Encoding the natural responses of primate retina

Nandita Bhaskhar[1], Ananth Saran Yalamarthy[2], and Arushi Arora[1]

{nanbhas, ananthy, arushi15}@stanford.edu

[1]Department of Electrical Engineering and [2]Department of Mechanical Engineering, Stanford University

1. Introduction

Can machine learning and neural networks help us model the eye's retina accurately? We explore 2 models: a ***Linear Nonlinear Poisson (LNP)*** model and a ***CNN*** to predict the neuronal response (spike trains) of 2 types of Retinal Ganglion cells (RGC) – ON & OFF RGCs, given visual stimuli.

2. Problem Formulation

For an RGC, given *m* training examples

$$\{x^{(i)}, y^{(i)}\},\ i \in \{1,2,\dots m\}$$

- Stimulus movie:
  $x^{(i)} \rightarrow$ frame size L x B, duration D (= N frames at refresh rate r)
- Observed spike train:
  $y^{(i)} = 1\{t_j \text{ observed a spike}\},\ j \in \{1, 2, \dots D/res\}$
- Predict:
  $\hat{y}$ for a test $x_{(test)}$ to minimise $err(y_{(test)}, \hat{y})$ for a reasonable error metric "*err*".

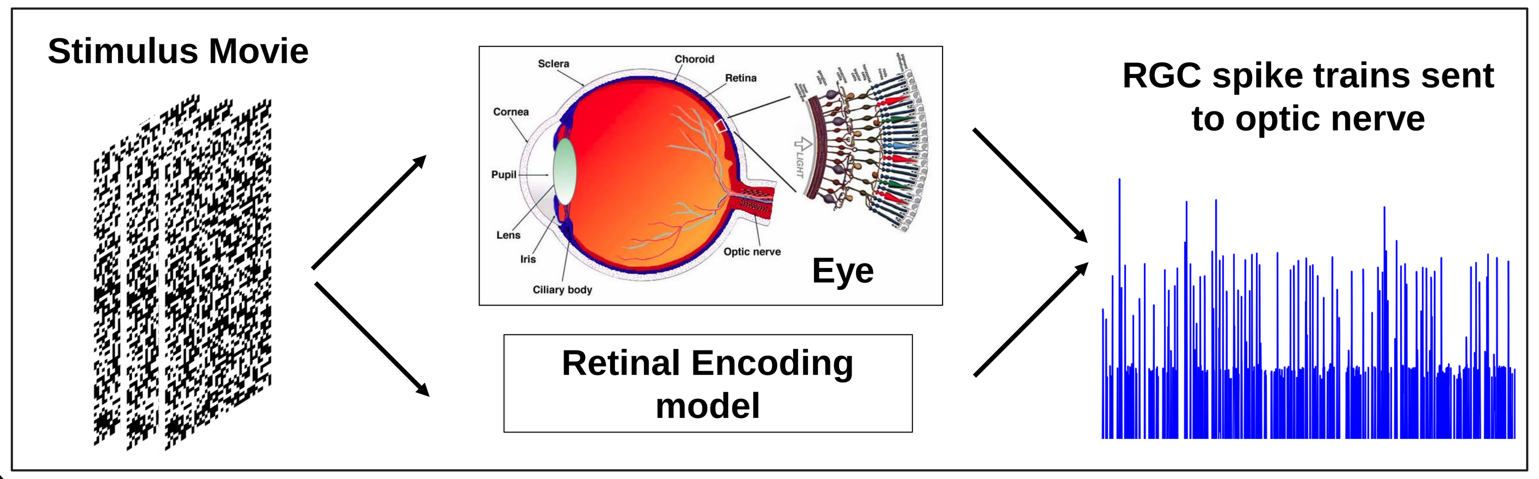

3. Data

- *Unique* multi-electrode recordings from large ensembles of RGCs in macaque monkey retina, visually stimulated with movies of white noise, shared generously by the ***Chichilnisky lab*** at Stanford.
- Gaussian binary white noise stimuli :– several 30s movies of (80 x 40) frame size with 120 Hz resolution .
- For every RGC exposed to this, there is an output vector containing the cell spike times
- A test movie stimulus is shown several times to obtain repeats of the cell responses. These are used to determine accuracy of prediction.

4. Linear Nonlinear Poisson Model

Consists of a linear block, followed by a non-linear estimate of the firing rate, subjected to a Poisson process for spike generation[1,2].

Stimulus → Linear → Generator → Nonlinear → Rate → Poisson → Spikes

5. Linear Block – STA

The Spike triggered average (STA) vector is given by

$$\omega = \left(\textstyle\sum_D s_t f_t\right) / \textstyle\sum_D f_t$$

- $s_t \rightarrow$ fixed fraction of the stimulus preceding time bin t
- $f_t \rightarrow$ the number of spikes observed in time bin t
- Spatial structure of ω → neuron's ***receptive field*** that isolates region of visual space that the neuron is active to.
- Temporal structure of ω → ***impulse response*** of the underlying linear summation

6. Nonlinear Block

Expected rate to the stimulus, is given by

$$R(s) = E[f \mid s] = \textstyle\sum_f f\, P(f|s)$$

- Simplifying assumption: R is a static non-linear function i.e. $R(s) = N(g)$ where $g = \omega\,.\,s$ is the *generator signal.*
- Estimation of N:
  - averaging f_t across time bins that give rise to the same generator signal g_t
  - By plotting $\bar{f}$ *vs* $\bar{g}$ for different values, the nonlinear function N is estimated as a ***sigmoid.***

7. Results

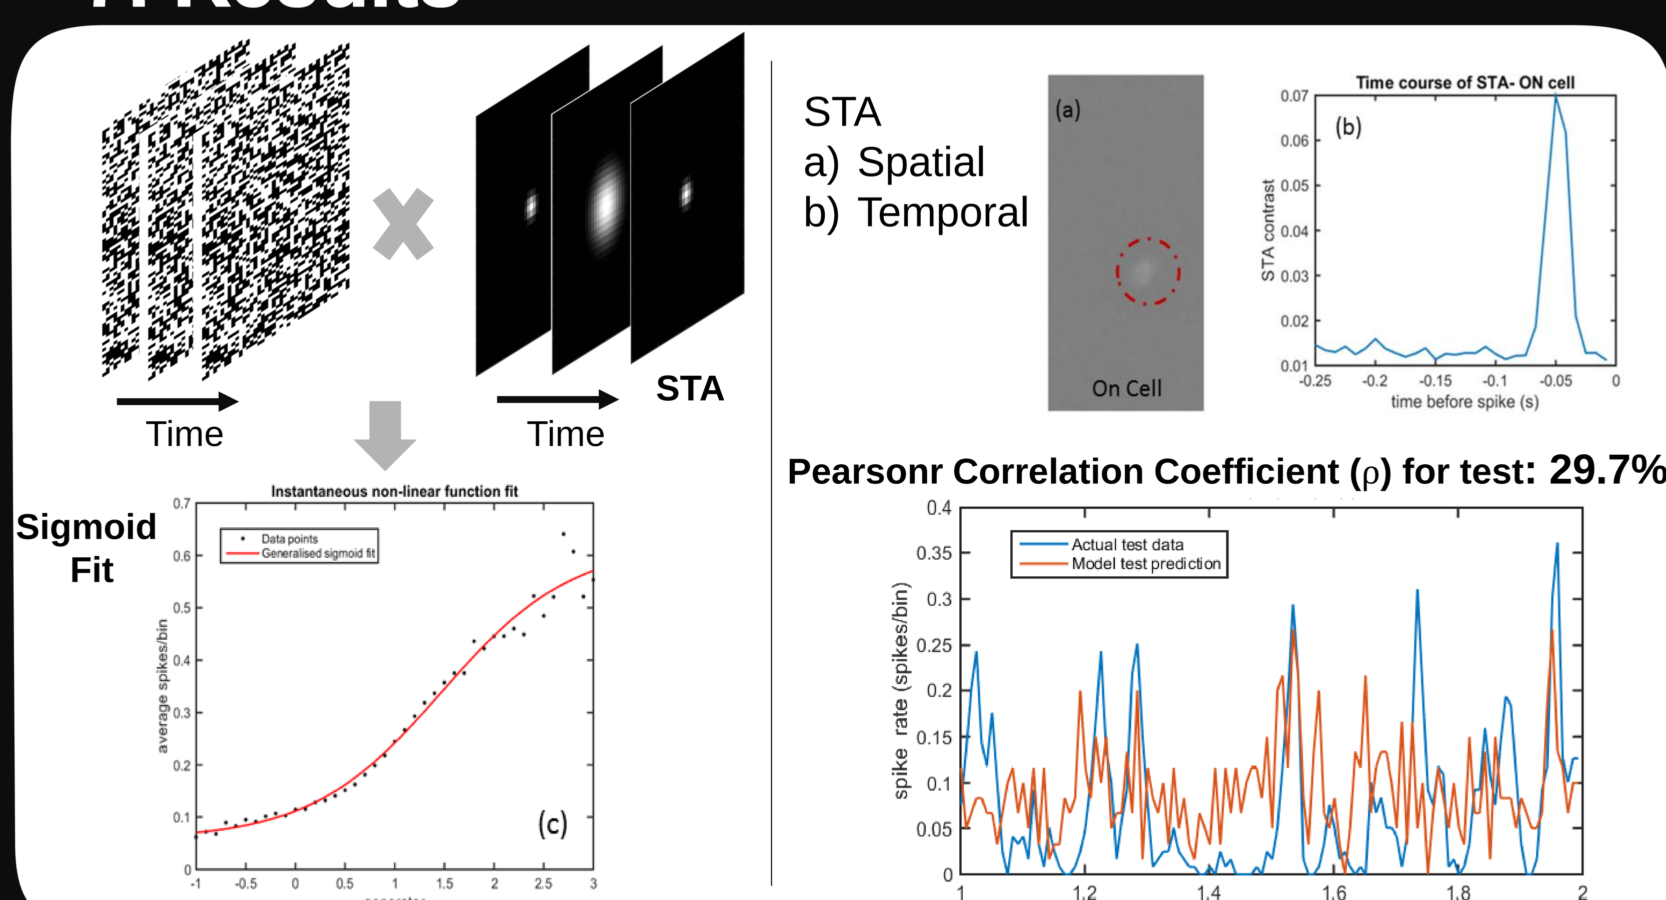

8. Deep Learning: CNN Architecture

- Experimented with RNNs & CNNs, chose CNNs.
- **Input:** 36k (6min) training examples, $80 \times 40 \times 10$ timesteps.
- **Output**: Single number that represents spiking probability, trained against Gaussian-convolved spike train (ground truth).

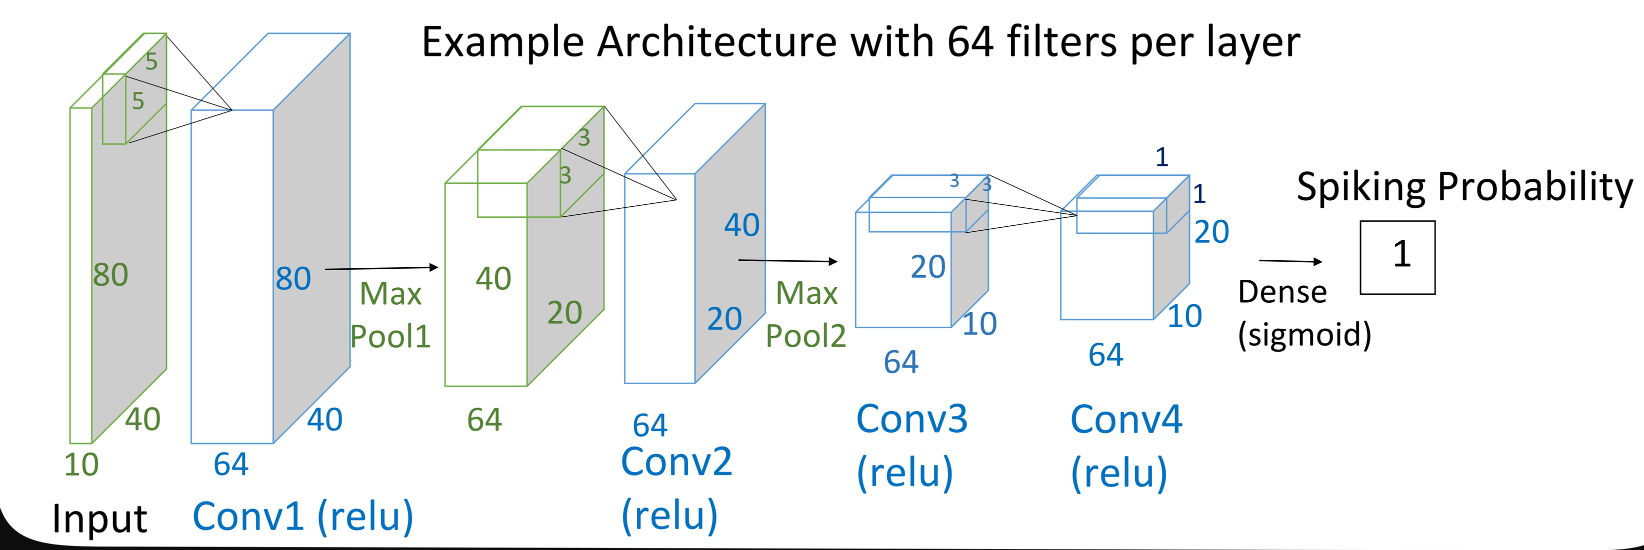

9. Deep Learning Results

Correlation coefficients for training data

Scale	Learning Rate (x1e-4) 10	5	1
0.2	0.98	0.993	0.44
	0.978	**0.996**	0.817
0.25	0.9795	0.991	0.478
	0.9845	0.935	0.355
0.3	0.989	0.947	0.368
	-0.012	-0.012	-0.012

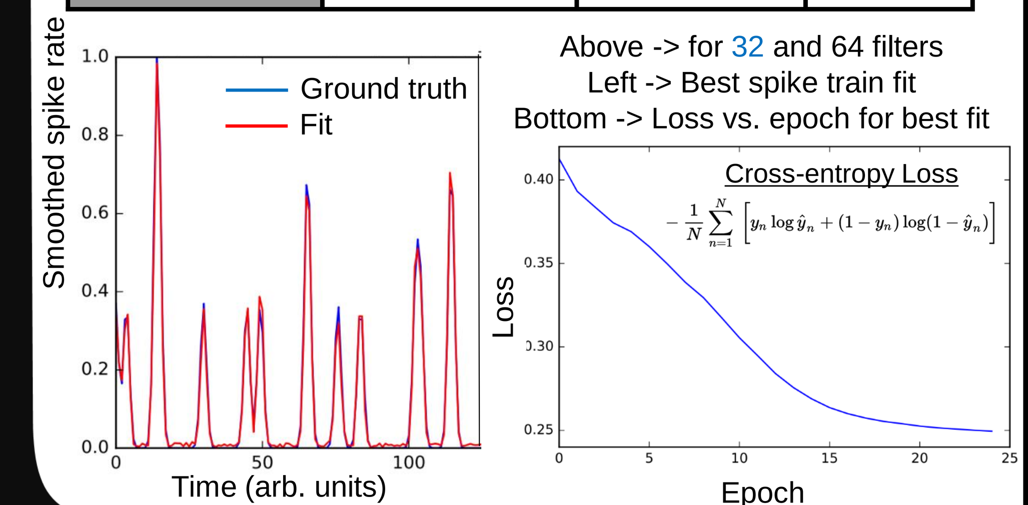

Above -> for 32 and 64 filters
Left -> Best spike train fit
Bottom -> Loss vs. epoch for best fit

Test data

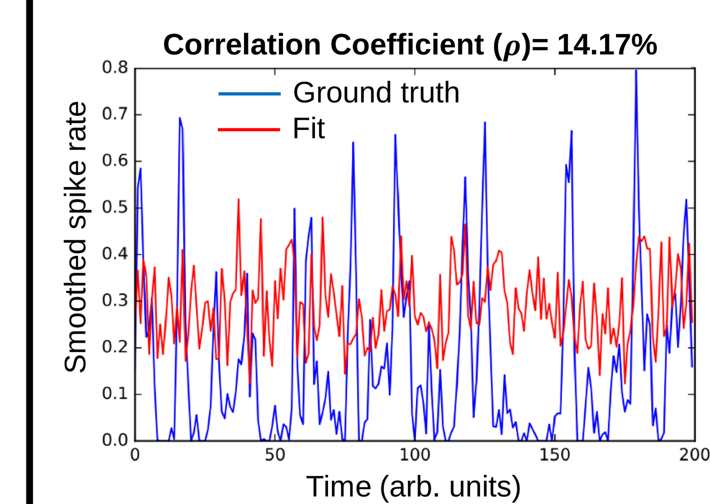

- Training data optimized hyper-parameters lead to ρ_{test}<10% due to overfitting.
- Test data optimized hyper-parameters gives maximum ρ of 14.2%

10. Discussion and Future Work

- The simplistic LNP model shows better test results compared to the CNN when tuned for minimum training error.
- CNN model needs to be tuned further to minimize test error.
- Possible architecture improvements:
  1. Add 3D convolution, including an LSTM layer in the end
  2. Explore RNNs further.
  3. Use a different RGC: literature shows[3] different RGCs give different ρ when trained on the same CNN model. It could be that we've used a neuron that has an intrinsically low ρ.

References and Acknowledgements

[1] Field GD, Chichilnisky EJ (2007) Information processing in the primate retina: circuitry and coding. Annu Rev Neurosci 30:1.f
[2] Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses, Comput. Neural Syst. 12, 199–213
[3] L. Macintosh, N. Maheshwaran, CS231n Stanford Project Report, Deep Learning Model of the Retina, Winter 2015, Stanford

We'd like to thank Prof EJ Chichilnisky for sharing the dataset with us and Nora Brackhill, Vishakh Hegde & Ashwin Paranjape for immensely useful discussions. Thanks to Bo (our TA) for feedback.