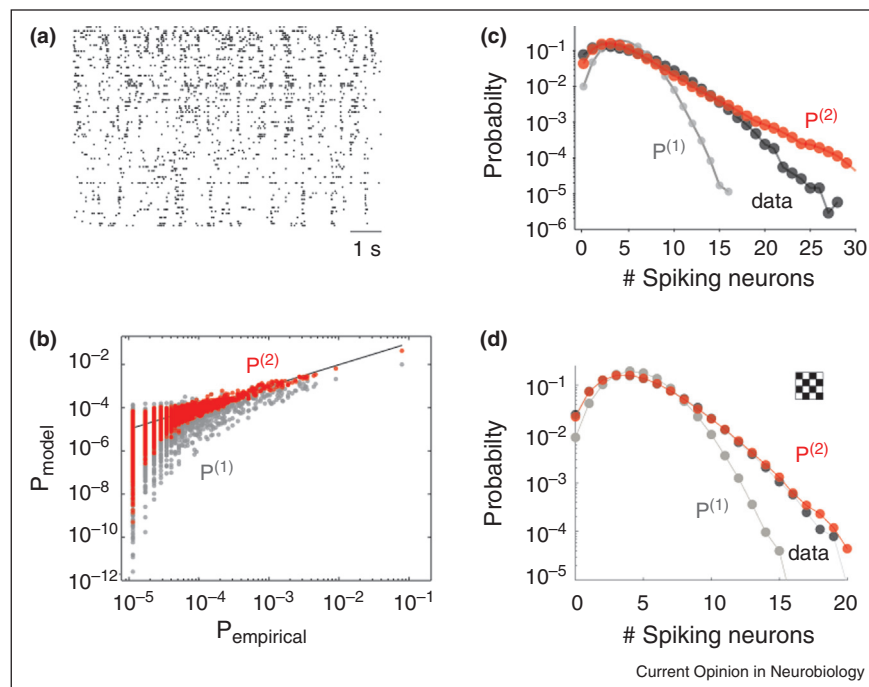
Mapping the semantic space of population codes, and a neural metric on stimulus space

The correlated activity of neuronal populations suggests that we should revise the classic view of the receptive field of a neuron with that of the collective response properties of populations [27,28,60[•]]. To understand information representation and computation in large populations, we must go beyond detailed analysis of combinatorial coding patterns [61–63] to characterize the 'semantic organization' of the spaces of stimuli and responses. Characterizing the noisy nature of mapping between stimuli and responses would enable deciphering the meaning of novel patterns that we have not seen before, and decode new stimuli that have been presented before. But how can we learn such neural dictionaries by generalizing from unavoidably severely undersampled set of observations of stimuli and responses, without making arbitrary assumptions?

Quantifying the similarity of population activity patterns $\{x_i\}$ and $\{x_j\}$ by their semantic overlap, namely, how similar are the stimuli these patterns are used to encode by comparing $P(s|\{x_i\})$ and $P(s|\{x_j\})$ — gives an assumption-free way to map the organization of the neural codebook (Figure 3a). In the vertebrate retina this revealed that the neural codebook of groups of 20 cells is organized in a relatively small number of semantic clusters that capture almost all the information carried by the population (Figure 3b and c). Importantly, metrics like Hamming distance, edit-distance, or other linear or bi-linear metrics fail to capture this organization of the space of population patterns [42[•]].

On the decoding side, mean-squared-error (MSE) has been commonly used to measure how well sensory stimuli

Figure 1

Example of the high accuracy of a pairwise maximum entropy models for spatial activity patterns of a large network responding to natural stimuli, and signs of higher order interactions. **(a)** A segment of the simultaneous action potential sequences of 99 ganglion cells in the salamander retina responding to a natural movie clip. Each line corresponds to a single neuron, each tick represents a single spike. **(b)** The probability of occurrence of each simultaneous (spatial) population activity pattern that appeared in the experiment as predicted if all cells were independent ($P^{(1)}$, grey), or by the 2nd order maximum entropy model, which takes into account pairwise correlations ($P^{(2)}$, red), are plotted against the measured rate. Note that while most rare patterns fall within the confidence region (not shown), frequently observed patterns are misestimated by the pairwise model. Adding a small number of high-order dependencies in this case gives a nearly perfect fit (not shown). **(c)** Probability distribution of synchronous spiking events in the 99 cell population in response to a long natural movie (black). The distribution of synchronous events for the same 99 cells predicted by the independent model (grey, $P^{(1)}$), and the synchrony distribution predicted from the 2nd order maximum entropy model $P^{(2)}$ (red) are shown **(d)**. Probability distribution of synchronous spiking events in the 99 cell population in response to artificial white noise stimuli (black). 2nd order model, and independent model, built as described in (c). Unlike the responses to natural movies, the pairwise model provides a very good fit to responses to white noise stimuli, reflecting a negligible role for higher order interactions in the population activity evoked by such stimuli.
Adapted from [53•].

or movements are reconstructed from neural activity. The similarity between two stimuli s_i and s_j, can be quantified by the overlap of the neural responses they elicit, $P(\{x\}|s_i)$ and $P(\{x\}|s_j)$ — again, without the need of arbitrary assumptions of similarity. For the vertebrate retina this showed that similarity as judged by the retina was very different than what we would expect from Euclidean metrics, and MSE in particular [40,41•] (Figure 2c).
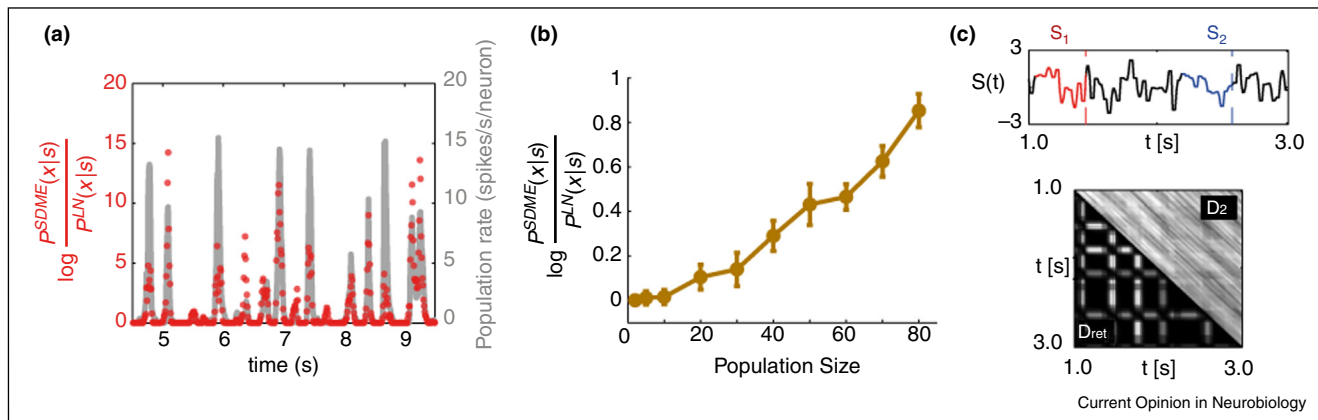
Thus, the brain's 'metrics' on stimulus space and on neural space are very different than our intuitive notion of similarity, and the clustered semantic organization of the space of neural activity patterns cannot be inferred from simple intuitive notions of what would constitute spiking similarities (see [42•]). Extending these ideas beyond sensory circuits and linking the brain metrics on stimuli and neural activity patterns would be necessary for our understanding of the code of high brain areas.

Neural population codes of larger populations and cognitive processing

Cognitive processing relies on computational architectures that are far more complicated than those of sensory circuits. Higher brain areas are usually built of many more cells, display a myriad of recurrent connections, and no apparent simple correlation structures. These imply intricate spatio-temporal processing that still empowers fast classification, decisions, and actions.

Experimental advances that would enable accurate recordings from thousands, millions, or all the neurons in an animal's brain at single spike resolution [18,22,64], would still suffer from the inability to quantify directly the full nature of correlations in large populations because of sampling issues. In many cases no specific population activity pattern would repeat itself over any experimental time window [53•]. Can smaller networks reveal simplifying

Figure 2

Stimulus dependent ME model captures noise correlations between cells, and allow for learning a metric on stimulus space, based on the neural population. **(a)** Pairwise SDME model predicts population activity patterns for $N = 100$ retinal ganglion cells responding to Gaussian white noise stimulus much better than conditionally independent model. The log-likelihood ratio of the population firing patterns under the SDME model and under the conditional independent model, shown as a function of time (red dots, scale at left) for an example stimulus repeat (models fitted on train data; example shown is for test data). For reference, the average time dependent population firing rate is shown in grey (scale at right). **(b)** The performance of the SDME model relative to conditionally independent LN models. The average log likelihood ratio between the SDME and the conditional independent models evaluated on the test set, as a function of the population size, N (error bars = std over 10 randomly chosen groups of neurons at that N). **(c)** An example of learning a metric on stimulus space from population neural responses: Stimulus segment with two 400 ms stimulus clips s_1 (red), s_2 (blue). Bottom: For every pair of time bins in the experiment, the Euclidean distance D_2 between the corresponding stimulus clips is shown in the upper diagonal part of the matrix, and the retinal neurons' distance D_{ret} in the lower part. Adapted from [40,41[•]].

principles for large networks? The organization of the code of a few hundred cells in sensory systems already diverges from the view at the level of tens of cells — in terms of higher-order interactions [53[•],54,65[•]], and the spatial decay of functional dependencies [52]— suggesting that there may be a typical functional module size. A promising direction has been given by an example of learning the codebook of a large population by combining many models of small overlapping subnetworks, suggesting a hierarchical organization of learnable high-order codes [53[•],66].

An especially intriguing feature of the vocabulary of neural populations is the distribution of frequencies of observed activity patterns, which resembles a power-law [53[•],54]. While some of this behavior might be expected from the nature of collective responses with hidden elements (be it the stimulus or other neurons) [53[•],59,67,68], it has been suggested to reflect design rather than epiphenomena of the code (see also [69]). Interestingly, this structure allows for learning the Reliable Interaction model (an approximate non-normalized ME model), which finds the dominant interactions of any order based on the frequent population patterns. In the retina, this model gave a very sparse high-order model with unparalleled accuracy in predicting the appearance of individual activity patterns in test data — showing that the code learnable from examples [53[•]]. More broadly,
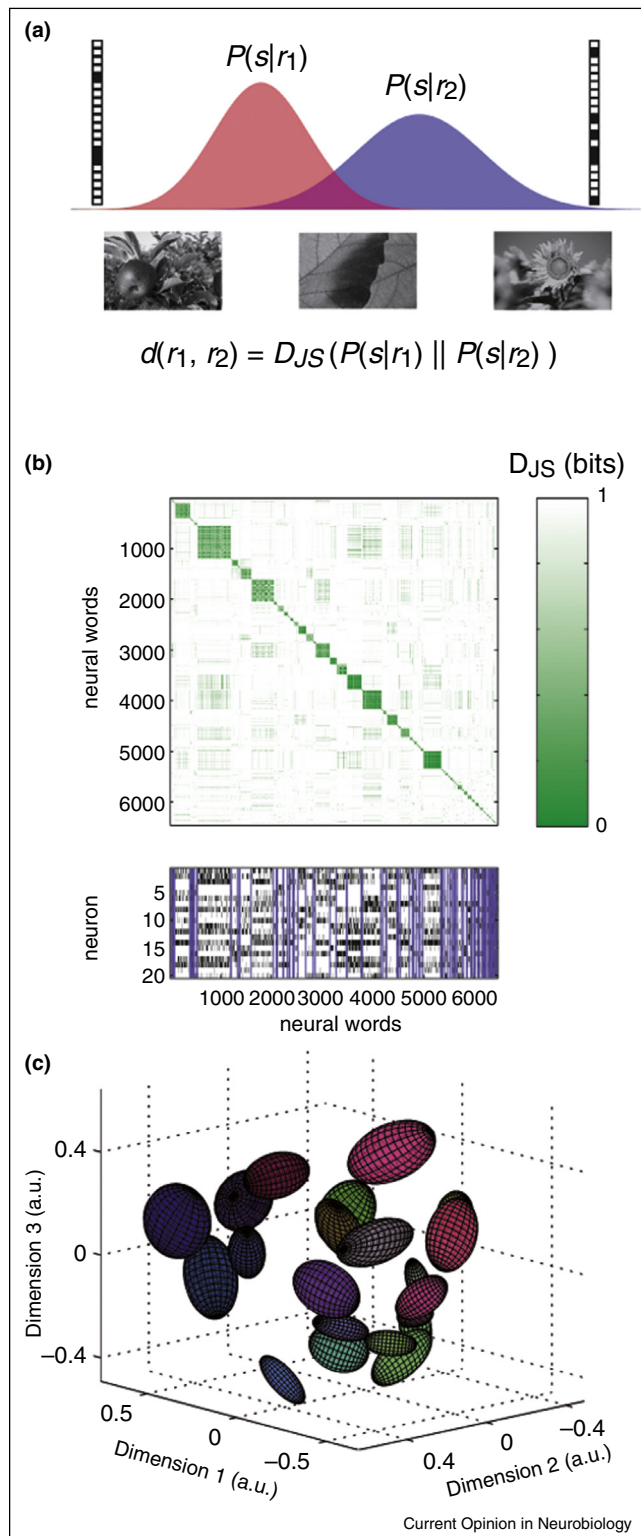
the Zipf-like distribution of population states is reminiscent of physical systems near critical behavior, congruent with the possibility of collective behavior states of the code of large groups of neurons that allow for a wide range of distinct responses that would be easily decodable [65[•],70]. Signatures of such critical behavior and its possible role as a feature of large networks in coding and adaptation have been presented in several different systems [65[•],70,71].

Merging these ideas of codebook generalization, sparse high-order interactions, semantic clusters, and the collective behavior of large populations, would be fundamental to our ability to understand the design principles of large networks and cognitive computations.

What could the nature of spatio-temporal codes reveal about the cognitive computations they perform?

The functional interactions between cells that underlie correlated codes set a form of a prior over population activity patterns, thus shaping the way stimuli would be perceived and encoded. Analysis of simple population models suggests that the interactions between cells that would maximize the information they carry about their stimuli, are the ones that would result in spontaneous activity that resembles the patterns that would be used to encode the stimuli [9,72]

Figure 3

$$d(r_1, r_2) = D_{JS}(P(s|r_1) \| P(s|r_2))$$

Current Opinion in Neurobiology

Mapping the semantic organization of population codebook, reveals a small set of semantic clusters. **(a)** A cartoon of the way the semantic similarity between neural activity patterns was estimated. For two neural words, we inferred the distribution over stimuli given the neural word, using an SDME model for the neural population response. Then the similarity between patterns was estimated using the Jensen–Shannon divergence between the distributions, d. **(b)** The neural codebook of the retina is comprised of distinct clustered responses with highly similar meaning. Example of the similarity matrix of the population responses of representative group of 20 neurons to a natural movie. Each entry in the matrix corresponds to the similarity d between two population responses ('words') observed in the test data. Matrix rows (and columns) are ordered according to the result of clustering the words using k-means. On the bottom we show the population responses corresponding to the entries in the matrix; the blue lines mark borders between different clusters. **(c)** A graphical representation of the organization of the codebook of these cells: The responses belonging to clusters that contain 30–300 patterns were embedded in 3D using a dimensionality reduction algorithm (The Euclidean distance in the plot approximates the similarity measure d). Each ellipse represents the 1 STD Gaussian fit to all responses belonging to a single cluster. The coordinates also correspond to the RGB value of each ellipse, thus nearby clusters share similar colors. Adapted from [42•].

— in accordance with experimental results of spontaneous (ongoing) cortical activity [37•,51,73,74,75,76•]. Moreover, rather than encoding specific stimuli, it has been suggested that neurons may be representing the likelihood of their inputs — under an internally learned model [77–79]. Hierarchical population codes may allow for neurons to learn and encode these likelihoods as biological plausible models suggest [80–83].

Extending the models we presented here, for the spatio-temporal dynamics of population activity in high brain areas, would allow us to explore the computation or algorithms that these circuits implement. Ultimately, mapping the way population codes change or adapt on short time scales [84,85•], and during learning, would delineate the dynamics of individual cells and of groups in terms of stimulus features or information they learn to represent and compute [86,87].

Finally, the ability to map the layout of connections in large networks [88] and to 'put numbers' on these connections [89,90•], and a model-based description of population activity patterns of the kind we presented here, would hopefully allow for understanding the relations between neural architecture and function and the development of neural circuits, and allow for tinkering and engineering of networks and neural computation.

Conflict of interest statement
Nothing declared.

Acknowledgements

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. Barlow HB: **Sensory Communication**. The MIT Press; 1962.

2. Atick JJ: **Could information theory provide an ecological theory of sensory processing?** *Netw Comput Neural Syst* 1992, **3**:213-251.

3. Atick JJ, Redlich AN: **What does the retina know about natural scenes?** *Neural Comput* 1992, **4**:196-210.

4. Segal IY, Giladi C, Gedalin M, Rucci M, Ben-Tov M, Kushinsky Y, Mokeichev A, Segev R: **Decorrelation of retinal response to natural scenes by fixational eye movements**. *Proc Natl Acad Sci U S A* 2015, **112**:3110-3115.

5. Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS: **Decorrelated neuronal firing in cortical microcircuits**. *Science* 2010, **327**:584-587.

6. Vinje WE, Gallant JL: **Sparse coding and decorrelation in primary visual cortex during natural vision**. *Science* 2000, **287**:1273-1276.

7. Dan Y, Atick JJ, Reid RC: **Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory**. *J Neurosci* 1996, **16**:3351-3362 May.

8. Barlow H: **Redundancy reduction revisited**. *Network* 2001, **12**:241-253.

9. Tkacik G, Prentice JS, Balasubramanian V, Schneidman E: **Optimal population coding by noisy spiking neurons**. *Proc Natl Acad Sci U S A* 2010, **107**:14419-14424.

10. Averbeck BB, Latham PE, Pouget A: **Neural correlations, population coding and computation**. *Nat Rev Neurosci* 2006, **7**:358-366.

11. Cohen MR, Kohn A: **Measuring and interpreting neuronal correlations**. *Nat Neurosci* 2011, **14**:811-819.

12. Panzeri S, Pola G, Petroni F, Young MP, Petersen RS: **A critical assessment of different measures of the information carried by correlated neuronal firing**. *Biosystems* 2002, **67**:177-185.

13. Schneidman E, Bialek W, Berry MJ: **Synergy, redundancy, and independence in population codes**. *J Neurosci* 2003, **23**:11539-11553.

14. Abbott LF, Dayan P: **The effect of correlated variability on the accuracy of a population code**. *Neural Comput* 1999, **11**:91-101.

15. Olshausen BA, Field DJ: **Emergence of simple-cell receptive field properties by learning a sparse code for natural images**. *Nature* 1996, **381**:607-609.

16. Kang K, Sompolinsky H: **Mutual information of population codes and distance measures in probability space**. *Phys Rev Lett* 2001, **86**:4958-4961.

17. Ahrens MB, Orger MB, Robson DN, Li JM, Keller PJ: **Whole-brain functional imaging at cellular resolution using light-sheet microscopy**. *Nat Methods* 2013, **10**:413-420.

18. Prevedel R, Yoon Y-G, Hoffmann M, Pak N, Wetzstein G, Kato S, Schrödel T, Raskar R, Zimmer M, Boyden ES, Vaziri A: **Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy**. *Nat Methods* 2014, **11**:727-730.

19. Szuts TA, Fadeyev V, Kachiguine S, Sher A, Grivich MV, Agrochão M, Hottowy P, Dabrowski W, Lubenov EV, Siapas AG, Uchida N, Litke AM, Meister M: **A wireless multi-channel neural amplifier for freely moving animals**. *Nat Neurosci* 2011, **14**:263-269.

20. Aronov D, Tank DW: **Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system**. *Neuron* 2014, **84**:442-456.

21. Berényi A, Somogyvári Z, Nagy AJ, Roux L, Long JD, Fujisawa S, Stark E, Leonardo A, Harris TD, Buzsáki G: **Large-scale, high-density (up to 512 channels) recording of local circuits in behaving animals**. *J Neurophysiol* 2014, **111**:1132-1149.

22. Hochbaum DR, Zhao Y, Farhi SL, Klapoetke N, Werley CA, Kapoor V, Zou P, Kralj JM, Maclaurin D, Smedemark-Margulies N, Saulnier JL, Boulting GL, Straub C, Cho YK, Melkonian M, Wong GK-S, Harrison DJ, Murthy VN, Sabatini BL, Boyden ES, Campbell RE, Cohen AE: **All-optical electrophysiology in mammalian neurons using engineered microbial rhodopsins**. *Nat Methods* 2014, **11**:825-833.

23. Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, El Gamal A, Schnitzer MJ: **Long-term dynamics of CA1 hippocampal place codes**. *Nat Neurosci* 2013, **16**:264-266.

24. Freeman J, Field GD, Li PH, Greschner M, Gunning DE, Mathieson K, Sher A, Litke AM, Paninski L, Simoncelli EP, Chichilnisky E: **Mapping nonlinear receptive field structure in primate retina at single cone resolution**. *Elife* 2015, **4**:e05241.

25. Schneidman E, Berry MJ, Segev R, Bialek W: **Weak pairwise correlations imply strongly correlated network states in a neural population**. *Nature* 2006, **440**:1007-1012.

26. Cunningham JP, Yu BM: **Dimensionality reduction for large-**
• **scale neural recordings**. *Nat Neurosci* 2014, **17**:1500-1509.
Reviews dimensionality reduction approaches to characterization of large neural population activity patterns.

27. Singer W: **Cortical dynamics revisited**. *Trends Cogn Sci* 2013, **17**:616-626.

28. Buzsáki G: **Neural syntax: cell assemblies, synapsembles, and readers**. *Neuron* 2010, **68**:362-385.

29. Smith MA, Kohn A: **Spatial and temporal scales of neuronal correlation in primary visual cortex**. *J Neurosci* 2008, **28**:12591-12603.

30. Puchalla JL, Schneidman E, Harris RA, Berry MJ: **Redundancy in the population code of the retina**. *Neuron* 2005, **46**:493-504.

31. Cafaro J, Rieke F: **Noise correlations improve response fidelity and stimulus encoding**. *Nature* 2010, **468**:964-967.

32. Pola G, Thiele A, Hoffmann KP, Panzeri S: **An exact method to quantify the information transmitted by different mechanisms of correlational coding**. *Network* 2003, **14**:35-60.

33. Reich DS, Mechler F, Victor JD: **Independent and redundant information in nearby cortical neurons**. *Science* 2001, **294**:2566-2568.

34. Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, Litke AM, Chichilnisky EJ: **The structure of multi-neuron firing patterns in primate retina**. *J Neurosci* 2006, **26**:8254-8266.

35. Shlens J, Field GD, Gauthier JL, Greschner M, Sher A, Litke AM, Chichilnisky EJ: **The structure of large-scale synchronized firing in primate retina**. *J Neurosci* 2009, **29**:5022-5031.

36. Marre O, El Boustani S, Frégnac Y, Destexhe A: **Prediction of spatiotemporal patterns of neural activity from pairwise correlations**. *Phys Rev Lett* 2009, **102**:138101.

37. Berkes P, Orbán G, Lengyel M, Fiser J: **Spontaneous cortical**
• **activity reveals hallmarks of an optimal internal model of the environment**. *Science* 2011, **331**:83-87.
Show that with age spontaneous and stimulus driven cortical population activity become more similar, suggesting spontaneous activity serving as a prior over neural responses.

38. Truccolo W, Hochberg LR, Donoghue JP: **Collective dynamics in**
• **human and monkey sensorimotor cortex: predicting single neuron spikes**. *Nat Neurosci* 2010, **13**:105-111.
Show that spiking history of small and randomly sampled ensembles of cortical neurons could predict subsequent single neuron spiking with high accuracy and better than by the ensemble's instantaneous state.

39. Tang A, Jackson D, Hobbs J, Chen W, Smith JL, Patel H, Prieto A, Petrusca D, Grivich MI, Sher A, Hottowy P, Dabrowski W, Litke AM, Beggs JM: **A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro**. *J Neurosci* 2008, **28**:505-518.

40. Granot-Atedgi E, Tkačik G, Segev R, Schneidman E: **Stimulus-dependent maximum entropy models of neural population codes**. *PLoS Comput Biol* 2013, **9**:e1002922.

41. Tkačik G, Granot-Atedgi E, Segev R, Schneidman E: **Retinal
•   metric: a stimulus distance measure derived from population neural responses**. *Phys Rev Lett* 2013, **110**:1-5.
Present a way to learn a neural based metric on stimulus space, and show it is significantly different from commonly used Euclidean metrics.

42. Ganmor E, Segev R, Schneidman E: **A thesaurus for a neural
•   population code**. *Elife* 2015, **4**:e06134.
Show that that population codebook is organized as a set of semantic clusters, which cannot be inferred from syntactic metrics on population patterns, and that most of the information carried by neuronal population is captured by the cluster a pattern belongs to.

43. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP: **Spatio-temporal correlations and visual signalling in a complete neuronal population**. *Nature* 2008, **454**:995-999.

44. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN: **A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects**. *J Neurophysiol* 2005, **93**:1074-1089.

45. Pillow JW, Ahmadian Y, Paninski L: **Model-based decoding,
•   information estimation, and change-point detection techniques for multineuron spike trains**. *Neural Comput* 2011, **23**:1-45.
Present efficiently learnable families of models that predict population responses to stimuli and used for stimulus decoding.

46. Pillow JW, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ: **Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model**. *J Neurosci* 2005, **25**:11003-11013.

47. Ahmadian Y, Pillow JW, Paninski L: **Efficient Markov chain Monte Carlo methods for decoding neural spike trains**. *Neural Comput* 2011, **23**:46-96.

48. Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA: **A Statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells**. *J Neurosci* 1998, **18**:7411-7425.

49. Brown EN, Kass RE, Mitra PP: **Multiple neural spike train data analysis: state-of-the-art and future challenges**. *Nat Neurosci* 2004, **7**:456-461.

50. Jaynes ET: **Information Theory and Statistical Mechanics**. *Phys Rev* 1957, **106**:620-630.

51. Marre O, Yger P, Davison AP, Frégnac Y: **Reliable recall of spontaneous activity patterns in cortical networks**. *J Neurosci* 2009, **29**:14596-14606.

52. Ganmor E, Segev R, Schneidman E: **The architecture of functional interaction networks in the retina**. *J Neurosci* 2011, **31**:3044-3054.

53. Ganmor E, Segev R, Schneidman E: **Sparse low-order
•   interaction network underlies a highly correlated and learnable neural population code**. *Proc Natl Acad Sci U S A* 2011, **108**:9679-9684.
Show that high-order interactions are needed to describe the vocabulary of large neural populations, and a way to learn a highly sparse set of these interactions, suggesting the code is learnable.

54. Tkačik G, Marre O, Amodei D, Schneidman E, Bialek W, Berry MJ: **Searching for collective behavior in a large network of sensory neurons**. *PLoS Comput Biol* 2014, **10**:e1003408.

55. Köster U, Sohl-Dickstein J, Gray CM, Olshausen BA: **Modeling
•   higher-order correlations within cortical microcolumns**. *PLoS Comput Biol* 2014, **10**:e1003684.
Compare different families of models to capture the high order structure of cortical population activity patterns.

56. Amari S-I, Nakahara H: **Correlation and independence in the neural code**. *Neural Comput* 2006, **18**:1259-1267.

57. Tkacik G, Marre O, Mora T, Amodei D, Berry MJ, Bialek W: **The simplest maximum entropy model for collective behavior in a neural network**. *J Stat* 2012:1-5.

58. Macke JH, Opper M, Bethge M: **Common input explains higher-order correlations and entropy in a simple model of neural population activity**. *Phys Rev Lett* 2011, **106**:208102.

59. Schneidman E, Still S, Berry MJ, Bialek W: **Network information and connected correlations**. *Phys Rev Lett* 2003, **91**:238701.

60. Yuste R: **From the neuron doctrine to neural networks**. *Nat Rev
•   Neurosci* 2015, **16**:487-497.
Reviews the need to change our view of neuronal function and encoding from single neurons to ensembles.

61. Osborne LC, Palmer SE, Lisberger SG, Bialek W: **The neural basis for combinatorial coding in a cortical population response**. *J Neurosci* 2008, **28**:13522-13531.

62. Schneidman E, Puchalla JL, Segev R, Harris RA, Bialek W, Berry MJ II: **Synergy from silence in a combinatorial neural code**. *J Neurosci* 2011, **31**:15732-15741.

63. Brenner N, Strong SP, Köberle R, Bialek W, van Steveninck RRdeR: **Synergy in a neural code**. *Neural Comput* 2000, **12**:1531-1552.

64. Keller PJ, Ahrens MB: **Visualizing whole-brain activity and development at the single-cell level using light-sheet microscopy**. *Neuron* 2015, **85**:462-483.

65. Tkačik G, Mora T, Marre O, Amodei D, Palmer SE, Berry MJ,
•   Bialek W: **Thermodynamics and signatures of criticality in a network of neurons**. *Proc Natl Acad Sci U S A* 2015, **112**:11508-11513.
Present the divergence of codebook properties for large neural populations, suggesting critical behavior of the code.

66. Soudry D, Keshri S, Stinson P, Oh M-H, Iyengar G, Paninski L: **Efficient 'shotgun' inference of neural connectivity from highly sub-sampled activity data**. *PLoS Comput Biol* 2015, **11**:e1004464.

67. Schwab DJ, Nemenman I, Mehta P: **Zipf's law and criticality in multivariate data without fine-tuning**. *Phys Rev Lett* 2014, **113**:1-5.

68. Aitchison L, Corradi N, Latham PE: *Zipf's law arises naturally in structured high-dimensional data*. 2014:. arXiv:1407.7135.

69. Buzsáki G, Mizuseki K: **The log-dynamic brain: how skewed distributions affect network operations**. *Nat Rev Neurosci* 2014, **15**:264-278.

70. Tetzlaff C, Okujeni S, Egert U, Wörgötter F, Butz M: **Self-organized criticality in developing neuronal networks**. *PLoS Comput Biol* 2010, **6**:e1001013.

71. Mora T, Deny S, Marre O: **Dynamical criticality in the collective activity of a population of retinal neurons**. *Phys Rev Lett* 2015, **114**:078105.

72. Berkes P, Orbán G, Lengyel M, Fiser J: **Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment**. *Science* 2011, **331**:83-87.

73. Arieli A, Sterkin A, Grinvald A, Aertsen A: **Dynamics of ongoing activity**. *Science* 1996, **273**:1868-1871.

74. Tsodyks M, Kenet T, Grinvald A, Arieli A: **Linking spontaneous activity of single cortical neurons and the underlying functional architecture**. *Science* 1999, **286**:1943-1946.

75. Luczak A, Barthó P, Harris KD: **Spontaneous events outline the realm of possible sensory responses in neocortical populations**. *Neuron* 2009, **62**:413-425.

76. Miller J-EK, Ayzenshtat I, Carrillo-Reid L, Yuste R: **Visual stimuli
•   recruit intrinsically generated cortical ensembles**. *Proc Natl Acad Sci U S A* 2014, **111**:E4053-E4061.
Present the relations between stimulus-induced ensemble responses and spontaneous population activity, and the dynamic nature of ensemble membership in stimulus encoding.

77. Zemel RS, Dayan P, Pouget A: **Probabilistic interpretation of population codes**. *Neural Comput* 1998, **10**:403-430.

78. Deneve S, Latham PE, Pouget A: **Efficient computation and cue integration with noisy population codes**. *Nat Neurosci* 2001, **4**:826-831.

79. Yang T, Shadlen MN: **Probabilistic reasoning by neurons**. *Nature* 2007, **447**:1075-1080.

80. Deneve S: **Bayesian spiking neurons I: inference**. *Neural Comput* 2008, **20**:91-117.

81. Deneve S: **Bayesian spiking neurons II: learning**. *Neural Comput* 2008, **20**:118-145.

82. Hoyer PO, Hyvärinen A: **Interpreting neural response variability as Monte Carlo sampling of the posterior**. *Advances in Neural Information Processing Systems*. 2003:293-300.

83. Buesing L, Bill J, Nessler B, Maass W: **Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons**. *PLoS Comput Biol* 2011, **7**:e1002211.

84. Gutnisky DA, Dragoi V: **Adaptive coding of visual information in neural populations**. *Nature* 2008, **452**:220-224.

85. Palmer SE, Marre O, Berry MJ, Bialek W: **Predictive information**
   • **in a sensory population**. *Proc Natl Acad Sci U S A* 2015, **112**:6908-6913.
Analyze the nature of predictive coding by population of cells based on stimulus statistics.

86. Goodman ND, Tenenbaum JB, Feldman J, Griffiths TL: **A rational analysis of rule-based concept learning**. *Cogn Sci* 2008, **32**:108-154.

87. Cohen Y, Schneidman E: **High-order feature-based mixture models of classification learning predict individual learning curves and enable personalized teaching**. *Proc Natl Acad Sci U S A* 2013, **110**:684-689.

88. Helmstaedter M, Briggman KL, Turaga SC, Jain V, Seung HS, Denk W: **Connectomic reconstruction of the inner plexiform layer in the mouse retina**. *Nature* 2013, **500**:168-174.

89. Cossell L, Iacaruso MF, Muir DR, Houlton R, Sader EN, Ko H, Hofer SB, Mrsic-Flogel TD: **Functional organization of excitatory synaptic strength in primary visual cortex**. *Nature* 2015, **518**:399-403.

90. Jiang X, Shen S, Cadwell CR, Berens P, Sinz F, Ecker AS, Patel S,
   • Tolias AS: **Principles of connectivity among morphologically defined cell types in adult neocortex**. *Science* 2015, **350** http://dx.doi.org/10.1126/science.aac9462.
Present large-scale, comprehensive profiling of neocortical neurons and their type-dependent characteristic input-output connectivity profile.