

RESEARCH ARTICLE

Randomly connected networks generate emergent selectivity and predict decoding properties of large populations of neurons

Audrey Sederberg^{1,2*}, Ilya Nemenman^{1,2,3}

1 Department of Physics, Emory University, Atlanta, Georgia, United States of America, **2** Initiative in Theory and Modeling of Living Systems, Emory University, Atlanta, Georgia, United States of America, **3** Department of Biology, Emory University, Atlanta, Georgia, United States of America

* audrey.j.sederberg@emory.edu



OPEN ACCESS

Citation: Sederberg A, Nemenman I (2020) Randomly connected networks generate emergent selectivity and predict decoding properties of large populations of neurons. PLoS Comput Biol 16(5): e1007875. <https://doi.org/10.1371/journal.pcbi.1007875>

Editor: Samuel J. Gershman, Harvard University, UNITED STATES

Received: January 30, 2020

Accepted: April 14, 2020

Published: May 7, 2020

Copyright: © 2020 Sederberg, Nemenman. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All code necessary to generate the simulations and to produce the figures in this paper is available at <https://github.com/ajsederberg/rnn-population-decoding>.

Funding: This work was supported by NIH Grants R01NS084844 (AS and IN), R01EB022872, and R01NS099375 (IN), and by NSF Grant BCS-1710452 (IN). Early conception of this project was supported by NIH/NINDS The Trial Version R01NS104528. The funders had no role in study design, data collection and

Abstract

Modern recording methods enable sampling of thousands of neurons during the performance of behavioral tasks, raising the question of how recorded activity relates to theoretical models. In the context of decision making, functional connectivity between choice-selective cortical neurons was recently reported. The straightforward interpretation of these data suggests the existence of selective pools of inhibitory and excitatory neurons. Computationally investigating an alternative mechanism for these experimental observations, we find that a randomly connected network of excitatory and inhibitory neurons generates single-cell selectivity, patterns of pairwise correlations, and the same ability of excitatory and inhibitory populations to predict choice, as in experimental observations. Further, we predict that, for this task, there are no anatomically defined subpopulations of neurons representing choice, and that choice preference of a particular neuron changes with the details of the task. We suggest that distributed stimulus selectivity and functional organization in population codes could be emergent properties of randomly connected networks.

Author summary

What do we learn about neural circuit organization and function from recordings of large populations of neurons? For example, in population recordings in the posterior parietal cortex of mice performing an evidence integration task, particular patterns of selectivity and correlations between cells were observed. One hypothesis for an underlying mechanism generating these patterns is that they follow from intricate rules of connectivity between specific neurons, but this raises the question of how such intricate patterns arise during learning or development. An alternative hypothesis, which we explore here, is that such patterns emerge from generic properties of certain random networks. We find that a random network model matches many features of experimental recordings, from single cells to populations. We suggest that such emergent selectivity could be an important principle in brain areas in which a broad distribution of selectivity is observed.



The Trial Version of pdfElement has some limitations. The funders had no role in study design, data collection and

analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Introduction

Experimental and data-processing advances have opened the floodgates of neural data: recordings from hundreds to thousands of neurons, in animals performing well-defined tasks, are becoming widespread in the literature [1–3]. Fundamentally, it is not clear what we expect to find in recordings of thousands of neurons. Without theory guiding expectations, one approach is to search for statistical differences across groups of cells and then to match these statistical trends in neural activity to the algorithm used for the task in question. Theoretical neuroscience has many such high-level models for computational algorithms underlying important functions, including making predictions [4], or accumulating evidence in making a decision [5, 6].

How are these functional models instantiated by real populations of neurons? For instance, suppose that we have a functional model in which pools of excitatory or inhibitory neurons selectively connect to each other, resulting in mutual suppression and patterns of activity that are selective for specific classes of input. The straightforward interpretation of this functional model is that sub-populations of excitatory and inhibitory neurons in the brain reflect the organization of the functional model. This interpretation raises the question of how such structural patterns of connectivity arose, perhaps through learning or development.

We explore an alternative mechanism for the emergence of such function. Our aim is to raise the question, for any high-throughput neuroscience experiment, of whether observed patterns could be explained by dynamics of a random network. Here we focus on a particular task, which has the advantage that it has been analyzed experimentally [1]. In this task, mice discriminate stimuli consisting of randomly timed impulses with different numbers of impulses per unit time (frequency), and input classes are distinguished by their average frequency. We hypothesize that a randomly connected network can produce patterns of activity that are sufficiently distinct to differentiate these classes of inputs. The eigenvalue spectrum of a randomly drawn connectivity matrix often has a tail of eigenvalues with large real parts [7], and this remains true even when networks follow Dale's principle of separate excitatory and inhibitory neurons [8, 9]. We reason that input modes overlapping with fast-growing modes are amplified by network dynamics, resulting in an activation pattern that depends on the temporal frequency of the input. Thus, the network would produce different patterns of activation as a function of the input frequency, generating emergent selectivity across the population. We are primarily interested in whether such a generic random network of firing rate units could generate the observed patterns of selectivity, functional connectivity, and cell-type-specific readout weight distributions simply through randomly arising heterogeneity in synaptic connectivity.

Our simulation results support this hypothesis. Our main findings are (i) that heterogeneity in connectivity generates differences in inputs to single cells that are dependent on stimulus frequency and (ii) that these differences are sufficient to distinguish between low- and high-frequency inputs. Our model reproduces experimental findings, including the distribution of single-neuron selectivity, patterns of pairwise noise correlations, the performance of a classifier, and the distribution of readout weights. Our theory makes the verifiable experimental predictions that, if the mechanism is through emergence rather than specific connectivity, then (1) there is no anatomical basis for sub-populations tuned to a particular choice and (2) when task parameters, such as input frequency are changed, neural selectivity also changes. These results suggest a mechanism for how cortical networks could exhibit functional organization without requiring specific patterns of cortical connectivity.

Results

Evidence accumulation tasks explore how the brain makes decisions based on the temporal integration of incoming sensory information. One class of models for performing this discrimination is that of attractor networks. In an attractor network model of decision making, pools of neurons fire selectively for a particular choice. A transient activation is prolonged through slow recurrent excitation within the pool while inhibiting other pools of neurons that are selective in their firing for other choices through non-specific inhibition [10]. The results of recent experiments in decision-related areas of rodent parietal cortex have called this model into question [1, 11, 12]. These experiments showed that, contrary to the predictions of the original models, inhibitory cells are also selective for choice, suggesting an alternative mechanism, in which pools of inhibitory neurons selectively inhibit neurons representing evidence for the opposite choice [1, 13]. However, it is unclear how the specific pattern of connectivity would be generated.

More specifically, in this task [1], a rodent is presented with an irregular train of either visual or auditory impulses and must determine whether the average frequency with which those pulses arrived is above or below some internally remembered threshold. Recordings of population activity in the posterior parietal cortex (PPC) during the task revealed weak choice selectivity in single cells, with a fraction of individual cells showing significant selectivity for one of the two choices. A linear classifier operating on the activity across the population decoded the choice with high accuracy. Both excitatory and inhibitory neurons were selective for choice, and noise correlations between pairs of neurons reflected whether stimulus selectivity was shared or opposing. A straightforward mechanism for these observations is that some specific pattern of connectivity exists in the cortical network that separates excitatory and inhibitory cells into “pools” that are selective for specific choices. In this paper, we explore an alternative mechanism for these experimental observations.

Simulation of a temporal evidence accumulation task

We study a randomly connected rate-based network of N_E excitatory cells and N_I inhibitory cells that performed the evidence accumulation task described above (Fig 1A). The firing rate of individual units is the sum of any external input, representing the stimulus, and synaptic inputs from the rest of the network, passed through a saturating non-linearity so that activity is between 0 and 1. Connectivity in the network is sparse, with a 20% probability of connection between two cells. Excitatory and inhibitory synaptic weights are drawn from truncated normal distributions ($\sim N(g_{E,I}, \sigma^2)$) (Fig 1B, Methods), and excitatory and inhibitory synaptic weights to neurons are balanced on average, such that $N_E g_E - N_I g_I = 0$. We fix g_E , g_I and σ , and we focus on a single combination of parameters, but the essential results are not dependent on making these particular choices (see S3 Fig for additional parameter choices). The stimuli consist of pulsed inputs arriving at random times (Fig 1C, see Methods) and with average frequency f . There is no spatial component of the task: inputs stimulate the same subset of excitatory input neurons in all trials. From the simulated activity of the network, we decode the frequency of the input, using either the full population or only excitatory cells or inhibitory cells.

This network was conceived as a model of the posterior parietal cortex (PPC) of rodents, which does not receive direct sensory inputs but rather receives inputs that have passed through upstream networks [14]. Moreover, the average population firing rates in PPC during such a task are not directly related to stimulus frequency [11]. To account for this effect in our simulation, we scaled the amplitude of inputs such that the network firing rate is, on average, equal for each frequency (Fig 1D and S1 Fig). From a computational standpoint, this choice




Fig 1. Overview: Simulation of an input frequency discrimination task in a recurrent network of excitatory and inhibitory neurons with random connectivity. A: Recurrent E-I network of randomly connected neurons. Synaptic weights are heterogeneous (line thickness), and connectivity in the network is random. Inputs to the network have the same pattern of spatial activation but different temporal features. Readouts of network activity (from all neurons, or restricted to excitatory or inhibitory neurons) discriminate between stimulus categories. We then analyze the distributions of readout weights. B: Distributions of non-zero inhibitory (shaded red) and excitatory (shaded blue) synaptic weights. 20% of the weights are non-zero. C: Inputs to the network are pulses arriving at random times with low (top row, yellow) or high (bottom row, red) average frequency. D: Histogram of population average network activity over trials for low- and high-frequency stimuli. Average network activity is not informative of input categories.

<https://doi.org/10.1371/journal.pcbi.1007875.g001>

makes the task of decoding from the network activity more difficult. We emphasize that selectivity in the network is not due to overall higher firing rates for higher-frequency stimuli, because average firing rates are controlled across stimuli.

We present the simulation results by following the experimental observations presented by Najafi and colleagues [1]. First, we analyze simulated choice selectivity at the level of single cells and pairs of cells. Given that we know the network connectivity, we analyze both cell-cell cross-correlations and the underlying connectivity pattern. Next, using the simulated population activity, we decode choice from the simulated population activity, and we compare the distribution of weights from the readouts of simulated activity to the distribution acquired for experimentally recorded activity. Finally, we simulate new conditions, in which we change the

frequencies being discriminated, and from this simulation, we predict how this changes the selectivity for choice, both in single neurons and across the population.

Emergent selectivity for stimulus category in single cells

We first examine choice selectivity in single cells from the network simulations. For this analysis, choice was defined to be the correct stimulus label (i.e., low vs. high frequency of input pulses) on each trial, which is equivalent to analyzing the correct trials only in a behavioral experiment. Single-neuron activity was averaged over time, yielding a single number per trial for each cell. An ideal observer analysis was used to discriminate between low- and high-frequency inputs (see [Methods](#)). Choice selectivity [1] was defined as the area under the receiver-operator curve (AUC, [15]), which is less (greater) than 0.5 when the cell is selective for the low-frequency (high-frequency) stimulus (see [S2 Fig](#) for examples of AUC values for single-cell response distributions). For each network realization, this generated a distribution of AUC values across all cells ([Fig 2A](#)). To assess the significance of a single AUC calculation, we computed the 95% confidence interval of AUC values obtained with shuffled trial labels. Values exceeding these bounds are significant. Approximately 30% of excitatory and 30% of inhibitory cells in the example network in ([Fig 2A](#)) had significant selectivity, based on the AUC analysis ([Fig 2B](#), network instance 1). Across different simulated networks, we observed proportions from 0.14 to 0.82 of single cells that had significant AUC values, and there was no difference between excitatory and inhibitory cells (Average fraction selective, $N = 14$ networks: 0.36 ± 0.15 (excitatory cells) and 0.35 ± 0.15 (inhibitory cells)). We also computed the average choice selectivity, defined as $2|AUC - 0.5|$ for each network ([Fig 2C](#)). Across networks, the average choice selectivity was 0.08 ± 0.02 in both excitatory and inhibitory cells ($N = 14$ networks, range 0.04 to 0.14), compared to the reported values 0.1 to 0.16 [1]. To summarize, in the model, single cells were weakly selective for choice, and excitatory and inhibitory cells exhibited similar levels of selectivity. Thus, we have set up the network such that the average population response is not strongly selective for choice, and in this regime, single cells across the population have selectivity values that are comparable to experimentally recorded values.

Cross-correlations reflect relative selectivity of pairs of cells

We next asked whether this simple, unstructured network also explained pairwise relationships observed in the experimental recordings. Specifically, pairs of cells in PPC that shared the same selectivity had higher correlations [1] than pairs that had opposite selectivity. To compare this to our simulation results, we computed correlations as the neuron-neuron cross-correlation of stimulus-specific responses, averaged across stimuli. Because the only source of noise in the simulation, the variable timing of inputs, is a shared input to all neurons, we expect correlation in the model to be higher in the model than in the data, but cells with the same selectivity are expected to have higher correlation. Restricting analysis to the cells that exceeded the significance criterion for AUC values, we categorized cells by selectivity and compared average correlation between pairs of cells with same and opposite preference. As observed experimentally, correlation in the simulation was higher between pairs of cells with the same selectivity than between cells with opposite selectivity ([Fig 2D](#)). Thus, organization of functional connectivity in the network emerged without setting up distinct clusters of connections in the network.

We further examined the patterns of connectivity between these sets of cells ([Fig 2E](#), [S4 Fig](#) and [S1 Table](#)). The overall probability of a synapse was set to 20% for all simulations. There was a nearly identical probability of connection from an excitatory cell to an inhibitory cell with the same selectivity, $20\% \pm 1\%$ (SD), as with opposite selectivity, $19\% \pm 3\%$ (SD). The




Fig 2. Weak selectivity in both excitatory and inhibitory cells. A. Distribution of AUC values for a single simulation, for excitatory (blue) and inhibitory (red) cells. Shaded regions are those AUC values exceeding significance bounds generated by shuffling trial labels, shown for excitatory cells only. Bounds for inhibitory cells are similar. B. The fraction of cells selective is the fraction of cells in the excitatory or inhibitory populations whose AUC exceeded the significance bounds. The fraction was variable across network realizations, but within each network, it was similar for excitatory and inhibitory cells. Error bars are derived from counting statistics. C. Average choice selectivity (defined as in [1] as the average of $2|AUC - 0.5|$). Averages are over all cells (selective and non-selective) and error bars are SEM. D. Correlations (average value) among cells that share the same selectivity is higher than among cells that have opposite selectivity. Error bars are SEM. Correlations are substantial (> 0.5) because there is no internal noise in the network (all noise is input-driven, and shared across cells). E. Probability of connection for four combinations of pairs of cell types and selectivity: excitatory to inhibitory, with the same selectivity; excitatory to inhibitory, with opposite selectivity; and inhibitory to excitatory with same and opposite selectivity. Error bars are derived from counting statistics.

<https://doi.org/10.1371/journal.pcbi.1007875.g002>

inhibitory to excitatory connection between cells with the same selectivity was similar as well, $18\% \pm 2\%$ (SD). Among cells with opposite selectivity, the probability of connection from an inhibitory to an excitatory cell was $24\% \pm 4\%$ (SD). Even this small amount of excess connectivity between inhibitory and excitatory cells of opposite selectivity was sufficient to reproduce the experimentally observed trends in functional connectivity and stimulus selectivity patterns. We emphasize that this bias in connectivity was not put in by hand, but rather was uncovered by the dynamics (which were shaped by the connectivity).

Decoding from population activity

To determine how accurately the simulated network represented choice, we trained linear classifiers to discriminate between low and high stimulus frequency. We fit a classifier using




Fig 3. Classifiers decode stimulus identity from recurrent network activity. A: Accuracy of a classifier over the course of the stimulus presentation. For each set of cells (all, black; 100 e-cells, blue; 100 i-cells, red), we fit the classifier at a single point in time (circle) and classified activity over the trial. B: Accuracy of classifier across network instances (same order as Fig 2). Error bars are ± 1 SE estimated by cross-validation (see Methods). C–E: Characterization of classifier weights. C: Weights for classifiers fit on the activity of all cells at each point in time in the stimulus presentation window. Weights from excitatory units are blue and from inhibitory units are red. D: Cumulative distribution of weights in each of the networks, for excitatory (blue) and inhibitory (red) cells. E: Average of CDFs shown in D. Distributions are overlapping for excitatory and inhibitory cells. For reference, we also plot the CDF of a normal distribution with standard deviation matched to that of weight distributions (black).

<https://doi.org/10.1371/journal.pcbi.1007875.g003>

activity near the end of the stimulus period (circles, Fig 3A) and tested the classifier over the full stimulus period on a set of reserved trials (Methods). In this simulation, classifier accuracy reached maximal performance within 5τ (time constant of the network) and decoded with $83\% \pm 2\%$ accuracy over the last half of the stimulus period. A classifier fit using only the activity of the inhibitory cells performed with $76\% \pm 2\%$ accuracy, and a randomly drawn subset of excitatory cells equal in number to the number of inhibitory cells decoded with $78\% \pm 2\%$ accuracy. Across all instances of the network, the accuracy of decoders was comparable for inhibitory cells and excitatory ones (Fig 3B). The range of performance we observed across randomly drawn networks ($69\% \pm 2\%$ to $86\% \pm 2\%$) was highly consistent with the population decoding accuracy observed in experiments (about 70% to 85%, [1]).

For the classifier built from the full population activity, we inspected the readout weights (Methods). Weights from inhibitory neurons and weights from excitatory neurons were not significantly different (Fig 3C). In all simulated networks, the distributions of weights for excitatory and inhibitory cells were overlapping (Fig 3D and 3E). The weight distributions are not normal (for all network, Lillifors test, $p < 0.001$). Thus, as was reported experimentally, we find that both excitatory and inhibitory cells contribute to stimulus decoding and that readout weights are not significantly different between the two.

Selectivity in a network with inputs to excitatory neurons only

In Figs 2 and 3, we analyzed selectivity in excitatory and inhibitory cells, all of which received inputs from a pool of excitatory cells, and we found equal selectivity for excitatory and inhibitory cells. In a second set of simulations, all excitatory cells in the network received direct inputs and inhibitory cells received inputs only through their recurrent connections. Connectivity in the network was determined probabilistically, with slightly stronger weights from (+2.5%) inhibitory cells to excitatory cells than from inhibitory cells to inhibitory cells, which offset the input current received by excitatory cells only. Again, we found that many cells in the network are selective (Fig 4A–4C), with 52% to 62% of cells showing statistically significant selectivity and the average selectivity index of 0.10 ± 0.01 for excitatory cells and 0.11 ± 0.01




Fig 4. Selectivity and decoding from random network with inputs to excitatory cells only. A. Distribution of AUC values for a single simulation, for excitatory (blue) and inhibitory (red) cells. Shaded regions are those AUC values exceeding significance bounds generated by shuffling trial labels, shown for excitatory cells only. Bounds for inhibitory cells are similar. B. The fraction of cells selective is the fraction of cells in the excitatory or inhibitory populations whose AUC exceeded the significance bounds. The fraction was variable across network realizations, but within each network, it was similar for excitatory and inhibitory cells. Error bars are derived from counting statistics. C. Average choice selectivity (defined as in [1] as the average of $|AUC - 0.5|$). Averages are over all cells (selective and non-selective) and error bars are SEM. D: Accuracy of a classifier over the course of the stimulus presentation. For each set of cells (all, black; 100 e-cells, blue; 100 i-cells, red), we fit the classifier at a single point in time (circle) and classified the stimulus over the trial. E: Accuracy of classifier across network instances (same order as Fig 2). Error bars are ± 1 SE estimated by cross-validation (see

for inhibitory cells. Additionally, the stimulus could be decoded from the population accurately (range: 83% to 95%) and equally well by excitatory (range: 77% to 94%) and inhibitory (range: 76% to 93%) sub-populations (Fig 4). Analysis of connectivity between pairs of neurons with the same selectivity and opposite connectivity showed that inhibitory cells were more likely to connect to excitatory cells of the opposite selectivity, as in the first set of simulations (S4 Fig and S1 Table). Overall, we find that whether the inputs drive only excitatory cells, or both excitatory and inhibitory cells—a rather large change to the structure of the problem—has negligible effect on the salient properties of the network dynamics. This suggests that our results may be largely independent of the many details of the brain networks, and hence are more generally applicable.

Eigenvalues of the linearized rate network and selectivity

We returned to our hypothesis that selectivity in this network emerges because different frequencies of inputs will drive different dynamics in the network, and potentially these could be predicted from the eigenvalue spectrum of the network. To explore this further, we linearized the rate equation about the zero-input fixed point and examined the relationship between eigenvectors in the linearized network and the vector of selectivity values across cells (Methods). Fig 5 shows the distribution of all eigenvalues across networks (Fig 5A) and the distribution of “selectivity-correlated eigenvalues,” the eigenvalues corresponding to eigenvectors whose real or imaginary parts had a significant linear correlation with the pattern of selectivity across the network (Fig 5B). In any network, there were multiple eigenvectors significantly correlated with the selectivity vector: on average, each network had 14 (range, 9 to 27, 6 networks, excitatory inputs only) pairs of complex eigenvectors and 1 (range 0–3; 6 networks, excitatory inputs only) real eigenvector with significant linear correlation ($p < 0.0001$, reflecting a Bonferroni correction for multiple comparisons with a single selectivity vector) with the selectivity vector. While significant, the values of these correlations were moderate in size: the average magnitude (of significant correlations) was 0.27 ± 0.02 (SE) across $N = 6$ networks, and the largest correlation value in each network was 0.52 ± 0.09 (SE, $N = 6$ networks). As we expected, selectivity-correlated eigenvalues were concentrated on the positive half of the plane, which were the modes expected to dominate network dynamics, and the likelihood of a particular eigenvector being highly correlated with the selectivity pattern increased with the size of the real part of its eigenvalue (Fig 5C). In summary, we confirmed our intuition for




Fig 5. Eigenvalue spectrum of the linearized rate network. A. Density plot of all eigenvalues in the linearized network. B. Density plot of eigenvalues corresponding to eigenvectors with significant correlation to the pattern of selectivity across the network. C. The fraction of eigenvectors that are significantly correlated with selectivity pattern, computed as the ratio of density in (B) to that in (A). For A and B, the radius of the distribution varied across networks so to average across networks we normalized by the maximum absolute value of eigenvalues a within each network ($a = 0.168 \pm 0.005$). Density is smoothed with a Gaussian kernel of width $\sigma = a/10$ and zeroed outside the unit circle.

<https://doi.org/10.1371/journal.pcbi.1007875.g005>




Fig 6. Changes in selectivity when input frequencies are changed. A: Density showing AUC (Fig 2A) for the 8 Hz vs. 16 Hz task against the 10 Hz vs. 20 Hz task, in the same network. Equality line is shown for reference. B: Cross-tabulation of selectivity, based on shuffle criterion. Some selective cells become non-selective, and non-selective cells become selective. C: Average cross-tabulation of selectivity across all networks.

<https://doi.org/10.1371/journal.pcbi.1007875.g006>

why a particular pattern of selectivity emerges in a network: these patterns are related to some of the fastest-growing modes of the linearized network in the zero-input fixed point. Predicting which specific modes drive selectivity for a specific task, and how many tasks such a network can perform, remains to be fully explored.

Selectivity in the network under different task conditions

Finally, we performed a new simulation in which the same network discriminated between different input frequencies. Originally, we set the average frequency of inputs to 8 Hz and 16 Hz. We now ask how selectivity changes when the details of the task change, by shifting the average frequency of inputs to 10 Hz and 20 Hz. In all cases, we used the normalization strategy for input amplitudes as before, so that average firing rate did not increase with average frequency. We then compared the single-cell AUC values for single cells on this new discrimination task to those on the original (8 Hz vs. 16 Hz) task. Fig 6A shows the shifts in the selectivity of single cells for an example network. For this simulated network, some cells that had low selectivity on the original task increased their selectivity on the new task, while a subset of selective cells lost selectivity on the new task. We calculated statistical significance for selectivity as in Fig 2. For this network, more cells were selective when the task was to distinguish 10 Hz from 20 Hz than 8 Hz vs. 16 Hz (Fig 6B), shown by the weight in the off-diagonal entries of the cross-tabulation of selectivity for the original and new task (10 Hz vs. 20 Hz). Across different realizations of the network, this was not a strong trend: a fraction of cells either gained or lost selectivity as the task parameters changed (Fig 6C). The fraction of all cells that changed selectivity (in either direction) varied across networks but was always greater than zero, averaging 22% of cells (+/- 9%, range 8% to 42%). To summarize, we predict that the set of selective cells depends on the temporal features of the input and will change if the task changes.

Discussion

We presented a set of simulation results that account for several key features of population recordings in PPC during an evidence accumulation task. This simulated network consists of excitatory and inhibitory neurons with random connectivity. It receives as inputs pulsed sensory signals, which have been filtered by sensory areas. The pulse times are random, with either low or high average rate in time. From the network patterns of activity, we measure single-cell selectivity for input rate, patterns of noise correlations between same and opposite selectivity cells, and readout performance and readout weight distributions. We find that these measures are highly consistent with those observed experimentally. Importantly, our simulations do not

include specific patterns of connectivity between excitatory and inhibitory neurons; any biases that emerge are the result of dynamics shaped by random heterogeneity in connectivity patterns. We suspect this is a fairly generic effect of a network with a broad spectrum of eigenvalues, such as the random networks that we studied, but the theoretical connection between the spectra of a connectivity matrix in a nonlinear network and the emergence of selectivity and distribution of readout weights among excitatory and inhibitory populations remains to be explored.

Experimental work has shown that PPC is integral to performing a sensory evidence integration task [11, 12]. Further, the examination of the representation of choice across PPC showed that both excitatory and inhibitory cells have choice selectivity, that cells with the same selectivity had higher noise correlations, and that decoders trained on the population activity patterns to read out choice had both positive and negative weights for both excitatory and inhibitory cells [1]. Based on these experimental observations, the authors suggested a model in which multiple pools of excitation and inhibition take as input some variable representing choice and, through specific patterns of connectivity, represent choice across the population, with weak individual selectivity, selective patterns of noise correlations, and zero-mean distributions of decoder weights. We showed that an unstructured network produces all of these effects as well.

Previous studies have argued that, while the posterior parietal cortex is involved in performing a visual sensory evidence integration task [11, 12] in rodents, PPC does not itself integrate evidence [6, 12]. One possibility, consistent with our simulations, is that the population response is subtly distinct across input frequencies, and these distinctions are learned through a reinforcement mechanism in some other area, which then feeds back to PPC to enhance the distinction between sensory inputs. This feedback mechanism may interact with other biases, such as motivation and trial history [16, 17], which modulate choice. Such a feedback loop could further enhance apparent selectivity in connectivity in PPC because it emphasizes dynamics that were shaped by heterogeneity in connectivity.

Finally, we found that stimulus information could be decoded very early in the stimulus window. The technical reason for this is that due to the input scaling, the first pulse of the input series carried stimulus information. Scaling was used to make the decoding task more challenging and to reproduce the experimental observation that PPC responses to low-frequency inputs are not lower than the PPC responses to high-frequency inputs. We did not implement more realistic adaptation dynamics, which could be the mechanism underlying such scaling. One would expect that such adaptation mechanisms (e.g., short-term depression or facilitation) could be transformed into a population code in a spiking network, readily decoded downstream [18], and it is an interesting question for future study whether spiking networks such as these replicate the experimentally observable quantities that we focused on here. Adding this feature to the model would slow the ramp-up in decoding accuracy, but would also require more choices about adaptation rates and tuning. Moreover, such elaboration of the model is superfluous to answering the question of whether a heterogeneous network of excitatory and inhibitory units can distinguish between inputs with different temporal frequencies and produce statistical features comparable to experimental observations.

Relationship to other network models

Recurrent network models have been used in other contexts to study how networks of neurons perform specific tasks or simulate neural activity [17, 19–24]. By comparison with such models, our model is exceedingly simple: it is a sparsely connected, random network of excitatory and inhibitory neurons with a firing rate non-linearity. In this network, temporal information

(about stimulus frequency) is transformed into a spatial representation. There are a number of ways to make this model more realistic. For instance, the emergent spatially distinct neural representations could trigger distinct neural trajectories [24], matching the spatio-temporal multineuronal dynamics on single trials more closely. We did not pursue this here, as our goal was to show that heterogeneity in network connectivity could explain many features of population recordings during a simple discrimination task.

One of our results is that selectivity can emerge for task parameters from an unstructured, random network. Several theoretical studies have previously examined the emergence of selectivity in random networks [23, 25–30]. As early as the 1970s, it was suggested that orientation selectivity in primary visual cortex could emerge from random projections from geniculate inputs [25, 26]. In balanced state networks, such selectivity would be robust due to the dynamic cancellation of non-selective inputs [23, 27–29]. Selectivity from random projections could be enhanced through learning mechanisms either in the feed-forward projections [25] or in the recurrent cortical network [23]. The mechanism for selectivity presented here adds to these by demonstrating selectivity for parameters that reflect temporal, rather than spatial, patterns. Additionally, our network matches specific experimental recordings quantitatively in several aspects, such as the accuracy of population activity classifiers and the average selectivity.

Our model is also related to reservoir computing. In a reservoir computing or liquid state machine [31, 32] framework, a recurrent network with fixed, random connections is driven by inputs. When the network is tuned with maximal eigenvalues close to the boundary of stability, the network generates a “reservoir” of temporal dynamics, from which a target behavior can be achieved by training readout weights appropriately. A reservoir model was suggested for capturing the function of prefrontal cortex in macaques performing a task that required switching between exploration and exploitation of probabilistic rewards [33]. There is a variation of the reservoir framework in which the recurrent network is in the chaotic regime, but particular dynamics are stabilized by adjusting either recurrent weights or feedback to the recurrent network. Such networks can generate stereotyped trajectories [21, 24, 34]. Our network has random connections between neurons, but it is not tuned to be a reservoir: in the linearized network, eigenvalues are well within the stable regime (see Fig 5, caption). In this simplified task setting, we have shown that modeling the network in posterior parietal cortex in this way explains experimental observations of single-cell selectivity, pairwise correlations, and population decoding properties. Moving forward, the most interesting question is not whether or not our network is a reservoir, but what reservoirs can do and what regular random networks can do, versus what the brain actually does. This is a topic that will take many years to fully develop, and what we have done here is a small step in that direction.

How realistic are our parameter choices?

Relative to a range of past modeling studies with a similar degree of realism in the model [1, 19, 24], the parameter choices made in this study are squarely in the middle of the pack. Yet, there are apparent discrepancies with experimentally observed connectivity statistics in cortical microcircuits. For instance, experimental measurements in sensory areas show that the probability of connection from inhibitory to excitatory cells (50 to 80%, depending on cortical area) is much higher than the probability of connection within the excitatory population (< 20%) [35–38]. These are distinct from the parameters used in our model, which had a uniform probability of connection of 20%. Our analysis of the relationship between selectivity and the eigenvalues of the linearized network (Fig 5) suggests that any sufficiently broad distribution of network eigenvalues will generate some selective modes, and the spectrum of a random

network depends on the synaptic weight distribution, not only the probability of a connection. Based on available data on synaptic weight distributions, quantified from post-synaptic potentials measured intracellularly in sensory areas [37, 38], the probability of a strong connection may be relatively low even when the overall probability of a connection is high. As more details of synaptic weight distributions are characterized across cortical areas in the mouse, it will be interesting to see to what extent the mechanism for selectivity explored here accounts for general features of dynamics and selectivity across the cortex.

Predictions for future experiments

Finally, our theory makes the following prediction, which should be verifiable experimentally. Suppose, the decision boundary for reporting low- versus high-frequency stimuli changed. If the network is structured as separate pools of excitatory and inhibitory neurons representing choice for one or the other stimulus category, then the representation of choice in PPC will not change with the task. If instead functional organization is generated by emergent network properties, when the task changes, the selectivity of individual cells will shift, as different pools of neurons represent the low- versus high-frequency stimuli. These pools of neurons would be overlapping, as the frequencies being discriminated change.

Conclusions

As ever larger populations of neurons are simultaneously recorded, and experiments frequently focus on a variants of well-controlled sensory discrimination tasks, we face a tremendous challenge in inferring mechanism from observations. Very generally, there are two mechanisms that could account for diverse and mixed selectivity along with patterned functional connectivity across a population of neurons engaged in some experimental task. The first is that the network is wired specifically to achieve this, and if that is the case, then one must also explain the developmental or learning process that produced such intricate wiring specificity in the network. The alternative, which we explored here, is that the observed patterns of selectivity can be explained as an emergent phenomena from simple patterns of statistical connectivity. In the particular case examined here, we were able to reproduce distributions of selectivity, functional connectivity, and population readout weight patterns that were experimentally observed. Even though we analyzed the emergence of selectivity in a specific experiment, we believe that similar conclusions could hold in other applications, in which a broad distribution of selectivity is observed.

Materials and methods

Network simulation

Random recurrent firing rate network. We simulated a recurrent neural network of $N = 500$ neurons ($N_E = 400$ excitatory neurons and $N_I = 100$ inhibitory neurons) using standard firing rate equations:

$$\tau \dot{\mathbf{x}} = J\mathbf{r} + \mathbf{ci}(t) - \mathbf{x}, \quad (1)$$

$$r_i = g(x_i), \quad (2)$$

$$g(x) = 0.5(1 + \tanh(x - b)), \quad (3)$$

where x_i is the “membrane potential” of neuron i and r_i is its firing rate, obtained from x_i through the nonlinear transfer function g . We set τ to 1 so all time is measured in units of the

unit time constant. The transfer function g scales activity between 0 and 1 for all cells. We included a bias term b , set to 2, and this ensured small (0.05 or less) spontaneous ($i(t) = 0$) firing rates in steady state. Neurons interacted through the matrix J , which was sparse (20% non-zero) and defined whether neurons are excitatory (80% of cells) or inhibitory. Weights originating from excitatory (inhibitory) neurons were drawn from a normal distribution with mean $g_E = 0.18$ ($g_I = -0.72$) and standard deviation 0.045 (both excitatory and inhibitory), which balanced excitatory and inhibitory synaptic inputs on average across the population. Synaptic delays are not modeled. Finally, the external stimulus was the product of a scalar function $i(t)$ capturing the impulses (described below) and the binary input vector $\mathbf{c}$. The vector $\mathbf{c}$ was 1 for the excitatory cells that received direct inputs (20% of cells, randomly selected) and otherwise 0, and this vector was fixed. In other words, the same subset of cells received inputs for all stimuli.

In a second set of simulations (Fig 4), all excitatory cells receive direct inputs and no inhibitory cells receive direct inputs. Connections were random with probability $p = 0.2$. Connection strengths between excitatory cells and inhibitory cells were balanced so that average inputs to either pool were close to zero, which meant that inhibitory to excitatory connections were slightly stronger than inhibitory to inhibitory connections in order to match the input-driven current. Specifically, the connection weights originating from the excitatory pools had parameter $g_{E \leftarrow E} = g_{I \leftarrow E} = 0.2$. Weights originating from the inhibitory pool had parameter $g_{E \leftarrow I} = -0.8$ and $g_{I \leftarrow I} = -0.82$.

The network was simulated using custom-written code in Matlab.

The frequency discrimination task. In the simulated task, the network was driven by stimuli consisting of irregular impulses with average frequency f (Fig 1A). The length of the sampling period in the simulated task was 50τ , corresponding to an effective τ of 20 ms. In the first set of simulations, either $f = 8$ Hz or $f = 16$ Hz. In simulations in Fig 6, we analyzed $f = 10$ Hz to $f = 20$ Hz.

For each trial with average frequency f , the input times (t_k) were selected randomly with uniform probability from the stimulus interval (50τ , or 1000ms) by drawing a fixed number of time points (e.g., 8). Temporal precision of impulse timing was 0.01τ , so impulse times were drawn from the integers 1 to 5000 without replacement. For simulated trials with the same impulse frequency, the trial-by-trial variation in input times was the only source of variability across trials.

We assumed that upstream sensory networks filtered the pulse inputs, so the overall input current to the network was described by

$$i(t) = \alpha \sum_{t_k < t} \frac{(t - t_k)^2}{a^2} \exp\left(-\frac{t - t_k}{a}\right), \quad (4)$$

where a is a filtering timescale of pre-processing networks and the summation was taken over all $t_k < t$. We set $a = 0.5$ for all simulations (full width at half max of a single pulse is 1.7τ).

Input currents were scaled by a frequency-dependent factor α to match the average firing rate across the network between conditions (Fig 1, S1 Fig). To implement this scaling, 25 sets of input parameters (frequency and amplitude) were simulated with input frequencies from 0.1 to 0.5 (per τ), and amplitudes from 0.3 to 15, a range spanning parameters that generated a range of firing rates in the network. From each of these 25 simulations, we computed the average network firing rate, and we interpolated this surface to find contours of equal firing rate. For the simulated experiment, we used combinations of frequency and amplitude that fall on a fixed contour. Across networks, the typical amplitude ratio between low- and high-frequency

inputs was 2.1. We verified post-simulation that the average firing rates match across frequency conditions (see, e.g., Fig 1D).

Simulation analysis

Single neuron selectivity. For a pair of stimuli (e.g., 8 Hz and 16 Hz), we used an ideal observer to determine selectivity in single neurons. For each cell, the area under the receiver-operator curve (AUC) was computed nonparametrically from the distribution of low-frequency responses and the distribution of high-frequency responses at the end of the stimulus period [15]. Significance bounds are the 2.5-97.5 percentiles of the trial-shuffled distributions. Single neurons were selective if their AUC value fell outside the significance bounds. The probability density plot of AUC values is smoothed with a Gaussian kernel of width $\sigma = 0.02$, which is the standard error of AUC values estimated by bootstrapping. Selectivity (Fig 2C) is reported as $2|AUC - 0.5|$, an index where 0 is non-selective and 1 is maximally selective.

Selectivity and eigenvector analysis. We computed, for each network, the linearized connectivity:

$$\mathcal{J}_{ij} = J_{ij} \frac{\partial g(x_j)}{\partial x_j} \Big|_{x_j(0)} \quad (5)$$

evaluated at the fixed point

$$x_i(0) = \sum_j J_{ij} g(x_j(0)), \quad (6)$$

which was found by evolving the network equations with zero input. Note that $\mathcal{J}$ is real but not symmetric, so to solve the linear equation we diagonalize $\mathcal{J} = \mathcal{V} \mathcal{D} \mathcal{V}^{-1}$, where the diagonal matrix $\mathcal{D}$ has entries $\mathcal{D}_{ii} = \lambda_i = a_i + i\omega_i$ and the columns of $\mathcal{V} = [v_1 v_2 \dots v_N]$ are the associated eigenvectors. The solutions to the linear equation are

$$x_i(t) = \sum_j v_j \int e^{(1-\lambda_j)(t-\tau)} \Theta(t-\tau) (\mathcal{V}^{-1} \mathbf{c}(\tau))_j d\tau. \quad (7)$$

In words, the activity in the (complex) eigenmodes is a convolution of the input function with $\exp(-(1-\lambda_i)t) = \exp((a_i-1)t)\exp(i\omega_i t)$. We expect that the pattern of selectivity across the network of N neurons is related to the eigenvectors of modes with eigenvalues with large real parts a_i and imaginary parts ω_i that pick up differences from different input frequencies. We computed the linear correlation between the real and imaginary parts of each eigenvector with the AUC for each neuron for all eigenvectors in each network. Significant correlations were those for which p-values are less than $0.05/N$. In Fig 5, we analyze eigenvalue distributions across networks. The radius of the distribution varied from network to network, so we normalized by the maximum absolute value of eigenvalues a within each network. Density is smoothed with a Gaussian kernel of width $\sigma = a/10$ and zeroed outside the unit circle.

Pairwise cross-correlations. We computed pairwise correlations from the activity at the end of the stimulus period by subtracting the stimulus-averaged response and then computing neuron-neuron correlation coefficients.

Classifier analysis. The goal of the classification was to discriminate between two input frequencies using the simulated activity patterns. Simulated network activity was temporally down-sampled by averaging over time windows of size 1 (τ). Neurons that received direct inputs were excluded from the decoding analysis, leaving 320 excitatory neurons and 100 inhibitory neurons. We fit classifiers separately on the full population ("all," 420 neurons), a

subset of excitatory neurons (“exc-sub”, 100) and a subset of inhibitory neurons (“inh”, 100). We split the 800 simulated (400 in each condition) into “test” and “training” sets. We trained a classifier (linear kernel, SVM) to predict the stimulus label based on the activity (in “all,” “exc-sub” and “inh”) at each time point. We trained the classifier on the z-scored activity from each cell [12]:

$$z_i = \frac{r_i - \bar{r}_i}{\sigma_{r_i}}. \quad (8)$$

The classifier finds a rule ξ, η

$$\xi * \mathbf{z} > \eta. \quad (9)$$

We also calculated the weights ($\mathbf{w}$) and bias (b) that operate on firing rates directly

$$\mathbf{w} * \mathbf{r} > b. \quad (10)$$

Classifier accuracy was calculated on the reserved test set. In Fig 3A, a classifier was fit at a single time point at the end of the stimulus window and tested at all other time points. Uncertainty in classifier accuracy was estimated by fitting the classifier using different cuts of the data: train/cross-val/test sets were drawn randomly, a classifier fit, and weights and accuracy on the test set recorded. This was repeated 50 times, and the error bar on classifier accuracy is the standard deviation across test set accuracy generated in this way.

Supporting information

S1 Fig. Setting the average network firing rate. A: Surface is a spline interpolation of average (over time and neurons) firing rate across 25 combinations of amplitude and frequency. Color on parula scale indicates firing rate (blue to yellow). Lines show contours at fixed average firing rate. For each firing rate (indicated by line color), we extract amplitudes and frequencies on the corresponding contour. B: Amplitudes for the high-frequency (freq2 = 2*freq1) input, plotted against the lower frequency (freq1). C: Amplitudes for the low-frequency input, plotted against the lower frequency (freq1). The highest firing rate contour that is defined over freq1 = 0.14 to 0.2 is simulated. Units of frequency are per τ ; multiplication by 50 converts to Hz.
(TIF)

S2 Fig. Single-cell response distributions and selectivity analysis. Example histograms of single-cell firing rate for the low- (yellow) and high-frequency (purple) inputs, with the area under the receiver-operator curve (AUC) statistic indicated. AUC values close to 0.5 are non-selective. A: (column) examples from one of the networks examined in Figs 1 to 3. B: (column) examples from a network receiving inputs to all excitatory cells (Fig 4).
(TIF)

S3 Fig. Additional parameter combinations. Fixing network topology (i.e., which elements of $\mathbf{J}$ are non-zero), we simulated three networks: with original weights (i), with all synaptic weights scaled by a factor of 1.5, (ii) and with homogeneous excitatory and homogeneous inhibitory synaptic weights (iii). A: Image of network connectivity for 25 (of 500 total) neurons showing that the topology was kept the same for each simulation. B: Distribution of non-zero excitatory and inhibitory weights in the network. Note that there are approximately four times as many excitatory weights, but they are on average a quarter of the strength of inhibitory weights. C: Histogram of stimulus 1 (yellow) and stimulus 2 (purple) population firing rates for each parameter scaling. Firing rates are matched for each simulation individually; these are

operating over a similar population firing rate range. D: Decoding accuracy of the full population in each network. E: Decoding accuracy of excitatory and inhibitory cells in each network. D and E show that decoding accuracy persists after a drastic parameter change, for this network topology. F: Decoding performance for all simulated networks. (TIF)

S4 Fig. Connectivity statistics for selectivity-defined pools. A: Average probability of connection between pairs of cells with the same selectivity (top row) and with opposite selectivity (bottom row) for pairs of excitatory cells (left column) and pairs of inhibitory cells (right column). B: Summary diagram of connectivity for the network when cells are divided into pools based on selectivity. Blue arrows: excitatory connections. Red end-stop lines: inhibitory connections. Weights label the probability of connection relative to the baseline (random chance) probability (1.0). Standard error across network simulations is 0.05 for inhibitory connections and 0.02 for excitatory ones. Connection probabilities are averaged across all networks in which there were cells in each of the four pools (excluding network instance 13, see A, lower right) and symmetrized across selectivity pools (i.e., the probability of connection from E1 to E2 is the same as the probability from E2 to E1, and from E1 to I1 is the same as from E2 to I2, and so on). Color of number (blue/red) matches the color of the connection drawn in the graph. C: Same diagram as in B, but for the network that only receives inputs to e-cells. Both cases show excess connectivity between pools I1 and I2 and from I1 to E2 (and from I2 to E1). (TIF)

S1 Table. Selectivity induces non-random connectivity in the network. For each network, pairs of neurons were randomly connected with probability $p = 0.2$. Following simulation, cells are divided into pools based on cell “type” (e-cells or i-cells) and selectivity (stimulus 1 and 2), yielding four pools: E_1 , E_2 , I_1 , and I_2 . The average probability of connection between each of the pools was calculated. Table entries are averages $\pm$ standard error across networks, normalized by the baseline probability of connection (0.2). Entries larger than 1 indicate more connections than the baseline expectation of 0.2. (TIF)

Author Contributions

Conceptualization: Audrey Sederberg, Ilya Nemenman.

Formal analysis: Audrey Sederberg.

Investigation: Audrey Sederberg, Ilya Nemenman.

Methodology: Audrey Sederberg, Ilya Nemenman.

Resources: Ilya Nemenman.

Software: Audrey Sederberg.

Visualization: Audrey Sederberg, Ilya Nemenman.

Writing – original draft: Audrey Sederberg.

Writing – review & editing: Audrey Sederberg, Ilya Nemenman.

References

1. Najafi F, Elsayed GF, Cao R, Pnevmatikakis E, Latham PE, Cunningham JP, et al. Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during

- Learning. *Neuron*. 2020 Jan; 105(1):165–179.e8. Available from: <https://linkinghub.elsevier.com/retrieve/pii/S0896627319308487> PMID: 31753580
2. Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*. 2019; 364 (6437). <https://doi.org/10.1126/science.aav7893> PMID: 31000656
 3. Gauthier JL, Tank DW. A Dedicated Population for Reward Coding in the Hippocampus. *Neuron*. 2018; 99(1):179–193.e7. Available from: <https://doi.org/10.1016/j.neuron.2018.06.008> PMID: 30008297
 4. Heeger DJ. Theory of cortical function. *Proceedings of the National Academy of Sciences of the United States of America*. 2017; 114(8):1773–1782. <https://doi.org/10.1073/pnas.1619788114> PMID: 28167793
 5. Najafi F, Churchland AK. Perceptual Decision-Making: A Field in the Midst of a Transformation. *Neuron*. 2018; 100(2):453–462. Available from: <https://doi.org/10.1016/j.neuron.2018.10.017> PMID: 30359608
 6. Brody CD, Hanks TD. Neural underpinnings of the evidence accumulator. *Current Opinion in Neurobiology*. 2016; 37:149–157. Available from: <http://dx.doi.org/10.1016/j.conb.2016.01.003> PMID: 26878969
 7. Mehta ML. *Random matrices*. vol. 142. Elsevier; 2004.
 8. Rajan K, Abbott LF. Eigenvalue spectra of random matrices for neural networks. *Physical Review Letters*. 2006; 97(18):2–5. <https://doi.org/10.1103/PhysRevLett.97.188104>
 9. Aljadeff J, Stern M, Sharpee T. Transition to chaos in random networks with cell-type-specific connectivity. *Physical Review Letters*. 2015; 114(8):1–5. <https://doi.org/10.1103/PhysRevLett.114.088101>
 10. Wang XJ. Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. *Neuron*. 2002 dec; 36(5):955–968. Available from: <http://linkinghub.elsevier.com/retrieve/pii/S0896627302010929> PMID: 12467598
 11. Raposo D, Kaufman MT, Churchland AK. A category-free neural population supports evolving demands during decision-making. *Nature Publishing Group*. 2014; 17(12):1784–1792. Available from: <http://dx.doi.org/10.1038/nn.3865>
 12. Licata AM, Kaufman MT, Raposo D, Ryan MB, Sheppard JP, Churchland AK. Posterior Parietal Cortex Guides Visual Decisions in Rats. *Journal of Neuroscience*. 2017; 37(19):4954–4966. <https://doi.org/10.1523/JNEUROSCI.0105-17.2017> PMID: 28408414
 13. Aksay E, Olasagasti I, Mensh BD, Baker R, Goldman MS, Tank DW. Functional dissection of circuitry in a neural integrator. *Nature Neuroscience*. 2007; 10(4):494–504. <https://doi.org/10.1038/nn1877> PMID: 17369822
 14. Whitlock JR. Posterior parietal cortex. *Current Biology*. 2017; 27(14):R691–R695. Available from: <http://dx.doi.org/10.1016/j.cub.2017.06.007> PMID: 28743011
 15. Green DM, Swets JA. *Signal Detection Theory and Psychophysics*. New York: John Wiley and Sons; 1966.
 16. Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Scholvinck ML, et al. The Detection of Visual Contrast in the Behaving Mouse. *Journal of Neuroscience*. 2011; 31(31):11351–11361. Available from: <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.6689-10.2011> PMID: 21813694
 17. Morcos AS, Harvey CD. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature Neuroscience*. 2016; 19(12):1672–1681. <https://doi.org/10.1038/nn.4403> PMID: 27694990
 18. Buonomano DV, Merzenich MM. Temporal Information Transformed into a Spatial Code by a Neural Network with Realistic Properties. *Science*. 1995; 267(February):1028–1030.
 19. Renart A, De La Rocha J, Bartho P, Hollender L, Parga N, Reyes A, et al. The asynchronous state in cortical circuits. *Science*. 2010; 327(5965):587–590. <https://doi.org/10.1126/science.1179850> PMID: 20110507
 20. Toyozumi T, Abbott LF. Beyond the edge of chaos: Amplification and temporal integration by recurrent networks in the chaotic regime. *Physical Review E* 2011; 84(5):1–8. <https://doi.org/10.1103/PhysRevE.84.051908>
 21. Laje R, Buonomano DV. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature Neuroscience*. 2013; 16(7):925–933. Available from: <http://dx.doi.org/10.1038/nn.3405>
 22. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503(7474):78–84. Available from: <http://dx.doi.org/10.1038/nature12742> PMID: 24201281
 23. Sadeh S, Clopath C, Rotter S. Emergence of Functional Specificity in Balanced Networks with Synaptic Plasticity. *PLoS Computational Biology*. 2015; 11(6):1–27. <https://doi.org/10.1371/journal.pcbi.1004307>

24. Rajan K, Harvey CD, Tank DW, Rajan K, Harvey CD, Tank DW. Recurrent Network Models of Sequence Generation and Memory Recurrent Network Models. *Neuron*. 2016; 90(1):128–142. Available from: <http://dx.doi.org/10.1016/j.neuron.2016.02.009> PMID: 26971945
25. von der Malsburg C. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*. 1973; 14(2):85–100. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/7621943> PMID: 4786750
26. Ringach DL. Haphazard wiring of simple receptive fields and orientation columns in visual cortex. *Journal of Neurophysiology*. 2004; 92(1):468–476. <https://doi.org/10.1152/jn.01202.2003> PMID: 14999045
27. Pattadkal JJ, Mato G, van Vreeswijk C, Priebe NJ, Hansel D. Emergent Orientation Selectivity from Random Networks in Mouse Visual Cortex. *Cell Reports*. 2018; 24(8):2042–2050.e6. Available from: <https://doi.org/10.1016/j.celrep.2018.07.054> PMID: 30134166
28. Hansel D, van Vreeswijk C. The Mechanism of Orientation Selectivity in Primary Visual Cortex without a Functional Map. *Journal of Neuroscience*. 2012; 32(12):4049–4064. Available from: <https://www.jneurosci.org/content/32/12/4049> PMID: 22442071
29. Pehlevan C, Sompolinsky H. Selectivity and sparseness in randomly connected balanced networks. *PLoS ONE*. 2014; 9(2). <https://doi.org/10.1371/journal.pone.0089992> PMID: 24587172
30. Natale JL, Hentschel HGE, Nemenman I. Precise Spatial Memory in Local Random Networks. *arXiv*. 2019; Available from: <https://arxiv.org/abs/1911.06921>.
31. Jaeger H, Haas H. Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication. *Science*. 2004; 304(5667):78–80. Available from: <https://science.sciencemag.org/content/304/5667/78> PMID: 15064413
32. Maass W, Natschläger T, Markram H. Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation*. 2002 nov; 14(11):2531–2560. Available from: <https://doi.org/10.1162/089976602760407955> PMID: 12433288
33. Enel P, Procyk E, Quilodran R, Dominey PF. Reservoir Computing Properties of Neural Dynamics in Prefrontal Cortex. *PLoS Computational Biology*. 2016; 12(6):1–35. <https://doi.org/10.1371/journal.pcbi.1004967>
34. Sussillo D, Abbott LF. Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*. 2009; 63(4):544–557. Available from: <http://dx.doi.org/10.1016/j.neuron.2009.07.018> PMID: 19709635
35. Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, Mrsic-Flogel TD. Functional specificity of local synaptic connections in neocortical networks. *Nature*. 2011; 473(7345):87–91. <https://doi.org/10.1038/nature09880> PMID: 21478872
36. Hofer SB, Ko H, Pichler B, Vogelstein J, Ros H, Zeng H, et al. Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature Neuroscience*. 2011; 14(8):1045–1052. <https://doi.org/10.1038/nn.2876> PMID: 21765421
37. Avermann M, Tomm C, Mateo C, Gerstner W, Petersen CCH. Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex. *Journal of Neurophysiology*. 2012 jun; 107(11):3116–3134. Available from: <https://www.physiology.org/doi/10.1152/jn.00917.2011> PMID: 22402650
38. Pala A, Petersen CCH. In Vivo Measurement of Cell-Type-Specific Synaptic Connectivity and Synaptic Transmission in Layer 2/3 Mouse Barrel Cortex. *Neuron*. 2014; 85(1):68–75. Available from: <http://linkinghub.elsevier.com/retrieve/pii/S0896627314010770> PMID: 25543458