Computational Approach to Identifying Contrast-Driven Retinal Ganglion Cells

Gault, R., Vance, P., McGinnity, T. M., Coleman, S., & Kerr, D. (2021). Computational Approach to Identifying Contrast-Driven Retinal Ganglion Cells. In *LNCS - Proceedings of the European Neural Network Society*

[Link to publication record in Ulster University Research Portal](#)

Published in:
LNCS - Proceedings of the European Neural Network Society

Publication Status:
Published (in print/issue): 14/09/2021

Document Version
Author Accepted version

General rights
Copyright for the publications made accessible via Ulster University's Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Ulster University's institutional repository that provides access to Ulster's research outputs. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact pure-support@ulster.ac.uk.

Computational Approach to Identifying Contrast-Driven Retinal Ganglion Cells^{*}

Richard Gault¹[0000–0001–6097–8981], Philip Vance²[0000–0002–2189–3608], T.M. McGinnity^{2,3}[0000–0002–9897–4748], Sonya Coleman²[0000–0002–4676–7640], and Dermot Kerr²[0000–0002–5077–0658]

¹ School of Electronics, Electrical Engineering and Computer Science, Queen’s University, Belfast, Northern Ireland, UK. richard.gault@qub.ac.uk

² Intelligent Systems Research Centre, Ulster University, Northern Ireland, UK. {p.vance, tm.mcginfinity, sa.coleman, d.kerr}@ulster.ac.uk

³ Dept. of Computer Science, Nottingham Trent University, Nottingham, UK martin.mcginfinity@ntu.ac.uk

Abstract. The retina acts as the primary stage for the encoding of visual stimuli in the central nervous system. It is comprised of numerous functionally distinct cells tuned to particular types of visual stimuli. This work presents an analytical approach to identifying contrast-driven retinal cells. Machine learning approaches as well as traditional regression models are used to represent the input-output behaviour of retinal ganglion cells. The findings of this work demonstrate that it is possible to separate the cells based on how they respond to changes in mean contrast upon presentation of single images. The separation allows us to identify retinal ganglion cells that are likely to have good model performance in a computationally inexpensive way.

Keywords: Retinal modelling · encoding natural images · identifying cell behaviour · visual modelling

1 Introduction

Visual perception begins with the encoding of visual stimuli into neuronal spikes by the retina. The retina receives few afferent signals from the central nervous system allowing the function of the retina to be studied as a black box with an input, i.e. light stimulus, and output, i.e. spiking activity of retinal ganglion cells (RGCs). The majority of RGCs can be classified as being either ON- and/or OFF-cells that are described as being transient or sustained [1]. ON-cells respond to an increase in light intensity while conversely OFF-cells respond to a

^{*} The authors would like to thank Jian Liu, Tim Gollisch and the “Sensory Processing in the Retina” research group at the Department of Ophthalmology, University of Göttingen who supplied the experimental data as part of the “VISUALISE” project funded under the European Union Seventh Framework Programme (FP7-ICT-2011.9.11); grant number [600954] (“VISUALISE”).

decrease in light intensity. A subdivision of each class can be made into transient and sustained responses. RGCs described as sustained produce continual firing in response to stimulus while transient cells respond only momentarily to temporal changes. For the investigation of how the retina encodes single images it is likely that sustained RGCs are of greatest interest to model the encoding using the corresponding spiking responses [1]. Many types of RGC have been documented for their selective response to particular types of stimuli, for example, direction specific motion, orientation or global environmental changes [1]. It is unlikely that cells with such diverse functionality are as important for single image processing as those with a sustained response.

The stimuli directly responsible for a RGC's activity are solely contained within the receptive field (RF) of the cell. Environmental changes or saccades of the eye mean that the stimuli that falls on the receptive field is constantly changing. It is therefore not surprising a significant proportion of RGC behaviour is directly coupled with temporal changes in environment. RGCs with transient responses are sensitive to subtle temporal changes in the RF but not stationary patterns [2]. In order to model the retinal encoding of single images it is important to identify RGCs whose behaviour is sustained over longer temporal periods. The ability to identify RGCs with sustained (or transient) responses to single images during the analysis stage would remove the need to use artificial stimuli which is known to not comprehensively probe the cell's functionality.

RGCs receive inputs from various networks of cells in the preceding layers of the retina which all contribute to forming the RF. Identification of the stimulus values that contribute to a cell eliciting a response is an essential stage in identifying the cells receptive field. Factorisation techniques have been used to infer potential sub-receptive fields (SRFs) that contribute to the overall RFs. Accurate identification of these sub-receptive fields and the underlying bipolar cells is essential in order to develop an accurate retinal model [3]. Models that consider the SRFs of the RGCs have been shown to provide improved accuracy when compared with approaches that consider the information contained in the entire RF as one single input [3]. SRFs allow detailed spatial information to be efficiently combined in a non-linear manner within RGCs when compared with single RF modelling approaches.

The aim of this work is to present an approach that can efficiently identify RGCs suitable for single image retinal models. Section 2 will outline the methodology used; this includes Sections 2.2 and 2.3 that outline the investigation of cell identification and the modelling approaches used for single image processing. The results are first presented in Section 3 before being discussed in detail in Section 4. Section 5 provides a discussion of the findings of this work and an outline of directions for future work.

2 Methods

Multi-electrode array recordings from isolated Tiger Salamander retinas were prepared and conducted at the Gollisch Lab, University Medical Center Göttingen

as outlined and described in [3–5]. The stimulus set comprised of a wide range of three hundred grayscale natural images from the “McGill Calibrated Colour Image Database” (<http://tabby.vision.mcgill.ca/html/browsedownload.html>) which were presented to the isolated retinas at a resolution of 256 x 256 pixels. Images were presented in a pseudo-random order for 200ms with an inter-stimulus interval of 800ms allowing sufficient time for no overlap in the cells response to different images. Spikes occurring within 300ms of stimulus onset were considered to be evoked by the stimulus whilst later activity was considered to be evoked by the removal of the stimulus and are ignored in this study.

The overall data analyses pipeline for the present study is comprised of identifying the RF of each cell and subsequently modelling and optimising SRFs for each cell. The input-output relationship that maps the information contained in the stimulus image to the resulting spiking activity is determined and modelled using the RF and SRFs. These components are described in detail through the remainder of Section 2.

The initial RF estimate was determined using a reverse correlation approach. Each retina was stimulated with spatio-temporal checker stimuli and the resultant spike-triggered average (STA) calculated. Singular Value Decomposition (SVD) is used to extract the spatial element of the STA over time [6]. The size, location and shape of a RGC’s RF is approximated by fitting a 2D Gaussian distribution to the extracted data. Artificial stimuli are known to not capture the complexity of natural images and it has been shown that RGCs produce a different RF response when stimulated with natural images [7, 8]. Therefore in this study the RF is refined using a subset of natural images based on the approach outlined in [9].

We then apply a number of machine learning methods to derive input-output models of the RGC using RFs that are comprised of different numbers of SRFs, and then compare the models’ responses against the actual real RGC neuron’s response. Results gathered from experiments involving several RGCs provide quantitative evidence on the benefits of considering SRFs of RGCs when deriving computational models rather than considering the complete RF as a summation of all its parts.

2.1 Modelling sub-receptive fields

Considering the stimulus values contained within the RF as a singular value, such as the mean contrast, results in the loss of much of the spatial information contained within this RF [10]. A better approach is to consider the RF as composed of a number of SRFs and corresponding singular values from each of these SRF regions. Here, the SRFs are characterised using two separate approaches. The first approach is a straightforward geometric approximation of the RF into several equal sectors. While analytically simple, this approach is in no way biologically plausible. Each SRF is obtained by segmenting the elliptical RF along its circumference and intersecting these points in the centre of the RF. Each of the resulting sectors that the elliptical RF is now composed is considered to be a SRF in such a way that there is approximately equal pixel coverage across each

Fig. 1: Illustration of a geometric segmentation of a RF in to 2,...,8 SRFs

of the SRFs. An illustration of this approach is shown in Figure 1. No emphasis is placed on the particular arrangement or orientation of each SRF as no prior knowledge of the underlying biological SRF is known.

The second approach to identify SRFs is aimed at approximating the receptive fields of the bipolar cells in a more biologically plausible way using non-negative matrix factorisation (NNMF) [3]. The pixels of 300 natural images (each image being 256×256) are restructured into a matrix, X , of dimension $256^2 \times 300$. The NNMF methodology allows for the dimension of the original problem to be reduced by approximating X as follows

$$X \approx FY \quad (1)$$

where F is a $256^2 \times n$ non-negative dictionary matrix of n factors and Y is the $n \times 300$ expansion matrix of weights. This technique is naturally suited for decomposing grayscale images as the original images and the corresponding factors contain only non-negative values. Each column vector of F can be restructured in to a 256×256 pixels image by reversing the process used to originally restructure the natural images into column vectors. The NNMF was carried out using the MATLAB function `nnmf` with alternating least-squares approach. The number of non-negative factors, n , equated to the number of SRFs being modelled.

Given the factors of each cell, now restructured to a resolution of 256×256 pixels, the SRFs are identified by fitting 2D-Gaussian distributions to each factor in a similar way to the original RF approximation. In the case of the geometric approach, 100% of the receptive field is covered by SRFs while the NNMF approach cannot guarantee such coverage. A Genetic Algorithm (GA) is used to optimise the size of the NNMF generated SRFs to maximise their coverage within the RF. The sizes of the SRFs are constrained such that at least 60% of the SRF must be located within the original RF and no two SRFs may overlap by more than 30% of their individual size. An example of these optimised SRFs (where the number being modelled is 1-8 SRFs) is shown in Figure 2 for one particular RGC.

2.2 Assessing Input-Output relationship

Information present in grayscale images can be represented in many different ways. In this study the mean contrast of each SRF contained within the RF is

Fig. 2: Example of a cell's optimised SRFs (red) of the full RF (green) for 1- 8 SRFs displayed relative to the 256×256 pixel stimuli size

considered in line with previous studies [9, 11]. The mean contrast is defined as

$$C_{RF} = \frac{M_{RF} - M_{gray}}{M_{gray}} \quad (2)$$

where M_{RF} is the mean intensity of the RF and M_{gray} is the mean intensity of the entire image. The output is considered to be the average spike count of each cell in response to each image. To assess the significance of the input-output relationships for each cell, the input-output pairs are fitted with a linear regression model. This provides an efficient approach to describe the significance of the input-output relationships. The gradient of the linear least-squares fit quantifies the rate of change in output with respect to change in the input quantity. Therefore, a large gradient (either positive or negative) indicates there is a significant change in the output behaviour relative to the input feature. Conversely, little or no gradient indicates that similar outputs are found for varying input values. An illustration of a large gradient is shown in the fitted data of Figure 3a whilst an example where there is a weak input-output gradient is shown in Figure 3b. It is clear from the example in Figure 3b that if varying input values create the same output that the input feature of the image does not represent the driving force behind the cell's behaviour.

Fig. 3: Illustration of the linear least-squares fit (gray) to an example of a strong input-output relationship (a) and a weak input-output relationship (b) for two different cells across all input images; in these cases the input property is the mean contrast of each image.

2.3 Modelling retinal behaviour

Inspired by previous modelling approaches for natural image processing [11–13] a number of different machine learning techniques are considered in this study; namely a Multi-Layer Perceptron (MLP) [14], Bayesian Regularised Neural Networks (BRNN) [15], Support Vector Regression (SVR) [16] and k-Nearest Neighbours (k-NN) [17]. Methodological details of each model can be found in the associated literature [18]. The data are randomly separated into training (80%) and testing (20%) sets which remained constant across all modelling approaches. Each model is subject to 5 fold cross-validation and evaluated on the test data in each fold. In a practical sense the MLP is implemented in MATLAB by first constructing a network with 10 neurons in the hidden layer (using `fitnet`) before training the network using the training function `trainlm`. The trained network is then evaluated using the unseen testing data. The BRNN model is implemented in a similar way to the MLP case with the exception that the training function used is `trainbr`. The SVR model is fitted to the training data using the function `fitrsvm` with a radial basis kernel function. The k-NN model is implemented using the `knnsearch` function with the distance metric being the Euclidean distance. The optimum value for k was determined through an exhaustive search of possible values for k between 1 and 100 with k chosen to minimise the model’s error.

Additionally, more classical modelling approaches are used as computationally efficient alternatives to the machine learning approaches outlined above; namely linear regression and non-linear (quadratic) modelling. The linear model is fitted using an ordinary least-squares approach utilising the `fitlm` function in

MATLAB. The quadratic model is fitted using the MATLAB function `fitnlm` using a unitary initial guess and a constant error based variance model.

3 Performance Evaluation

To evaluate the effectiveness of each model in representing the behaviour of the RGCs in response to natural images, the coefficient of determination, or R^2 measure, is used. Given the model's predictions, p , compared with the true RGC output, t , over the 300 images the R^2 measure defined as

$$R^2 = 1 - \frac{\sum(t - p)^2}{\sum(t - \bar{t})^2} \quad (3)$$

where $\bar{t}$ denotes the mean of the true spike rates t over all images. The model is considered to perfectly match the real world behaviour when $R^2 = 1$ and the R^2 metric is bounded below by 0.

Figure 4 shows the average performance across all cell models when a threshold is placed on the absolute value of the input-output least-squares gradient (described in Section II-B). The input is considered to be the mean contrast for the entire receptive field of each cell. The y -axis indicates the average model performance across all cells with an absolute input-output least-squares gradient of at least the threshold indicated on the x -axis.

Figure 4 illustrates that when using the mean contrast as model input we can observe a positive relationship between the gradient of the least-squares fit and the average R^2 across all the models. A gradient threshold may be used to divide the cells into two groups for further analysis. This removes the need for individual analysis of each cell to consider those where spiking response is unlikely to be correlated to differing mean contrasts within the RF. The threshold is chosen to be indicative of a significant input-output gradient. Therefore, for the remainder of this analysis the upper-quartile of the input-output gradients is selected as the threshold (this represents the median value of the upper portion of the gradients).

The second stage of the analysis is concerned with modelling performance when considering additional spatial information through the use of SRFs. The 138 cells are now separated into two groups. Firstly the absolute input-output gradient of all cells is considered. The upper-quartile of these data is taken to be the differentiation point. Group 1 consists of those cells with a minimum absolute input-output gradient of the upper-quartile (3.61) and Group 2 consists of cells with an absolute input-output gradient below this threshold. When considering cells from Group 1 (35 cells) all models perform significantly better (at least $R^2 = 0.4$) compared with the inclusion of all cells (at most $R^2 = 0.28$) as seen in Figures 5 and 4 respectively. Conversely, when the cells in Group 2 (103) are considered in isolation (Figure 6) the cell model performance is greatly reduced ($R^2 \in [0.1, 0.3]$) compared with the performance observed when modelling Group 1 cells ($R^2 \in [0.4, 0.72]$). This is true irrespective of the modelling technique or the number of SRFs considered. Thus, it has been possible to identify RGCs

Fig. 4: Average performance for each model given cells with a minimum input-output gradient listed along the x -axis.

Fig. 5: The average model performance across all Group 1 cells for varying numbers of geometrically defined SRFs with an absolute input-output gradient of at least 3.61.

suitable for the processing contrast-driven natural images through the division of the entire cell dataset into two groups prior to modelling.

Fig. 6: The average model performance across all Group 2 cells for varying numbers of geometrically defined SRFs with an input-output gradient of less than 3.61.

As the geometric approach to SRF modelling does not yield biologically realistic SRF of the cells, the NMF approach to modelling SRF is considered on those cells in Group 1. The results are illustrated in Figure 7. The results show again that including SRF models improves modelling performance compared with modelling only a single RF irrespective of the modelling approach used. However, comparing each model in turn leads in general to the unexpected finding that NMF defined SRFs (results shown in Figure 7) did not provide better performance than the geometrically defined counterparts (results shown in Figure 5). One possible explanation could be that SRFs derived through NMF did not cover all of the RF and thus may omit some spatial information in contrast to the geometric approach to SRF modelling.

4 Discussion

The primary aim of this work was to investigate whether it is possible to identify relevant RGCs for processing natural images i.e cells that predominantly respond to variations in mean contrast. Analysis indicated that model performance was improved with an increase in the absolute gradient of the input-output least-squares fit (Figure 4). The simple process of fitting a linear least-squares line to the input-output data provided the basis for data to be segmented into two groups. It was determined that cells with a large absolute gradient were most likely to provide the best model performance. It is appropriate to separate RGCs by functionality as different cell types are known to have disjoint pathways. We

Fig. 7: The average model performance across all cells for varying numbers of NNMF defined SRFs with an absolute input-output gradient of at least 3.61.

postulate that the proposed approach is a step towards an analytical approach for classifying cell functionality.

It is well known that modelling a RGC Receptive Field using a singular value results in a loss of spatial information. Vance et al. [10] have shown that using finer grained spatial information led to improved model performance compared with single RF models (Figures 5 - 7). The NNMF constructed SRFs provide improved modelling performance over single RFs models (Figure 7); however, surprisingly the results from of the NNMF approach did not, in general, improve upon the results from the geometrically defined SRFs (Figure 5). Comparing the geometric and NNMF approaches, it can be deduced that it is possible to produce models with at least the accuracy of the bio-inspired NNMF SRF methodology with a relatively efficient and simple geometric approach. Further work is required to ascertain the reason why the geometric approach produced the best results for SRF modelling despite having no immediately apparent biological connection to the SRFs of the RGCs.

A number of machine learning and regression modelling approaches were considered in this work. Figures 5 to 7 show the regression approaches (illustrated with dotted lines) performed similar to the machine learning approaches. The quadratic model performed better than the linear model irrespective of the number of SRFs taken into account. This is unsurprising as the RGCs are known to respond in a non-linear way to stimuli [3]. In the case of NNMF derived SRF models (Figure 7) the quadratic model outperformed all other modelling approaches irrespective of the number of RFs modelled. Using the geometrically derived SRFs the quadratic model provided the overall best performance com-

pared with all other models with the maximum number of SRFs. Amongst the machine learning approaches SVR performed, in general, the strongest across different numbers of SRFs whilst MLP consistently performed poorly. It should be noted that all modelling approaches, with the possible exception of the linear model, performed similarly and showed that an accurate model of single image processing can be achieved once appropriate cells are identified and at least 6 SRFs are considered (Figures 5 and 7).

5 Conclusion

The encoding of visual information is carried out by a variety of functionally diverse cells whose response is driven by a particular characteristic of the stimuli. To create bio-inspired models of retinal processing of visual information this work aimed to identify a method for generating a subset of functionally similar retinal cells appropriate for single image processing. A computationally efficient linear fit of the input-output relationship of a stimulus attribute, such as mean contrast, and the RGC response, namely the mean firing rate, was used to identify cells whose behaviour is proportional to the stimulus attribute. It was possible to divide the cells into two groups using this information prior to modelling the RGC behaviour. Separating cells in this way allows us to identify those cells that could be modelled accurately. Therefore, it is postulated that the identified cells are appropriate for single image processing. Increasing the number of SRFs modelled in general led to increased model performance irrespective of the modelling approach used. The machine learning models performed comparably with the more computationally efficient quadratic model.

Future work is required to identify additional attributes of input stimuli that can be used to construct functionally homogeneous subgroups of RGCs such as those that respond to transient stimuli rather than stationary stimuli. Accurate modelling of RGC behaviour with distinct functionality could provide the building blocks to construct a complete representation of the encoding of visual information by the retina. Future work will also need to consider a greater repertoire of cell functionality including modelling the temporal encoding by RGCs compared with the firing rate models considered in the present work.

References

1. Onkar S Dhande, Benjamin K Stafford, Jung-Hwan A Lim, and Andrew D Huberman. Contributions of retinal ganglion cells to subcortical visual processing and behaviors. *Annual Review of Vision Science*, 1:291–328, 2015.
2. Tim Gollisch and Markus Meister. Eye smarter than scientists believed: neural computations in circuits of the retina. *Neuron*, 65(2):150–164, 2010.
3. Jian K Liu, Helene M Schreyer, Arno Onken, Fernando Rozenblit, Mohammad H Khani, Vidhyasankar Krishnamoorthy, Stefano Panzeri, and Tim Gollisch. Inference of neuronal functional circuitry with spike-triggered non-negative matrix factorization. *Nature Communications*, 8(1):149, 2017.

4. Liu, Jian K and Gollisch, Tim. Spike-triggered covariance analysis reveals phenomenological diversity of contrast adaptation in the retina. *PLoS computational biology*, 11(7):e1004435, 2015.
5. Arno Onken, Jian K Liu, PP Chamanthi R Karunasekara, Ioannis Delis, Tim Gollisch, and Stefano Panzeri. Using matrix and tensor factorizations for the single-trial analysis of population spike trains. *PLoS computational biology*, 12(11):e1005189, 2016.
6. Jeffrey L Gauthier, Greg D Field, Alexander Sher, Martin Greschner, Jonathon Shlens, Alan M Litke, and EJ Chichilnisky. Receptive fields in primate retina are coordinated to sample visual space more uniformly. *PLoS biology*, 7(4):e1000063, 2009.
7. Joaquin Rapela, Jerry M Mendel, and Norberto M Grzywacz. Estimating nonlinear receptive fields from natural images. *Journal of Vision*, 6(4):11–11, 2006.
8. Jon Touryan, Gidon Felsen, and Yang Dan. Spatial structure of complex cell receptive fields measured with natural images. *Neuron*, 45(5):781–791, 2005.
9. P. J. Vance, G. P. Das, D. Kerr, S. A. Coleman, T. M. McGinnity. Refining Receptive Field Estimates using Natural Images for Retinal Ganglion Cells. *IARIA, COGNITIVE 2016*, 77–82, 2016.
10. P. J. Vance, G. P. Das, D. Kerr, S. A. Coleman, T. M. McGinnity, T. Gollisch, and J. K. Liu. Bioinspired approach to modeling retinal ganglion cells using system identification techniques. *IEEE Transactions on Neural Networks and Learning Systems*, PP(99):1–13, 2017.
11. P. J. Vance, G. P. Das, S. A. Coleman, D. Kerr, E. P. Kerr, and T. M. McGinnity. Investigation into Sub-Receptive Fields of Retinal Ganglion Cells with Natural Images. *2018 International Joint Conference on Neural Networks (IJCNN)*, 1–8, 2018.
12. D Kerr, SA Coleman, and TM McGinnity. Modelling and analysis of retinal ganglion cells with neural networks. *Irish Machine Vision and Image Processing 2014*, 95–100, 2014.
13. Gautham Das, Philip Vance, Dermot Kerr, SA Coleman, and T Martin McGinnity. Modelling retinal ganglion cells stimulated with static natural images. *COGNITIVE 2016: The Eighth International Conference on Advanced Cognitive Technologies and Applications, Rome, Italy. IARIA 2016*.
14. Simon Haykin and Neural Network. A comprehensive foundation. *Neural Networks*, 2(2004):41, 2004.
15. David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992.
16. Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
17. Naomi S Altman. An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician*, 46(3):175–185, 1992.
18. Hastie, T. and Tibshirani, R. and Friedman, J.. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. *Springer Series in Statistics*, Springer New York, 2013.
19. Jonathan W Pillow, Liam Paninski, Valerie J Uzzell, Eero P Simoncelli, EJ Chichilnisky. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *Journal of Neuroscience*, 25(47):11003–11013, 2005.