

# What can neuronal populations tell us about cognition?

Iñigo Arandia-Romero<sup>1</sup>, Ramon Nogueira<sup>1</sup>, Gabriela Mochol<sup>1</sup>, and Rubén Moreno-Bote<sup>1,2</sup>

<sup>1</sup> Center for Brain and Cognition & Department of Information and Communications Technologies, University Pompeu Fabra, 08018 Barcelona, Spain

<sup>2</sup>Serra Húnter Fellow Programme, 08018 Barcelona, Spain

Corresponding author:

Rubén Moreno-Bote

Email: [ruben.moreno@upf.edu](mailto:ruben.moreno@upf.edu)

Contributions: All authors discussed and wrote the paper

Acknowledgements: This work has been supported by the Spanish PSI2013-44811-P and SlowDyn FLAGERA-PCIN-2015-162-C02-02 grants from MINECO to R. M. B., and IJCI-2014-21937 grant from MINECO to G. M.

## **Highlights**

- **Neuronal population (NP) level analysis is slowly pervading systems neuroscience**
- **NPs provide both unprecedented temporal and spatial resolution to study decision-making**
- **However, large NPs come with challenging data analysis and interpretative problems**
- **Challenges can be alleviated by using models of NP activity**

## **Abstract**

**Nowadays, it is possible to record the activity of hundreds of cells at the same time in behaving animals. However, these data are often treated and analyzed as if they consisted of many independently recorded neurons. How can neuronal populations be uniquely used to learn about cognition? We describe recent work that shows that populations of simultaneously recorded neurons are fundamental to understand the basis of decision-making, including processes such as ongoing deliberations and decision confidence, which generally fall outside the reach of single-cell analysis. Thus, neuronal population data allow addressing novel questions, but they also come with so far unsolved challenges.**

## **Introduction**

Single-neuron electrophysiology has provided golden ages in neuroscience passing from the discoveries of Lord Adrian [1] to those by Hubel and Wiesel [2]. This approach relies on finding mappings between the activity of a single neuron and external world, body or internal variables, most often summarized with the so-called ‘tuning curve’, that is, the mean neuron response as a function of the relevant variable. As such, single-cell electrophysiology has proven to be a central pillar over which many theories of the brain rest [2-8].

In the last 20 years, there has been a surge of new technologies that allow recording the activity of hundreds of cells within and across brain areas, either with dense electrodes or with imaging techniques [9-16] and pioneering work has shown the benefits of these neuronal population (NP) recordings [17-22]. New emerging technologies even promise several orders of magnitude improvement in the number of recorded cells, leaping from hundreds to thousands of cells [23, 24]. However, these promises come with new problems. The most prominent challenges arise in experimental settings that involve animal behavior, where the number of trials is limited to avoid animal satiation and exhaustion, and to minimize overtraining to better study naturalistic behaviors. Considering the activity of each neuron as one dimension in the neuronal state space, how should we analyze and model multi-dimensional neuronal activity with just a few tens of trials per condition? How can we benefit from the many dimensions that NP offers? And, most importantly, what are the new questions that can be addressed with NPs about cognitive processes that cannot, or at least hardly, be asked with single neurons?

We review recent literature that can help to answer the above questions with the focus on decision-making. We provide examples of a better understanding of the transformation from stimulus to choices by using the activity of simultaneously recorded cells. We argue about the difficulty, and even impossibility, of drawing qualitatively similar conclusions using single-cell electrophysiology. Then, we talk about the challenges that we face with large NPs and limited number of trials, as those expected from experiments that combine both animal behavior and NPs recordings. Finally, we point to past and more recent work that could benefit from using NP recordings and discuss possible promising future directions.

## **Advantages of NPs**

The success of NP approaches has been clearly exemplified in brain-machine interface applications [25-28]. Recent studies have shown how it is possible to restore hand control in a

quadriplegic patient [29] or aid non-human primates in walking [30] by reading out the population activity of motor cortex and converting these signals into motor commands bypassing injured nerves. ‘Reading out’ or ‘decoding’ means computing an online estimate of the value of an external world or internal variable from the activity of the recorded neurons [31-33]. For instance, decoding the population activity of V1 neurons with linear classifiers can predict with high accuracy the orientation of a stimulus drifting grating in a trial-by-trial basis [20, 34, 35]. Recent literature tends to favor approaches based on decoding over information theory [19, 20, 34, 36-43], mostly because decoding techniques can be applied with relatively small number of trials while information-theoretic methods require computing probability distributions over as many dimensions as neurons in the NP, a problem daunted by the curse of dimensionality [33]. In addition, decoding techniques provide trial-by-trial estimates of external and internal variables that can be directly interpreted and compared to their actual values [44].

Despite the success of NP approaches to understand the basics of attention [36], working memory [45] and decision-making [37, 38, 46], they are sometimes contemplated with suspicion. Indeed, in some applications it could be argued that NPs do not offer radically new knowledge compared to single-neuron approaches. For instance, while decoding analysis can provide better estimates of information content in a NP by taking into account the trial-by-trial correlated noise among cells, rough estimates can be obtained by extrapolating the information of singles neurons to small NPs of a few tens of cells [47] (but not for populations larger than 100-1000 cells due to the presence of ‘differential correlations’ [48]). Moreover, research based on NP recordings sometimes ends up using traditional single-neuron-based analysis [49, 50]. What are the unique opportunities that NPs can offer? Here we give examples of how NPs can provide new knowledge that is very hard - and in some cases impossible - to acquire from single-neuron-based analysis.

There are at least two broad families of problems that cannot be addressed with single-neuron approaches: 1) internal neuronal dynamics that is not time-locked to observable stimulus or body variables, and 2) variables that are encoded at the population level and not at the single-neuron level. The first category relates to neuronal processing that is internally or externally triggered without experimental control. These phenomena might include trial-by-trial fluctuations of attention that are not cued experimentally, mind wandering, and any other neuronal processing that is not time-locked to the stimulus or body variables [34, 36, 37, 39, 41, 45, 51-63]. Because of their lack of significant correlation to observable experimental variables, these phenomena cannot be fully studied using single-neuron

approaches. The second category relates to neuronal processing that requires coordination of neurons in a NP. For instance, stable representation of a world variable, such as world-centered body location of mice [64], can occur despite unstable representation of that variable at the single-cell level (Fig. 1a). Similarly, theoretical work shows that stable representation of working memory signals can be achieved in NPs despite unstable single-neuron representations [65]. Thus, observing single neurons can lead to the erroneous conclusion that the encoding of the variable is unreliable, while the reality is different, namely, the variable is reliably encoded at the NP level. Further, theoretical studies suggest that the encoding of variables can involve temporal delays between neurons without concomitant changes of activity in any single cell [66] (Fig. 1b), and conclusions on information content drawn from combining single cells recordings can be misleading if there is shared response variability among cells in the NP [32, 48] (Fig. 1c).

## **Opportunities from NPs to understand decision-making**

NPs can provide new ways to study the processes underlying decision-making and cognition. These novel opportunities with NP arise because crucial decision-making variables might not be generally representable by single neurons. . An example of such a variable is ‘decision confidence’ [67-71], that is, the confidence of a decision maker about her choice being correct based on accumulated noisy evidence [68]. A large family of mathematical models for two-alternative decision-making tasks called ‘race models’ requires at least two types of neurons racing to generate the choice [67, 68]. Here each racing neuron represents at each time point the evidence for each option accumulated so far (Fig. 2). For this family of models, decision confidence is encoded in the difference between the activity of two racing neurons (in addition to elapsed time), the so-called ‘balance of evidence’, a quantity that cannot be represented by the activity of neither of the neurons alone [68, 72]. The activity of any single racing neuron has little information about confidence. For example, it is possible to have no change in firing rate with confidence (Fig. 2a,b, cell 1) or even increase in firing rate (Fig.2, cell 2) for lower confidence (lower panels of Fig 2a vs Fig. 2b). More realistic models based on probabilistic population codes [73], which incorporate the Poisson-like nature of neuronal responses, also represent both choices and confidence simultaneously in the difference in activity of different neuronal pools [74].

Recent experimental work has shed new light into how changes of mind might be encoded in neuronal populations during decision-making, a process that can also guide

further studies in the encoding of decision confidence on NPs. Kiani and colleagues recorded the activity of large NPs in area 8Ar of prearcuate gyrus while macaque monkeys performed a motion discrimination task [37]. The authors used decoding techniques to find a hyperplane in multidimensional activity space that separated well animals' choices on a trial-by-trial basis (Fig. 3a). Importantly, the distance of the population activity vector from the hyperplane could be used to infer in which trials the animal 'changed its mind' (Fig 3b). This method allows studying the ongoing process of choice at unprecedented levels of temporal resolution. Decision confidence theories predict that confidence should be related to the same NP activity distance, but these predictions remain to be tested.

A recent paper by Rich and Wallis has studied the deliberation process in economic-based choices at a surprising level of temporal resolution using up to 16 electrodes in areas 11 and 13 in the monkey orbitofrontal cortex [38]. Following a smart experimental design, a linear decoder learned to discriminate between different offers presented on forced-choice single-offer trials, showing how offers were encoded in the NP when they were presented separately. This decoder was then used without retraining to predict the internal deliberation process in free-choice trials where two offers were simultaneously available, uncovering alternations over time between the states representing the available options (Fig. 3c). The untested assumption in this work is that the representation of offers is the same in forced and free choice trials. However, consistently with this assumption, longer total duration of the state corresponding to the chosen offer was observed compared to the unchosen one. Thus, the activity of NPs reveals a rich dynamic representation of choice options, which would be very hard to obtain with single neurons. Similarly complex dynamics has been also reported in other tasks and areas during decision-making [41, 45].

## **Challenges from NPs**

The above examples illustrate the importance of NP analysis. But what is the NP size that we ought to record from to address a given question? A priori, there is no limit to the desired size: the larger, the better, although in some cases with just a bunch of simultaneously recorded cells it is possible to address questions that are unfeasible with just a single cell, as described above for relevant theories of decision confidence. As a rule of thumb, the larger the NP, the more precisely the state of the neuronal network can be assessed [34, 53, 75], the more accurately stimulus or body variables can be decoded [20, 34, 35], and the more variables can be studied simultaneously [29, 30].

However, the promises of population-based analysis do not come free of challenges. The most important challenge is model specification given large number of parameters and limited number of trials. Consider, for instance, a population of 100 neurons with 200 trials per distinct condition and a binary classification task (400 trials in total). Using a quadratic discriminant for this task requires estimating a number of parameters that scales as the square of the number of neurons, around 5 thousand parameters, which is outside the reach of this limited amount of data. A linear classifier could be still used in this condition because 400 trials is in general sufficient to estimate its 101 free parameters. However, consider now a population of 1000 neurons with again 200 trials per condition and the same binary classification task. In this case, even for a linear classifier, we need to estimate 1001 parameters, which is above the number of data points collected. Roughly speaking and based on our experience, we would need at least around 5 times more trials per class and per neuron recorded to train a linear classifier ( $\sim 5 \times \text{neurons} \times \text{classes}$ ), which for large NPs will certainly be difficult to achieve. The problem is even bigger for more complex decoders with many more free parameters, such as deep feedforward networks [76]. Deep learning approaches that have boosted machine learning in the last few years can be hardly suitable for neuronal data based on a few trials per neuron and condition [76]. Novel approaches that would work with standard ‘few-trials’ data need to be developed.

What approaches are then suitable for ‘few-trials’ data? This is a domain ripe for research, and here we discuss two recently used methods. The first one is to use linear decoders together with regularization techniques [37, 52, 77]. In this way, parameters that provide none or little explanatory power are pushed to zero or very small values, therefore reducing the complexity of the decoder and thus dampening overfitting. Regularization has been shown to be useful for NPs approximately as large as 150 neurons with more than 500 trials per condition [37]. However, a possible danger is that regularization implicitly introduces priors, which constraint the family of solutions, but might turn to be wrong. The second method consists of building models with increasing complexity starting from basic principles and performing model comparison using cross-validation techniques [34, 78, 79]. Basic principles refer here to the knowledge that has been acquired across many years of research. For instance, knowledge about translational invariance of real-world objects incorporated into convolutional networks allows to significantly reduce their number of free parameters, thus improving dramatically their classification performance [76]. Another example of prior knowledge that can be included in decoders is the Poisson-like nature of

neuronal responses [4, 73, 80-83] and the gain modulations induced by global fluctuations in population activity [34, 51, 54, 75, 78, 84]. For instance, a recent model of NP activity with Poisson-like firing and both heterogeneous multiplicative and additive gains has been shown to explain better macaque V1 responses than alternative models ('multi-gain' model [34]) (Fig. 4). The first important advance is that this analysis directly exploits NP data, as otherwise it is impossible to study the modulation of single-neuron activity with global gain. Secondly, the resulting model has culminated research that has shown that global gain modulations can be either purely additive [51, 75], purely multiplicative [84], homogeneous multiplicative and additive [54], or a mixture of homogeneous multiplicative and heterogeneous additive gains [78] (Fig. 4).

Finally, it is important to avoid some potential pitfalls when using NPs. First, as larger NPs are available, computing 'noise correlation' matrices (correlations between all pairs of neurons at fixed stimulus conditions, [32, 85, 86]) will become unfeasible, because the covariance matrix becomes singular when the number of neurons exceeds the number of trials [87], unless strong priors are used [88]. There is also a tendency to build large surrogate NPs from independently recorded single-neurons [63, 89, 90]. Surrogate NPs can be useful in some special circumstances [63, 89-91], but overusing them could lead to erroneous conclusions. This can be the case when neuronal populations contain 'differential correlations', which can be very small and yet have a huge impact on information and computations, and are absent in surrogate populations [48].

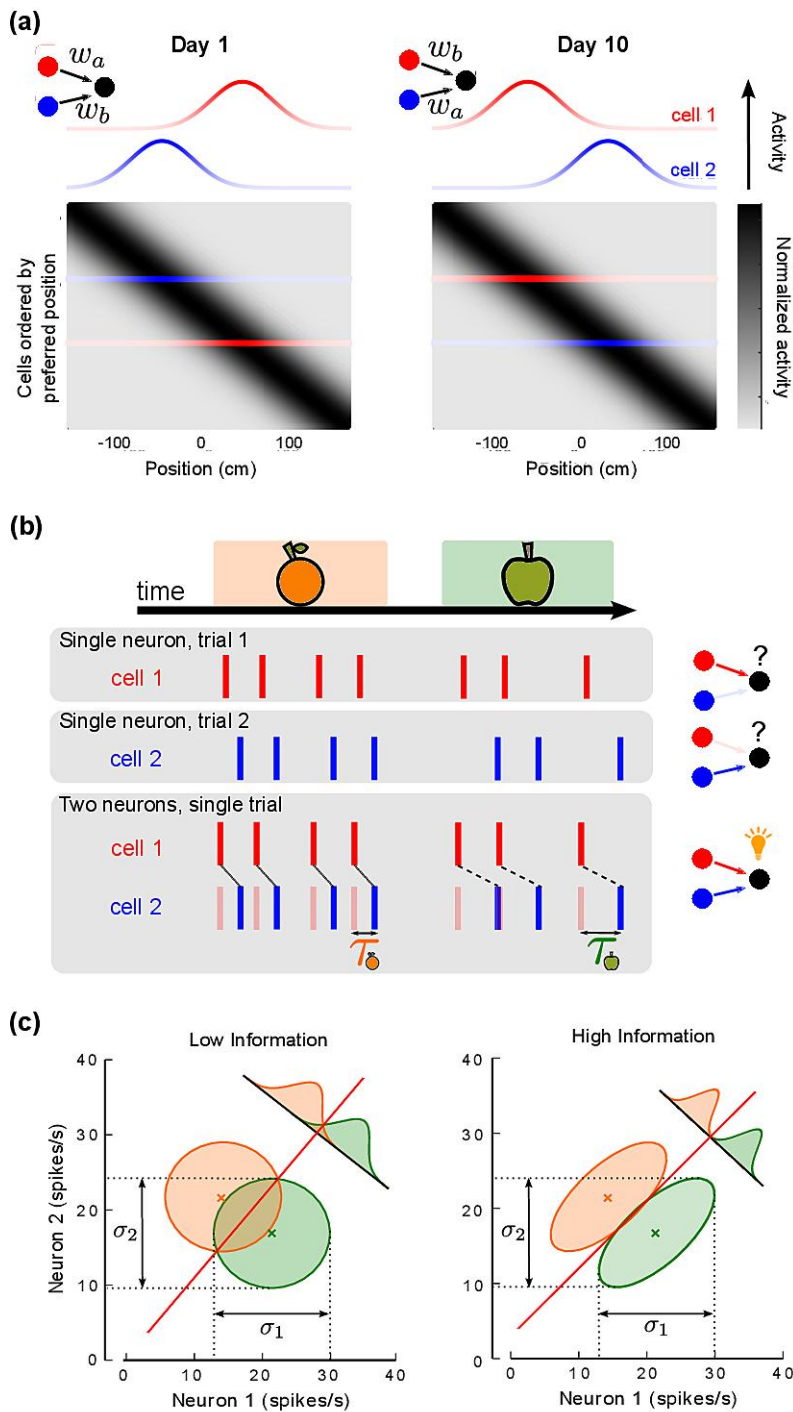
## **Conclusions: what can NPs offer in the future?**

We have shown that NP analysis is slowly pervading neuroscience, but fulminating strides are expected to occur soon. First, ongoing debates on the role of certain brain areas in decision-making and on its neuronal mechanisms will dramatically benefit from measuring the activity of very large NPs [92-95]. Second, decision-making is characterized by complex behavior that in real-life conditions involve a myriad of variables, such as complex stimuli and body variables. The many-fold dimensions of neuronal activity will allow studying the simultaneous encoding of those variables in the brain. Dimensionality reduction techniques [96], including decoders, will leave room to new methods (e.g., [97]) more suitable when the complexity of the experimental design does not permit for such reductionist, low dimensional, view of neuronal activity, as computational neural network models predict [98-100]. Finally, it is also expected that new-generation decoders and other approaches suitable



for ‘few-trials’ datasets, typical for behavioral and systems neuroscience, will allow precise estimates of sensory information from increasingly large NPs [48, 86].

## Figures and legends




**Figure 1:** Neuronal populations (NPs) are necessary to understand the neuronal code.

(a) NPs can represent stable information while the tuning of individual neurons changes [64]. The activity of cell 1 (red) and cell 2 (blue) for different positions can change from the first day (top-left) to the tenth day (top-right), while the information represented by the NP

remains constant (bottom). The ordering of the cells by preferred position differs from the first (bottom-left) to the tenth (bottom-right) day due to changes in individual tuning, but jointly they represent the same map. To recover the information represented by the NP a readout neuron (black circle at the top) adapts the weights ( $w_i$ -s) to the tuning of each neuron (red and blue circles) at a given day. The information recovered across days is the same.

(b) Stimulus information in theory can be encoded in the relative spike timing between neurons [66]. Vertical bars represent the spiking activity of cell 1 (red) and cell 2 (blue) in three cases while two stimuli (orange and apple) are presented consecutively (top). The first two cases show the activity recorded from single cells. Here, the readout neuron (black circle on the right) is not able to recognize the stimulus that has been presented. In the third case, in contrast, the readout neuron can measure the difference in the timing of spikes from cell 1 and cell 2, and recover the stimulus by using the relative spike latencies.

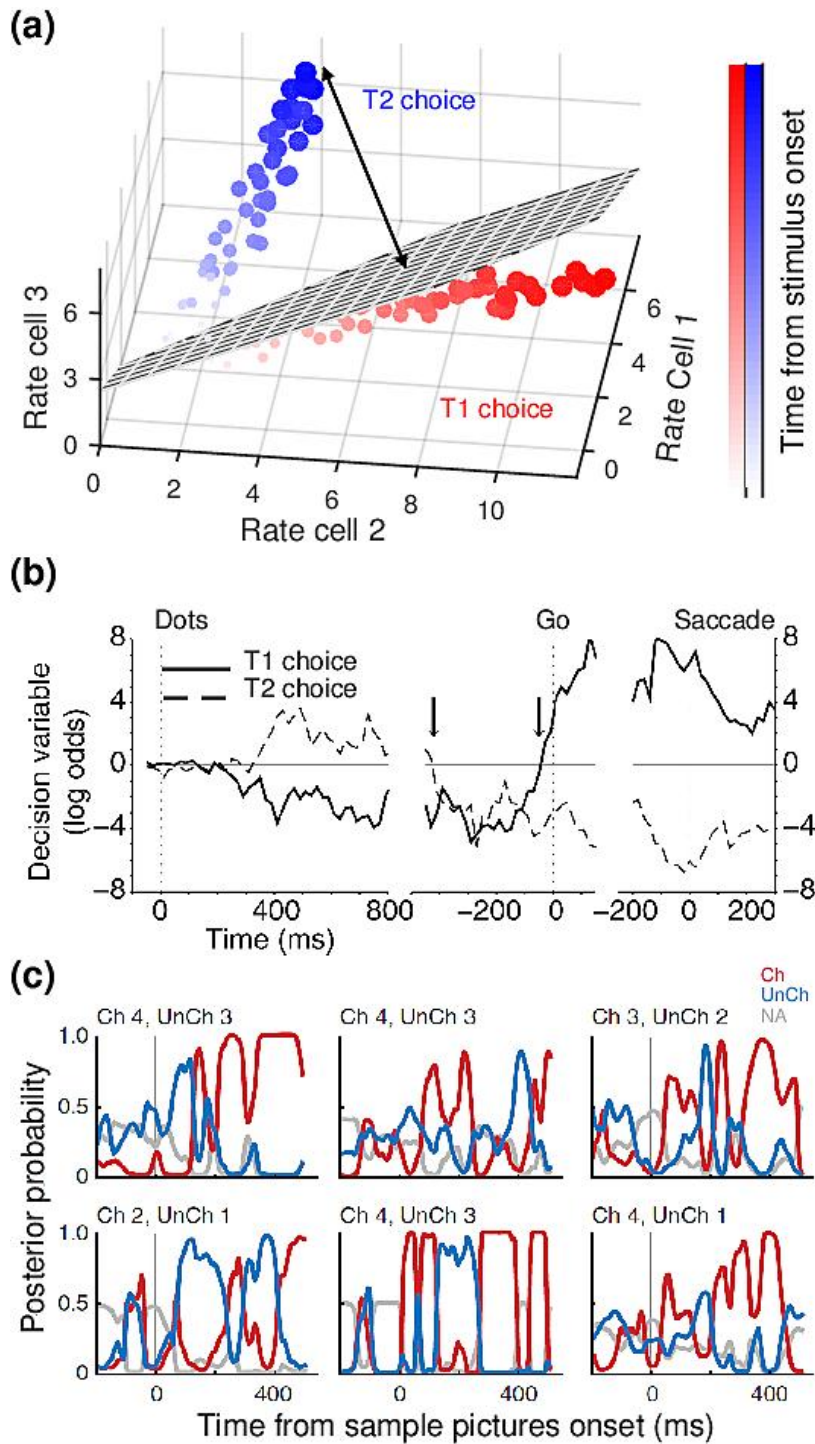
(c) Pair-wise correlations can help decoding [32]. Each plot shows the mean activity (crosses) and the variability in the response (standard deviation represented by clouds) of two neurons when two different stimuli are presented (orange and green). The responses in the left panel have circular clouds, meaning that the variability is the same in all directions, and therefore the correlation between the two neurons is zero for both stimuli. The panel in the right shows two correlated neurons with the same amount of variability as in the left panel, but spread along the same direction (ellipses). The red line marks the decoding boundary, dividing the space between the area which corresponds to the orange stimulus and to the green one. The clouds overlapping in the left panel mean that in some trials where the green stimulus has been presented the decoder will predict the orange one, and vice-versa. However, for the correlated activity case (right) most of the trials are correctly classified, as the clouds do not overlap. Therefore, predicting the performance of a neuronal code depends on neuronal correlations, a problem that lies beyond single-neuron electrophysiology.



**Figure 2:** Decision confidence in a 'race model'.

(a) A decision maker chooses between two offers (green apple and orange), whose values are represented by the activity of two neurons (red and green lines), growing over time as more evidence is accumulated. The decision rule is based on what option depicts the highest activity at decision time (dashed vertical line). In this case, the choice was easy because the values of the offers were very distinct, and therefore the confidence on the decision (represented by the distance between the activity of the neurons) was high [68].

(b) When the two offers have similar values (red vs green apple), the choice is more difficult, and thus decision confidence (distance between the races) is lower. The estimation of confidence would not be possible if the two neurons were not recorded simultaneously. Comparing the high and low confidence cases (a-b) the activity of the neuron preferring the green apple (cell 1) is identical, and thus this neuron alone does not correlate with decision confidence at the time of the decision. Although the neuron that prefers the lower value offer changes its activity in the two cases, the direction of this change is opposite to the change in confidence (cell 2 has higher rate for the low confidence trial) and cannot give a good approximation of the latter. nor provide information about the state of the other racing neuron, making confidence estimation impossible from single neuron recordings.



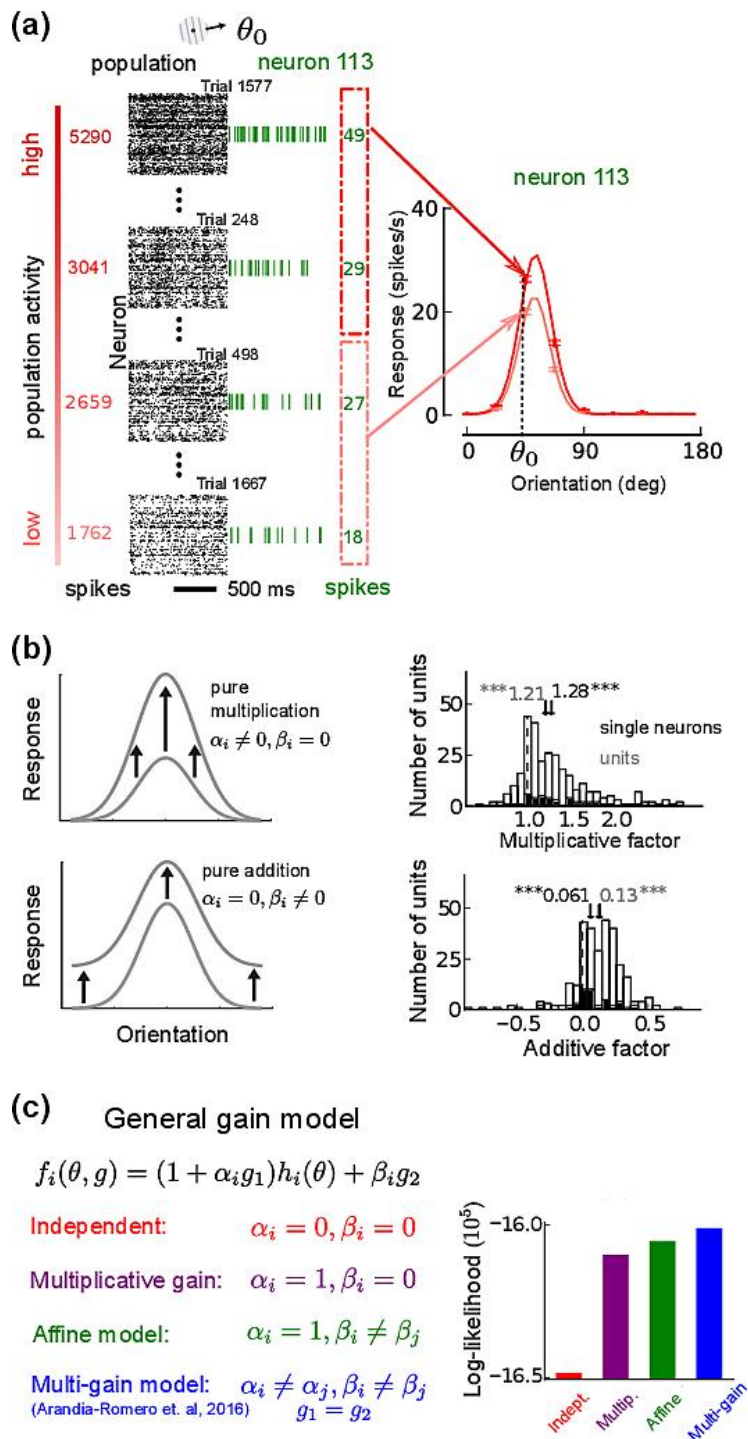
**Figure 3:** Decision making in neuronal response space.

(a) The dynamics of decision making of an animal performing a two-alternative task can be visualized in the neuronal response space, that is, the space where each direction corresponds to the activity of a recorded neuron. At a given time window the population activity is a single point in this space (blue and red dots). A linear decoder is trained to find a hyperplane (in this case, a plane in three dimensions) which best separates the points belonging to the

two classes (red for T1 choice, and blue for T2 choice). At the beginning of each trial (faded colors) the population activity provides weak evidence about stimulus conditions, and thus the points are close to the hyperplane. As more evidence is accumulated, the population activity diverges away from the hyperplane, making easier to predict the animal's choice. The distance between population activity and the classification hyperplane defines the decision variable (DV, black line).

**(b)** The temporal dynamics of DV reveals 'changes of mind'. In a given trial, the DV crossed the ambiguity point equal to zero (indicated by an arrow), suggesting the presence of a change of mind, that is, an instance in which the hypothetical intended choice flips (two trials are shown). Tracking these changes of mind would not be possible with single-neuron analysis. Reproduced from [37] with the permission of the publisher.

**(c)** Ongoing deliberation inferred NP activity in monkey orbitofrontal cortex. A linear decoder, trained in forced-choice trials, was used to predict choices between simultaneously presented offers in free-choice trials. The posterior probability derived from the decoder for chosen (red), unchosen (blue) and unavailable options (gray, average of both unavailable options), are shown in six different trials. This analysis allows tracking the ongoing deliberation processes at unprecedented temporal resolution. Reproduced from [38] with the permission of the publisher



**Figure 4:** Population activity models based on Poisson-like firing and global gain modulation. A large variety of models have been recently used to characterize the activity of NP including the effect of the population activity on single neuronal activity in primary visual cortex (V1) [34, 75, 78, 84].

(a) Population activity fluctuates across trials for the same stimulus (left panels). The tuning of a V1 neuron (right) is modulated with population activity (called here global gain), which is estimated as the sum of spikes across all neurons excluding the activity from the neuron for



which tuning is being characterized; thus, this analysis cannot be done with single neurons, it requires NPs. Trials were ranked from high (top left) to low (bottom left) global gain for each stimulus condition. The activity of the selected neuron (green spike trains) was averaged across either the top (red box) or bottom (pink box) 50th percentile of trials, and the averages were plotted as a function of stimulus orientation (rightmost panel). The tuning was modulated with global gain (red vs. pink lines). Points and error bars are mean responses and s.e.m., respectively; lines are von Mises fits. Reproduced from [32] with the permission of the publisher.

(b) Global gains can affect the tuning of V1 cells in a multiplicative (top-left) and/or an additive (bottom-left) way. The modulatory effect of the population activity on the tuning curves is heterogeneous, with different multiplicative (top-right) and additive factors (bottom-right) for different neurons. \*\*\*:  $p < 0.001$ . Reproduced from [32] with the permission of the publisher.

(c) Models of NP activity including global gains. Neurons are assumed to fire as independent Poisson processes (or variations of this, see [78, 84]) conditioned to gain factors  $g_1$  and/or  $g_2$ , which reflect modulatory factors in the NP. The mean firing rate of each cell  $i$  in the NP is given by the function  $f_i(\theta, g)$ . The function  $h_i(\theta)$  determines the tuning of the cells when the gain factors are zero. The gains can produce multiplicative effects on the tuning curves if  $\alpha_i$  is non-zero, and additive effects if  $\beta_i$  is non-zero (see panel **b left**). A recent model for NP activity suggests heterogeneous multiplicative and additive modulations (multi-gain model [34]). The multiplicative gain model assumes that there is not additive gain and that the multiplicative gain is identical for all neurons ( $\alpha_i = 1, \beta_i = 0$  for all  $i$ ). The affine model allows for heterogeneous additive gains, but the multiplicative gain is still uniform for all neurons ( $\alpha_i = 1$  for all  $i$ ). The multi-gain model permits heterogeneous multiplicative and additive gains (it also uses  $g_1 = g_2$ ). The multi-gain model provides better fits of V1 population activity than other models. Cross-validated log-likelihood is shown.



## References

1. Adrian, E.D. and Y. Zotterman, *The impulses produced by sensory nerve endings: Part 3. Impulses set up by Touch and Pressure*. J Physiol, 1926. **61**(4): p. 465-83.
2. Hubel, D.H. and T.N. Wiesel, *Receptive fields, binocular interaction and functional architecture in the cat visual cortex*. Journal of Physiology London, 1962. **160**: p. 106-154.
3. Wilson, F.A., S.P. Scalaidhe, and P.S. Goldman-Rakic, *Dissociation of object and spatial processing domains in primate prefrontal cortex*. Science, 1993. **260**(5116): p. 1955-8.
4. Shadlen, M.N. and W.T. Newsome, *The variable discharge of cortical neurons: implications for connectivity, computation, and information coding*. J Neurosci, 1998. **18**(10): p. 3870-96.
5. Desimone, R., *Face-selective cells in the temporal cortex of monkeys*. J Cogn Neurosci, 1991. **3**(1): p. 1-8.
6. Rizzolatti, G. and L. Craighero, *The mirror-neuron system*. Annu Rev Neurosci, 2004. **27**: p. 169-92.
7. O'Keefe, J. and J. Dostrovsky, *The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat*. Brain Res, 1971. **34**(1): p. 171-5.
8. Miyashita, Y., *Inferior temporal cortex: where visual perception meets memory*. Annu Rev Neurosci, 1993. **16**: p. 245-63.
9. Maynard, E.M., C.T. Nordhausen, and R.A. Normann, *The Utah intracortical Electrode Array: a recording structure for potential brain-computer interfaces*. Electroencephalogr Clin Neurophysiol, 1997. **102**(3): p. 228-39.
10. Kelly, R.C., et al., *Comparison of recordings from microelectrode arrays and single electrodes in the visual cortex*. J Neurosci, 2007. **27**(2): p. 261-4.
11. Blanche, T.J., et al., *Polytrodes: high-density silicon electrode arrays for large-scale multiunit recording*. J Neurophysiol, 2005. **93**(5): p. 2987-3000.
12. Csicsvari, J., et al., *Massively parallel recording of unit and local field potentials with silicon-based electrodes*. J Neurophysiol, 2003. **90**(2): p. 1314-23.
13. Stosiek, C., et al., *In vivo two-photon calcium imaging of neuronal networks*. Proc Natl Acad Sci U S A, 2003. **100**(12): p. 7319-24.
14. Shoham, D., et al., *Imaging cortical dynamics at high spatial and temporal resolution with novel blue voltage-sensitive dyes*. Neuron, 1999. **24**(4): p. 791-802.
15. Buzsaki, G., *Large-scale recording of neuronal ensembles*. Nat Neurosci, 2004. **7**(5): p. 446-51.
16. Kerr, J.N. and W. Denk, *Imaging in vivo: watching the brain in action*. Nat Rev Neurosci, 2008. **9**(3): p. 195-205.
17. Schoenbaum, G. and H. Eichenbaum, *Information coding in the rodent prefrontal cortex. II. Ensemble activity in orbitofrontal cortex*. J Neurophysiol, 1995. **74**(2): p. 751-62.
18. Warland, D.K., P. Reinagel, and M. Meister, *Decoding visual information from a population of retinal ganglion cells*. J Neurophysiol, 1997. **78**(5): p. 2336-50.
19. Nikolich, D., et al., *Distributed fading memory for stimulus properties in the primary visual cortex*. PLoS Biol, 2009. **7**(12): p. e1000260.
20. Graf, A.B., et al., *Decoding the activity of neuronal populations in macaque primary visual cortex*. Nat Neurosci, 2011. **14**(2): p. 239-45.
21. Aertsen, A.M., et al., *Dynamics of neuronal firing correlation: modulation of "effective connectivity"*. J Neurophysiol, 1989. **61**(5): p. 900-17.

22. Riehle, A., et al., *Spike synchronization and rate modulation differentially involved in motor cortical function*. *Science*, 1997. **278**(5345): p. 1950-3.
23. Ahrens, M.B., et al., *Brain-wide neuronal dynamics during motor adaptation in zebrafish*. *Nature*, 2012. **485**(7399): p. 471-7.
24. Kim, T.H., et al., *Long-Term Optical Access to an Estimated One Million Neurons in the Live Mouse Cortex*. *Cell Rep*, 2016. **17**(12): p. 3385-3394.
25. Lebedev, M.A. and M.A. Nicolelis, *Brain-machine interfaces: past, present and future*. *Trends Neurosci*, 2006. **29**(9): p. 536-46.
26. Andersen, R.A., et al., *Cognitive neural prosthetics*. *Trends Cogn Sci*, 2004. **8**(11): p. 486-93.
27. Collinger, J.L., et al., *High-performance neuroprosthetic control by an individual with tetraplegia*. *Lancet*, 2013. **381**(9866): p. 557-64.
28. Sadtler, P.T., et al., *Neural constraints on learning*. *Nature*, 2014. **512**(7515): p. 423-6.
29. Bouton, C.E., et al., *Restoring cortical control of functional movement in a human with quadriplegia*. *Nature*, 2016. **533**(7602): p. 247-50.
30. Capogrosso, M., et al., *A brain-spine interface alleviating gait deficits after spinal cord injury in primates*. *Nature*, 2016. **539**(7628): p. 284-288.
31. Abbott, L.F., *Decoding neuronal firing and modelling neural networks*. *Q Rev Biophys*, 1994. **27**(3): p. 291-331.
32. Averbeck, B.B., P.E. Latham, and A. Pouget, *Neural correlations, population coding and computation*. *Nat Rev Neurosci*, 2006. **7**(5): p. 358-66.
33. Quiñero, R. and S. Panzeri, *Extracting information from neuronal populations: information theory and decoding approaches*. *Nat Rev Neurosci*, 2009. **10**(3): p. 173-85.
- \*\*34. Arandia-Romero, I., et al., *Multiplicative and Additive Modulation of Neuronal Tuning with Population Activity Affects Encoded Information*. *Neuron*, 2016. **89**(6): p. 1305-16.  
This work demonstrates the influence of the population activity on the tuning of individual monkey V1 neurons. Population activity enhances or reduces information in small neuronal ensembles, depending on whether neurons are multiplicatively or additively modulated by population activity.
35. Berens, P., et al., *A fast and simple population code for orientation in primate V1*. *J Neurosci*, 2012. **32**(31): p. 10618-26.
- \*36. Cohen, M.R. and J.H. Maunsell, *A neuronal population measure of attention predicts behavioral performance on individual trials*. *J Neurosci*, 2010. **30**(45): p. 15241-53.
- \*\*37. Kiani, R., et al., *Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials*. *Curr Biol*, 2014. **24**(13): p. 1542-7.
38. Rich, E.L. and J.D. Wallis, *Decoding subjective decisions from orbitofrontal cortex*. *Nat Neurosci*, 2016. **19**(7): p. 973-80.
- \*39. Tremblay, S., et al., *Attentional filtering of visual information by neuronal ensembles in the primate lateral prefrontal cortex*. *Neuron*, 2014. **85**(1): p. 202-15.
40. Adibi, M., et al., *Population decoding in rat barrel cortex: optimizing the linear readout of correlated population responses*. *PLoS Comput Biol*, 2014. **10**(1): p. e1003415.
- \*\*41. Morcos, A.S. and C.D. Harvey, *History-dependent variability in population dynamics during evidence accumulation in cortex*. *Nat Neurosci*, 2016. **19**(12): p. 1672-1681.
42. Sussillo, D., et al., *Making brain-machine interfaces robust to future neural variability*. *Nat Commun*, 2016. **7**: p. 13749.

43. Montijn, J.S., et al., *Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective*. Cell Rep, 2016. **16**(9): p. 2486-98.
44. Churchland, A.K. and R. Kiani, *Three challenges for connecting model to mechanism in decision-making*. Curr Opin Behav Sci, 2016. **11**: p. 74-80.
- \*45. Stokes, M.G., et al., *Dynamic coding for cognitive control in prefrontal cortex*. Neuron, 2013. **78**(2): p. 364-75.
46. Nogueira, R., et al., *Lateral orbitofrontal cortex anticipates choices and integrates prior with current information*. Nat Commun, 2017. **8**: p. 14823.
47. Ecker, A.S., et al., *Decorrelated neuronal firing in cortical microcircuits*. Science, 2010. **327**(5965): p. 584-7.
- \*\*48. Moreno-Bote, R., et al., *Information-limiting correlations*. Nat Neurosci, 2014. **17**(10): p. 1410-7.  
 In this study, it is characterized the type of neuronal pairwise correlations that limit information. A particular pattern of trial-by-trial shared variability is generated when the input layer is noisy itself, the so-called differential correlations, which make information saturate for large neuronal ensembles.
49. Hosokawa, T., et al., *Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex*. J Neurosci, 2013. **33**(44): p. 17385-97.
50. Sul, J.H., et al., *Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making*. Neuron, 2010. **66**(3): p. 449-60.
51. Arieli, A., et al., *Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses*. Science, 1996. **273**(5283): p. 1868-71.
52. Kiani, R., et al., *Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex*. Neuron, 2015. **85**(6): p. 1359-73.
53. Okun, M., et al., *Diverse coupling of neurons to populations in sensory cortex*. Nature, 2015. **521**(7553): p. 511-5.
- \*54. Mochol, G., et al., *Stochastic transitions into silence cause noise correlations in cortical circuits*. Proc Natl Acad Sci U S A, 2015. **112**(11): p. 3529-34.
55. Scholvinck, M.L., et al., *Cortical state determines global variability and correlations in visual cortex*. J Neurosci, 2015. **35**(1): p. 170-8.
56. Jezzini, A., et al., *Processing of hedonic and chemosensory features of taste in medial prefrontal and insular networks*. J Neurosci, 2013. **33**(48): p. 18966-78.
57. Ponce-Alvarez, A., et al., *Dynamics of cortical neuronal ensembles transit from decision making to storage for later report*. J Neurosci, 2012. **32**(35): p. 11956-69.
- \*\*58. Engel, T.A., et al., *Selective modulation of cortical state during spatial attention*. Science, 2016. **354**(6316): p. 1140-1144.  
 This work demonstrates the influence of the population activity on the tuning of individual monkey V1 neurons. Population activity enhances or reduces information in small neuronal ensembles, depending on whether neurons are multiplicatively or additively modulated by population activity.
59. Abeles, M., et al., *Cortical activity flips among quasi-stationary states*. Proc Natl Acad Sci U S A, 1995. **92**(19): p. 8616-20.
60. Seidemann, E., et al., *Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task*. J Neurosci, 1996. **16**(2): p. 752-68.
61. Kaufman, M.T., et al., *Cortical activity in the null space: permitting preparation without movement*. Nat Neurosci, 2014. **17**(3): p. 440-8.
- \*62. Kaufman, M.T., et al., *Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex*. Elife, 2015. **4**: p. e04677.

- \*63. Mante, V., et al., *Context-dependent computation by recurrent dynamics in prefrontal cortex*. Nature, 2013. **503**(7474): p. 78-84.
- \*64. Ziv, Y., et al., *Long-term dynamics of CA1 hippocampal place codes*. Nat Neurosci, 2013. **16**(3): p. 264-6.
- 65. Druckmann, S. and D.B. Chklovskii, *Neuronal circuits underlying persistent representations despite time varying activity*. Curr Biol, 2012. **22**(22): p. 2095-103.
- 66. VanRullen, R., R. Guyonneau, and S.J. Thorpe, *Spike times make sense*. Trends Neurosci, 2005. **28**(1): p. 1-4.
- 67. Vickers, D., *Decision processes in visual perception*. 1979, New York: London Academic Press. 406.
- \*68. Moreno-Bote, R., *Decision Confidence and Uncertainty in Diffusion Models with Partially Correlated Neuronal Integrators*. Neural Computation, 2010. **22**(7): p. 1786-1811.
- \*69. Kiani, R. and M.N. Shadlen, *Representation of confidence associated with a decision by neurons in the parietal cortex*. Science, 2009. **324**(5928): p. 759-64.
- \*70. Kepecs, A., et al., *Neural correlates, computation and behavioural impact of decision confidence*. Nature, 2008. **455**(7210): p. 227-31.
- 71. Drugowitsch, J., et al., *The cost of accumulating evidence in perceptual decision making*. J Neurosci, 2012. **32**(11): p. 3612-28.
- 72. Drugowitsch, J., R. Moreno-Bote, and A. Pouget, *Relation between belief and performance in perceptual decision making*. PLoS One, 2014. **9**(5): p. e96511.
- 73. Ma, W.J., et al., *Bayesian inference with probabilistic population codes*. Nat Neurosci, 2006. **9**(11): p. 1432-8.
- 74. Beck, J.M., et al., *Probabilistic population codes for Bayesian decision making*. Neuron, 2008. **60**(6): p. 1142-52.
- \*75. Ecker, A.S., et al., *State dependence of noise correlations in macaque primary visual cortex*. Neuron, 2014. **82**(1): p. 235-48.
- 76. Goodfellow, I., Y. Bengio, and A. Courville, *Deep Learning*. 2016: MIT Press.
- 77. Bishop, C.M., *Pattern recognition and machine learning*. 2006, Singapore: Springer.
- \*78. Lin, I.C., et al., *The Nature of Shared Cortical Variability*. Neuron, 2015. **87**(3): p. 644-56.
- 79. Rabinowitz, N.C., et al., *Attention stabilizes the shared gain of V4 populations*. Elife, 2015. **4**: p. e08998.
- 80. Softky, W.R. and C. Koch, *The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs*. J Neurosci, 1993. **13**(1): p. 334-50.
- 81. Tolhurst, D.J., J.A. Movshon, and A.F. Dean, *The statistical reliability of signals in single neurons in cat and monkey visual cortex*. Vision Res, 1983. **23**(8): p. 775-85.
- 82. Geisler, W.S. and D.G. Albrecht, *Visual cortex neurons in monkeys and cats: detection, discrimination, and identification*. Vis Neurosci, 1997. **14**(5): p. 897-919.
- \*83. Moreno-Bote, R., *Poisson-like spiking in circuits with probabilistic synapses*. PLoS Comput Biol, 2014. **10**(7): p. e1003522.
- \*84. Goris, R.L., J.A. Movshon, and E.P. Simoncelli, *Partitioning neuronal variability*. Nat Neurosci, 2014. **17**(6): p. 858-65.
- 85. Cohen, M.R. and A. Kohn, *Measuring and interpreting neuronal correlations*. Nat Neurosci, 2011. **14**(7): p. 811-9.
- 86. Kohn, A., et al., *Correlations and Neuronal Population Information*. Annu Rev Neurosci, 2016. **39**: p. 237-56.
- 87. Kanitscheider, I., et al., *Measuring Fisher information accurately in correlated neural populations*. PLoS Comput Biol, 2015. **11**(6): p. e1004218.



88. Yatsenko, D., et al., *Improved estimation and interpretation of correlations in neural circuits*. PLoS Comput Biol, 2015. **11**(3): p. e1004083.
- \*89. Rigotti, M., et al., *The importance of mixed selectivity in complex cognitive tasks*. Nature, 2013. **497**(7451): p. 585-90.
- \*90. Ince, R.A., S. Panzeri, and C. Kayser, *Neural codes formed by small and temporally precise populations in auditory cortex*. J Neurosci, 2013. **33**(46): p. 18277-87.
91. Panzeri, S., et al., *Neural population coding: combining insights from microscopic and mass signals*. Trends Cogn Sci, 2015. **19**(3): p. 162-72.
92. Latimer, K.W., et al., *Single-trial spike trains in parietal cortex reveal discrete steps during decision-making*. Science, 2015. **349**(6244): p. 184-7.
93. Latimer, K.W., et al., *Response to Comment on "Single-trial spike trains in parietal cortex reveal discrete steps during decision-making"*. Science, 2016. **351**(6280): p. 1406.
94. Shadlen, M.N., et al., *Comment on "Single-trial spike trains in parietal cortex reveal discrete steps during decision-making"*. Science, 2016. **351**(6280): p. 1406.
95. Bollimunta, A., D. Totten, and J. Ditterich, *Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex*. J Neurosci, 2012. **32**(37): p. 12684-701.
96. Cunningham, J.P. and B.M. Yu, *Dimensionality reduction for large-scale neural recordings*. Nat Neurosci, 2014. **17**(11): p. 1500-9.
97. Kobak, D., et al., *Demixed principal component analysis of neural population data*. Elife, 2016. **5**.
- \*98. Boerlin, M., C.K. Machens, and S. Deneve, *Predictive coding of dynamical variables in balanced spiking networks*. PLoS Comput Biol, 2013. **9**(11): p. e1003258.
99. Moreno-Bote, R. and J. Drugowitsch, *Causal Inference and Explaining Away in a Spiking Network*. Sci Rep, 2015. **5**: p. 17531.
100. Grabska-Barwinska, A., et al., *A probabilistic approach to demixing odors*. Nat Neurosci, 2016. **20**(1): p. 98-106.