inferring retinal cone locations and functional connectivity from ganglion cell recordings

Kolia Sadeghi [1], Greg D. Field [2], Jeffrey L. Gauthier [2], Alexander Sher [3], Martin Greschner [2], Timothy Machado [1], Lauren H. Jepson [2], Jonathon Shlens [4], Deborah E. Gunning [5], Keith Mathieson [5], Alan M. Litke [3], and E.J. Chichilnisky [2], Liam Paninski [1]

Columbia University [1], Salk Institute for Biological Studies [2], Santa Cruz Institute for Particle Physics [3], Google [4], University of Glasgow [5]

introduction

Recently, it has been shown [3] that it is possible to get a handle on both the input and the output cells of macaque retinas, and the functional connections between them: if sufficiently fine-grained stimuli are used to excite the retina, the Spike Triggered Average receptive field of ganglion cells appear to be composed of small islands of light sensitivity, which are in fact the receptive fields of individual cones. Here, we address the problem of identifying the number, locations and types of cones in a way that provides information on how certain we can be of our inference. This is done using Markov Chain Monte Carlo (MCMC) on a familiar encoding model of ganglion cells where the functional weights have been integrated out. We obtain inferences of higher quality than with the greedy method used in [3].

stimulus

- Spatio-temporal 'white noise' : binary RGB
- High-resolution pixels: 5-6 micrometers wide
- 15-30Hz frames
- 30-240 min. experiments

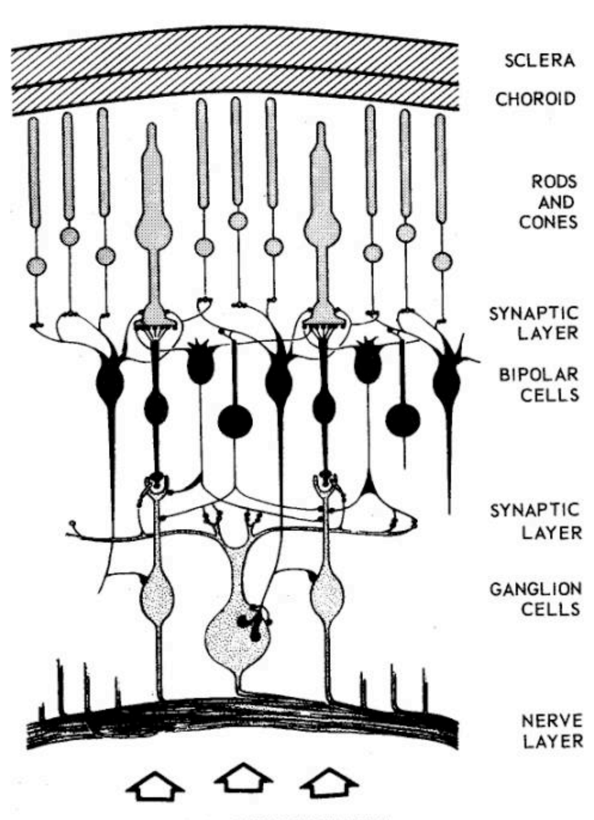

responses : 200 - 50,000 spikes / ganglion cell

Spike Triggered Averages

We assume STAs are separable in space and time. Consider the spatial components of a few STAs.

Low res. (downsampled) High resolution

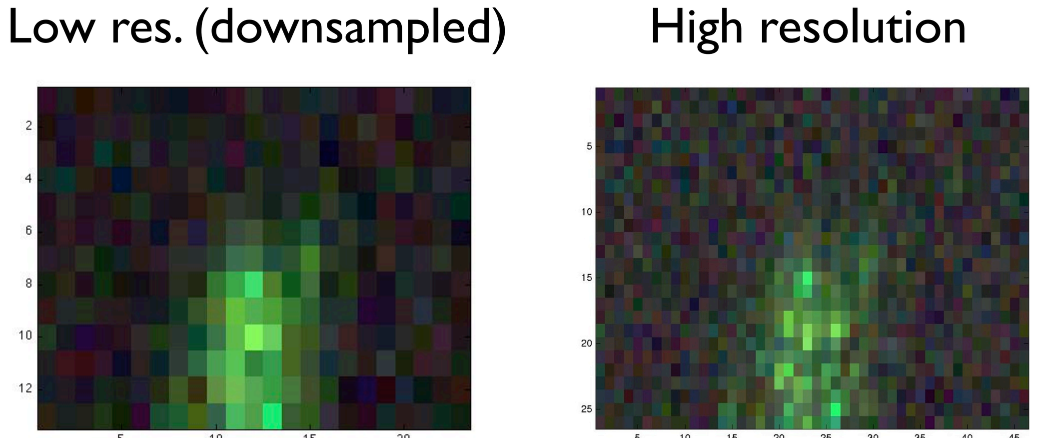

Small islands of light sensitivity are visible within the receptive field center, only in the high resolution STA.

[*G. Field, J. Gauthier, A. Sher, M. Greschner, T. Machado, L. Jepson, J. Shlens, D. Gunning, K. Mathieson, L. Paninski, A. Litke, and E.J. Chichilnisky , Nature 2010*] shows we are seeing cones.

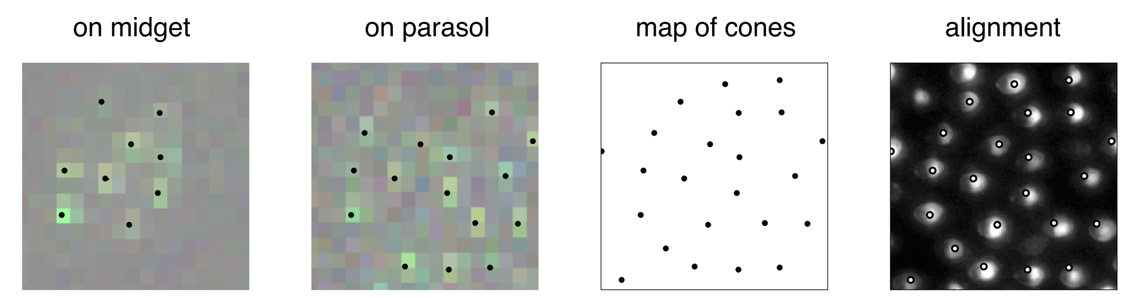

For cells with fewer spikes, are these cones or noise?

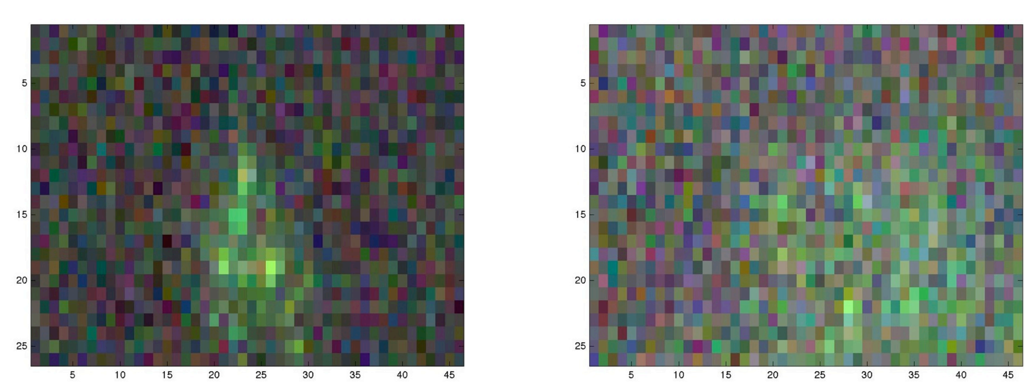

LNP model

$$n(t) \sim Poiss[\, e^{b+\mathbf{k}.\mathbf{s}_t}\, dt \,]$$

$n(t)$: spike train of a ganglion cell

$\mathbf{s}_t$: stimulus used to excite cells.

$b \in \mathbb{R}$: offset parameter.

$\mathbf{k}$: a linear filter acting on the stimulus $\mathbf{s}_t$.

$$\begin{aligned}\log p(\ n(t) \mid b,\ \mathbf{k},\ \mathbf{s}_t\ ) &= \sum_t n(t)(b+\mathbf{k}.\mathbf{s}_t) - e^{b+\mathbf{k}.\mathbf{s}_t} dt \ + const \\ &= N_{spikes}\ (b+\mathbf{k}.\mathbf{STA}) - \sum_t e^{b+\mathbf{k}.\mathbf{s}_t} dt \ + const\end{aligned}$$

gaussian stimulus approximation:

$$\begin{aligned}&\approx N_{spikes}\ (b+\mathbf{k}.\mathbf{STA}) \ - \ T \int e^{b+\mathbf{k}.\mathbf{s}} p(s) ds \ + \ const \\ &\approx N_{spikes}\ (b+\mathbf{k}.\mathbf{STA}) \ - \ T\ \exp\left(b + \frac{\sigma^2}{2}||\mathbf{k}||^2\right) \ + \ const\end{aligned}$$

Maximum likelihood b and $||\mathbf{k}||^2$: $k = \alpha\ STA$

$$\log p(\ n(t) \mid b, \alpha, \mathbf{s}_t\ ) \approx N_{spikes}(b + \alpha ||\mathbf{STA}||^2) \ - \ T\ \exp\left(b + \frac{\sigma^2\alpha^2}{2}||\mathbf{STA}||^2\right) \ + \ const$$

maximum-likelihood : $\begin{cases}\alpha = 1/\sigma^2 \\ b = \log\left(\frac{N_{spikes}}{T}\right) - \frac{||\mathbf{STA}||^2}{2\sigma^2}\end{cases}$

taylor expansion around max. likelihood value of $||\mathbf{k}||^2$:

$$\begin{aligned}\frac{1}{N_{spikes}} \log p(\ n(t) \mid \mathbf{k},\ \mathbf{STA}\ ) &\approx \mathbf{k}.\mathbf{STA} \ - \ \exp\left(-\frac{||\mathbf{STA}||^2}{2\sigma^2} + \frac{\sigma^2}{2}||\mathbf{k}||^2\right) \ + \ const. \\ &\approx \mathbf{k}.\mathbf{STA} \ - \ \frac{\sigma^2}{2}\,||\mathbf{k}||^2 \ + \ const.\end{aligned}$$

cones $\boxed{\mathbf{k} \ = \ \mathbf{W}(cones)\ \mathbf{a}}$

a : vector of functional weights between cones and ganglion cells

W : each column is a cone receptive field

both a and W depend on the set of cone locations and colors

log-likelihood

$$\boxed{\frac{1}{N_{spikes}} \log p(\ n(t) \mid \mathbf{W},\ \mathbf{a},\ \mathbf{STA}\ ) \approx \ \mathbf{STA}^T\,\mathbf{W}\,\mathbf{a} \ - \ \frac{\sigma^2}{2}\,\mathbf{a}^T\mathbf{W}^T\mathbf{W}\mathbf{a} \ + \ const.}$$

prior on weights

$$p(\,\mathbf{a} \mid cones\,) = \frac{1}{\sqrt{|2\pi(g\mathbf{W}^T\mathbf{W})^{-1}|}} \exp\left(-\frac{1}{2}\mathbf{a}^T g \mathbf{W}^T\mathbf{W}\mathbf{a}\right) \qquad g \ \text{determined by} \ ||\mathbf{STA}||/\sigma^2$$

integrating out the weights

$$\begin{aligned}p(\,data \mid cones\,) &= \int d\mathbf{a}\ p(\,\mathbf{a} \mid cones\,)\ p(\,data \mid cones\,, \mathbf{a}\,) \\ &\propto \int d\mathbf{a}\ \exp\left[\,N_{spikes}\left(\mathbf{STA}^T\,\mathbf{W}\,\mathbf{a} \ - \ \sigma^2\,\mathbf{a}^T\mathbf{W}^T\mathbf{W}\mathbf{a}\,/\,2\right) \ - \ g\,\mathbf{a}^T\mathbf{W}^T\mathbf{W}\mathbf{a}\,/\,2\,\right]\end{aligned}$$

$$\begin{aligned}2\ \log p(\,data \mid cones\,) &= \frac{N_{spikes}^2}{N_{spikes}\,\sigma^2 + g}\ \mathbf{STA}^T\,\mathbf{W}\,(\mathbf{W}^T\mathbf{W})^{-1}\,\mathbf{W}^T\,\mathbf{STA} \\ &\quad - \ \log\left|\frac{2\pi}{g}(\mathbf{W}^T\mathbf{W})^{-1}\right| \ + \ \log\left|\frac{2\pi}{N_{spikes}\,\sigma^2 + g}(\mathbf{W}^T\mathbf{W})^{-1}\right| \\ &= \frac{N_{spikes}^2}{N_{spikes}\,\sigma^2 + g}\ \mathbf{STA}^T\,\mathbf{W}\,(\mathbf{W}^T\mathbf{W})^{-1}\,\mathbf{W}^T\,\mathbf{STA} \ + \ N_{cones}\ \log\left(\frac{g}{N_{spikes}\,\sigma^2 + g}\right)\end{aligned}$$

$$\log p(\,data \mid cones\,) = \frac{1}{2}\sum_i \frac{N_{spikes_i}^2}{N_{spikes_i}\,\sigma^2 + g}\ \mathbf{STA}_i^T\,\mathbf{W}\,(\mathbf{W}^T\mathbf{W})^{-1}\,\mathbf{W}^T\,\mathbf{STA}_i \ + \ \frac{N_{cones}}{2}\ \log\left(\frac{g}{N_{spikes_i}\,\sigma^2 + g}\right)$$

$$\mathbb{E}\,(\,\mathbf{a}_i \mid cones\,,\,data\,) = \left[(N_{spikes_i}\,\sigma^2 + g)\mathbf{W}^T\mathbf{W}\right]^{-1}\ N_{spikes_i}\,\sigma^2\ \mathbf{W}^T\ \mathbf{STA}$$

$$\mathbf{Cov}\,(\,\mathbf{a}_i \mid cones\,,\,data\,) = \left[(N_{spikes_i}\,\sigma^2 + g)\mathbf{W}^T\mathbf{W}\right]^{-1}$$

hard cone exclusion prior

Cones cannot overlap in space: we place a hard exclusion prior on cone locations.

visualizing the evidence

Ignoring overlaps between cones due to pixelization gives: $W^TW = I$

$$\sum_i \frac{N_i^2}{N_i\,\sigma^2 + g}\,\mathbf{STA}_i^T\,\mathbf{W}\mathbf{W}^T\,\mathbf{STA}_i \longrightarrow$$

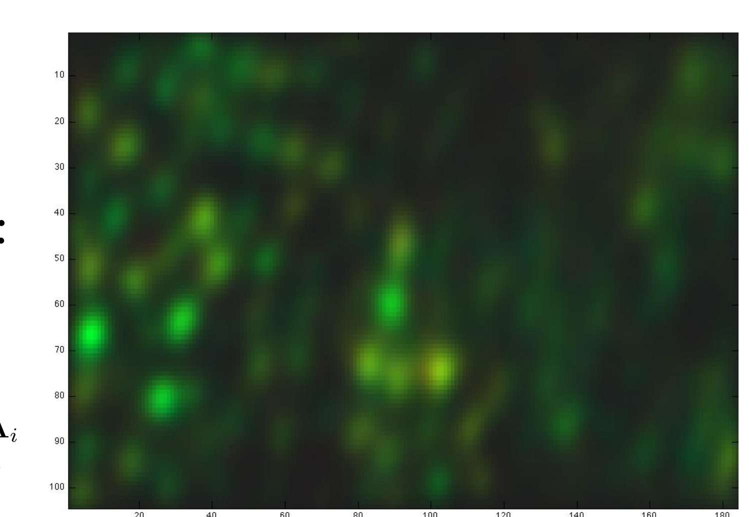

MCMC moves

- cone addition and deletion
- change of cone color
- shift of cone locations $\longrightarrow$

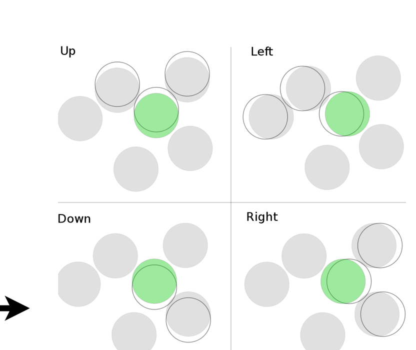

local maxima are a problem

Due to the hard cone exclusion prior and the strength of the evidence pooled across ganglion cells, MCMC with these moves rapidly gets stuck in local maxima.

Parallel Tempering

In order to overcome local maxima, we want to flatten the log-likelihood landscape, while sampling only from the true log-likelihood. This can be done by doing MCMC on a sequence of coupled MCMC instances with progressively increasing 'temperatures' β, δ .

$$\log p(\,data \mid cones\,,\,\beta,\delta\,) = \frac{1}{2}\sum_i \frac{N_{spikes_i}^2}{\beta(N_{spikes_i}\,\sigma^2 + g)}\left(\mathbf{STA}_i^T\,\mathbf{W}\,(\mathbf{W}^T\mathbf{W})^{-1}\,\mathbf{W}^T\,\mathbf{STA}_i\right)^{\delta} + \frac{N_{cones}}{2\beta}\ \log\left(\frac{g}{N_{spikes_i}\,\sigma^2 + g}\right)$$

Regular MCMC moves within each instance are alternated with swap moves which exchange a cluster of cones between neighboring configurations:

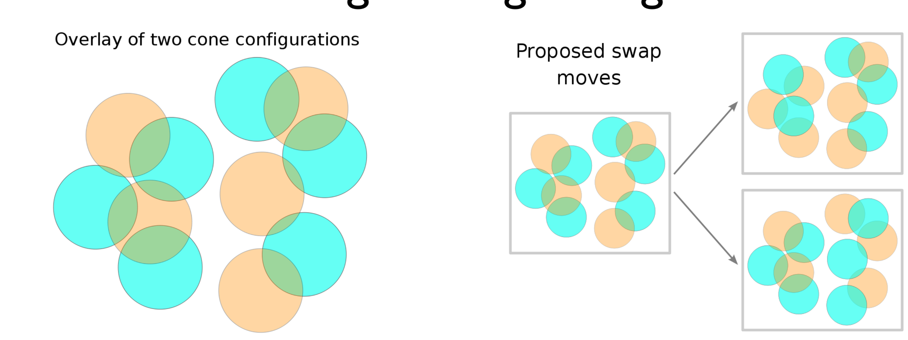

results

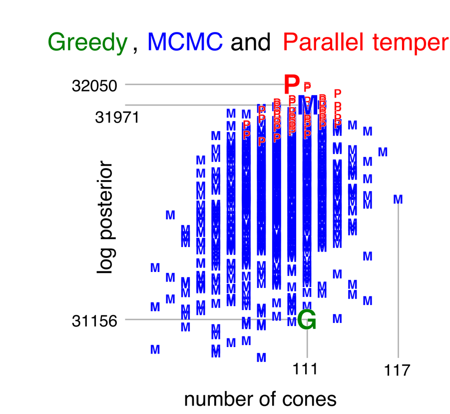

MCMC with Parallel tempering achieves higher likelihoods than regular MCMC and than the greedy method used previously.

The cones found by MCMC avoid some pitfalls of greedy optimization.

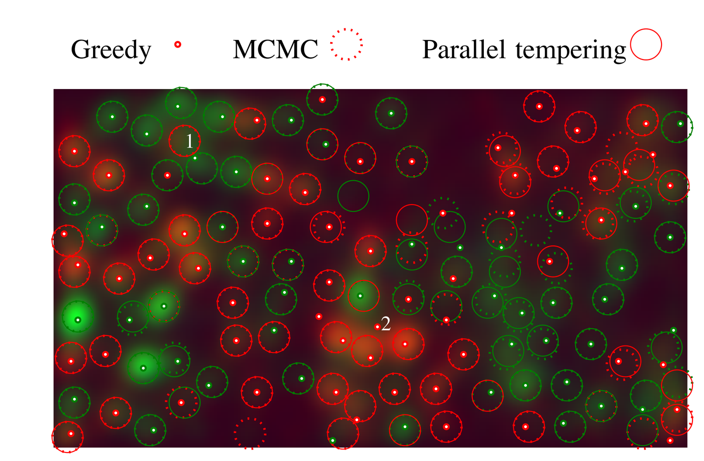

denoised STAs

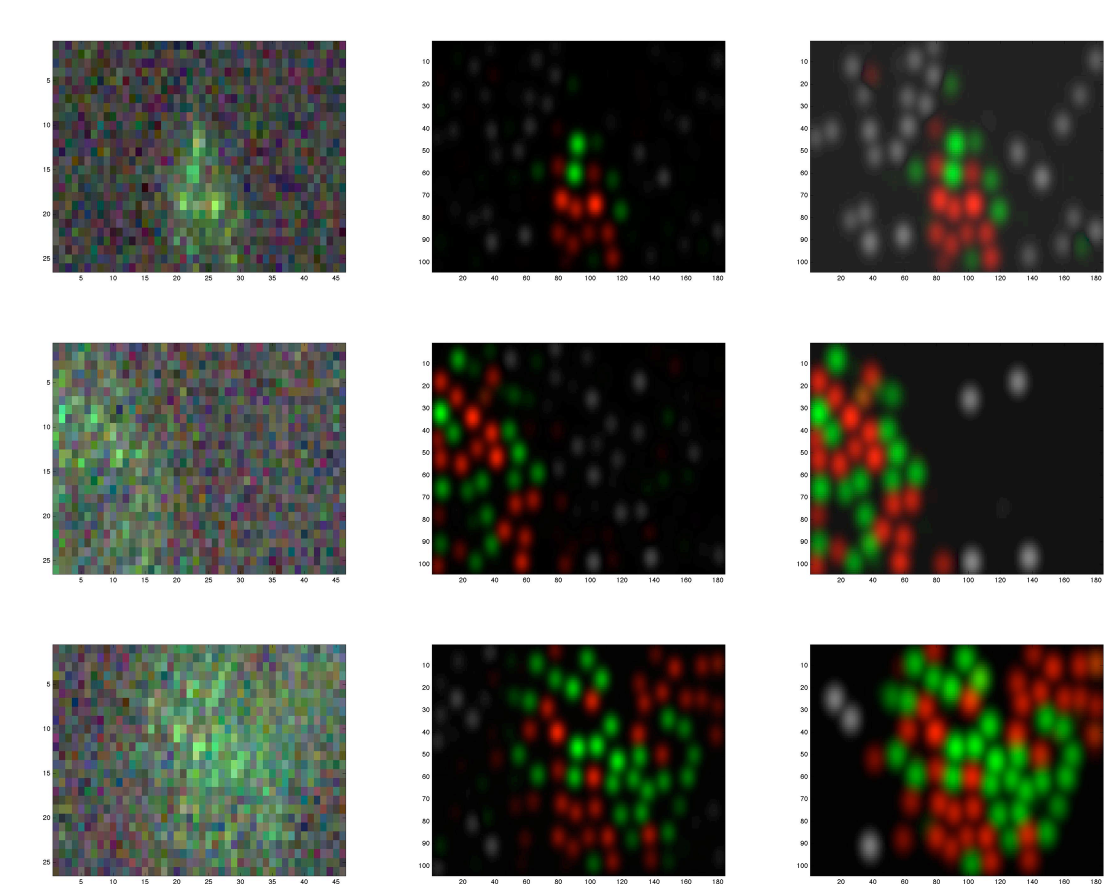

STA denoised STA w/ 2 color maps