

# Population Coding in Retinal Ganglion Cells

*Reza Abbasi Asl*

Electrical Engineering and Computer Sciences  
University of California at Berkeley

Technical Report No. UCB/EECS-2018-23

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2018/EECS-2018-23.html>

May 1, 2018



Copyright © 2018, by the author(s).  
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

### Acknowledgement

Parts of the the results presented in this report is a joint work with Mitya Chklovskii, Cengiz Pehlevan, and Bin Yu and has been published as a paper in the IEEE Proceedings of 50th Asilomar Conference on Signals, Systems and Computers. The authors would like to thank Michael Berry, Jason Prentice and Mark Ioffe for providing the dataset and related instructions. Bin Yu acknowledges partial research support from National Science Foundation (NSF) Grant CDS/E-MSS 1228246, and the Center for Science of Information, an NSF Science and Technology Center, under Grant Agreement CCF-0939370.

# Population Coding in Retinal Ganglion Cells

by

Reza Abbasi Asl

A research project submitted in partial satisfaction of the

requirements for the degree of

Master of Science

in the department of

Electrical Engineering and Computer Sciences,

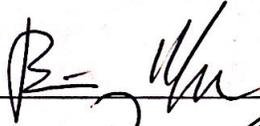
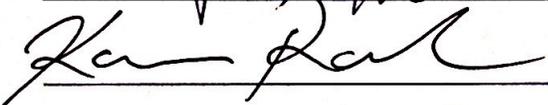
University of California, Berkeley

Committee in charge:

Professor Bin Yu, Chair  
Professor Kannan Ramchandran

Spring 2018

The dissertation of Reza Abbasi Asl, titled Population Coding in Retinal Ganglion Cells, is approved:

Chair   


Date 4/19/18

Date 4/26/18

University of California, Berkeley

## Abstract

Population Coding in Retinal Ganglion Cells

by

Reza Abbasi Asl

Master of Science in Electrical Engineering and Computer Sciences,

University of California, Berkeley

Professor Bin Yu, Chair

Several theories of early sensory processing suggest that it whitens sensory stimuli. Here, we test three key predictions of the whitening theory using recordings from 152 ganglion cells in salamander retina responding to natural movies. We confirm the previous finding that firing rates of ganglion cells are less correlated compared to the pixels in naturalistic stimulus, although significant correlations remain. We show that while the power spectrum of ganglion cells decays less steeply than that of natural scenes, it is not completely flattened. Finally, retinal ganglion cells reduce the dimensionality of the photoreceptor signal, as evidenced by their significantly lower numbers. Whitening theory predicts that only the top principal components of the visual stimulus should be transmitted. By comparing the right singular vectors of stimulus and firing rate covariance matrices, we find evidence that supports this prediction.

# Contents

<b>Contents</b>	<b>i</b>
<b>List of Figures</b>	<b>ii</b>
<b>1 Introduction and problem setup</b>	<b>1</b>
<b>2 Experimental results</b>	<b>3</b>
2.1 Dataset and notation . . . . .	3
2.2 Retinal ganglion cells outputs are not perfectly uncorrelated . . . . .	4
2.3 Power spectrum of retinal ganglion cell firing rate . . . . .	6
2.4 Retinal output transmits dynamics in the principal subspace of visual scenes . . . . .	6
<b>3 Conclusion</b>	<b>13</b>
<b>Bibliography</b>	<b>14</b>

# List of Figures

- 2.1 Output of ganglion cells are decorrelated compared to natural scenes. A. A sample frame from stimulus. B. Estimated receptive field for an On and an Off ganglion cell. C. Pearson correlation coefficient of firing rates for all ganglion cell-type pairs as a function of the distance between their receptive field centers. The curves corresponds to median correlation coefficient in a distance window of 1 pixel. Similar correlation coefficient and curve are shown for stimulus. . . . . 5
- 2.2 Power spectrum of ganglion cells are not flattened. A. Eigenvalues of stimulus and ganglion cells firing rate covariance matrices. Eigenvalues are normalized by the largest eigenvalue and are shown in log scale. The index is also normalized by the largest value. A line can be fitted on each plot with slope -4.7 and -5.3, respectively. B. Scatter plot of top 152 eigenvalue pairs. A line can be fitted to this plot with slope 1.09. . . 7

- 2.3 Input-output singular vector pair identification. A. Red curve shows correlation coefficient between the input-output right singular vector pair with the highest peak value among other pairs. Other curves show correlation coefficients between right singular vectors for 100 repeats of randomly time-permuted data (see text for more details). In this plot, the output singular vector is shifted by -30 to +30 time samples. B. Heatmap of peak correlation coefficients for all of the possible input-output pairs. The singular vectors of input and output belonging to the top 50 matches are ordered differently to maximize the matched pair correlation coefficients. For details of the sorting procedure see text. C. Same as B for randomly permuted data. D. Red: Histogram of correlation coefficients for 152 identified pairs from original data. Blue: Histogram of highest peak correlation coefficient for 100 randomly permuted trials. A Gaussian function can be fitted on this histogram. The threshold is set to be three standard deviation from mean of these 100 points. This leads to the selection of 31 pairs with significant peak correlation coefficient. . . . . 9
- 2.4 A. Scatterplot of top selected eigenvalues and fitted linear function. The slope of the fitted line is 0.08. B. Probability of selecting each right singular vector by the selection process (see text for more details) for stimulus matrix. Singular vectors are ordered by values of the corresponding singular values (high-to-low). C. Same as B for ganglion cell firing rate matrix. . . . . 10

# Chapter 1

## Introduction and problem setup

Efficient coding theories have been extremely influential in understanding early sensory processing in general [3, 4] and, visual processing in the retina in particular [16, 9, 2, 1, 17, 8, 6]. Many efficient coding theories predict that the retina should whiten sensory stimuli [7]. For instance, whitening is optimal from the predictive coding theory perspective, which postulates that only the unpredicted part of a signal should be transmitted. An optimal linear predictive coding filter is a whitening filter, and when applied to the retina this theory predicts that the receptive fields of retinal ganglion cells whiten visual stimuli [16, 9]. Redundancy reduction theory proposes that redundant information transmission at the retinal output should be minimized by decorrelated firing of ganglion cells [2]. Information theory, too, favors whitening. A linear filter that maximizes mutual information between noisy inputs and noisy outputs is a whitening filter [17, 8, 10]. An influential take on this line of thinking proposes that redundant information transmission at the retinal output should be avoided [4] by decorrelating the visual stimulus [2]. Experimental tests have so far been limited to one central prediction of these theories, that retinal ganglion cell outputs should be decorrelated at high-contrast. It was found that retinal ganglion cells are indeed decorrelated compared to strong spatial correlations of natural scenes [13], but significant redundancy remains [15]. Finally, whitening was proposed by another normative theory [11], whose starting point was efficient dimensionality reduction through similarity matching [12].

Here, using electrophysiological recordings from 152 ganglion cells in salamander

retina responding to natural movies, we test three key predictions of the whitening theory. 1) Retinal ganglion cells should be uncorrelated. We confirm previous tests of this prediction finding that retinal ganglion cells are indeed decorrelated compared to strong spatial correlations of natural scenes [13], but significant redundancy remains [15]. 2) Output power spectrum should be flat. Whereas we observe flattening of the output compared to the input, it is not completely flat. 3) If the number of channels are reduced from the input to the output, the top input principal components are transmitted.

We start by detailing the dataset we used for this work. Then, we present our results, which test the three predictions. Our results show that perfect whitening is not the whole story, consistent with findings of [15, 13].

# Chapter 2

## Experimental results

### 2.1 Dataset and notation

The dataset was collected by Michael Berry’s laboratory at Princeton University. It consists of recorded responses from 152 larval tiger salamander retinal ganglion cells to a 7-minute natural movie stimulus. The stimulus consists of gray scale movie of leaves and branches blowing in the wind (figure 2.1-A) and is projected onto the array from a CRT monitor at 60 Hz. The frame size of the stimulus is  $512 \times 512$ . Details of recording procedure, including the spike-sorting algorithm, are described in [14].

In this thesis,  $\mathbf{X}$  and  $\mathbf{Y}$  denote the stimulus and the firing rate matrices respectively.  $\mathbf{X}$  is a  $n \times T$  matrix, where  $n$  is the total number of pixels in each frame (reshaped to form a vector) and  $T$  is the number of time samples. In our setting, for computational purposes, stimulus frames are down-sampled to  $n = 128 \times 128 = 16,384$  and the 7-minute movie shown at 60 Hz corresponds to  $T = 25,200$ . Moreover,  $\mathbf{Y}$  is the  $k \times T$  firing rate matrix, where  $k = 152$  is the number of ganglion cells that data is recorded from. The firing rates are reported based on spike counts in sliding 80 ms time windows. Both  $\mathbf{X}$  and  $\mathbf{Y}$  are centered in time.

## 2.2 Retinal ganglion cells outputs are not perfectly uncorrelated

We start by testing whether output of ganglion cells are indeed decorrelated. This question was to a large extent answered by previous studies [15, 13], and here we confirm their findings.

Our goal is to study the correlation between firing rates of ganglion cells as a function of distance between their receptive field center, therefore, it is necessary to first estimate the receptive field of each cell. We use regularized spike-triggered average also known as regularized reverse correlation technique [5] to estimate the spatial receptive field of each cell. In this technique, firing rate matrix is assumed to be a linear function of stimulus matrix,  $\mathbf{Y} = \mathbf{R}\mathbf{X}$ . Here,  $\mathbf{R}$  is the  $k \times n$  receptive field matrix for ganglion cells.  $\mathbf{R}$  can be estimated using ridge regression which has the following closed form:

$$\mathbf{R} = \mathbf{Y}\mathbf{X}^T(\mathbf{X}\mathbf{X}^T + \alpha\mathbf{I})^{-1}.$$

Here,  $\alpha$  is regularization parameter which we chose with 10-fold cross validation and  $\mathbf{I}$  is identity matrix. Examples of the estimated On and Off receptive fields are shown in figure 2.1-B. The center of receptive field for each cell is simply identified as the location with maximum (for On cells) or minimum (for Off cells) value of estimated receptive field.

Figure 2.1-C shows the Pearson correlation coefficient (CC) of firing rates for all ganglion cell-type pairs as a function of the distance between their receptive field centers. This is depicted for On-On, Off-Off, and On-Off pairs. To show the trend of each group, the median correlation coefficient in a distance window of 1 pixel is also shown. We compare this correlation with the spatial correlation in the stimulus. Consistent with Pitkow and Meister's results [13], firing rates of ganglion cells are less correlated than the pixels in naturalistic stimulus. However, significant correlations remain, questioning the validity of the whitening theory (see also [15]).

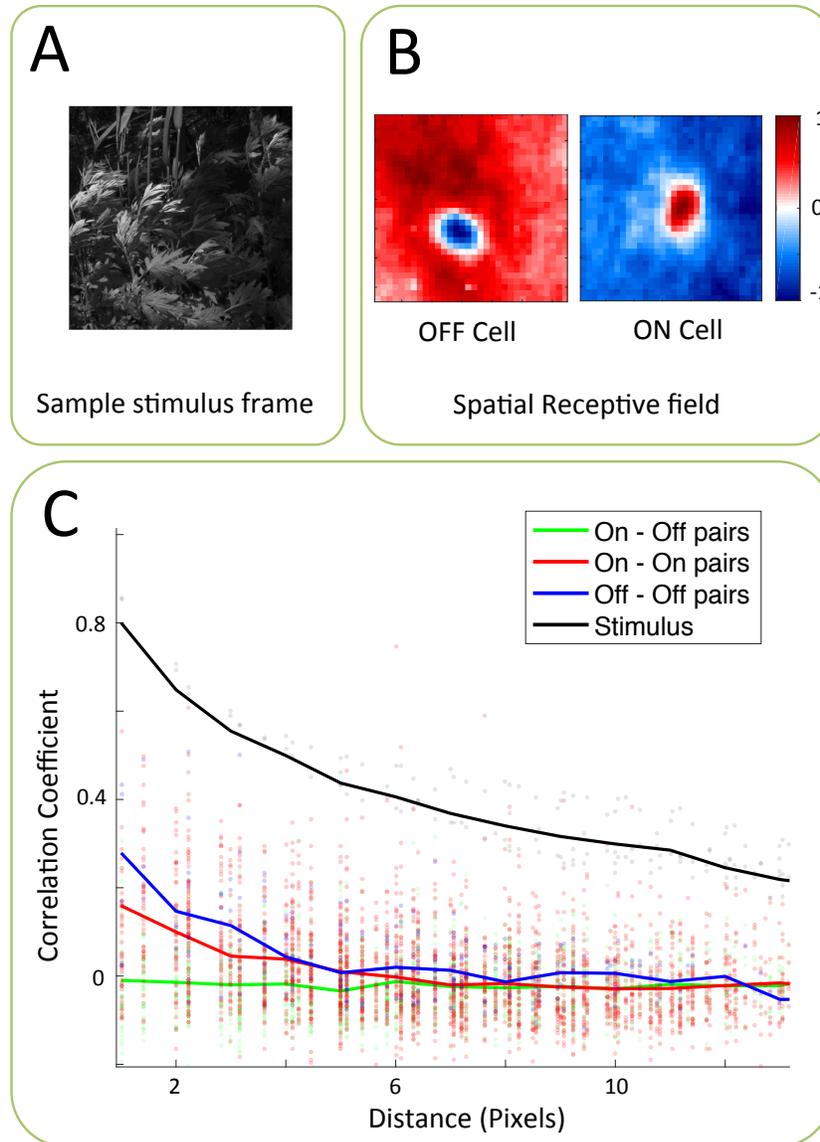


Figure 2.1: Output of ganglion cells are decorrelated compared to natural scenes. A. A sample frame from stimulus. B. Estimated receptive field for an On and an Off ganglion cell. C. Pearson correlation coefficient of firing rates for all ganglion cell-type pairs as a function of the distance between their receptive field centers. The curves corresponds to median correlation coefficient in a distance window of 1 pixel. Similar correlation coefficient and curve are shown for stimulus.

## 2.3 Power spectrum of retinal ganglion cell firing rate

Next we test whether the power spectrum of ganglion cells is flattened. In Figure 2.2-A, we show the eigenvalues of the ganglion cell firing rate covariance matrix, normalized by the highest eigenvalue. We also show similar plot for the visual stimuli covariance matrix for comparison. We observe a decay in the power spectrum of the firing rates.

To compare the rate of power spectrum decay, we plot the sorted eigenvalues of the firing rate covariance matrix against top 152 eigenvalues of natural stimuli covariance matrix, Figure 2.2-B, and observe that a line can be fitted to this plot with slope 1.09, suggesting again that the decay is similar (A perfectly equalized output spectrum would lead to a slope of 0). This analysis assumed that in the scatter plot of Figure 2.2-B, the correct correspondence between stimulus and firing-rate covariance eigenvalues is in order of their magnitude. It is possible that due to noise, we misestimate the rank-order of eigenvalues. We explore this possibility in more detail below.

## 2.4 Retinal output transmits dynamics in the principal subspace of visual scenes

Finally we test whether retinal ganglion cell outputs represent a projection of natural scenes to their principal subspace. To be mathematically precise, we want to test if

$$\mathbf{Y} = \mathbf{O}\mathbf{D}\mathbf{P}\mathbf{X},$$

where  $\mathbf{P}$  is a  $k \times n$  matrix with  $k < n$ , rows of  $\mathbf{P}$  are the principal eigenvectors of the stimulus covariance matrix,  $\mathbf{D}$  is a diagonal, nonnegative,  $k \times k$  matrix, which could be rescaling power in each component for e.g. whitening, and  $\mathbf{O}$  is an arbitrary  $k$ -dimensional orthogonal matrix. Since we do not know  $\mathbf{O}\mathbf{D}$ , we do not have a direct prediction for what  $\mathbf{Y}$  should be.

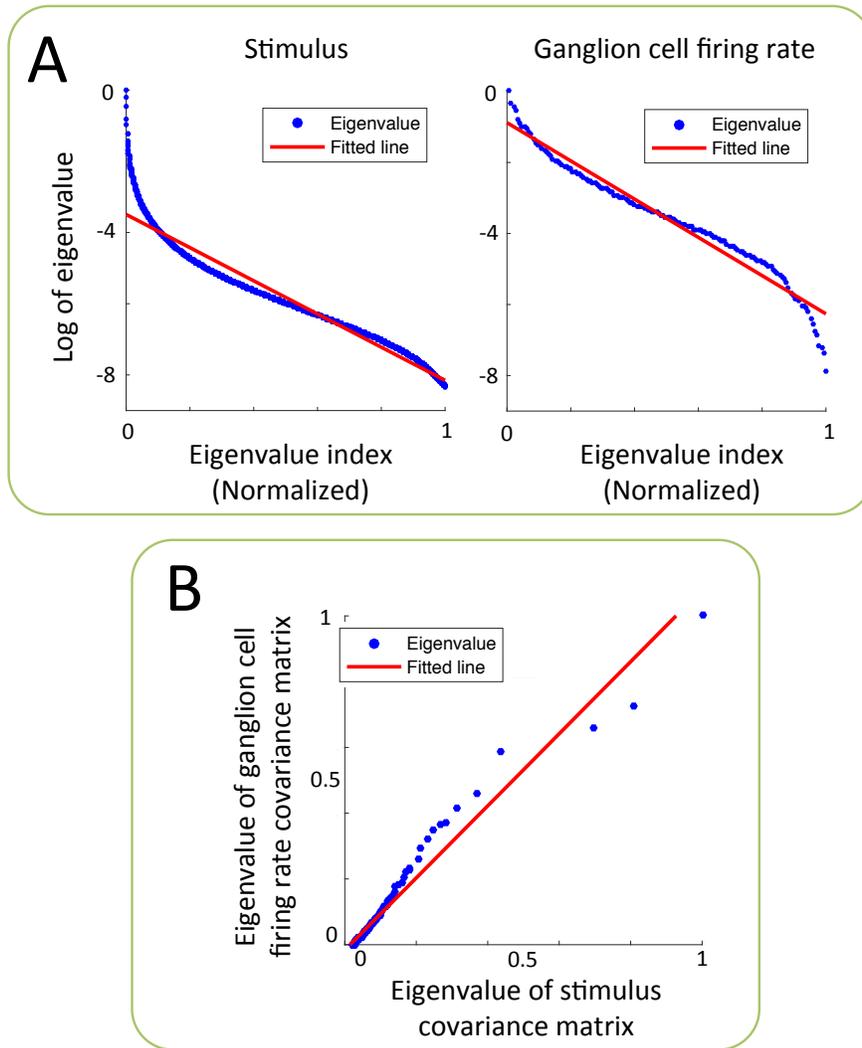


Figure 2.2: Power spectrum of ganglion cells are not flattened. A. Eigenvalues of stimulus and ganglion cells firing rate covariance matrices. Eigenvalues are normalized by the largest eigenvalue and are shown in log scale. The index is also normalized by the largest value. A line can be fitted on each plot with slope  $-4.7$  and  $-5.3$ , respectively. B. Scatter plot of top 152 eigenvalue pairs. A line can be fitted to this plot with slope  $1.09$ .

However, a prediction can be made for the right-singular vectors of  $\mathbf{Y}$ . Suppose a singular-value decomposition (SVD) for  $\mathbf{X} = \mathbf{U}_X \mathbf{\Lambda}_X \mathbf{V}_X^T$ , where singular values are

ordered in decreasing order. Then,

$$\mathbf{Y} = \mathbf{ODPU}_X \Lambda_X \mathbf{V}_X^T.$$

But, rows of  $\mathbf{P}$  are the first  $k$  columns of  $\mathbf{U}_X$ , and hence

$$\Lambda_Y := \mathbf{DPU}_X \Lambda_X$$

is a  $k \times N$  diagonal matrix. Then,  $\mathbf{Y} = \mathbf{R}\Lambda_Y \mathbf{V}_X^T$  is an SVD, and we can claim that the top  $k$  right singular vectors of  $\mathbf{X}$  and  $\mathbf{Y}$  (columns of  $\mathbf{V}_X$ ) should match.

Of course the match will not be perfect due to noise, which may effect our analysis in two ways. First, our estimate of the ordering of singular vectors will be imperfect. Second, the singular vector estimates will be imperfect.

Therefore, we came up with the following procedure. We calculate the SVD of  $\mathbf{X} = \mathbf{U}_X \Lambda_X \mathbf{V}_X^T$  and  $\mathbf{Y} = \mathbf{U}_Y \Lambda_Y \mathbf{V}_Y^T$  from data. We match right singular vectors of  $\mathbf{X}$  and  $\mathbf{Y}$ , columns of  $\mathbf{V}_X$  and  $\mathbf{V}_Y$ , using CC as a similarity measure between these vectors. We compute CC for all possible pairs of input-output singular vectors and rank each pair based on this similarity measure.

To address the delay of neural response to stimulus, CC is computed for several time shifts between stimuli and firing rates and the peak value is considered as the similarity. Fig. 2.3-A (red curve) shows the CC between the pair of singular vectors with the highest peak value. There is a considerable correlation between the two vectors, when compared to peak CC values for randomly permuted data (see also below).

Finally we use a greedy sorting procedure to identify top pairs. We select the pair with the highest rank as the first pick. This pair corresponds to certain index of input and output singular vectors. We then remove all the other pairs coming from these indexes and repeat a similar process for the remaining pairs. Fig. 2.3-B shows the heatmap of similarity measures for sorted pairs. The diagonal structure implies that there is an approximate one-to-one mapping between the top sorted singular vectors.

To test the significance of our matching between singular values, we provide a reference by breaking the temporal relationship between stimuli and firing rates and redoing the similarity analysis. We randomly permute columns in  $\mathbf{Y}$  (neural

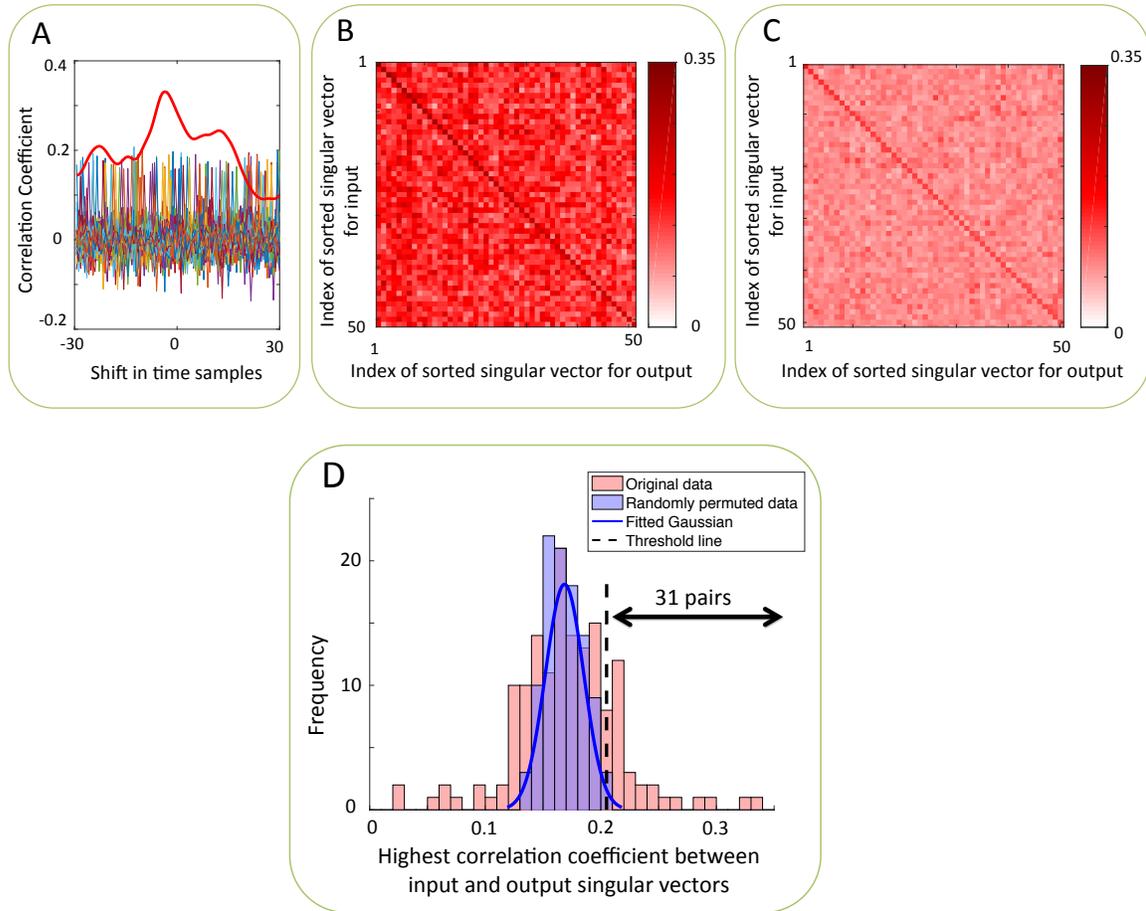


Figure 2.3: Input-output singular vector pair identification. A. Red curve shows correlation coefficient between the input-output right singular vector pair with the highest peak value among other pairs. Other curves show correlation coefficients between right singular vectors for 100 repeats of randomly time-permuted data (see text for more details). In this plot, the output singular vector is shifted by -30 to +30 time samples. B. Heatmap of peak correlation coefficients for all of the possible input-output pairs. The singular vectors of input and output belonging to the top 50 matches are ordered differently to maximize the matched pair correlation coefficients. For details of the sorting procedure see text. C. Same as B for randomly permuted data. D. Red: Histogram of correlation coefficients for 152 identified pairs from original data. Blue: Histogram of highest peak correlation coefficient for 100 randomly permuted trials. A Gaussian function can be fitted on this histogram. The threshold is set to be three standard deviation from mean of these 100 points. This leads to the selection of 31 pairs with significant peak correlation coefficient.

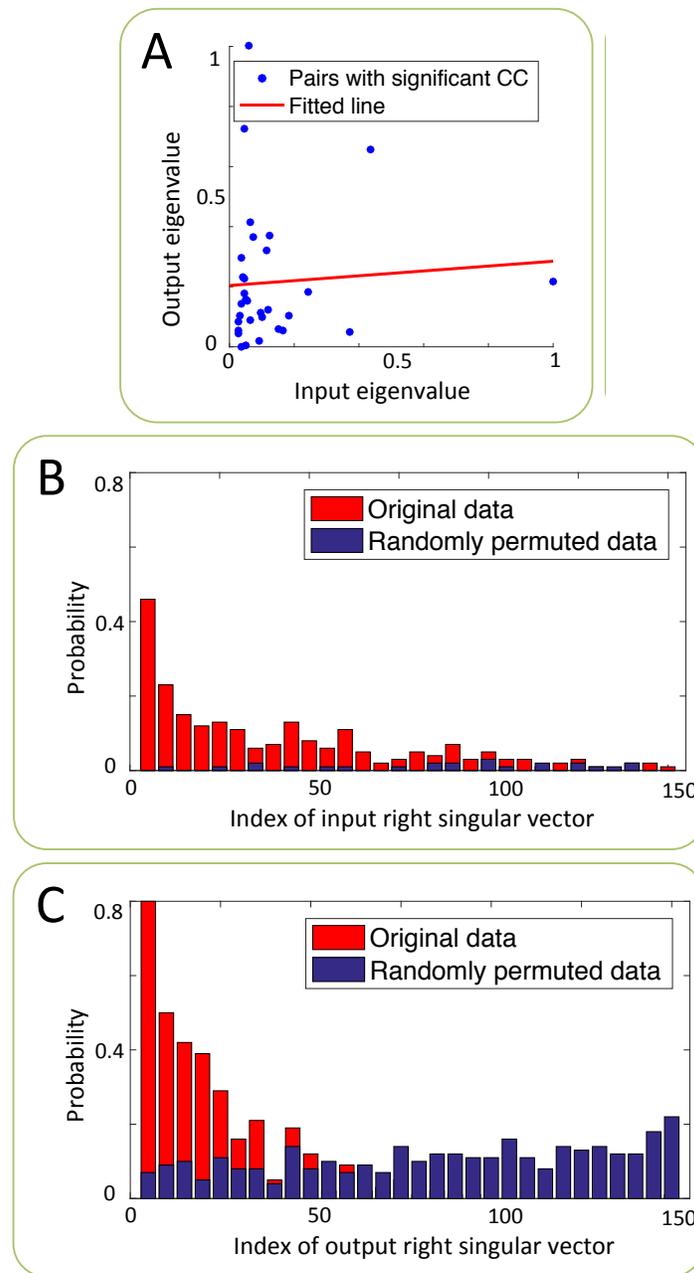


Figure 2.4: A. Scatterplot of top selected eigenvalues and fitted linear function. The slope of the fitted line is 0.08. B. Probability of selecting each right singular vector by the selection process (see text for more details) for stimulus matrix. Singular vectors are ordered by values of the corresponding singular values (high-to-low). C. Same as B for ganglion cell firing rate matrix.

response matrix) while we keep  $\mathbf{X}$  (stimulus matrix) the same. This dataset is called the *randomly permuted* data in this thesis. We redo the similarity analysis for 100 trials of this random permutation. Fig. 2.3-A shows 100 CC curves with the highest peak for each of these randomly permuted data. The red curve which corresponds to the original data has higher peak CC compared to other curves. Fig. 2.3-C shows the heatmap of peak CC values for one repeat of randomly permuted data. Comparing this figure to fig. 2.3-B, we observe that correlations in the original data are higher compared to the randomly permuted case.

To identify pairs with significantly higher peak CC in original data, it is necessary to find a threshold CC based on randomly permuted data. Figure 2.3-D shows the process to find this threshold. We first take the highest peak CC from 100 randomly permuted trials. Then, the threshold is set to be three standard deviations from mean of these 100 points. This leads to the selection of 31 pairs with significant peak CC. Out of these 31 input-output pairs, 12 was among the top 31 input singular values and 20 was among the top 31 output singular values.

To summarize our findings in this section up until now, 1) we found that there is a significant correlation between some pairs of stimulus and firing rate right singular vectors, 2) we can use these correlations to come up with a one-to-one mapping between the singular vectors, and 3) the one-to-one mapping we obtained suggests that the top principal components of the input is represented at the output, consistent with whitening theory.

After identifying the top input-output pairs of singular vectors, which also gives a mapping between covariance matrix eigenvalues, we go back to the question we asked in the previous section: does the power spectrum of retinal ganglion cells decay less steeply than the power spectrum of natural scenes? In Fig. 2.4-A, we show the scatterplot of the normalized covariance eigenvalues with our new matching and the fitted linear function. This should be compared to Figure 2.2-B. We found that the slope is lower in this case, 0.08, suggesting a flattening effect at the output.

To test the validity of our results, we perform further analyses. We estimate the probability of a singular vector being selected by our selection process (described above). First, we divide out the data into 20 non-overlapping segments and repeat the pair identification process for these subsets of data. The probability of selecting

each singular value is estimated by dividing the total number of selection for that singular value by 20. To provide a reference, we randomly permute columns in each of the neural response subset matrices while keeping stimulus subset matrix the same and re-estimate the probabilities. Figure 2.4-B and 2.4-C show the probability of selection for top 150 right singular vectors of the stimulus and firing rate matrices, respectively. We find that the top singular vectors, both for stimuli and firing rates, have a significantly higher than chance level probability of selection. These figures suggest that top right singular vectors of the stimulus matrix are highly expected to have a significant CC with the top right singular values of the ganglion cell firing rate matrix. This observation is consistent with the theory that ganglion cell activity represents a projection of visual stimulus onto its principal subspace.

## Chapter 3

### Conclusion

In this thesis, we put three key predictions of whitening theory to test. We found that 1) retinal ganglion cells outputs are not perfectly uncorrelated (also see[15, 13]), 2) the power spectrum of ganglion cell firing is not completely flat, but rather slopes downwards and 3) there are significant correlations between top right singular vectors of visual stimuli and retinal ganglion cell firing rates, suggesting that retinal output transmits dynamics in the principal subspace of visual scenes. Overall, our findings suggest that whitening alone cannot fully explain response properties of retinal ganglion cells.

# Bibliography

- [1] Joseph J Atick. “Could information theory provide an ecological theory of sensory processing?” In: *Network: Computation in neural systems* 3.2 (1992), pp. 213–251.
- [2] Joseph J Atick and A Norman Redlich. “What does the retina know about natural scenes?” In: *Neural computation* 4.2 (1992), pp. 196–210.
- [3] Fred Attneave. “Some informational aspects of visual perception.” In: *Psychological review* 61.3 (1954), p. 183.
- [4] Horace B Barlow. “Possible principles underlying the transformations of sensory messages”. In: *Sensory Communication*. Ed. by W.A. Rosenblith. Cambridge, MA: MIT press, 1961, pp. 217–234.
- [5] EJ Chichilnisky. “A simple white noise analysis of neuronal light responses”. In: *Network: Computation in Neural Systems* 12.2 (2001), pp. 199–213.
- [6] Eizaburo Doi et al. “Efficient coding of spatial information in the primate retina”. In: *The Journal of Neuroscience* 32.46 (2012), pp. 16256–16264.
- [7] Daniel J Graham, Damon M Chandler, and David J Field. “Can the theory of “whitening” explain the center-surround properties of retinal ganglion cell receptive fields?” In: *Vision research* 46.18 (2006), pp. 2901–2913.
- [8] M Haft and JL Van Hemmen. “Theory and implementation of infomax filters for the retina”. In: *Network: Computation in Neural Systems* 9.1 (1998), pp. 39–71.
- [9] Yanping Huang and Rajesh PN Rao. “Predictive coding”. In: *Wiley Interdisciplinary Reviews: Cognitive Science* 2.5 (2011), pp. 580–593.

- [10] Ralph Linsker. “Self-organization in a perceptual network”. In: *Computer* 21.3 (1988), pp. 105–117.
- [11] Cengiz Pehlevan and Dmitri Chklovskii. “A normative theory of adaptive dimensionality reduction in neural networks”. In: *Advances in Neural Information Processing Systems*. 2015, pp. 2269–2277.
- [12] Cengiz Pehlevan, Tao Hu, and Dmitri B Chklovskii. “A hebbian/anti-hebbian neural network for linear subspace learning: A derivation from multidimensional scaling of streaming data”. In: *Neural computation* 27 (2015), pp. 1461–1495.
- [13] Xaq Pitkow and Markus Meister. “Decorrelation and efficient coding by retinal ganglion cells”. In: *Nature neuroscience* 15.4 (2012), pp. 628–635.
- [14] Jason S. Prentice et al. “Error-Robust Modes of the Retinal Population Code”. In: *PLOS Computational Biology* 12.11 (Nov. 2016), pp. 1–34. DOI: 10.1371/journal.pcbi.1005148.
- [15] Jason L Puchalla et al. “Redundancy in the population code of the retina”. In: *Neuron* 46.3 (2005), pp. 493–504.
- [16] Mandyam V Srinivasan, Simon B Laughlin, and Andreas Dubs. “Predictive coding: a fresh view of inhibition in the retina”. In: *Proceedings of the Royal Society of London B: Biological Sciences* 216.1205 (1982), pp. 427–459.
- [17] JH Van Hateren. “A theory of maximizing sensory information”. In: *Biological cybernetics* 68.1 (1992), pp. 23–29.