Untuned but not irrelevant: A role for untuned neurons in sensory information coding

Joel Zylberberg[1,2,3]

[1]Department of Physiology and Biophysics, Center for Neuroscience, and Computational Bioscience Program, University of Colorado School of Medicine, Aurora, CO 80045

[2]Department of Applied Mathematics, University of Colorado, Boulder, CO 80309

[3]Learning in Machines and Brains Program, Canadian Institute For Advanced Research, Toronto, ON M5G1Z8

May 6, 2017

Abstract

In the sensory systems, most neurons' firing rates are tuned to at least one aspect of the stimulus. Other neurons are untuned, meaning that their firing rates appear not to depend on the stimulus. Previous work on information coding in neural populations has ignored the untuned neurons, based on the tacit assumption that they are unimportant. Using theoretical calculations and analyses of *in vivo* neural data, I show that untuned neurons can contribute significantly to the population code. Ignoring untuned neurons can lead to severe underestimates of the amount of stimulus information encoded, and in some cases population codes can be made more informative by replacing tuned neurons with untuned ones.

Introduction

When you look at a picture, signals from your eyes travel along the optic nerve to your brain, where they evoke activity in neurons in the thalamus and visual cortex. As sensory systems neuroscientists, we ask how these patterns of stimulus-evoked brain activity reflect the outside world – in this case, the picture at which you are looking. Other related work asks how patterns of activity in different parts of the brain reflect motor commands sent to the muscles. Answers to these questions are important both for basic science, and for brain-machine interface technologies that either decode

brain activity to control prosthetic limbs or other devices [1, 2, 3], or stimulate the brain to alleviate sensory deficits [4, 5].

For decades, researchers have addressed these information coding questions by recording neural activity patterns in animals while they are being presented with different stimuli, or performing different motor tasks. That work revealed that many neurons in the relevant brain areas show firing rates that depend systematically on the stimulus presented to the individual, or on the motor task. This neural "tuning" underlies the ability of these neural circuits to encode information about the stimulus and/or behavior. At the same time, many neurons appear to be untuned, thus showing little or no systematic variation in their firing rates as the stimulus (or behavior) is changed [6]. These untuned neurons are typically ignored in studies of neural information coding because it is presumed that they do not contribute [7]. Instead, data collection and analysis are typically restricted to the tuned neurons (for example, consider the selection criteria used by [8, 9]). Here I question that approach, by asking whether and how untuned neurons can contribute to neural information coding.

To address this question, I used theoretical calculations, and then verified the predictions from those calculations by analyzing 2-photon imaging data collected in the visual cortex of mice that were shown drifting grating stimuli [10]. For the theoretical calculations, I used a common mathematical model of the neural population responses to sensory stimulation [11, 12, 13, 14, 15, 16, 17, 9, 18, 19, 20, 17, 21]. This model describes key features of sensory neural responses: the stimulus tuning (or lack thereof) of individual neurons; the trial-by-trial deviations (or "noise") in the neural responses [9, 22, 23, 24, 25, 26]; and the potential for that noise to be correlated between neurons [27, 28, 9, 29, 30, 27, 31, 32, 33, 34, 35, 36]. For different conditions – for example, including vs. excluding untuned neurons – I computed the amount of information about the stimulus that is encoded in the population firing patterns. By comparing the information across conditions, I characterized the impact that untuned neurons can have on the neural population code.

My main finding is that ignoring untuned neurons can lead to substantial underestimates of the information in a neural population code. This means that our understanding of sensory information coding could be incomplete unless we measure from, and consider the effects of, untuned neurons. Moreover, for brain-machine interface technologies, performance could be improved by using decoders that take into account the activities of both tuned and untuned neurons. A secondary finding is that neural population codes can sometimes be made more informative by replacing some of the tuned neurons with untuned ones. This provides a potential explanation for why these neurons exist: despite their lack of explicit stimulus tuning, they can improve the brain's ability to encode useful information about the outside world.

Results

I first study a theoretical model of information coding in neural populations, to understand whether and how untuned neurons contribute to information coding. I then validate the main prediction from the theory by analyzing data collected in mouse visual cortex.

Theoretical study of how untuned neurons impact sensory information coding

To investigate the role of untuned neurons in sensory information coding, I studied populations of neurons that encode information about the motion direction of a visual stimulus via their randomly shaped and located tuning curves (Fig. 1A). Many different population sizes were considered. For each population, 70% of the neurons were tuned, and the other 30% were untuned. (These numbers match the fraction of well-tuned neurons selected for analysis in a recent population imaging study [36], and are comparable to the fraction of tuned neurons in the experimental data that we study below. Choosing larger or smaller fractions of untuned neurons does not qualitatively affect the results – see Fig. S1.) The untuned neurons had flat tuning curves that did not depend on the stimulus – see the dashed lines in Fig. 1A.

The neurons had Poisson-like variability: for each cell, the variance over repeats of a given stimulus was equal to the mean response to that stimulus. This mimics the experimentally observed relation between means and variances of neural activities [23, 21]. The variability was correlated between cells, and the correlation coefficients were chosen to follow the "limited-range" structure reported experimentally [28, 16, 37, 38, 39], and used in previous theoretical studies [11, 12, 13, 18]. With this structure, the correlation coefficients were large for neurons with similar preferred directions, and smaller for neurons with very different preferred directions (see Methods and Fig. 1B).

For each population, I computed the Fisher information (Fig. 1D, black curve), which quantifies how well an observer – like a downstream neural circuit – can estimate the stimulus direction angle from the neural activities (see Methods). I compared that with the Fisher information obtained from only the tuned subset of neurons – in other words, the information that would be obtained if the untuned cells were ignored (Fig. 1D, red curve). The difference was stark. Ignoring the untuned neurons leads to a dramatic underestimate of the encoded stimulus information. This emphasizes that, despite their lack of stimulus dependence, the untuned neurons can still contribute significantly to the population code.

Does this effect arise because there is some special property to the untuned neurons, or because ignoring *any* neurons (tuned or not) sacrifices information? To address this question, I computed the information for random subsets of 70% of the neurons in each population (i.e., subsets the same size as are obtained when untuned neurons are ignored). This yielded less information than did the full populations (Fig. 1D, blue curve), and showed more information than when only untuned neurons were ignored. This result was surprising because these random subsets contain both tuned and untuned neurons. The fact that they contain more information than subsets of the same size, but with only tuned neurons, implies that an untuned neuron can sometimes contribute more to a population code than does a tuned one.

These results suggest that, at least under some conditions, the brain's population code could be made more informative by replacing some tuned neurons with untuned ones. To test this hypothesis, I repeated the calculations from Fig. 1, but with different fractions of the population untuned. For each population size and tuning fraction, I computed the Fisher information in the entire population.

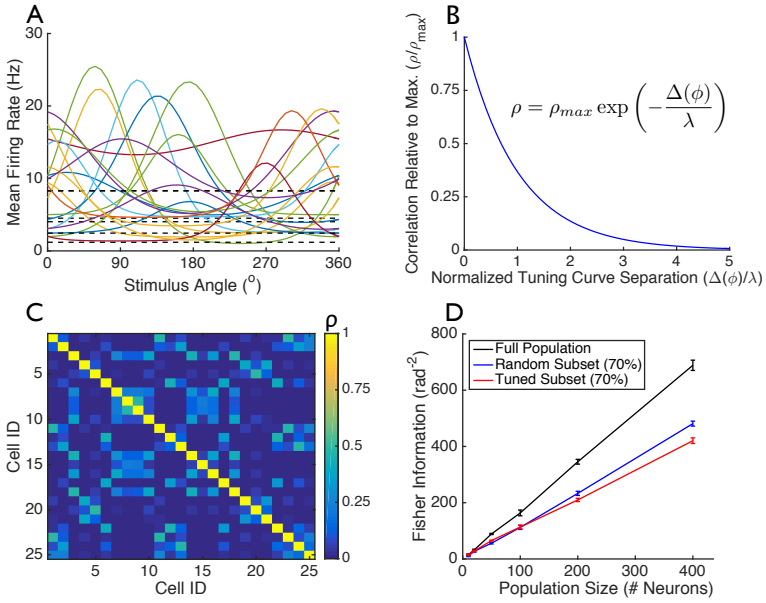

Figure 1: **Untuned neurons can play an important role in sensory information coding.** I considered populations of neurons with randomly shaped and located tuning curves (A). Of those neurons, 70% were tuned to the stimulus, whereas 30% were untuned – their mean firing rates do not depend on the stimulus (dashed black lines in panel A). The neurons' trial to trial variability was Poisson-like and correlated between neurons. These correlations followed the "limited-range" structure (B) with $\rho_{max} = 0.5$ and $\lambda = 0.5$ radians (29^o). The mean correlation coefficients (averaged over neurons) were 0.08, which is comparable to values reported in primary visual cortex [28]. (Modifying these values did not qualitatively change our results – see Fig. S2). The correlation matrix for a small population is shown in panel (C). For different sized populations, I computed the Fisher information, which quantifies how well the stimulus can be estimated from the neural population activities (D). The different lines correspond to: the Fisher information for the full neural populations (black); the Fisher information for the tuned 70% of the populations (red); or the Fisher information for random subsets of 70% of the populations (including both tuned and untuned cells: blue). Data points shown in (D) are mean ± S.E.M., computed over 5 different random draws of the tuning curves.

These calculations show that in large neural populations, having at least some untuned neurons can lead to better information coding than having only tuned neurons (Fig. S1). While this effect was modest in magnitude, a larger effect was observed in populations with a different correlation structure (Fig. 3B). These findings suggest a functional reason for why there are untuned neurons in the sensory areas.

How do untuned neurons contribute to neural information coding?

In Fig. 1, and in reality, the noise is correlated between neurons. This means that, while the untuned neurons' activities do not directly reflect the stimulus, they do reflect the trial-specific noise in the other neurons' activities. Accordingly, a downstream readout – like the circuit receiving these neural spikes – can obtain a less noisy estimate of the stimulus by using the untuned neurons' activities to estimate the noise in the activities of the tuned neurons, and subtracting that noise estimate from the observed firing rates. Ignoring untuned neurons leads to the loss of the information available through this "de-noising".

To illustrate this point, I considered a pair of neurons, one of which is tuned to the stimulus (Fig. 2A). In response to stimulation, the neurons give noisy responses, and that noise is correlated between the two cells. When plotted in the space of the two cells' firing rates, the distributions of neural responses to each stimulus are defined by ellipses, shown in Fig. 2B. (These are the 1 standard deviation probability contours.) The correlation between cells is reflected in the fact that these ellipses are diagonally oriented. These ellipses are well separated, meaning that the neural responses to the different stimuli do not overlap, and so it is relatively unambiguous to infer from the neural firing rates which stimulus was presented. For contrast, consider the neural activities observed when the untuned neuron is ignored. In that case, only the tuned neuron is observed, and the distributions in its responses to the different stimuli overlap substantially (Fig. 2B, right vertical axis). This means that, based on only observations of the tuned cell, the stimulus cannot reliably be determined. Ignoring the untuned neuron leads to a loss of stimulus information.

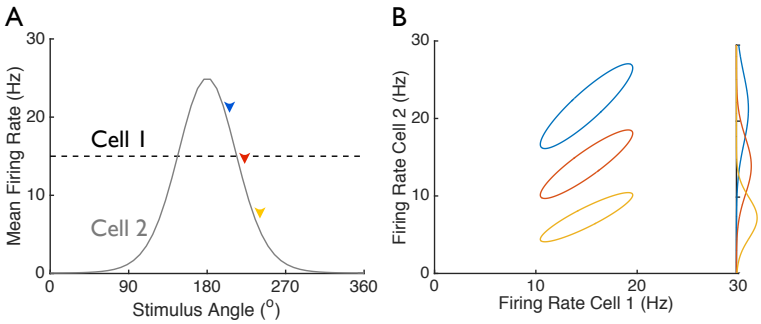

Figure 2: **Untuned neurons can shape noise, improving the population code.** Two neurons' tuning curves are shown in panel (A). Cell 1 is untuned. In response to stimulation, they give noisy responses. That noise is correlated between the two neurons. The distribution of noisy responses to each stimulus is described by an ellipse in the space of the two neurons' firing rates (B). The stimulus values are indicated by arrows in panel (A). The ellipses are well separated, meaning that the stimuli can be readily discriminated based on the two cells' firing rates. If the untuned cell is ignored, then only the tuned cell is observed. The distribution of the tuned cell's firing rate in response to each stimulus is shown along the right vertical of panel (B). Because those distributions overlap substantially, the stimulus cannot be readily discriminated based only on the firing rate of the tuned cell.

When do untuned neurons improve the population code?

Because they have no stimulus tuning, untuned neurons do not, in isolation, convey information about the stimulus. Instead, their contributions to the population code are indirect, and stem from the fact that their activities reflect the noise in the activities of the tuned cells. Accordingly, the untuned neurons' activities give information that can correct for the noise, thereby making the population code more informative. Consequently, untuned neurons contribute to the population code when their trial-specific activities are correlated with those of the tuned neurons.

To demonstrate this point, I considered neural populations with two different correlation structures. As in Fig. 1, I computed the stimulus information available in either the full population (black curves), the tuned subset of the population (red curves), or a random subset the same size as the tuned subset, but including both tuned and untuned neurons (blue curves). In the first of these populations, the untuned neurons are independent of the tuned ones, and the untuned neurons do not contribute to the population code (Fig. 3A: red and black curves overlap in the center panel). In the second population, the untuned neurons were correlated with the tuned ones, and the untuned neurons contribute substantially to the population code (Fig. 3B: red curve is lower than black curve in the center panel). These findings emphasize that untuned neurons enhance population codes when they are correlated with the tuned cells.

Moreover, the random subsets of 70% of the neurons in Fig. 3B contain more information than do subsets of the same size, but containing only tuned neurons (blue curve is above red curve). In other words, populations containing untuned neurons can have more stimulus information than do populations of the same size, but containing only tuned cells.

To generate these examples, I used the same random tuning curve shapes from Fig. 1A, but different covariance matrices. (In both cases, the neurons had Poisson-like variability, as is seen experimentally). In Fig. 3A, the neurons had the *differential correlations* studied by [17, 21] (see Methods for details). These correlations are such that the shared (correlated) part of the population noise mimics the changes in neural firing pattern induced by changes in the stimulus, thereby causing the distributions of responses to different stimuli to overlap substantially (Fig. 3A, lower). As a a result of that overlap in the stimulus-evoked response distributions, the noise substantially hinders the population code. Because changes in the stimulus do not change the mean firing rates of the untuned neurons – and the differential correlations mean the correlated noise mimics the stimulus-evoked changes in firing rates – the untuned neurons are unaffected by the correlated noise. This means that the noise in the untuned neurons is independent from the noise in the tuned ones. Consequently, ignoring the untuned neurons causes no loss of information (Fig. 3A: full population and tuned subset have the same information values).

To make Fig. 3B, I modified the differential correlation structure such that the untuned neurons were correlated with the tuned ones (see Methods). For the tuned subset of the population, the noise structure was identical to Fig. 3A (both contain differential correlations), and thus the tuned subsets of neurons have the same information in both cases. (Red data points in Fig. 3A have the same values as do points on the red curve in Fig. 3B). Different from Fig. 3A, the full population contained much more information than did the tuned subset, indicating that the untuned neurons do contribute to

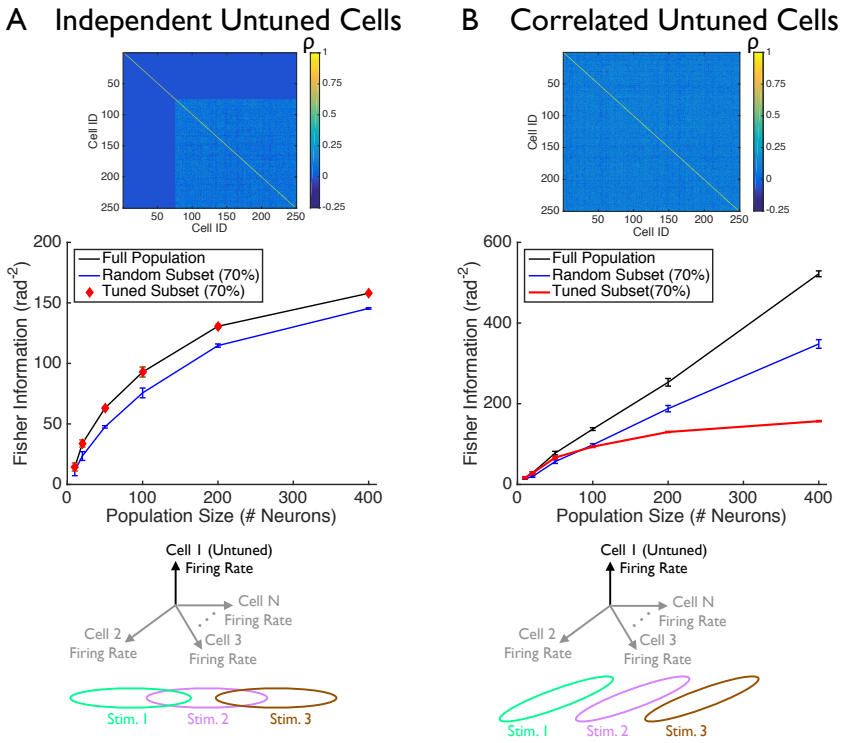

Figure 3: **Untuned neurons improve population coding when they are correlated with the tuned neurons.** I considered neural populations with tuning curves as in Fig. 1, and where the untuned neurons were either independent of the tuned ones (A), or where the untuned neurons were correlated with the tuned ones (B). 70% of the neurons in each population were tuned to the stimulus, and 30% were untuned. Upper panels show correlation matrices for 250-cell populations: cells 1 through 75 are untuned, whereas the remainder were tuned. Center panels show the Fisher information for the full neural populations (black), for the tuned subsets of neurons (red), and for random subsets of 70% of the neurons in each population (blue). (Data points shown are mean ± S.E.M., computed over 5 different random draws of the tuning curves). The cartoons in the lower panels illustrate why these two different correlation structures lead to untuned neurons having such different effects on the population code (see text). The cartoons show the space of neural firing patterns: each axis is the firing rate of a different neuron. The vertical axis is the firing rate of an untuned neuron. The other axes are the firing rates of tuned cells. Ellipses represent the 1 standard deviation probability contours of the neural population responses to the 3 different stimuli.

the population code in this case.

This contribution of untuned neurons to the population code can be understood via the cartoon in Fig. 3B (lower), which shows the distribution of population responses to 3 different stimuli. In the cartoon, cell 1 is untuned, whereas the rest of the cells are tuned. This means that, as the stimulus changes, the mean responses change along the plane orthogonal to the cell 1 axis. Because the untuned

neuron is correlated with the tuned ones, the noise distributions are tilted along the vertical axis. In this configuration, the distributions do not overlap very much. If, however, the untuned neuron is made independent from the tuned ones (as in Fig. 3A), the vertical tilt goes away, causing much more overlap in the response distributions. In other words, when the untuned neurons are correlated with the tuned ones, they improve the population code by separating the responses to different stimuli. This effect disappears when the untuned neurons are independent of the tuned ones.

Decoding *in vivo* neural activities with untuned neurons either included or excluded from the analysis

The theoretical work in the preceding Section makes one key prediction: the ability to decode a stimulus from the evoked neural population activities could be improved if untuned neurons are included in those populations, as opposed to being ignored. To test that prediction, I analyzed data from 2-photon Ca^{2+} imaging recordings done in primary visual cortex of awake mice (data from [10]) whose neurons expressed the genetically encoded calcium indicator GCaMP6f. The mice were presented with stimuli consisting of gratings drifting in 8 different directions, and the fluorescence levels of $\mathcal{O}(100)$ neurons were observed in each experiment. I analyzed the data from 46 such experiments.

For each stimulus presentation and neuron, I extracted the mean fluorescence change during the stimulus presentation, relative to the fluorescence in the period before the stimulus presentation: this $\Delta F/F$ value measures the stimulus-induced change in neural activity. I then computed the neurons' tuning curves by averaging these $\Delta F/F$ values over all trials in which the stimulus drifted in each direction. Some of the neurons had well-defined direction tuning curves (Fig. 4A), whereas others were relatively untuned (Fig. 4B). Following [10], I categorized these cells as tuned or untuned based on their direction selectivity indices (see Methods). Between the 46 experiments, $2973/4495 \approx 66\%$ of the neurons were tuned.

Along with the tuning, I measured the correlations in the cells' trial-to-trial variability over repeats of each stimulus. These "noise correlations" are shown for one of the experiments in Figs. 4C and D. The correlation coefficients were similar between pairs of tuned neurons ("TT"), pairs of untuned neurons ("UU"), and mixed pairs consisting of one tuned and one untuned neuron ("TU"). Because there were correlations between the tuned and untuned neurons, the theory predicts that stimulus decoding could be improved by including the untuned neurons, as opposed to ignoring them.

To test this prediction, I used the k-nearest neighbors (KNN) decoder to estimate the stimulus direction corresponding to the population activity pattern observed on each trial (Fig. 5A). KNN implements a simplified form of maximum likelihood estimation, works with even modest amounts of data, and has previously been used to study neural population coding [9, 25]. To estimate the stimulus corresponding to a given activity pattern (like the question mark in Fig. 5A), the classifier identifies the k most similar activity patterns in the dataset (similarity measured by Euclidean distance between data points; $k = 5$ in Fig. 5A). The classifier then takes a majority vote over the stimulus directions associated with those activity patterns, to estimate the stimulus that is responsible for the test point. (Note that, for decoding each data point, the KNN decoder is constructed from *all other* data points.

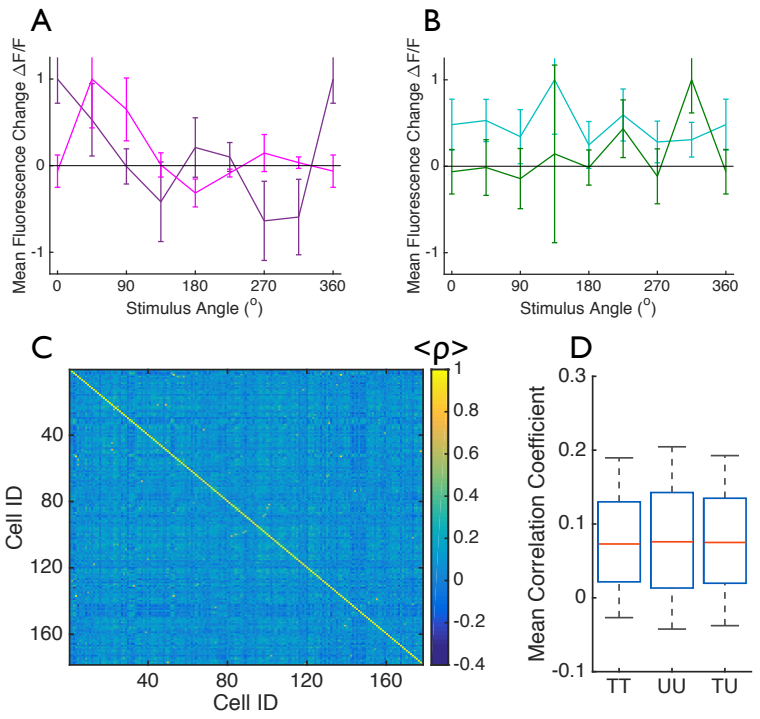

Figure 4: **Tuned and untuned neurons are correlated *in vivo*.** Neurons' responses to drifting grating stimuli were measured using 2-photon Ca^{2+} imaging. Example tuning curves for two direction tuned neurons are shown in panel (A), and for two untuned neurons in panel (B). Markers show mean $\Delta F/F \pm$ S.E.M, calculated over 75 trials of each stimulus direction. I measured the correlations in the responses of the different cells over repeats of each stimulus. These are shown in panel (C): correlation coefficients were averaged over all stimuli. The distributions of these mean correlation coefficients are shown in panel (D) for cell pairs of different types: where both cells were direction tuned ("TT"; $n = 7875$ pairs); where both cells were untuned ("UU"; $n = 1326$ pairs); and where one cell was tuned and one was untuned ("TU"; $n = 6552$ pairs). Each box plot shows the median, the range (maximum and minimum indicated by black bars), and the boundaries of the 25^{th} and 75^{th} percentiles (blue box) of the distributions.

This means that the test point is held-out from the decoder's construction. This is important, because otherwise the test point would be used to decode itself, which could yield erroneously high performance values.)

For one of the experiments (the same one shown in Fig. 4), I performed the KNN decoding for different values of k, and computed the fraction of trials on which the stimulus direction was correctly identified. This performance measure did not depend strongly on k (Fig. 5B), so I chose $k = 10$ for the subsequent analyses. Next, I performed the KNN decoding on the neural populations from each of the 46 different recordings. I separately performed the decoding on the full populations (including both tuned and untuned neurons), or on the subsets of tuned neurons in each recording. In most of the

experiments, the stimulus could be decoded substantially better by including the untuned neurons, as opposed to ignoring them (Fig. 5C: $p = 5 \times 10^{-11}$; one-sided paired sample t-test; $t = -8.4$ with 45 degrees of freedom). On average, decoding performance was $36 \pm 6\%$ (mean $\pm$ S.E.M.) better using the full populations vs just the tuned subsets. These results are consistent with the theoretical work presented above, and indicate that untuned neurons can contribute to sensory information coding, and that in some cases, their contributions can be sizable.

Next, I asked whether – as in the theoretical calculations – populations that include both tuned and untuned neurons could yield better decoding vs populations of the same size but containing only tuned cells. For each population, I extracted a random subset of the neurons that was the same size as the set of tuned neurons. I then performed the KNN decoding on these random subsets, and compared the performance with that which was obtained on only the tuned subsets (Fig. 5D). On average, the decoding performance was $21 \pm 5\%$ (mean $\pm$ S.E.M.) better using the random subsets vs the fully tuned ones, a statistically significant difference ($p = 2 \times 10^{-6}$, single-sided paired sample t-test; $t = -5.3$ with 45 degrees of freedom).

The findings on the population imaging experiments validate the theoretical results from Figs. 1,3. Namely, they show that untuned neurons can enhance neural population coding (Fig. 5C), and that mixed populations of tuned and untuned neurons can sometimes yield better information coding than can populations of the same size but containing only tuned neurons (Fig. 5D).

Discussion

I showed that, when the variability in neural responses to stimulation is correlated between cells, untuned neurons can contribute to sensory information coding. This effect was observed in both a theoretical model (Figs. 1-3), and in large population recordings from mouse visual cortex (Fig. 5). Moreover, in at least some cases (Figs. 1D, S1, and 3B, 5D), populations with both tuned and untuned neurons can convey more information about the stimulus than do populations of the same size but containing only tuned neurons.

These results have three main implications. First, our understanding of how the sensory systems encode information about the outside world is likely to be incomplete unless it includes the contributions of untuned neurons. This means that current practices, in which untuned neurons are ignored during data collection and analysis, might be hindering progress.

Second, because adding untuned neurons can increase the stimulus information, there might be a functional reason why these neurons exist in the sensory areas. (At the same time, it is always possible that neurons that appear untuned are in fact tuned, but only to aspects of the stimulus that were not varied in the experiment.) It is important to note, however, that no brain area can encode more stimulus information than it received from its inputs [40, 21]. This is the *data-processing inequality*, and it implies that there is not a limitless increase in information to be obtained by adding large numbers of untuned neurons to neural circuits.

Third, because information in a neural circuit can be lost when untuned neurons are ignored, it may be possible to improve the decoding of brain signals to control prosthetic limbs via brain-

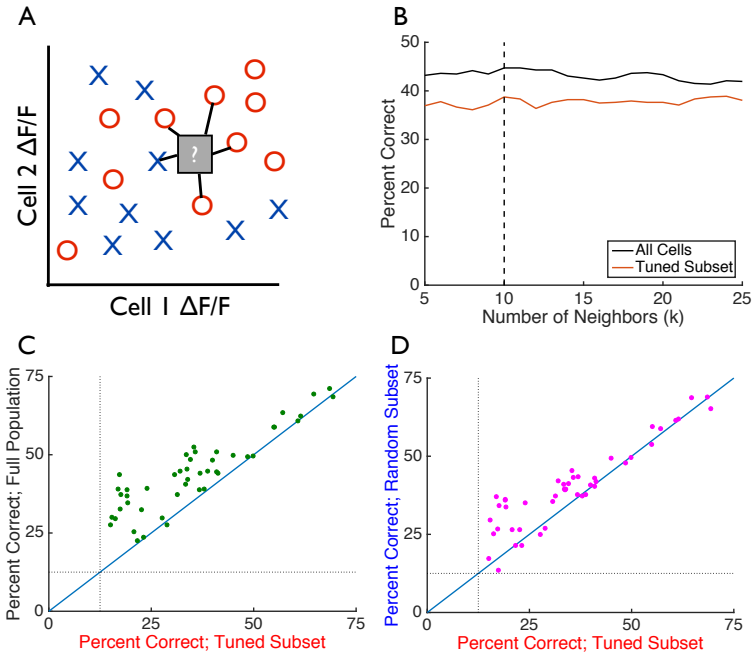

Figure 5: **K-nearest neighbor decoding reveals that untuned neurons can enhance information coding in *in vivo* neural populations.** The k-nearest neighbors decoder is illustrated in panel (A), which shows a cartoon of 2 neurons' activities, as measured by the $\Delta F/F$. The activities observed on each trial are shown, with the symbol type indicating which stimulus was presented. To decode a data point – like the one indicated by the question mark – the k-nearest data points are identified ($k = 5$ in this cartoon). A majority vote is taken over those data points' stimulus values to classify the test point. In the cartoon, the decoder would determine that the question mark data point was caused by the red circle stimulus. For different values of k, I applied this decoder to the population activities from one of the experiments. The percentage of trials in which the stimulus was correctly identified by the decoder is shown in panel (B), for decoding either the full population of 178 neurons (black curve), or for decoding only the 119 neurons with strong direction tuning (red curve). For $k = 10$, I applied the decoder to 46 different population recordings. Panel (C) shows the fraction of trials on which the stimulus was correctly identified using either the full population recordings, or the tuned subsets. Panel (D) shows the decoding performance when decoding either the tuned subsets of the populations, or random subsets the same size as the tuned subsets, but containing both tuned and untuned neurons. Chance performance for these decoding tasks is $1/8 = 12.5\%$. Diagonal lines in panels C and D denote equality.

machine interface devices, by including the activities of untuned neurons in those signals. This effect is demonstrated most strongly in Fig. 5C.

Observations related to those presented here have also been made by Insanally and colleagues (Insanally, Carcea, Albanna, and Froemke, Cosyne 2017 abstract), based on analyses of data from

awake behaving rats. There, as in the analysis of mouse data presented here, it is hard to distinguish weakly tuned neurons from purely untuned ones, and thus difficult to be certain that the coding benefits of putatively untuned neurons do not arise from non-zero tuning, that is nonetheless under the chosen threshold. This complication highlights the value of the theoretical work presented here (Figs. 1,3): in the model, the untuned neurons really have no stimulus dependence, enabling us to be pinpoint the role of untuned neurons in sensory information coding.

For large neural populations, an astronomically large number of different correlation patterns are possible (and this problem is confounded when one includes correlations of higher order than the pairwise ones considered here [41, 42]). Accordingly, it was not possible to simulate all possible correlation patterns in the theoretical study. Thus, it is natural to ask how general the results are over different correlation structures. Here, the fact that I made similar observations with two very different correlation structures is encouraging (Figs. 1 and 3). Moreover, I saw a qualitatively similar effect in the experimental data (Fig. 5) as in the theoretical model with limited range correlations (Fig. 1), which further argues for the generality and applicability of the findings.

Adding neurons to a population can never decrease the amount of encoded stimulus information: because a downstream read-out could always choose to ignore the added cells, those cells can at worst contribute zero information. Consequently, untuned neurons can never *hinder* the population code. (However, decoding based on observations with small numbers of trials is subject to overfitting. In this case, adding more cells can hinder the decoding because the decoder might be inaccurate). This means that the potential effects of untuned neurons on population coding range between no contribution (Fig. 3A), and positive contributions at least as large as those seen in Figs. 3B, 1D and S2. (I.e., at least 200% increase in information available by including vs. ignoring untuned neurons). There may be other cases, not explored here, in which the positive contributions of untuned neurons are even larger.

It is important not to interpret the results presented here as implying that neural tuning is not essential to sensory information coding. If there are no tuned neurons, there is no information in the neural population (Fig. S1: information approaches zero as the fraction of untuned neurons approaches 1). However, if there are some tuned neurons, then the untuned neurons can serve to make them more informative, thereby improving the population code overall. Thus, untuned neurons are not irrelevant for sensory information coding.

Methods

I first discuss the theoretical calculations, and then the analysis of experimental data.

Theoretical Calculations

Modeling the stimulus-evoked neural responses, and the information encoded

I considered for simplicity a 1-dimensional stimulus s (for example, the direction of motion of a drifting grating). In response to the stimulus presentation, the neural population displays firing rates $\vec{r}_i$, where the index i denotes the trial. (Each element of the vector $\vec{r}_i$ is the firing rate of a single neuron). These

responses have two components. The first, $\vec{f}(s)$, is the mean (trial-averaged) response to stimulus s, whereas the second component, $\vec{\epsilon}_i$, represents the trial-by-trial fluctuations, or "noise" in the neural firing rates.

$$\vec{r}_i = \vec{f}(s) + \vec{\epsilon}_i \tag{1}$$

The tuning curves were chosen to be either Von Mises functions (as in [13, 15, 21]), or, in the case of untuned neurons, to be constants (Fig. 1A). The parameters of the tuning curves were randomly drawn, using the same distributions as in [21].

The neurons' noise variances were chosen to match the mean responses, in accordance with experimental observations of Poisson-like variability. I considered different patterns of inter-neural correlation, as described below.

For each set of tuning curves and correlations, I used the typical linear Fisher information measure, $I(s)$, to quantify the ability of downstream circuits to determine the stimulus, s, from the noisy neural responses on each trial $\vec{r}_i$ [11, 12, 13, 14, 15, 16, 18, 20, 17, 19, 21]:

$$I(s) = \vec{f}'^{T}(s)\,[C(s)]^{-1}\,\vec{f}'(s), \tag{2}$$

where the prime denotes a derivative with respect to the stimulus, the supserscript T denotes the transpose operation, and $C(s) = \mathrm{cov}\,(\vec{\epsilon}_i|\,s)$ is the covariance matrix of the noise in the neural responses to stimulus s.

To compute the information for a subset of a neural population, I extracted the block of the covariance matrix, and the elements of the vector $\vec{f}'(s)$, that correspond to the neurons in that subset. I then used those values in Eq. 2.

For all of the information values presented here, I computed the information for each of 50 different stimulus values, evenly spaced over $[0^o, 360^o]$. The reported values are averages over these 50 stimuli. This accounts for the fact that Fisher information $I(s)$ is a *local* quantity which varies from stimulus to stimulus. By averaging over many stimuli, I avoid the possibility that the reported information values might be atypical, and affected by the specific stimulus at which the information was calculated.

Limited-range correlations, and Fig. 1

The elements of covariance matrix $C(s)$ were $C_{ij}(s) = \sqrt{f_i(s) f_j(s)}\rho_{ij}$, where ρ_{ij} is the correlation between cells i and j. The factor of $\sqrt{f_i(s) f_j(s)}$ ensures that the neurons have Poisson variability (variance of noise is equal to mean firing rate, meaning that standard deviation of noise is equal to square root of mean firing rate).

The correlation coefficients ρ_{ij} were calculated from the equation in Fig. 1B. The tuning curve separation $\Delta(\phi)$ for each cell pair was computed as $\Delta(\phi) = |\arccos[\cos(\phi_i - \phi_j)]|$, where ϕ_i and ϕ_j are the cells' preferred direction angles (the locations of their tuning curve peaks). This formula accounts for the fact that angles "wrap" around the circle: so values of 10^o and 350^o have a separation of 20^o (and not 340^o).

For the untuned neurons, their preferred stimulus angles were randomly assigned, uniformly over the range $[0^o, 360^o]$.

Differential correlations, "modified" differential correlations, and Fig. 3

For Fig. 3A (with differential correlations), the covariance matrix was given by

$$C_A(s) = C_o + \epsilon \vec{f}'(s) \vec{f}'^T(s), \tag{3}$$

where ϵ is a (small) scalar parameter that sets the strength of the differential correlations, and C_o is a diagonal matrix with entries equal to the mean firing rates given by $f(s)$. For $\epsilon = 0$, this covariance matrix describes independent neurons with Poisson variability. For the results in Fig. 3A, I chose $\epsilon = 5 \times 10^{-3}$, corresponding to weak but non-zero differential correlations.

Because untuned neurons have zero slope to their tuning curves – meaning the corresponding elements in $\vec{f}'(s)$ are zeros – the differential correlations formula (Eq. 3) ensures that the untuned neurons are independent from the tuned ones (as in Fig. 3A).

For Fig. 3B, I modified Eq. 3 in a fashion that kept the tuned subset of the population unchanged, but made the untuned subset correlated with the tuned one. To do this, I used the formula

$$C_B(s) = C_o + \epsilon \vec{g}(s) \vec{g}^T(s), \tag{4}$$

where $g_i(s) = f_i'(s)$ if the neuron is tuned, and $g_i(s) \sim \mathcal{N}(0, 5)$ if the neuron is untuned. If only the subset of tuned neurons is considered, $C_A = C_B$. For the untuned neurons, the corresponding rows and columns of C_A are all zeros, whereas for C_B, they are randomly generated non-zero values. Fig 3B used the same value of ϵ as Fig. 3A.

Analysis of *in vivo* neural recordings

Overview of the experiment

The full description of the experiment is given by [10], and so I briefly summarize here. GCaMP6f was expressed in the excitatory neurons of the forebrain of mice. 2-photon imaging was used to measure the fluorescence of neurons in visual cortex through a cranial window. The mice were presented with drifting grating stimuli. The stimuli could move in any of 8 different directions, and at 5 different temporal frequencies. The stimuli were presented for 2 seconds each, followed by a 1 second gray screen before the next stimulus was presented. Each combination of direction and frequency was presented repeatedly (either 15 or 30 times each, depending on the temporal frequency).

Data access and analysis

Following the example Jupyter notebook provided by [10] – which provides a template for accessing the experimental data – I retrieved the following data: average $\Delta F/F$ values for each neuron on each trial, the stimulus direction for each trial, the neurons' direction selectivity indices, and indicators of whether or not each neuron was visually responsive. Only visually responsive neurons were analyzed, and data were retrieved for 46 different experiments.

I calculated the tuning curves (Figs. 4A and B) by averaging the $\Delta F/F$ values for all trials of each direction: this *marginalizes* over the different temporal frequencies. The noise correlations coefficients

(Figs. 4C and D) were computed over repeats of the same stimulus (same orientation and temporal frequency), and then averaged over all stimuli.

For the decoding analyses, I used the k-nearest neighbors method (Fig. 5A) on the population activity vectors observed on each trial. These vectors had as elements the $\Delta F/F$ values for all of the visually responsive neurons (black curves in Figs. 5B, and vertical axis of Fig. 5C), for just the direction selective cells (red curve in Figs. 5B, and horizontal axes of Figs. 5 C and D), or for random subsets of visually responsive neurons that were the same size as the groups of direction selective cells (Fig. 5D, vertical axis).

For each analysis, I iteratively considered each single-trial activity vector as a "test" data point, and identified the k most similar other data points (smallest Euclidean distance) to the test point. I then took a majority vote over the stimulus directions of these k neighboring points, to guess the most likely stimulus direction for the test point. This was repeated for each test point, and I measured the neural coding performance as the percentage of trials on which the estimated stimulus direction matched the stimulus direction associated with the test point. (Note that, for decoding each data point, the KNN decoder is constructed from *all other* data points. This means that the test point is held-out from the decoder's construction. This is important, because otherwise the test point would be used to decode itself, which could yield erroneously high performance values.)

Following the example analysis of [10], direction selective cells were defined as those having a direction selectivity index (DSI) greater than 0.5. DSI values were obtained from [10], and computed as follows. Each cell's preferred direction was determined (direction yielding maximum neural response). The DSI was then defined as

$$DSI = \frac{R_{pref} - R_{null}}{R_{pref} + R_{null}}, \tag{5}$$

where R_{pref} is the mean response to the preferred direction, and R_{null} is the mean response to the "null" direction, which is the direction 180^o opposite to the preferred direction.

Statistical tests

After performing the KNN decoding separately on the full recorded populations, the tuned subsets of neurons, and random subsets of neurons, I obtained 3 coding performance values for each of the 46 experiments. I first tested the distributions of these coding performance values for Gaussianity, using the Kolmogorov-Smirnov (KS) test. All 3 distributions had KS statistics – which measure their deviation from a Gaussian distribution with mean equal to the sample mean, and standard deviation equal to the sample standard deviation – of less than 1, and p values of greater than 0.05, indicating that they do not significantly differ from the Gaussian. Accordingly, it was appropriate to test the significance of the diferences in coding performance using paired 1-sided t-tests.

These comparisons were done the coding performance using only tuned neurons and: a) the coding performance using all neurons; or b) the coding performance using random subsets of the neural population that were the same size as the tuned subset.

Code and data availability

All experimental data used here was accessed from the Allen Institute's database, and are freely available there [10]. The custom Jupyter and Matlab scripts used for the theory study, and analysis of experimental data, are available from the author upon request.

Acknowledgments

Thanks to Eric Shea-Brown and Alex Cayco-Gajic for helpful discussions, and to the Allen Institute for sharing their *in vivo* datasets. JZ gratefully acknowledges the following funding: Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholar Award, Sloan Research Fellowship, and Google Faculty Research Award.

References

[1] Lebedev MA, Nicolelis MA (2006) Brain–machine interfaces: past, present and future. Trends Neurosci 29: 536–546.

[2] Santhanam G, Ryu SI, Byron MY, Afshar A, Shenoy KV (2006) A high-performance brain–computer interface. Nature 442: 195–198.

[3] Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, et al. (2014) Neural constraints on learning. Nature 512: 423–426.

[4] Bensmaia SJ (2015) Biological and bionic hands: natural neural coding and artificial perception. Phil Trans R Soc B 370: 20140209.

[5] Delhaye BP, Saal HP, Bensmaia SJ (2016) Key considerations in designing a somatosensory neuroprosthesis. J Physiol - Paris (in press) .

[6] Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque v1: diversity and laminar dependence. J Neurosci 22: 5639–5651.

[7] Olshausen BA, Field D (2006) What is the other 85 percent of v1 doing? In: van Hemmen L, Sejnowski T, editors, Problems in Systems Neuroscience, Oxford Press. pp. 182–211.

[8] Graf A, Kohn A, Jazayeri M, Movshon J (2011) Decoding the activity of neuronal populations in macaque primary visual cortex. Nat Neurosci 14: 239–245.

[9] Zylberberg J, Cafaro J, Turner M, Shea-Brown E, Rieke F (2016) Direction-selective circuits shape noise to ensure a precise population code. Neuron (89): 369-383.

[10] Allen Brain Observatory (2016). `http://observatory.brain-map.org/visualcoding/`. Accessed: 2017-04-17.

[11] Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. Neural Comput 11: 91–101.

[12] Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci 7: 358–366.

[13] Ecker AS, Berens P, Tolias AS, Bethge M (2011) The Effect of Noise Correlations in Populations of Diversely Tuned Neurons. J Neurosci 31: 14272–14283.

[14] da Silveira RA, Berry MJ (2014) High-Fidelity Coding with Correlated Neurons. PLoS Comput Biol 10: e1003970.

[15] Hu Y, Zylberberg J, Shea-Brown E (2014) The sign rule and beyond: Boundary effects, flexibility, and noise correlations in neural population codes. PLoS Comput Biol 10: e1003469.

[16] Shamir M (2014) Emerging principles of population coding: in search for the neural code. Curr Opin Neurobiol 25: 140-148.

[17] Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, et al. (2014) Information-limiting correlations. Nat Neurosci 17: 1410-1417.

[18] Sompolinsky H, Yoon H, Kang K, Shamir M (2001) Population coding in neuronal systems with correlated noise. Phys Rev E 64: 051904.

[19] Averbeck BB, Lee D (2006) Effects of noise correlations on information encoding and decoding. J Neurophys 95: 3633–3644.

[20] Josić K, Shea-Brown E, Doiron B, de la Rocha J (2009) Stimulus-dependent correlations and population codes. Neural Comput 21: 2774–2804.

[21] Zylberberg J, Pouget A, Latham P, Shea-Brown E (2017) Robust information propagation through noisy neural circuits. PLoS Comput Biol 13: e1005497.

[22] Britten K, Shadlen M, Newsome W, Movshon J (1993) Responses of neurons in macaque MT to stochastic motion signals. Visual Neurosci 10: 1157-1169.

[23] Churchland M, Yu B, Cunningham J, Sugrue L, Cohen M, et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci 13: 369-378.

[24] Franke F, Fiscella M, Sevelev M, Roska B, Hierlemann A, et al. (2016) Structure of neural correlation and how they favor coding. Neuron 89: 409-422.

[25] Zylberberg J, Hyde R, Strowbridge B (2016) Dynamics of robust pattern separability in the hippocampal dentate gyrus. Hippocampus (29): 623-632.

[26] Faisal A, Selen L, Wolpert D (2008) Noise in the nervous system. Nat Rev Neurosci 9: 292-303.

[27] Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 370: 140–143.

[28] Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. Nat Neurosci 14: 811–819.

[29] Lampl I, Reichova I, Ferster D (1999) Synchronous membrane potential fluctuations in neurons of the cat visual cortex. Neuron 22: 361-374.

[30] Alonso J, Usrey W, Reid R (1996) Precisely correlated firing of cells in the lateral geniculate nucleus. Nature 383: 815–819.

[31] Goris R, Movshon J, Simoncelli E (2014) Partitioning neuronal variability. Nat Neurosci 17: 858-865.

[32] Smith M, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. J Neurosci 28: 12591-12603.

[33] Ecker A, et al (2014) State dependence of noise correlations in macaque primary visual cortex. Neuron 82: 235-248.

[34] Scholvinck M, Saleem A, Benucci A, Harris K, Carandini M (2015) Cortical state determines global variability and correlations in visual cortex. J Neurosci 35: 170-178.

[35] Lin IC, Okun M, Carandini M, Harris K (2015) The nature of shared cortical variability. Neuron 87: 644-656.

[36] Montijn JS, Meijer GT, Lansink CS, Pennartz CM (2016) Population-level neural codes are robust to single-neuron variability from a multidimensional coding perspective. Cell Rep 16: 2486–2498.

[37] Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. J Neurosci 21: 1676–1697.

[38] Reich DS, Mechler F, Victor JD (2001) Independent and redundant information in nearby cortical neurons. Science 294: 2566–2568.

[39] Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? J Neurosci 13: 2758–2771.

[40] Beck JM, Ma W, Pitkow X, Latham PE, Pouget A (2012) Not noisy, just wrong: the role of suboptimal inference in behavioral variability. Neuron 74: 30–39.

[41] Zylberberg J, Shea-Brown E (2015) Input nonlinearities can shape beyond-pairwise correlations and improve information transmission by neural populations. Phys Rev E 92: 062707.

[42] Cayco-Gajic A, Zylberberg J, Shea-Brown E (2015) Triplet correlations among similarly-tuned cells impact population coding. Front Comput Neurosci 9: 57.

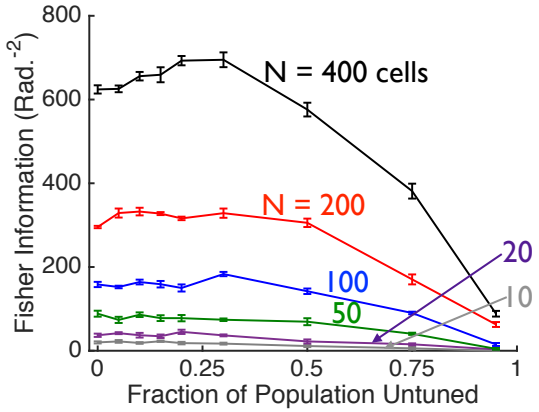

Supplementary Figure 1: **Populations with untuned neurons can encode more information than ones with only tuned neurons – dependence on tuning fraction.** (Related to Fig. 1.) I repeated the calculations from Fig. 1, but with different fractions of neurons left untuned in each population. For each such population, I computed the Fisher information. Error bars are the S.E.M. over 5 random sets of different tuning curves.

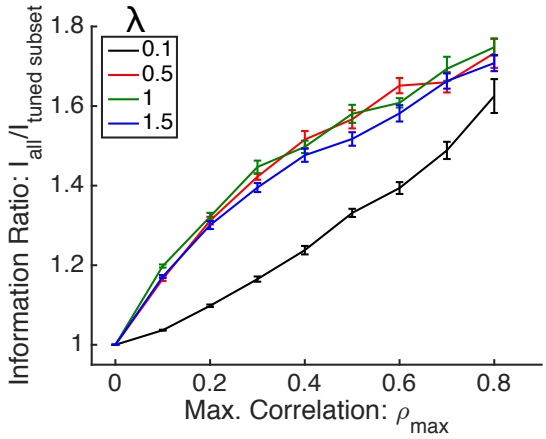

Supplementary Figure 2: **Dependence of information on limited range correlation parameters.** (Related to Fig. 1.) I repeated the calculations from Fig. 1, in all cases for populations of 200 neurons. I repeated the calculations for different values of ρ_{max} and λ, the parameters that define the limited-range correlations. For each set of parameters, I computed the ratio of Fisher information in the full population of 200 neurons, vs. the Fisher information in just the tuned subset of (70% of) the population. Error bars are the S.E.M. over 10 random sets of different tuning curves.