Contrast Invariance of Cell Populations in V1

Daniel Manson

January 2012

Abstract

Cells in mammalian primary visual cortex respond preferentially to bars of lightness or darkness on the retina. For a long time it has been known that the extent of this sensitivity does not depend on the contrast of the stimulus. It has recently been shown that this invariance also exists in the shape of the time-averaged population response. Here we examine whether invariance of the population holds over much shorter time periods, and tentatively conclude that it does.

Contents

1 Background

1.1 From photon to V1

[1] Each human retina contains approximately 100 million light sensitive neurons, known as rods and cones. Whenever a photon is absorbed by one of these cells, it triggers a change in the electrical potential across the cell's membrane. Usually these fluctuations would give rise to action potentials, but in the case of rods and cones they do not, and are instead relayed directly to bipolar cells via gap junctions. Each bipolar cell receives input from a circular group of either rods or cones, with some of the inputs having an excitatory effect and others an inhibitory effect. In some of the bipolar cells the excitatory region is in the center of the circle, and in other cells it is in the area surrounding the center. Thus bipolar cells can be considered either on-center or off-center, responding maximally when only the excitatory region is illuminated (see Figure 1).

Bipolar cells, like rods and cones, do not have synapses or produce action potentials, instead they are connected via gap junctions to amacrine and retinal ganglion cells. They are also connected recurrently via horizontal cells. This connectivity is depicted in Figure 2.

Amacrine cells and retinal ganglion cells do threshold their inputs, producing action potentials that leave the retina and travel along axons to the lateral geniculate nucleus (LGN) in the thalami. On route to the thalami, the axons pass through the optic chiasm where half the axons from each eye

[1] The content of this section can be found in many textbooks including [12].

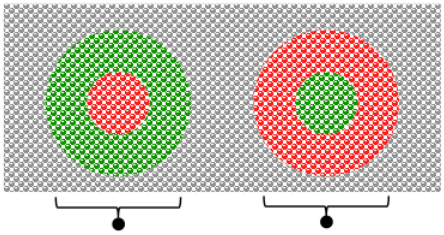

Figure 1: A rectangle representing a small section of light sensitive cells on the retina, with the two black circles below representing bipolar cells connected to the cells above. Green cells have an excitatory effect on one of the bipolar cell, red cells an inhibitory effect, and grey cells are not connected to either of the bipolar cells shown. The bipolar cell on the left is off-centre, and the one on the right is on-centre.

cross over from one side of the brain to the other in order that axons from the left and right visual fields are afferent on the right and left LGNs respectively (contralaterally).

LGN cells, and indeed many other thalamic nuclei, can be thought of as performing a role similar to that performed by the registers in a microprocessor: neocortex 'reads' from the LGN, performs a computation, and 'writes' an output back. Although this analogy is somewhat helpful, it should not be taken too literally since many cortical regions are connected to each other directly; this is certainly true of the visual areas.

In this work we focus on the first stage of neocortical visual processing, which takes place in an area known simply as V1.

1.2 Cell response types in V1

Output from LGNs cells, which carries roughly the same information as that from bipolar cells, gives rise to a variety of cell types in V1, some examples of which are given in Figures 3 and 4. The one thing that all the cells have in common is that they respond to bars of lightness or darkness falling on well defined receptive fields in the retina, and in all cases there is a specific orientation at which the bar produces a maximum response. Although the distribution of these different cell types does not appear to conform to any significant spatial pattern (ignoring small differences between layers), V1 does exhibit several other important spatial patterns:

- The layout of V1 is said to be retinotopic, meaning that areas close together in the retina connect to cells close together in V1.
- Cells receiving their dominant inputs from different eyes are segregated into bands, known as ocular dominance columns.
- The orientation selectivity of cells is organized into 'pinwheel' shapes so that neighboring cells respond preferentially to very similar orientations of bar. See Figure 5.

1.3 Levels of electrophysiological analysis

Most of the information provided so far has been acquired as a result of experiments in which monkeys or cats have had very fine electrodes implanted into their brains, before being shown different light

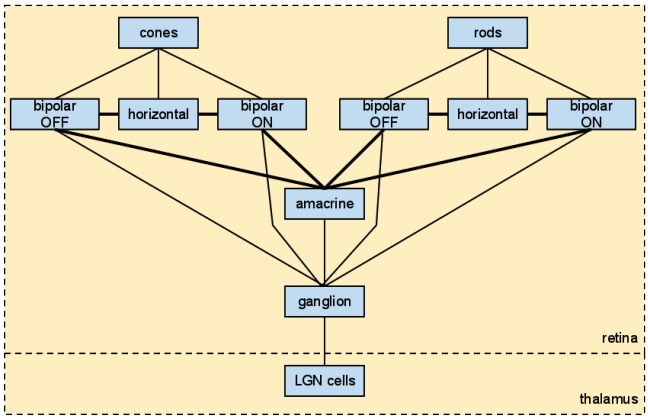

Figure 2: Start of the visual pathway. Signals flow along the thin black lines from top to bottom, and along the thicker lines in both directions. Note that most of the cell types shown can be further subdivided into classes showing only a subset of the connections depicted here.

patterns, all while under anesthetic [11]. Recording signals from brains in this way is an experimental technique from the field of electrophysiology, and can be done at several different scales:

Membrane potential. The sub-threshold fluctuations of the membrane potential can be measured for individual cells.[2]

Single unit. Spiking rates can be recorded for individual cells.

Multi unit. The aggregate spiking rates of several dozen [14] neighboring cells can be recorded. Commonly [6], a spike is said to have occurred when voltage exceeds 4 s.d. from the mean. However these events may also correspond to coincident pairs (or triples etc.) of true action potentials within the unit, implying a more complex relationship to the true total rate of the multiunit. Clearly multiunit analysis is only useful if neighboring cells are known to exhibit similar characteristics, which is the case in V1.

Population. Spiking rates of many single units or many multi units can be recorded simultaneously (separately not aggregated). We can then analyze this data to see how cells with different selectivity preferences are responding to stimuli. Recordings of this kind are done with arrays of approximately 100 electrodes (Utah arrays).

Local Field Potential Weakly localized temporal fluctuations of electric field can be recorded. This activity is thought to correspond to population-wide sub-threshold fluctuations rather than mean spiking rates. Interestingly, it has recently [13] been shown that LFP may be localized to within aproximately $250\mu m$. However, analysis of LFP data can only give us a coarse view of how much activity is occurring in a particular region of the brain. In many cases we will be able to analyze

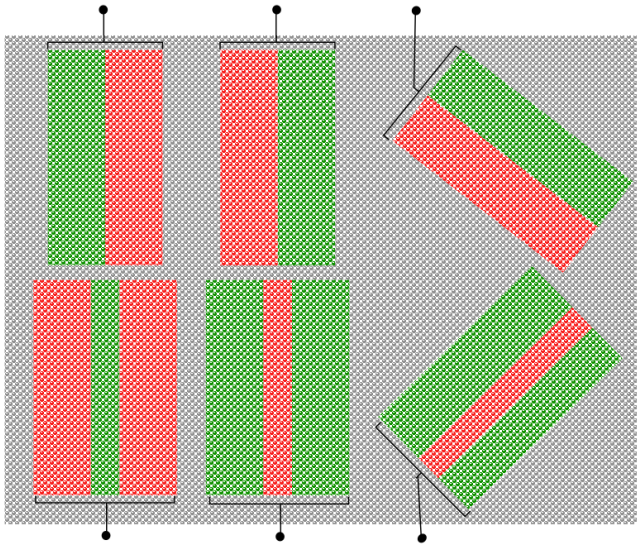

Figure 3: V1 cells are known as 'simple' if they have well defined excitatory and inhibitory regions on the retina. The meaning of the symbols here is as in Figure 1, except that here the black circles represent V1 cells not bipolar cells, and their connections to the light sensitive cells are not explicitly known. Note the two main types of simple cell are edge detectors (in the **top** row), and line detectors (in the bottom row). There are a few other known patterns, but these are the best recognized ones. Note that of course the axis can be at any angle.

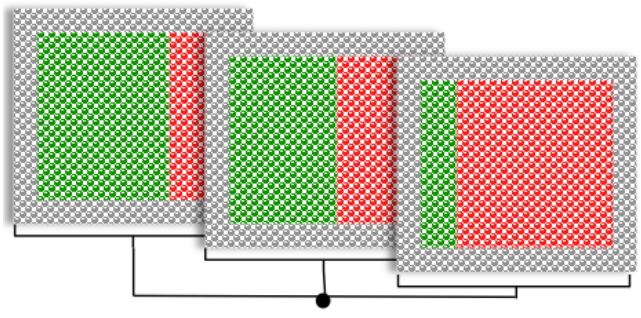

Figure 4: V1 cells are known as 'complex' if they do not have specific excitatory and inhibitory regions, but instead respond to bars placed anywhere in their field. Here the grey rectangles represents the same patch of retina three times over, giving a different boundary between the excitatory and inhibitory regions in each case. Some complex cells are particularly responsive to moving bars.

[2]This is done with a pipette rather than a wire, and is known as patch clamping.

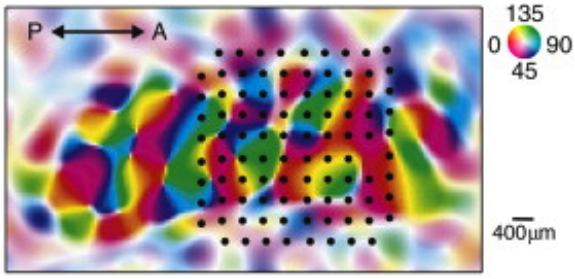

Figure 5: The spatial distribution of orientation selectivity in V1. Data for the image was obtained using voltage sensitive dye and optical imaging. Colors indicate preferred orientation and saturation indicates the strength of tuning. The black dots show, to scale, the position of the electrodes of a Utah array. Figure taken from [13].

the magnitude and frequency of oscillations at characteristic frequency bands, e.g. oscillations in the 5-10Hz band are common and referred to as theta waves. (This is less relevant in V1, except perhaps during development.)

Each of these levels of analysis has its shortcomings as well as its benefits - see Figure 7 for a discussion of one problem with population analysis. At this point it is important to emphasize that each of these levels of recording and analysis is distinct. We need to be especially careful in V1 because the mathematics and statistics can seem deceptively similar at all scales. Let us not forget though, that although we may be analyzing at the population level, ultimately we wish to elucidate the cognitive algorithm that that is being implemented at the single unit level.

1.4 Motivation for studying V1

The pioneering work of Hubel and Wiesel in the 1960s, some of which has been presented here, led Marr and others to make significant advances in the theory of visual processing. Over the course of several papers, and in his book, Vision [15], published posthumously in 1982, Marr described in some detail how V1-like cells might form the first stage in a complex chain of processing steps leading

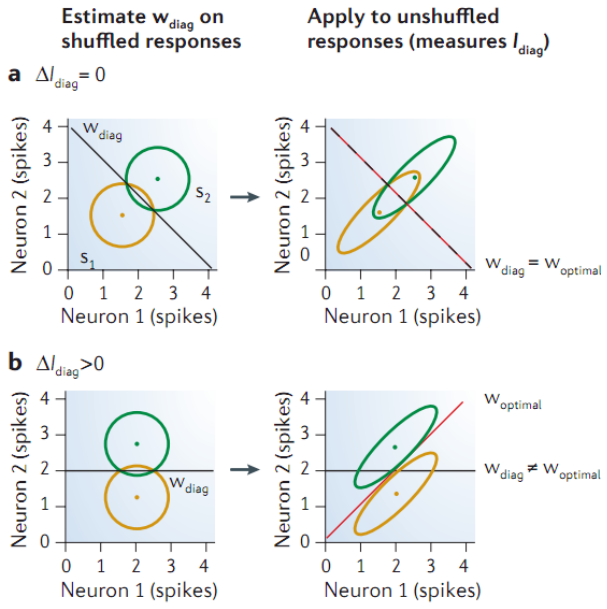

Figure 7: Current electrophysiological techniques only permit us to record from about two or three neighboring single units simultaneously. If the cells are tuned to similar stimuli their responses will obviously be correlated, i.e. when Neuron 1 has a high rate it is likely that Neuron 2 will also. In addition to the signal, these rates have an element of stochasticity, which sometimes exhibits a weak correlation across neurons. These two types of correlation interact and can either increase or reduce the amount of information theoretically extractable from the signal.

In each of the panels in the figure, the two circles/ellipses represent a range of responses to two distinct stimuli; the black line, w_{diag}, indicates the optimum division between the two responses if the noise is uncorrelated, and on the right the red line, $w_{optimal}$, indicates the optimum given the noise correlation depicted.
In **a** the signals are positively correlated; on the **left** the noise is uncorrelated and on the **right** there is a positive correlation. Here the correlation of the noise acts to reduce the information in the two neurons, which is shown by the large area of each ellipse on the incorrect side of w_{diag}. In this example, knowledge of the noise's correlation does not help us improve the amount of information we can extract from the pair of signals, which is shown by the equivalence of the red and black lines.
In **b** the signals are not correlated; again on the left the noise is uncorrelated, and on the right there is a positive correlation. Here the correlation of the noise acts to increase the information, which is shown by the fact that the red line now perfectly separates the two response regimes. However it is only possible to benefit from this if we know what gradient to give the line, i.e. what is the correlation of the noise. Simply using w_{diag} would actually gives us less information than in the case on the left.
Importantly, a weak effect of this kind between pairs of neurons - which may be very difficult to identify experimentally - will result in large changes in the information content of large populations of neurons.
Figure taken from [4]. See [17] and [9] for some experimental evidence of correlations in the retina and V1 respectively, and [1] for details of the mathematics.

from retinal data through what he called the raw primal sketch, full primal sketch, two-and-a-half dimensional sketch and ultimately to a full three dimensional model of the exterior world. Since the publication of this work, significant progress has been made in computer vision, but this has had little bearing on our understanding of visual cognition. The apparent lack of progress is due to the fact that until recently, the tools available for experimenting on live animals at the cellular level had changed little since the 1960s[3], and even now it is still not possible to record simultaneously from large populations of single neurons. As a result, due to its spatial patterns and simple responses, V1 remains one of the only brain regions that can be studied in any detail. Refining our understanding of V1 is no doubt a perquisite for studying V2 and beyond, but it will also serve to teach us the sorts of mechanisms that the brain employs in order to accomplish the tasks required of it.

2 Modeling V1 orientation selectivity

2.1 Single units

At the single unit level, a cell which is maximally selective for orientation θ will have a spiking rate $r_\theta(\varphi, c)$ when exposed to a stimulus of orientation φ and contrast c. The nature of this function r is well described by the following [18, 2]:

$$r_\theta(\varphi, c) = r_{\max} \frac{c^n}{c_{50}^n + c^n} \frac{e^{\kappa \cos(\varphi - \theta)}}{I_0(\kappa)}$$

The nice thing is that this function consists of two separate factors: a contrast factor and an orientation factor. This separability is known as *contrast invariance*. The contrast factor is a hyperbolic ratio parameterized by two numbers, n and c_{50}. And the orientation factor is a circular Gaussian parameterized by θ, describing the location of the peak; and κ or ${}^1/_\kappa$, describing the spread. I_0 is a modified Bessel function of order zero, this and $r_{\max}$ are only required for normalization.[4]

The observed rate (i.e. spike count in a short interval) resulting from this idealized rate is reasonably well modeled as a Poisson process with parameter $\lambda = r_\theta(\varphi, c)$, this is despite the fact that the probability of an action potential is dependent on the recent history of the neuron.

2.2 Noise and the iceberg effect

As noted in the 1970s and still discussed [8], contrast invariance at the single unit level is difficult to explain with Hubel and Wiesel's original wiring model [11], which proposed that V1 simple cells sum the responses from a set of collinear LGN on-center (or off-center) cells and then apply a threshold to give a spiking rate. The problem is that low contrast stimuli produce weaker retinal signals and thus cannot drive the V1 cells over their threshold - the 'iceberg effect'. Or rather, the orientation selectivity should decrease notably with contrast - which it does not. As it happens, low contrast signals are significantly amplified in the retina and LGN (this is termed 'gain control' and is remarkably consistent) but it still remains to explain why r is smooth across the domain of θ.

Further to this, two other mechanisms may be involved: perhaps recurrent inhibitory connections in V1 act to reduce firing at non-preferred orientations, thus permitting the use of a lower threshold; alternatively [3] sub-threshold 'noise' in V1 could play an important role, turning what was a fixed threshold into a stochastic boundary. As stated in [8], recent evidence strongly supports the later mechanism, although there is also substantial evidence of significant lateral connectivity in V1. And it is noted that noisy thresholds may be an important computational feature widely implement throughout the brain.

2.3 Multiunits

We now define a rate $R_V(\varphi, c)$, where $V \subset W_\theta$: $\{i \mid \text{neuron } i \text{ is roughly selective for orientation } \theta\}$:

$$R_V(\varphi, c) = \sum_{i \in V} r_\theta^i(\varphi, c)$$

[3] Recent advances in optical microscopy and fluorescence techniques hold the promise of offering higher resolution data.

[4] $I_0(\kappa) = \sum_{a=0}^{\infty} \frac{1}{a!(a+1)!} \left(\frac{\kappa^2}{4}\right)^a$

For which we note that the neurons in V all have the same value of θ, but different values of κ, n and c_{50}. To the experimenter, these three values can be treated as random variables K, N, and C_{50}, sampled independently from three distributions of some unknown shape (see [6, 16, 2] for experimentally obtained distributions of these variables and evidence of independence). Using this idea we can express an expectation for R as:

$$\mathbb{E}[R_V(\varphi,c)] = \mathbb{E}\Big[\sum_{i\in V} r^i_{\theta,\varphi,c}(K^i,N^i,C^i_{50})\Big] = \sum_{i\in V}\mathbb{E}\big[r^i_{\theta,\varphi,c}(K^i,N^i,C^i_{50})\big]$$

$$= m\mathbb{E}\big[r^i_{\theta,\varphi,c}(K,N,C_{50})\big]$$

where m is the number of neurons in V. Then, since the three variables are independent for individual cells, we can further expand the expectation:

$$= m\mathbb{E}\Big[\frac{c^N}{C_{50}^N+c^N}\Big]\mathbb{E}\Big[\frac{e^{K\cos(\varphi-\theta)}}{I_0(K)}\Big]$$

This expectation will be a good approximation only if m is large,which is probably the case (we noted earlier that it is of the order of several dozen; Figure 8 shows what can happen when $m = 2$). As with the case of single units, this formula is separable, and thus contrast invariant. However the shapes of the contrast function and the orientation function are no longer known - Jensen's Inequality in its weakest form says that in general $\mathbb{E}[f(X)] \neq f(\mathbb{E}[X])$.[5] As it happens, it is reasonable to approximate $R_V(\varphi,c)$ with an identical equation to that of $r_\theta(\varphi,c)$, but we should be clear that no analytical link has been demonstrated.

As before it is possible to model the observed rate as a Poisson process,[6] this time despite the possible complex internal correlations within the unit and the complexities of the definition of multi-unit spiking.

2.4 Population

Having stated our definition for R, we can now define the population response:

$$P_{\varphi,c}(\theta) = \frac{1}{n}\sum_{V\subset W(\theta)} R_V(\varphi,c) = \frac{1}{n}\sum_{V\subset W(\theta)}\sum_{i\in V} r^i_\theta(\varphi,c)$$

where n is the number of multiunits with preferred orientation roughly equal to θ. This double sum is equivalent to the single sum in the previous section, but now with a higher value of m. Thus we have effectively improved upon the approximation to the expectation. We can further increase n by expressing P as a function of $\varphi-\theta$ rather than θ, and aggregating over responses at different values of φ.

A slightly confusing point worth mentioning is that whether or not m is large, R will be a continuous function of φ since it is the sum of continuos functions. However P may have discontinuities in θ if the number of neurons in each $W(\theta)$ is too small. And to reiterate the point made earlier, no amount of neurons will allow us to claim that P or R will be of the same form as r, though experimental evidence suggest this is the case [6].

Observations of $P_c(\varphi-\theta)$ are observations of a distribution (or vector) not of a scalar, so we must leave the domain of simple Poission statistics. However it is crucial that we understand the statistics of P, since it is the instantaneous state of the population which is the input to downstream neurons,

[5] To see this, imagine taking the expectation of a simple gaussian $G = e^{-ax^2}$ where the parameter a has a uniform distribution on the interval $[1,4]$. We get $\mathbb{E}[G(A)] = \int_1^4 p(A)e^{-Ax^2}dA$ which is $\frac{1}{3x^2}[e^{-x^2} - e^{-4x^2}]$, i.e. it has a different form from G. Note that here we are deterministically summing Gaussian functions, i.e. we are not summing random variables drawn from Gaussian distributions, which would produce a further Gaussian.

[6] There is some debate as to whether this is true: in some experiments [7] responses appear almost Poisson, but not in others [10]. It is thought that higher variance may be due to fluctuations in the animal's state of awareness during the experiment.

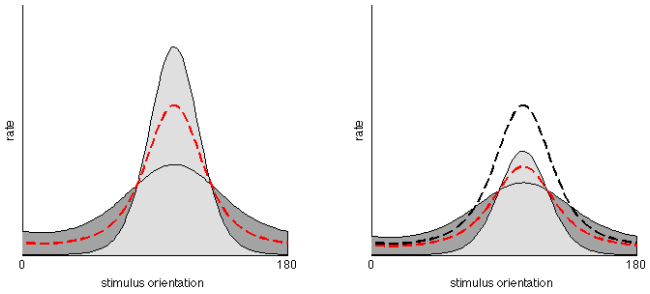

Figure 8: The mean/total rate of two contrast invariant neurons, with the same preferred orientation, in general is not itself contrast invariant [6]. The light-shaded Gaussian and dark-shaded Gaussians are the same in both panels except their magnitudes are smaller in the second panel. In other words the panel on the **left** represents response to a high contrast stimulus, and that on the **right** represents the response to a lower contrast stimulus. In both panels the red line shows the mean of the two Gaussians. For ease of comparison, a dotted black line is shown on the right in the same location as the red line in the left panel.

as well as likely impacting on V1 itself via recurrent connections. This brings us to the problem we wish to address in this work: how can we test whether the population response is contrast invariant over a short (one or two-second) time window? If it proves not to be invariant, we may have to adjust the model. Either way, we will improve our understanding of how sub-threshold 'noise' is regulated during changes in contrast.

Before considering this question it is temping to ask a more fundamental one: how exactly do downstream neurons 'interpret' (or 'decode') the population response? It is often assumed that this decoding is simply a vector averaging of the population, but there is some evidence [10] that individual neurons may be reliable enough to act as the population signal. Elsewhere, [9] a range of possibilities have been discussed and tested and it has been concluded that the optimal decoding technique (which uses units' covariance) is probably not being implemented, though with current technology it is difficult to make any strong statements. Another approach has been [6] to probe the population with two super-imposed stimuli of different orientation and contrast, and then fit the population response to various candidate models.[7] The data show that as the difference in contrast between the two stimuli is increased, the population transitions from representing both stimuli, to a winner-takes-all scenario in which the weaker contrast stimulus is underrepresented by the population.

However, without knowing the decoding algorithm, it is not clear whether we are characterizing the population response appropriately, or in other words we cannot be sure that we are looking at the true information content of the population.

3 Data analysis

3.1 Multiunits

The data presented in [6] were used again here.[8] The data consisted of multiunit spike times recorded from layers 2 and 3 of V1 cortex in cats. In each experiment approximately 95 multiunits were successfully recorded. The stimuli were sinusoidal gratings like those in Figure 9. Spatial frequency was chosen to maximize the response of the multiunits.

The first stage of analysis described in [6] can be broken into several steps:

Figure 9: Example stimuli used in experiment. Each stimulus has a fixed angle and a maximum contrast and all stimuli have the same spatial frequency. Contrast is temporally modulated by a 4Hz sine wave.

[7] The first model was the weighted sum of two $r(\varphi, c)$ functions. It provided good fits to the data, but had a large number of parameters and was not predictive. The second model took the form of a 'contrast normalization':

$$R_{1+2}(c_1, c_2) = r_{\max} \frac{c_1^n G_1 + c_2^n G_2}{c_{50}^n + (c_1^2 + c_2^2)^{\frac{n}{2}}}$$

where G_1 and G_2 were the Gaussian responses to the stimuli when presented separately. This model has fewer parameters and still provides a good fit to the data.

[8] Approximately half the stimuli were not suitable for this analysis due to having two gratings rather than one.

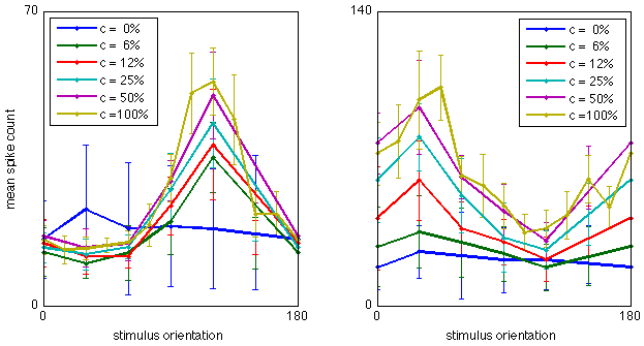

Figure 10: Example responses of two multiunits taken from CATZ082 series 6 experiment 3. Note that the response of the unit in the **left** panel drops significantly below the background rate, whereas the response of the multiunit on the **right** does not. Error bars show $\pm 1SE$.

1. Calculated a mean rate for each unit for each presentation of the stimuli. A more rigorous, though less successful, approach would be to extract the power in the 4Hz (or 8Hz) band of the signal, since this corresponds to the stimuli at maximum contrast. In the past, attempts were made to use different frequency components of the signal to distinguish between simple and complex cells (using a modulation ratio) but this was later found not to be possible - see [16].

2. Calculate the vector average for each unit over all orientations of the full-contrast stimuli:

$$R = \frac{\sum_k r_k e^{i2\theta_k}}{\sum_k r_k}$$

   where r_k is the mean rate of the unit for stimulus orientation θ_k. Define the unit's preferred orientation, φ, as $Arg[R]\frac{1}{2}$.

3. The circular variance, given by $1 - |R|$, indicates of how well tuned the unit is. Discard units with variance greater than 0.85.

4. Subtract each unit's mean background firing rate (obtained when no stimulus is present) from all its responses. Zero any negative rates that result. It is likely that this background rate not only corresponds to the baseline firing rates of cells comprising the multiunit [16], but also to the unrelated firing of cells that we do not wish to include in the multiunit. Figure 10 gives examples where one or other of these sources of background may be dominating, though there are a number of other possible explanations.

5. Normalize each unit's firing rate by dividing by the mean response to the maximum contrast stimuli).

There are a couple of problems with this approach. Firstly, the variance of the rates is roughly proportional to their means, which weakens the case for using the naive vector average. Secondly, it may be wrong to subtract out both sources of background: our definition of $r_\theta(\varphi, c)$ does not permit negative values (since rates cannot be negative), so perhaps we should treat sub-baseline responses as fractional rates rather than negative rates.

A possible solution to both these problems is to parametrically fit $r_\theta(\varphi)$ to the data, using an exhaustive search for the least squares solution. In addition to κ and θ, we parameterized the function

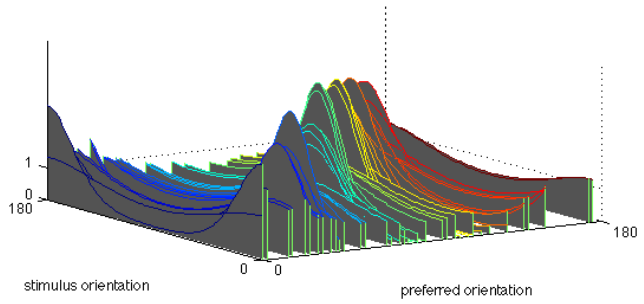

Figure 11: Parametrically fitted curves for the 52 well-tuned units in series 6 experiment 3 for CATZ082. Colors are purely to aid visualization. Curves are all mean normalized to one. For the 52 units pictured, the weighted-MSE of the fit is 0.17 ± 0.12. In the end these fits were not used in the analysis.

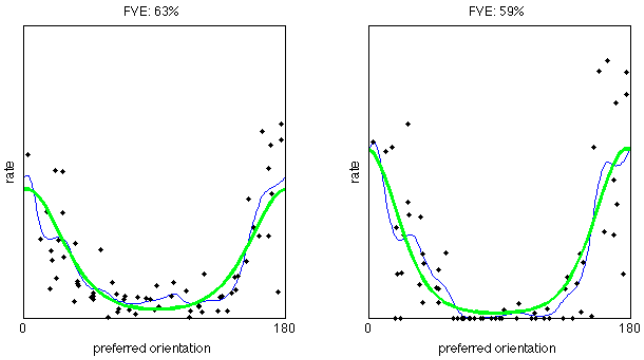

Figure 12: Example population response to the presentation of a stimulus at 0 degrees with 12% contrast. The **left** panel shows the response of the population averaged over repeats, the **right** panel shows the values for a single presentation of the stimulus. The black diamonds show the rate of individual multiunits and the blue line shows the binned, smoothed rates. The green curve is the result of fitting a circular Gaussian to the binned values. FVE values give the fraction of variance explained for the green curve applied to the black points. Data from CATZ084 series 12 experiment 16.

by a background level offset and a magnitude.[9] At the multiunit level this approach gave reasonable fits, however it proved difficult to use the results at the population level. So, in the end this method was abandoned - see Figure 11 for an example set of fits.

3.2 Population

The standard approach when analyzing the time-averaged population response is to bin units into 12 groups according to their preferred orientation and then find the mean rate in each group [6]. As mentioned in the modeling section, it is also possible to aggregate over all orientations of stimuli by binning according to the difference between the stimulus orientation and the preferred orientation, rather than the absolute stimulus orientation. After this processing step it is relatively easy to fit circular Gaussians to the population response.

Here, however, we do not have the luxury of averaging over either repeats or stimuli. Referring back to the modeling section, this means that not only will the observed values of R have a greater error, but $P(\theta)$ may also have some discontinuities.

In response to these issues it was decided to discard two out of the eight experiments, leaving only those with a greater number of well-tuned units (32, 32, 45, 51, 52, and 66). In addition, instead of using 12 bins, the population was binned into 180 bins and then smoothed with a gaussian kernel ($\sigma^2 = 36$ degrees^2). The mean values were obtained from this by dividing by analogously binned and smoothed unit counts. It was hoped that this would make P more Gaussian-like. To account for the greater Poisson-type nature of the data, during fitting it was decided that the squared errors should be weighted by an indicator of variance. To ensure this weight was neither too high nor too low, and reflected both the sample variance and the expected variance of a Poisson process, the weights were defined as:

$$W = \frac{2}{1 + v/\max_\theta(v)} \text{ where } v = \sqrt{P(\theta)\sigma(\theta)}$$

where $P(\theta)$ is the binned-smoothed mean and $\sigma(\theta)$ is the smoothed variance.[10]

Figure 12 shows an example of this fitting process applied to a low contrast stimulus, and the full set of results are given in Figure 13. Inevitably the fitting improves with contrast which makes it difficult to compare values of kappa, especially using the data in this aggregate form. Qualitatively however, kappa appears not to depend on contrast.

It might have been better to calculate ranked correlations from each experiment separately, and then summarize the results using p-values and an appropriate false discovery rate correction. Alternatively, we could calculate two FVE statistics for each trial, firstly using the fit to the individual trials and then using the fit to the average (as in Figure 12). A comparison between between these two values would indicate the strength of the fits obtained.

[9] The fit is performed on the mean normalized data, so the magnitude should be close to unity. Also note that when adjusting the offset and magnitude we have to take care to normalize the goodness of fit appropriately.

[10] Smooth variance was calculated by finding the squared error of each unit's response relative to $P(\theta)$ and then binning, smoothing, and normalizing as was done with P itself.

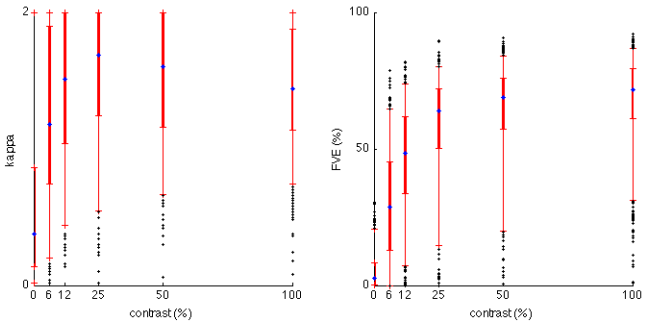

Figure 13: Box plots showing the results of fitting for all trials of all experiments. Outliers lie beyond the 5th and 95th percentiles. As in Figure 12, FVE stands for fraction of variance explained by fit.

3.3 Fluctuations of population response

We now consider a different approach. As has already been stated, the population response appears contrast invariant when averaged over enough trials. So if we wish to show that invariance exists for single trials, all we need to do is compare across trials of the same stimulus and demonstrate that the response is 'the same'.

The simplest way for the response to be different across repeats is for the peak rate to fluctuate around the value of the true stimulus orientation – see top two panels of Figure 14. If this is the case, we might be able to detect these fluctuations by measuring the angular variance of the population mean.

Figure 15 shows the results of such an analysis. Although there is some drop in the variance as contrast is increased, the relationship appears weak, and is difficult to interpret.

Perhaps a more powerful approach would be to examine correlations between pairs of cells across repeats of the same stimulus. In some regions we will expect a positive correlation and in other regions a negative correlation - see bottom panel of Figure 14. If on the other hand the population response is constant, there will not be any detectable correlations across repeats of a single stimulus.

An analysis of this type was explored using the following method: for each stimulus we can calculate the Spearman's rank correlation coefficient, ρ_{ab}, for all possible pairs of units a and b (e.g. with 32

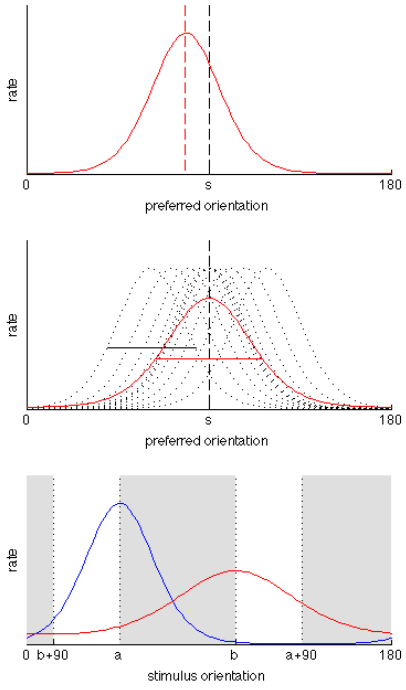

Figure 14: When presented with a stimulus of orientation s, perhaps the population recognizes it with some error, shifting the population response by a small angle relative to s. This is depicted in the **top** panel, with the center of the red curve being slightly offset from the true stimulus orientation. Over multiple repeats of the same stimulus, this error will presumably have a roughly normal distribution. This is depicted in the **centre** panel, with the dotted black lines representing the population response on individual trials, and the red line representing the averaged population response. Note that the width of the averaged response (red) is larger than the width of individual trials (black).
In the **bottom** panel, the response curves of a pair of neurons, a and b, are depicted (mean normalized). If the population response fluctuates around the true value as shown in the top two panels, we should expect a correlation between a and b across repeats of a single stimulus. The correlation will be negative in the shaded regions, and positive in the unshaded regions. If, on the other hand, the population response is always centered on the true stimulus orientation, we should not expect any correlations.

units we have $32 \times 31 \times \frac{1}{2} = 496$ pairs).[11] Then for each pair we identify in which region of Figure 14 each stimulus lies. From this we calculate weights, $w_{\theta ab}$, for each stimulus, where the sign of w is given by the sign of the expected correlation, and the magnitude is defined as one over the number of stimuli in the region (positive or negative region). Thus the weights in the negative region sum to -1, and in the positive region to $+1$. Once we have calculated weights and correlations, we can calculate a metric, M_{ab}, for each pair of cells at each contrast, c:

$$M_{ab}(c) = \sum_{\theta} w_{\theta ab} \rho_{\theta ab}$$

Strongly positive values of M correspond to differences in the population response across trials, and values of M_{ab} close to zero correspond to consistent population responses. However computing all the values of M is not enough to distinguish between the two hypotheses; we need to compute the distribution of M values under shuffles of the data and calculate significance values.

This approach does not require us to subtract background or normalize the units' rates, which is good. Also, in theory it should be relatively agnostic of the distribution of tuning widths within the population. We cannot say it is truly agnostic because values of M_{ab} will tend to be low when a or b has either a very wide or very narrow tuning width. The value and interpretation of M_{ab} will depend upon the number of repeats of the stimulus, making it non-trivial to compare across experiments with different protocols. More importantly, we should note that the variance of M_{ab} will depend on the number of stimuli in each region. For example if a and b have very similar preferred orientations, all

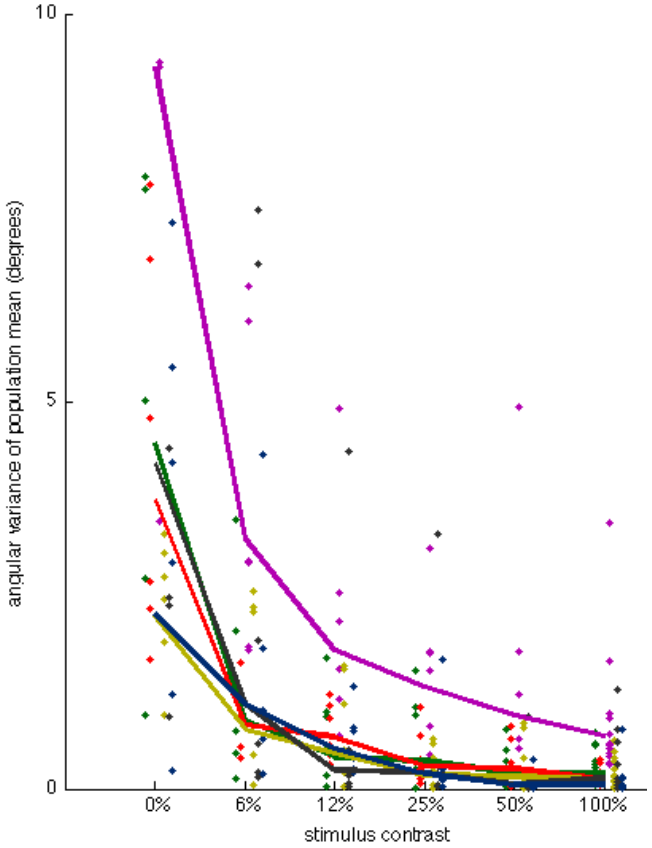

Figure 15: In the figure, each color corresponds to a separate experiment. The dots show the angular variance of the population mean across repeats of the same stimulus. The lines connect the average (geometric mean) of the angular variance at each contrast. Given the multiple levels of summary statistics used to produce this data, it was considered more appropriate to display it graphically rather than summarize it further and test for correlations.

[11] Spearman's rank correlation coefficient is given by:

$$p_{ab} = \frac{\sum (a_i - \bar{a})(b_i - \bar{b})}{\sum (a_i - \bar{a}) \sum (b_i - \bar{b})}$$

where a_i and b_i are the ranked values of rates for units a and b for the ith presentation of a stimulus, and $\bar{a}$ and $\bar{b}$ their means. Tied values are assigned the mean of the tied ranks.

stimuli will fall in the positive region and the variance of M will be related to $\frac{1}{n}$, whereas if a and b have perpendicular preferred orientations the weights will be equal to $\pm\frac{2}{n}$ and the variance of M_{ab} will be related to $\frac{2}{n} + \frac{2}{n} = \frac{4}{n}$. It may be possible to adjust the definition of M_{ab} to account for this, but it was hoped that by comparing against appropriate shuffles of the data, the analysis could be made insensitive to this problem.

This method is reasonably complex, and as such it remains to be fully implemented (no results are shown here).

4 Concluding remarks

We have described the link between averaged population responses and single unit responses, and have presented some initial analysis to test whether contrast invariance exists over short time scales. The initial results suggest invariance holds, but they are not conclusive. Further potential analytical approaches have been discussed. Ideally we would like to consider shorter and shorter time frames, but with the experimental protocol used it is not clear that this is possible.

References

[1] LF Abbott and P. Dayan. The effect of correlated variability on the accuracy of a population code. *Neural Computation*, 11(1):91–101, 1999.

[2] D.G. Albrecht and D.B. Hamilton. Striate cortex of monkey and cat: Contrast response function. *Journal of Neurophysiology*, 1982.

[3] J.S. Anderson, I. Lampl, D.C. Gillespie, and D. Ferster. The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science*, 290(5498):1968, 2000.

[4] B.B. Averbeck, P.E. Latham, and A. Pouget. Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5):358–366, 2006.

[5] P. Berens. Circstat: a matlab toolbox for circular statistics. *Journal of Statistical Software*, 31(10):1–21, 2009.

[6] L. Busse, A.R. Wade, and M. Carandini. Representation of concurrent stimuli by population activity in visual cortex. *Neuron*, 64(6):931–942, 2009.

[7] M. Carandini. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS biology*, 2(9):e264, 2004.

[8] M. Carandini. Melting the iceberg: contrast invariance in visual cortex. *Neuron*, 54(1):11–13, 2007.

[9] Y. Chen, W.S. Geisler, and E. Seidemann. Optimal decoding of correlated neural population responses in the primate visual cortex. *Nature neuroscience*, 9(11):1412–1420, 2006.

[10] M. Gur and D.M. Snodderly. High response reliability of neurons in primary visual cortex (v1) of alert, trained monkeys. *Cerebral Cortex*, 16(6):888, 2006.

[11] D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1):106, 1962.

[12] E.R. Kandel, J.H. Schwartz, T.M. Jessell, et al. *Principles of neural science*, volume 4. McGraw-Hill New York, 2000.

[13] S. Katzner, I. Nauhaus, A. Benucci, V. Bonin, D.L. Ringach, and M. Carandini. Local origin of field potentials in visual cortex. *Neuron*, 61(1):35–41, 2009.

[14] J.H. Macke, J.P. Cunningham, M.Y. Byron, K.V. Shenoy, and M. Sahani. Empirical models of spiking in neural populations. 2011.

[15] D. Marr. *Vision: a computational investigation into the human representation and processing of visual information*. W. H. Freeman, San Francisco, 1982.

[16] D.L. Ringach, R.M. Shapley, and M.J. Hawken. Orientation selectivity in macaque v1: diversity and laminar dependence. *The Journal of neuroscience*, 22(13):5639–5651, 2002.

[17] E. Schneidman, M.J. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.

[18] G. Sclar, RD Freeman, et al. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale*, 46(3):457, 1982.

[19] B.C. Skottun, A. Bradley, G. Sclar, I. Ohzawa, and R.D. Freeman. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *Journal of Neurophysiology*, 57(3):773–786, 1987.

Acknowledgement. The author is grateful to Matteo Carandini for making his data available, as well as more generally for providing guidance and help with this work.