bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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

June 29, 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

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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]).

Recently, researchers have begun to question that assumption: analyses of neural data in the prefrontal cortex [10] and auditory cortex (Insanally et al., 2017 cosyne abstract) show that even neurons with no obvious stimulus tuning can nevertheless contribute to the population code. These findings are intriguing, because they suggest that – by virtue of our ignoring the untuned neurons – our understanding of neural population coding might be incomplete. At the same time, several deep questions remain unanswered: Are the impacts of the putatively untuned neurons on population coding due to weak tuning that is nevertheless below the threshold the experimenters set for calling neurons tuned (vs untuned)? And why are there untuned neurons in the brain? Do mixed populations of tuned and untuned neurons have a functional advantage over populations containing only tuned neurons?

To answer these questions, I used theoretical calculations, and then verified the predictions from those calculations by analyzing 2-photon imaging data collected in the visual cortices of mice that were shown drifting grating stimuli [11]. For the theoretical calculations, I used a common mathematical model of the neural population responses to sensory stimulation [12, 13, 14, 15, 16, 17, 18, 9, 19, 20, 21, 18, 22]. 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, 23, 24, 25, 26, 27]; and the potential for that noise to be correlated between neurons [28, 29, 9, 30, 31, 28, 32, 33, 34, 35, 36, 37]. 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.

Because the untuned neurons in the theoretical model really have no stimulus tuning, these calculations enabled me to demonstrate conclusively that strictly untuned neurons really can contribute to population coding: I provide a geometrical explanation for this phenomenon. Moreover, by studying the information coding of neural populations containing different fractions of tuned vs untuned neurons, I demonstrated that mixed populations can encode stimulus information better than populations containing only tuned neurons. This provides a functional explanation for why the brain contains untuned neurons, and explains when untuned neurons improve brain function. Finally, I used decoding

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

analyses applied to data collected in the visual cortices of awake mice to validate the key predictions of the theory: excluding putatively untuned neurons hinders decoding; this effect depends on the noise-shaping mechanism described by the theory; and decoding random groups of neurons (both tuned and untuned) yields better performance than does decoding neural populations of the same size, but containing only tuned neurons.

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 predictions from the theory by analyzing data collected in mouse visual cortex.

Theoretical analysis

A role for untuned neurons in 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 [37], and are comparable to the fraction of tuned neurons in the experimental data that I analyzed. I later consider populations with different fractions of untuned neurons.) 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 [24, 22]. The variability was correlated between cells, and the correlation coefficients were chosen to follow the "limited-range" structure reported experimentally [29, 17, 38, 39, 40], and used in previous theoretical studies [12, 13, 14, 19]. 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. 1C, blue 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. 1C, 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.

Because the correlation coefficients in Fig. 1 did not depend on the stimulus, it is not the case that the untuned neurons themselves encode information indirectly, through their second-order statistics

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

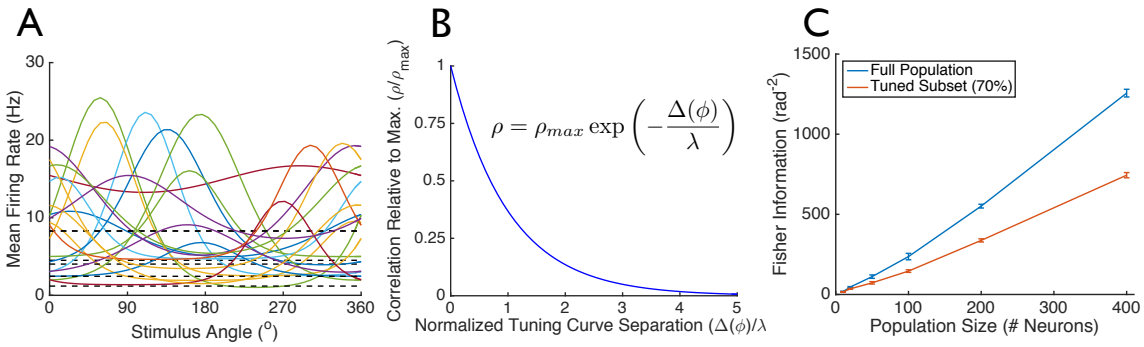

Figure 1: **Untuned neurons can play an important role in sensory information coding.** (A) I considered populations of neurons with randomly shaped and located tuning curves. Of those neurons, 70% were tuned to the stimulus, whereas 30% were untuned – their mean firing rates do not depend on the stimulus (dashed black lines in panel A). The neurons' trial to trial variability was Poisson-like and correlated between neurons. (B) These correlations followed the "limited-range" structure with $\rho_{max} = 0.75$ and $\lambda = 0.5$ radians (29^o). The mean correlation coefficients (averaged over neurons) were 0.12, which is comparable to values reported in primary visual cortex [29]. (Modifying these values did not qualitatively change the results – see Fig. S1). (C) For different sized populations, I computed the Fisher information, which quantifies how well the stimulus can be estimated from the neural population activities. The different lines correspond to: the Fisher information for the full neural populations (blue); and the Fisher information for the tuned 70% of the populations (red). Data points are mean ± S.E.M., computed over 5 different random draws of the tuning curves.

(as was the case in the theoretical model of [41]). This point is emphasized in Fig. 4, where the information in the population goes to zero as the fraction of untuned neurons approaches 100%. This suggests the question of how untuned neurons contribute to neural information coding. While the untuned neurons' activities do not reflect the stimulus, they do reflect the trial-specific noise in the tuned neurons' activities (because they are correlated). 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 relatively disjoint, meaning that the neural responses to the different stimuli have little overlap, and so it is relatively unambiguous to infer from the neural firing

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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.

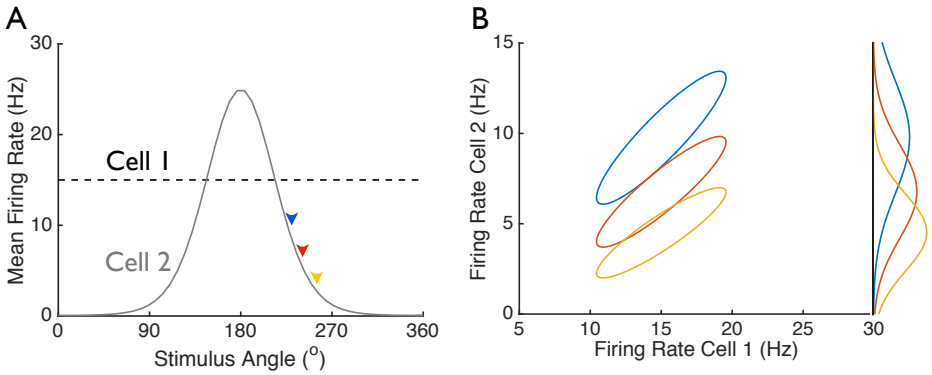

Figure 2: **Untuned neurons can shape noise, improving the population code.** (A) Two neurons' tuning curves are shown; cell 1 is untuned. In response to stimulation, the cells give noisy responses. That noise is correlated between the two neurons, with a correlation coefficient of 0.9. (B) The distribution of noisy responses to each stimulus is described by an ellipse in the space of the two neurons' firing rates. 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.

Because the untuned neurons' contribution to the population code relies on their activities reflecting the single-trial noise in the activities of the tuned cells, the untuned neurons do not contribute to population coding if they are independent of the tuned neurons. To demonstrate this point, I repeated the analysis from Fig. 1 (above), but made the untuned neurons uncorrelated from each other and from the tuned neurons. In that case, the untuned neurons do not contribute to the population code: the full population and the tuned subset both have the same amount of stimulus information (Fig. 3A).

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

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

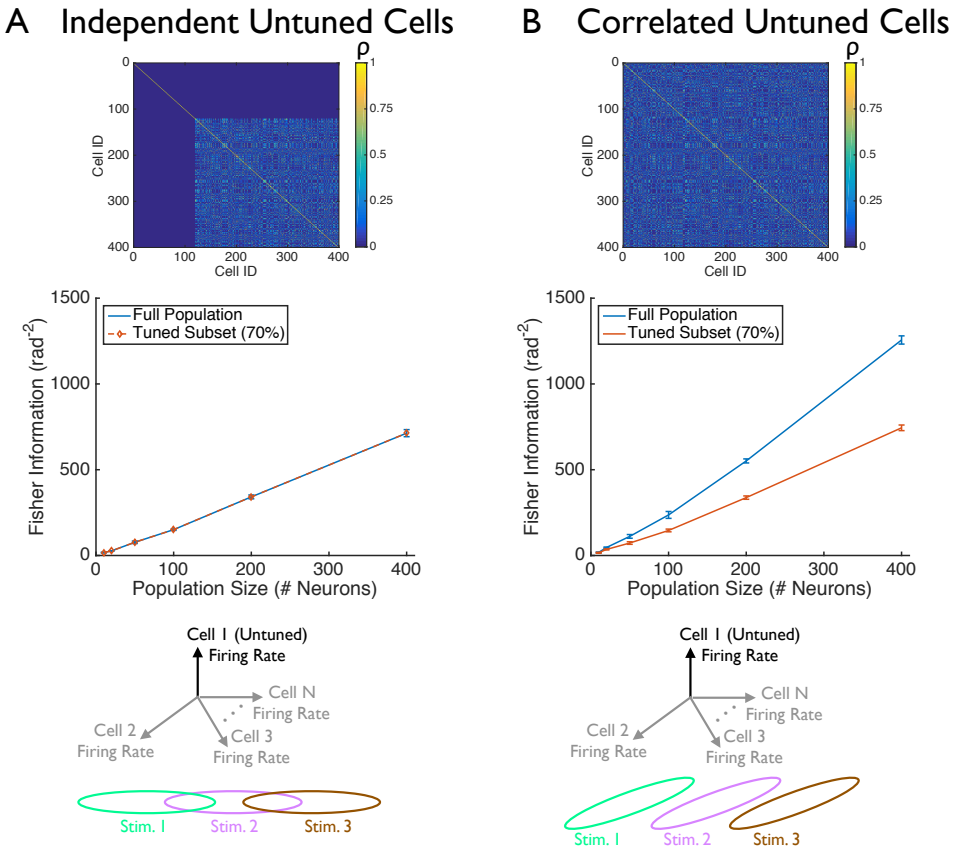

Figure 3: **Untuned neurons improve population coding when they are correlated with the tuned neurons.** (A) I repeated the analysis from Fig. 1C, and modified it so that the untuned neurons were independent of each other and of the tuned neurons. (B) For comparison, I also show again the results from Fig. 1C. As in Fig. 1, 70% of the neurons in each population were tuned to the stimulus, and 30% were untuned. Upper panels show correlation matrices for 400-cell populations: cells 1 through 120 are untuned, whereas the remainder were tuned. Center panels show the Fisher information for the full neural populations (blue), or for the tuned subsets of neurons in each population (red). (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.

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.

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

This effect disappears when the untuned neurons are independent of the tuned ones.

Mixed populations of tuned and untuned neurons can encode stimulus information more effectively than populations containing only tuned neurons

In the preceding analyses, I showed that neurons with no stimulus tuning can contribute to the population code: ignoring them entails a loss of stimulus information. Here, I turn to the question of why the brain contains those neurons at all. In other words, is there a functional benefit to including untuned neurons in a population vs having only tuned neurons?

To answer this question, I repeated the analysis from Fig. 1, but altered the fraction of untuned neurons in each population.The maximum information values were obtained with around 30% of neurons being untuned; this effect was larger in larger populations (Fig. 4A). Because the maximum information does not occur when all of the neurons are tuned (corresponding to an untuned neuron fraction of 0), this analysis shows that neural populations can be made more informative by replacing tuned neurons with untuned ones. This suggests that there may be a functional reason why the brain contains untuned neurons.

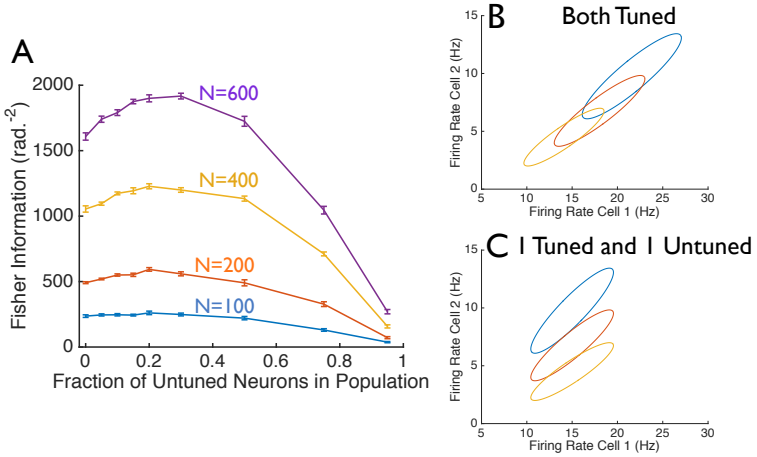

Figure 4: **Populations containing some untuned neurons can encode more information than ones with only tuned neurons.** (A) I repeated the calculations from Fig. 1, but with different fractions of untuned neurons in each population. For several different population sizes (indicated on the plot), the Fisher information is shown as a function of the fraction of untuned neurons in the population. Error bars are the S.E.M. over 5 random sets of different tuning curves. (B and C) In response to 3 different stimuli, I show the 1 standard-deviation probability contours in the responses of a pair of neurons. In all cases, the neurons have Poisson variability, and noise correlation coefficient of 0.9. In panel (B), both cells are tuned to the stimulus, whereas in panel (C), cell 2 is tuned to the stimulus, and cell 1 is untuned.

How (and when) are mixed populations of tuned and untuned neurons better at encoding information than populations of the same size but containing only tuned cells? The plots in Figs. 4B and

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

C provide some intuition. In both cases, the distributions of firing rates of two neurons are shown, in response to 3 different stimuli (similar to Fig. 2B): ellipses indicate 1 standard-deviation probability contours. In both panels, the neurons have Poisson variability in their firing rates, and the two cells are correlated (with noise correlation of 0.9 in all cases). In Fig. 4B, both cells are tuned to the stimulus, and the centers of the ellipses are correspondingly displaced relative to each other along both the vertical, and the horizontal, axes of the plot. With this geometrical configuration, the ellipses corresponding to different stimulus-evoked responses overlap substantially: that overlap means that there is ambiguity in determining the stimulus from the neural responses, and so the population code has relatively low information. Fig. 4C differs from Fig. 4B only in the tuning of cell 1: in Fig. 4C, cell 1 is untuned, whereas in Fig. 4B, it was tuned. This means that, in Fig. 4C (where only one of the cells is tuned to the stimulus), the different stimulus-evoked response distributions are displaced relative to each other only in the vertical direction, and not the horizontal one. Owing to the diagonal orientation of the ellipses, there is less overlap between the different response distributions in Fig. 4C than 4B. Consequently, the pair of neurons in Fig. 4C (one of which is untuned) is better at encoding stimulus information than the pair of neurons in Fig. 4B (both of which are tuned to the stimulus).

The examples in Figs. 4BC show how the presence of untuned neurons can improve the population code: including untuned neurons modifies the *signal correlation* structure (the correlation between neurons in the stimulus-evoked mean responses) relative to the case where both neurons are tuned. And because the relationship between the signal and noise correlations determines the population coding efficacy [16, 15, 13], this modification can improve the population code overall.

Analysis of *in vivo* neural activities

The theoretical work in the preceding Section makes a 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 this prediction, I analyzed data from 2-photon Ca^{2+} imaging recordings done in primary visual cortex of awake mice (data from [11]) 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. 5A), whereas others were relatively untuned (Fig. 5B). Following [42, 9], I categorized these cells as tuned or putatively untuned (hereafter referred to simply as *untuned*) based on their direction selectivity indices (see Methods). Between the 46 experiments, $5379/8943 \approx 60\%$ of the neurons were classified as being tuned for direction.

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 all pairs of simultaneously observed neurons

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

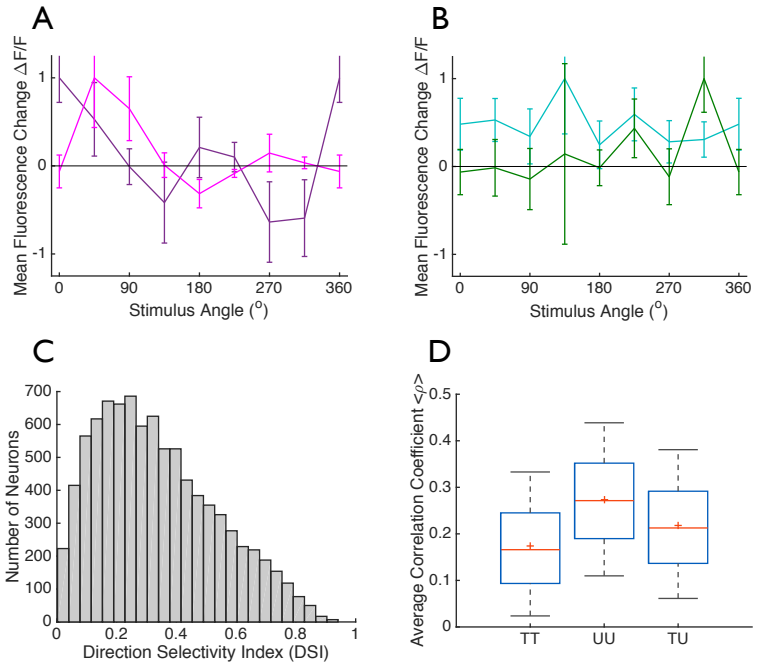

Figure 5: **Tuned and untuned neurons are correlated *in vivo*.** Neurons' responses to drifting grating stimuli were measured using 2-photon Ca^{2+} imaging. (A) Tuning curves for two direction tuned neurons. (B) Tuning curves of two untuned neurons. Markers show mean $\Delta F/F$ $\pm$ S.E.M, calculated over 75 trials of each stimulus direction. (C) Direction selectivity indices for the 8493 neurons whose stimulus-evoked responses were measured. (D) The distributions of correlation coefficients for cell pairs of different types: where both cells were direction tuned ("TT"; $n = 391833$ pairs); where both cells were untuned ("UU"; $n = 150801$ pairs); and where one cell was tuned and one was untuned ("TU"; $n = 434752$ 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.

(Fig. 5D). The correlation coefficients were typically positive for 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. 6A). KNN decoding works with even modest amounts of data, and has previously been used to study neural population coding [26]. To estimate the stimulus corresponding to a given activity pattern (like the question mark in Fig. 6A), the classifier identifies the k most similar activity patterns in the dataset (similarity measured by Euclidean distance between data points; $k = 5$ in Fig. 6A). The classifier then takes a majority vote over the stimulus directions associated with those activity patterns, to

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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. 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.) As a result of this procedure, KNN implements a simplified form of maximum likelihood estimation: to classify each test point, the decoder identifies the stimulus labels that occur with highest frequency in the region around the test point. Accordingly, it uses the empirical data distributions to sample the likelihood of response points similar to the test point, given each stimulus label, and classifies the test point with the stimulus that yields the highest likelihood.

For each experiment, 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. 6B), so I chose $k = 10$ for the subsequent analyses. Next, I separately performed the KNN 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. 6C: $p = 2.6 \times 10^{-7}$; one-sided paired sample t-test). On average, decoding performance was $24 \pm 5\%$ (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 (Figs. 1-3), and indicate that untuned neurons can contribute to sensory information coding, and that their contributions can be sizable.

Does the benefit of including putatively untuned neurons in the population code (vs excluding them) seen in Fig. 6C arise because of the noise-shaping effects described by the theory (Figs. 1-3), or because these putatively untuned neurons have weak stimulus tuning that, while below the chosen threshold, is nevertheless non-zero? To answer this question, I first noted that the noise-shaping mechanism described by the theory relies on *noise correlations*: the correlations in the neurons' trial-to-trial variability (Fig. 3). A trial-shuffling procedure that removes the noise correlations [9] abolishes this noise-shaping effect.

I thus repeated the analysis from Fig. 6C using trial-shuffled (uncorrelated) neural responses and I found that the full neural populations (including both tuned and untuned neurons) yielded $20 \pm 4\%$ (mean $\pm$ S.E.M over all 46 populations) better decoding performance than did the subsets of only tuned neurons. This shows that, even absent correlations, there is still more information in the full population than in the putatively tuned subset: this difference is due to the weak yet sub-threshold tuning discussed above. This weak tuning cannot, however, fully account for the results in Fig. 6C: with the correlated neural responses, the full populations yielded $24 \pm 5\%$ better decoding than did the tuned subsets. A comparison between the results with, and without, correlations (Fig. 6D), shows that the results in Fig. 6C arise from a combination of effects: first, even the putatively untuned neurons have some (weak) tuning, and second, the untuned neurons contribute to the population code by noise shaping, consistent with the theory in Figs. 1-3. This second effect is statistically significant: a paired one-sided t-test between the effect sizes in the raw data vs the trial-shuffled (uncorrelated) data yielded a value of $p = 0.020$.

It is also important to check that the results in Fig. 6C do not depend on the specific criterion

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

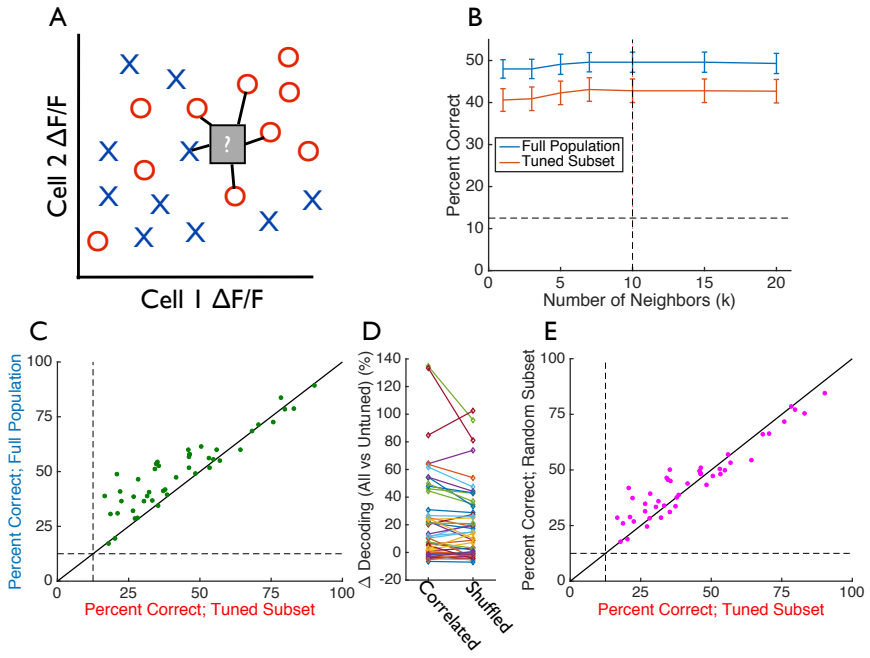

Figure 6: **Untuned neurons can enhance information coding in *in vivo* neural populations.** (A) k-nearest neighbors decoder. Cartoon shows 2 neurons' activities, measured by the $\Delta F/F$. Activities observed single trials are shown; symbol type indicates which stimulus was presented. To decode 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. (B) For different values of k, I applied this decoder to the population activities from all 46 the experiments. The percentage of trials in which the stimulus was correctly identified by the decoder is shown, for decoding either the full population (blue curve), or for decoding only the neurons with strong direction tuning (red curve). Data points show the mean ± S.E.M. over the 46 different experiments. (C) Fraction of trials on which the stimulus was correctly identified using either the full population recordings, or the tuned subsets; each data point corresponds to one of the 46 experiments. (D) Difference in decoding performance when all neurons are decoded vs just the tuned ones, expressed as percentage change relative to the accuracy when only the tuned neurons are decoded. For each experiment the difference in decoding accuracy is shown for the raw data ("Correlated"), and for uncorrelated surrogate data obtained via a trial shuffling procedure ("Shuffled"). (E) 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 E denote equality.

used to distinguish tuned from putatively untuned neurons. Consequently, I repeated the analysis from Fig. 6C with several different criteria (see Methods and Fig. S2). These results are all in

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

qualitative agreement with those of Fig. 6C: regardless of the specific criterion that is used, the putatively untuned neurons contribute to the population code.

Next, I asked whether – as in the theoretical calculations in Fig. 4 – populations that include both tuned and untuned neurons could yield better decoding vs populations of the same size but containing only tuned cells. To answer this question, I extracted a random subset of the neurons from each population, that was the same size as the set of tuned neurons within that population. I then performed the KNN decoding on these random subsets, and compared the performance with that which was obtained on the tuned subsets (Fig. 6E). On average, the decoding performance was $10 \pm 4\%$ (mean $\pm$ S.E.M.) better using the random subsets vs the fully tuned ones, a modest but statistically significant difference ($p = 0.028$, single-sided paired sample t-test).

Finally, it is important to verify that the results shown above apply to other decoders of the neural activity, and not just the KNN decoder. To perform this verification, I used the logistic regression method of [43] to identify the stimuli corresponding to the different population response vectors (see Supplemental Information). This analysis yielded results very similar to those obtained with the KNN decoder in Fig. 6: the full populations can be significantly better decoded to recover the stimulus than can the tuned subsets of the neurons (Fig. S3A), and this effect is diminished when the noise correlations are removed via a trial-shuffled procedure (Fig. S3B). Moreover, random subsets including both tuned and untuned neurons yield better decoding than do the subsets of only tuned neurons (Fig. S3C).

The findings on the population imaging experiments validate the theoretical results from Figs. 1-4. Namely, they show that untuned neurons can enhance neural population coding (Fig. 6C), that this effect depends on the correlations between neurons (Fig. 6D), and that mixed populations of tuned and untuned neurons can yield better information coding than populations of the same size but containing only tuned neurons (Fig. 6E).

Discussion

I showed that, when the variability in neural responses to stimulation is correlated between cells, untuned neurons can contribute to sensory information coding. Moreover, in at least some cases, 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 effects were observed in both a theoretical model (Figs. 1-4), in and in large population recordings from mouse visual cortex (Fig. 6). Moreover, the findings were not sensitive to the specific decoder used for the neural activities (Fig. S3), or to the specific criterion used to identify tuned vs untuned neurons (Fig. S2).

These results have two 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. Moreover, because there is not always a clear distinction between tuned and untuned neurons (Fig. 5C: histogram is unimodal) – and this effect is confounded by noise in the experimental measurements – selection criteria are largely

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

arbitrary. The results shown here suggest that, rather than discarding neurons with low selectivity indices, it may be better to simply include all the neurons in the analysis (Figs. 6C, S3A). This is especially true for the decoding of brain signals to control brain-machine interface devices: better decoding can be obtained by including the activities of untuned neurons in those signals.

Second, because adding untuned neurons can increase the stimulus information (Figs. 4, 6E, S3C), there might be a functional benefit to having some untuned neurons in a population. This is related to previous observations that heterogeneous tuning curves could confer advantages on the population code [44, 14]. Those previous studies did not, however, consider the role of untuned neurons in the population code. It is important to note, however, that no brain area can encode more stimulus information than it received from its inputs [45, 22]. 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.

Observations related to those presented here have also been made by Insanally and colleagues (Cosyne 2017 abstract), and by [10] based on analyses of *in vivo* neural data. 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 isolate the coding benefits of putatively untuned neurons due to noise shaping, vs those due to non-zero tuning, that is nonetheless under the chosen threshold. (However, the trial-shuffle analysis in Figs. 6D, S3B does help make this distinction, as does the fact that the results are not sensitive to the specific criterion uses to label neurons as tuned vs untuned: Fig. S2). This complication highlights the value of the theoretical work presented here (Figs. 1-4): 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 [46, 47]). 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 saw consistent effects in the experimental data (Fig. 6) as in the theoretical model with limited range correlations (Fig. 1), 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 Fig. 1C (i.e., at least 70% 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. Whereas the theoretical model of [41] can encode stimulus

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

information via changes in the correlations between neurons, that effect is not responsible for the results shown here. Notably, for the theoretical calculations in Figs. 1-4, the correlations do not depend on the stimulus, yet nevertheless the untuned neurons contribute to the population code. This is because of the noise-shaping effects shown in Figs. 2-4. Underscoring this point is the fact that, if there are no tuned neurons in our models, there is no stimulus information (Fig. 4: information approaches zero as the fraction of untuned neurons approaches 1). Moreover, the fact that a linear decoder can identify the stimuli presented to the mouse, based on the visual cortical activity patterns (Fig. S3), suggests that the information is being encoded in the firing rates and not the correlation patterns. Given that the neurally plausible decoder [43] used in Fig. S3 is would not be able to extract information that depended only on the stimulus-dependence of the correlations between neurons, this is an important distinction.

We conclude by noting that, even when untuned neurons do not by themselves encode information about the stimulus, they can shape the noise in the population responses, 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 [14, 16, 22]), 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 [22].

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 $\vec{r}_i$ [12, 13, 14, 15, 16, 17, 19, 21, 18, 20, 22]:

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

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

where the prime denotes a derivative with respect to the stimulus, the supserscript T denotes the transpose operation, and $C(s) = \mathrm{cov}\left(\vec{\epsilon}_i | \, s\right)$ is the covariance matrix of the noise in the neural responses to stimulus s. For all calculations, I checked that the correlation (and covariance) matrices were positive semi-definite (thus being physically realizable) before performing the Fisher information calculations.

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

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 $\Delta(\phi) = |\arccos\left[\cos(\phi_i - \phi_j)\right]|$, 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 $[0^o, 360^o]$.

Analysis of *in vivo* neural recordings

Overview of the experiment

The full description of the experiment is given by [11], 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 6 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).

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Data access

Following the example Jupyter notebook provided by [11] – 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, and the stimulus direction for each trial. I analyzed all of the neurons observed in each experiment, and not only those that were labelled as visually responsive.

Tuning curves

I calculated the tuning curves (Figs. 5A 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 (Fig. 5D) were computed over repeats of the same stimulus (same orientation and temporal frequency), and then averaged over all stimuli.

K-nearest (KNN) neighbors decoding

For these decoding analyses, I used the k-nearest neighbors method (Fig. 6A) on the population activity vectors observed on each trial. These vectors had as elements the $\Delta F/F$ values for all of the neurons (blue curves in Figs. 6B, and vertical axis of Fig. 6C), for just the direction selective cells (red curve in Fig. 6B, and horizontal axes of Figs. 6C and E), or for random subsets of neurons that were the same size as the groups of direction selective cells (Fig. 6E, 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.)

Identifying tuned vs untuned neurons

Following [42, 9], direction selective cells were identified via their *circular variance*, with the direction selectivity index (DSI) computed for each neuron as follows. For each stimulus direction, I computed the two-dimensional direction vector $d(\theta) = [\cos(\theta), \sin(\theta)]$, and multiplied that by the neuron's mean response to this stimulus $r(\theta)$ (i.e., the tuning curve value for that stimulus). This yielded a vector $v(\theta) = r(\theta)d(\theta)$ that points in the direction of the stimulus, with the length determined by the cell's mean response to the stimulus. I then averaged this vector over all stimulus directions. If the neuron gave equal responses to all stimuli, the horizontal and vertical components of $v(\theta)$ would average out to zero over all the stimuli, whereas if the neuron responds selectively to one stimulus direction, this cancellation would not occur. Consequently, the DSI is measured by the length of $\langle v(\theta) \rangle$, relative to

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

the neuron's mean response (averaged over all stimuli):

$$DSI = \left| \frac{\langle v(\theta) \rangle}{\langle r(\theta) \rangle} \right|. \tag{3}$$

If the neuron responds strongly to only one stimulus direction, the DSI can be as large as 1, and the DSI can be as small as 0 for neurons that respond equally to all stimuli.

To identify tuned (vs "untuned") neurons, I chose a cutoff of $DSI > 0.25$. This matches the smallest DSI of the direction-selective retinal ganglion cells studied by [9]. I also repeated the decoding analysis from Fig. 6 with different cutoffs on the DSI, and found qualitatively similar results: the precise value of the cutoff is unimportant (Fig. S2).

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.

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

[10] Leavitt ML, Pieper F, Sachs AJ, Martinez-Trujillo JC (2017) Correlated variability modifies working memory fidelity in primate prefrontal neuronal ensembles. Proc Natl Acad Sci USA 114: E2494–E2503.

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

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

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

[14] 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.

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

[16] 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.

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

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

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

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

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

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

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

[24] 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.

[25] 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.

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

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

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

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

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

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

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

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

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

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

[35] 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.

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

[37] 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.

[38] 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.

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

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

[41] Burge J, Geisler WS (2015) Optimal speed estimation in natural image movies predicts human performance. Nat Comm 6.

[42] Mazurek M, Kager M, Van Hooser SD (2014) Robust quantification of orientation selectivity and direction selectivity. Frontier Neural Circuits 8: 92.

[43] Berens P, Ecker A, Cotton R, Ma W, Bethge M, et al. (2012) A fast and simple population code for orientation in primate v1. J Neurosci 32: 10618–10626.

[44] Shamir M, Sompolinsky H (2006) Implications of neuronal diversity on population coding. Neural Comput 18: 1951–1986.

[45] 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.

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

[46] 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.

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

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Supplemental Information

Supplemental Figures

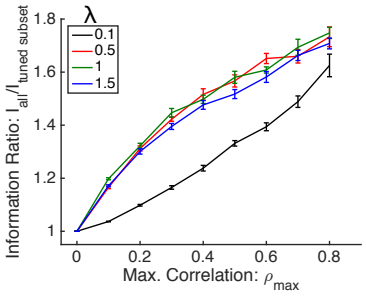

Supplementary Figure 1: **Dependence of information on limited range correlation parameters.** (Related to Fig. 1.) I repeated the calculations from Fig. 1, in all cases for populations of 200 neurons. I repeated the calculations for different values of ρ_{max} and λ, the parameters that define the limited-range 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.

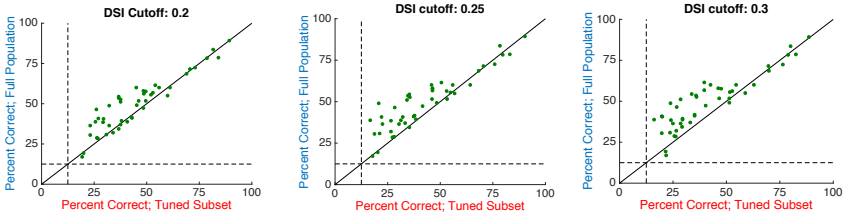

Supplementary Figure 2: **Dependence of the role of untuned neurons on the DSI cutoff.** (Related to Fig. 6C.) I repeated the calculations from Fig. 6C for 3 different values of the DSI cutoff that distinguishes tuned from putatively untuned neurons. For a lower DSI cutoff (vs the value of 0.25 for the results in the main paper) of 0.2 (left panel), the full populations can be decoded $15 \pm 3\%$ better to recover the stimulus than can the tuned subsets ($p = 7.4 \times 10^{-6}$; paired one-sided t-test). For a higher DSI cutoff of 0.3 (right panel), the full populations yield $29 \pm 5\%$ better decoding accuracy than the tuned subsets ($p = 3.6 \times 10^{-8}$; paired one-sided t-test). Reported values are mean $\pm$ S.E.M. over the 46 populations. Chance performance for this task is $1/8 = 12.5\%$. Diagonal line denotes equality.

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Logistic regression analysis

To confirm that the results from Fig. 6CDE are not specific to the KNN decoder, I repeated the same analysis using another decoder. I used the logistic regression method of [43] to take in vectors of neural activity recorded in response to one of 2 different stimuli, and to return a label ("0" or "1") that indicates which of the two stimuli was presented. I randomly divided the data into a training set (75% of the data) that was used to fit the weights of the classifier, and a test set (25% of the data) that was used to measure the performance. After training on the training data, I applied the classifier to the neural responses from the test data set, yielding an output value for each response vector. Values above 0.5 indicated that the stimulus was most likely to be stimulus "1", whereas values less than 0.5 were taken to indicate that response was most likely generated by stimulus "0". I then computed the fraction of these test trials on which this classifier correctly identified the stimulus that caused the neural response. This analysis was separately done for all $(8 \times 7)/2 = 28$ different stimulus pairings: reported performance values are averages over all such stimulus pairings.

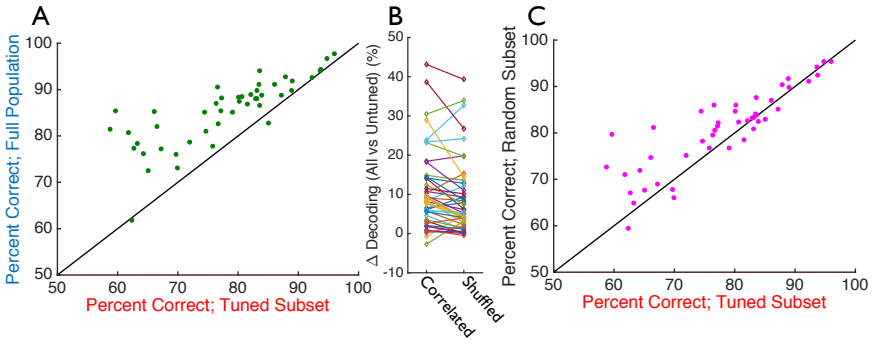

Supplementary Figure 3: **Role for untuned neurons – logistic regression decoder.** (Related to Figs. 6CDE.) I used logistic regression to perform pairwise discrimination on the population response vectors, to determine which of 2 different stimuli caused each response. I repeated this analysis for all possible pairs of stimuli: reported values are the percentage of trials for which the stimulus was correctly identified, averaged over all possible pairings (there is one data point per experiment). (A) Decoding accuracy when the full population response vectors were decoded (vertical axis) vs. when only the tuned subsets of the neurons are seen by the decoder (horizontal axis). (B) Difference in decoding performance when all neurons are decoded vs just the tuned ones, expressed as percentage change relative to the accuracy when only the tuned neurons are decoded. For each experiment the difference in decoding accuracy is shown for the raw data ("Correlated"), and for uncorrelated surrogate data obtained via a trial shuffling procedure ("Shuffled"). (C) Decoding accuracy when random subsets of the neurons (of the same size as the tuned subset, but containing both tuned and untuned neurons) are input to the decoder (vertical axis) vs. when only the tuned subsets of the neurons are seen by the decoder (horizontal axis). Chance performance for this binary discrimination task is 50%. Diagonal line denotes equality.

bioRxiv preprint doi: https://doi.org/10.1101/134379; this version posted June 29, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

I first performed this analysis on the full populations of recorded neurons – including the tuned and untuned ones – and compared this to the decoding performance when only tuned neurons (having $DSI > 0.25$) were used by the decoder (Fig. S3A). Using all of the neurons resulted in $11 \pm 1\%$ better decoding performance (mean $\pm$ S.E.M. $p = 5.3 \times 10^{-11}$; paired one-sided t-test).

Next, I asked whether the role of untuned neurons in population coding relies on the correlations between cells (as predicted by the "noise shaping" theory). To to this, I repeated the analysis from Fig. S3A with trial-shuffled data, in which the correlations between neurons were removed. On these trial-shuffled data, the full populations yielded $9 \pm 1\%$ (mean $\pm$ S.E.M.) better decoding performance than did the tuned subsets of neurons. A comparison, for each experiment, of the improvement in decoding performance by including putatively untuned neurons vs excluding them (Fig. S3B) shows that the effect of untuned neurons is larger with the correlated data vs the uncorrelated control ($p = 0.005$; paired single-sided t-test). Thus, similar to the results in the main paper (obtained with the KNN decoder), the impact of untuned neurons on the population code depends on the correlations between neurons, as is predicted by the theory (Figs. 1-3 of the main paper).

Finally, I used the logistic regression decoder to ask whether populations containing both tuned and untuned neurons could be better decoded than could populations the same size but containing only tuned neurons. To do this, I compared the decoding performance on the tuned neurons in each population (Fig. S3C; horizontal axis), and compared that with the decoding performance obtained on random subsets of each population that were the same size as the tuned subset (Fig. S3C; vertical axis). Using randomly chosen neurons resulted in $4 \pm 1\%$ (mean $\pm$ S.E.M.) better decoding performance than was obtained by selecting only tuned neurons ($p = 2.0 \times 10^{-4}$; paired single-sided t-test). Thus, the mixed populations of tuned and untuned neurons were better at encoding the stimulus than were populations of the same size but containing only tuned neurons.