Predicting Retinal Ganglion Cell Receptive Fields

based on material by Chris Williams & Mark van Rossum

Neural Information Processing
School of Informatics, University of Edinburgh

February 2018

Normative Modelling of the Visual System

Book: HHH [Hyvärinen et al., 2009] (free online) *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*, Springer 2009, chapter 1

- Normative vs Descriptive Theories: how *should* the system behave?
- Of course, this makes most sense if evolution has optimized the natural system. Effect of constraints
- "Statistical-ecological" approach

Chapter 10 of Dayan and Abbott (2001) is also useful.

Statistical-ecological approach

(HHH, p 21)

1. Different sets of features are good for different kinds of data.
2. The images that our eyes receive have certain statistical properties (regularities).
3. The visual system has learned a model of these statistical properties.
4. The model of the statistical properties enables (close to) optimal statistical inference.
5. The model of the statistical properties is reflected in the measurable properties of the visual system (e.g. receptive fields of the neurons)

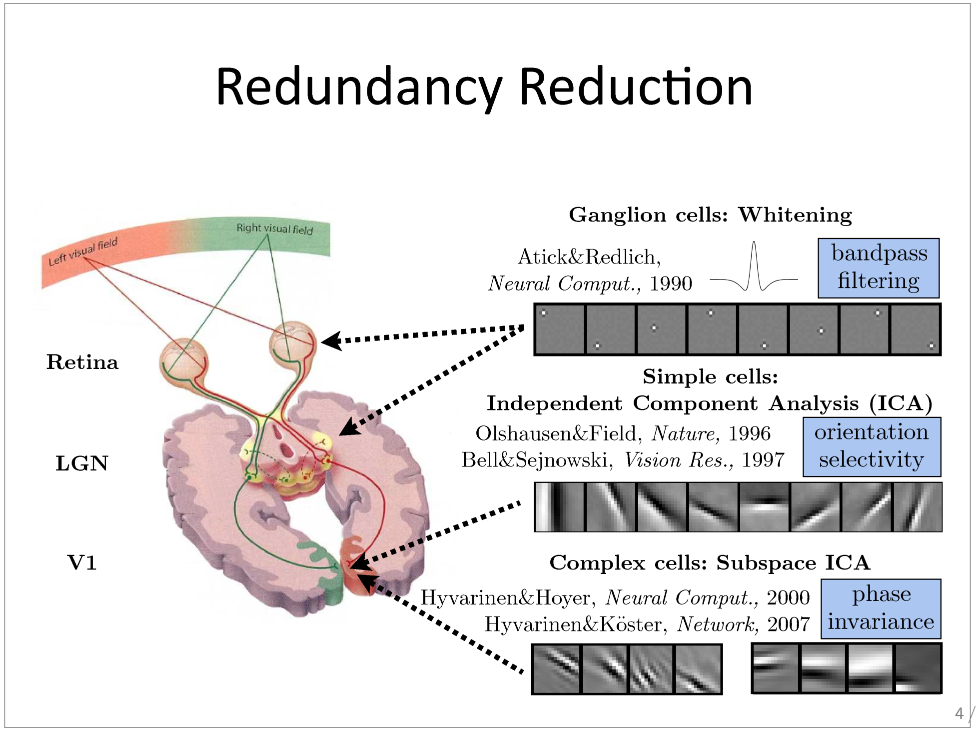

Mutual Informaton and Populations of Neurons

$$H(\mathbf{R}) = -\int p(\mathbf{r}) \log_2 p(\mathbf{r}) d\mathbf{r} - N \log_2 \Delta r$$

and

$$H(R_a) = -\int p(r_a) \log_2 p(r_a) d\mathbf{r} - \log_2 \Delta r$$

We have

$$H(\mathbf{R}) \leq \sum_a H(R_a)$$

(proof, consider KL divergence)

Recall that

$$I(\mathbf{R};\mathbf{S}) = H(\mathbf{R}) - H(\mathbf{R}|\mathbf{S})$$

so if noise entropy $H(\mathbf{R}|\mathbf{S})$ is independent of the transformation $S \to R$, we can maximize mutual information by maximizing $H(\mathbf{R})$ under given constraints

Factorial Coding

- Maximization of population response entropy is achieved by
  1. factorial coding $p(\mathbf{r}) = \prod_a p(r_a)$
  2. each response distribution must be optimized wrt the imposed constraints
- If all neurons have the same constraints $\Rightarrow$ probability equalization. This does not mean that each variable responds identically!
- Exact factorization and probability equalization are difficult to achieve
- A more modest goal is decorrelation (whitening)

$$\langle (\mathbf{r} - \langle \mathbf{r} \rangle)(\mathbf{r} - \langle \mathbf{r} \rangle)^T \rangle = \sigma_r^2 I$$

Second order statistics

- First order image statistics $\langle s(x,t) \rangle$
- Second order, correlation $Q(x,x',t,t') = \langle s(x,t)s(x',t') \rangle$
- By Wiener-Kinchin specifying Q is equivalent to specifying $PSD = |\tilde{s}(f)|^2$ (Wiener-Kinchin)
- Gaussian approximation $\Leftrightarrow Q(x,x') \Leftrightarrow PSD$
- Higher order statistics, e.g. $\langle s(x,t)s(x',t')s(x'',t'') \rangle$ will be discussed later

Principal Component Analysis

- Want $\langle \mathbf{r}\mathbf{r}^T \rangle = I$
- Subtract mean of $\mathbf{s}$. Linear model (!): $\mathbf{r} = W\mathbf{s}$
- One solution for W: PCA. Find the eigenvectors of $\mathrm{cov}(\mathbf{s}) = \langle \mathbf{s}\mathbf{s}^T \rangle = Q_{ss}$ and scale
- Write $Q_{ss} = U\Lambda U^T$ (where $U^T U = I$ and Λ is diagonal). Set $W = \Lambda^{-1/2} U^T$, then $\langle \mathbf{r}\mathbf{r}^T \rangle = I$
- First PC maximizes $\mathrm{var}(\mathbf{w}_1 \cdot \mathbf{s})$ subject to $|\mathbf{w}_1|^2 = 1$
- Subsequent components: subtract previous ones and repeat procedure
- Can also be used for dimensionality reduction by removing modes with lowest eigenvalues.

PCA on Natural Image Patches

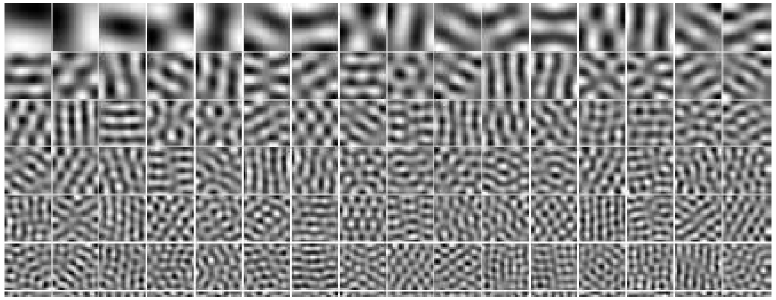

Figure: Hyvärinen, Hurri and Hoyer (2009)

If translation invariant covariance matrix, $C_{ij} = f(|i-j|)$: eigenvectors are periodic (proof: e.g. HHH p.125).
So PCA = Fourier analysis.

Whitening with PCA

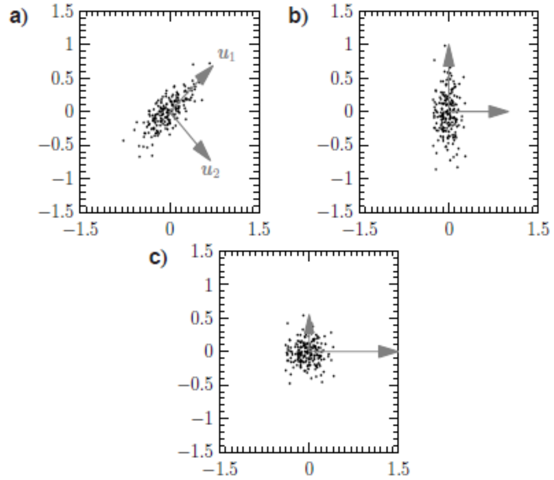

[Hyvärinen et al., 2009]
To whiten:1) do PCA projections 2) scale components with inverse variance.

Generative model with PCA

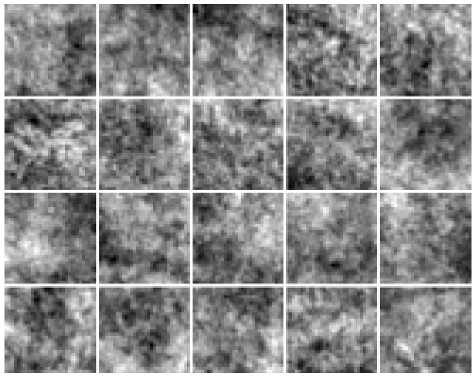

[Hyvärinen et al., 2009]
$\mathbf{s} = \sum_k \mathbf{w}_k r_k$
$P(\mathbf{r}) = \prod_k P(r_k) = \prod_k N(0, \sigma_k^2)$
Gaussian mix of principal components

Importance of Fourier Phase Infomation

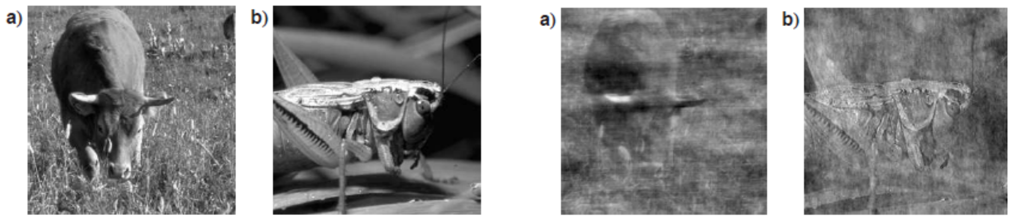

Figure: Hyvärinen, Hurri and Hoyer (2009)

- Left: sample images.
  Right: a) phase of (a) + amplitude of (b), b) v.v.
  (Method: Fourier transform image, split into magnitude and phase, mix, inverse transform)
- PSD contains no phase information, so second order stats miss important information ... tbc.

Retinal Ganglion Cell Receptive Fields

Continuous-space version of the above calculation.
Spatial part of the calculation only. [Atick and Redlich, 1990], also Dayan and Abbott §4.2 Find filter $D(x)$.

$$r(\mathbf{a}) = \int D(\mathbf{x}-\mathbf{a})s(\mathbf{x})d\mathbf{x}$$

$$Q_{rr}(\mathbf{a},\mathbf{b}) = \int\int D(\mathbf{x}-\mathbf{a})D(\mathbf{y}-\mathbf{b})\langle s(\mathbf{x})s(\mathbf{y})\rangle d\mathbf{x}d\mathbf{y}$$

For decorrelation we require

$$Q_{rr}(\mathbf{a},\mathbf{b}) = \sigma_r^2\delta(\mathbf{a},\mathbf{b})$$

Do calculations in the Fourier basis

$$\tilde{D}(\boldsymbol{\kappa}) = \int D(\mathbf{x})\exp(i\boldsymbol{\kappa}\cdot\mathbf{x})d\mathbf{x}$$

$$D(\mathbf{x}) = \frac{1}{4\pi^2}\int \tilde{D}(\boldsymbol{\kappa})\exp(-i\boldsymbol{\kappa}\cdot\mathbf{x})d\boldsymbol{\kappa}$$

to obtain

$$|\tilde{D}(\boldsymbol{\kappa})|^2\tilde{Q}_{ss} = \sigma_r^2 \qquad \Rightarrow \qquad |\tilde{D}(\boldsymbol{\kappa})| = \frac{\sigma_r}{\sqrt{\tilde{Q}_{ss}}}$$

- Whitening filter
- Notice that only $|\tilde{D}(\boldsymbol{\kappa})|$ is specified. Decorrelation and variance equalization do not fully specify kernel

- For natural scenes $\tilde{Q}_{ss}(\boldsymbol{\kappa}) \propto (\kappa_0^2 + |\boldsymbol{\kappa}|^2)^{-1}$ (Field, 1987)
- Filtering in the eye adds extra factor so that

$$\tilde{Q}_{ss}(\boldsymbol{\kappa}) = \frac{\exp(-\alpha|\kappa|)}{\kappa_0^2 + |\boldsymbol{\kappa}|^2}$$

- Implies that $|\tilde{D}(\boldsymbol{\kappa})|$ grows exponentially for large $|\boldsymbol{\kappa}|$.
- Whitening filter boosts the high frequency components (that have low power in $\tilde{Q}_{ss}$)

Filtering Input Noise

- Total input is $s(\mathbf{x}) + \eta(\mathbf{x})$, where $\eta(\mathbf{x})$ is noise, reflecting image distortion, photoreceptor noise etc
- Optimal least-squares filter is the Wiener filter with

$$\tilde{D}_\eta(\boldsymbol{\kappa}) = \frac{\tilde{Q}_{ss}(\boldsymbol{\kappa})}{\tilde{Q}_{ss}(\boldsymbol{\kappa}) + \tilde{Q}_{\eta\eta}(\boldsymbol{\kappa})}$$

Thus

$$\tilde{D}_s(\boldsymbol{\kappa}) = \tilde{D}(\boldsymbol{\kappa})\tilde{D}_\eta(\boldsymbol{\kappa})$$

$$|\tilde{D}_s(\boldsymbol{\kappa})| = \frac{\sigma_r\sqrt{\tilde{Q}_{ss}(\boldsymbol{\kappa})}}{\tilde{Q}_{ss}(\boldsymbol{\kappa}) + \tilde{Q}_{\eta\eta}(\boldsymbol{\kappa})}$$

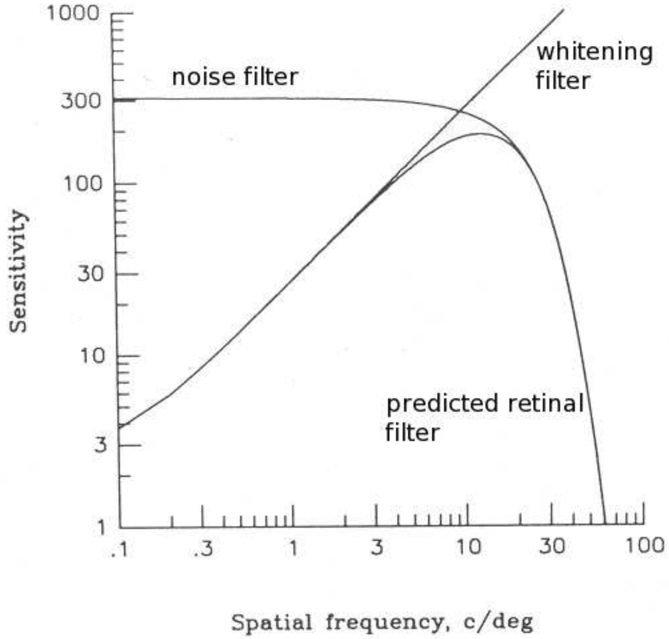

[Atick and Redlich, 1992]

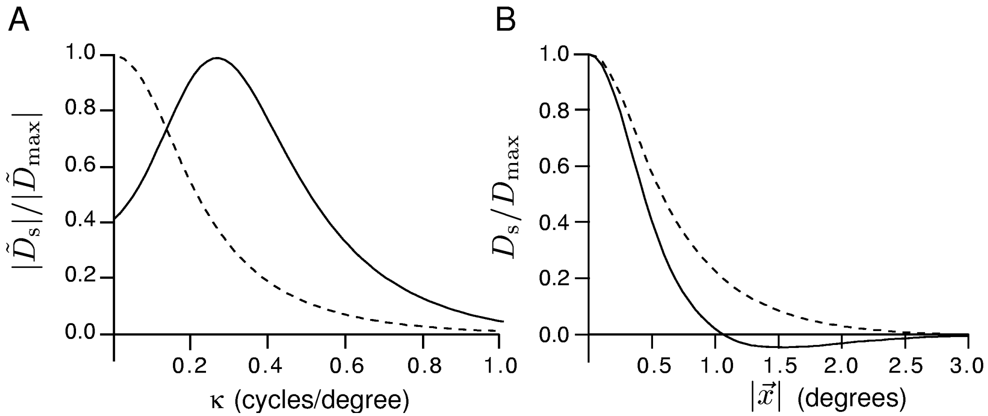

Figure: [Dayan and Abbott 2001]

Solid curve, low noise; dashed curve, high noise
Choose local, rotationally symmetric solution

- For low noise the kernel has a bandpass character, and the predicted receptive field has a centre-surround structure
- This eliminates one major source of redundancy arising from strong similarity of neighbouring inputs
- For high noise the structure of the optimal filter is low-pass, and the RF loses its surround
- This averages over neighbouring inputs to extract the signal which is obscured by noise
- Result is not simple PCA as we have enforced spatial invariance on the filter
- In the retina, low light levels ≡ high noise. The predicted change matches observations [Van Nes and Bouman, 1967]

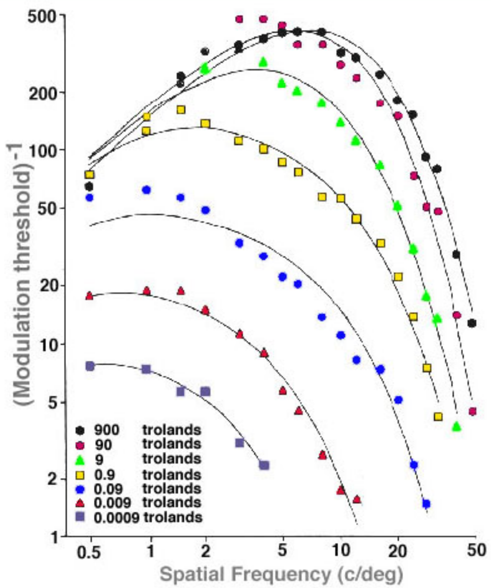

Figure 24. Contrast sensitivity function showing a change in shape from low pass at low luminances and bandpass at high luminances. van Ness' data from Lamming D., Contrast Sensitivity. Chapter 5. In: Cronly-Dillon, J., Vision and Visual Dysfunction, Vol 5. London: Macmillan Press, 1991.

Contribution of Spiking to de-correlation

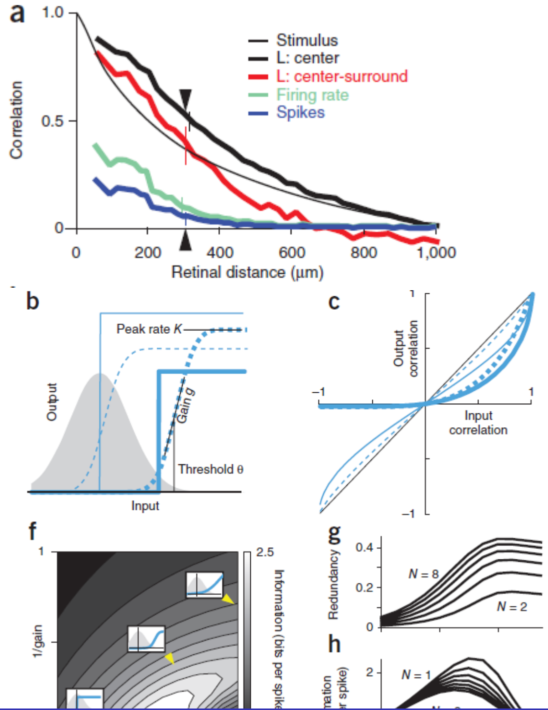

Further Decorrelation Analyses

- Spatio-temporal coding (Dong and Atick, 1995; Li, 1996). Power spectrum is $1/f^2$ but non-separable
- Colour opponency: red centre, green surround (and vice versa) [Atick et al., 1993]

Caveats for the Information Maximization Approach

- Information maximization sets limited goals and requires strong assumptions
- Analyzes representational properties but ignores computational goals e.g. object recognition, target tracking
- Cortical processing of visual signals requires analysis beyond information transfer. V1 can have no more information about the visual signal than the LGN, but it has many more neurons
- However, information transfer analysis does help understand mutual selectivities: RFs with preference for high spatial frequencies are low-pass temporal filters, and RFs with selectivity for low spatial frequency act as bandpass temporal filters

References I

Atick, J. J., Li, Z., and Redlich, A. N. (1993).
What does post-adaptation color appearance reveal about cortical color representation?
Vision Res, 33(1):123–129.

Atick, J. J. and Redlich, A. N. (1990).
Towards a Theory of Early Visual Processing.
Neural Comput, 2(3):308–320.

Atick, J. J. and Redlich, A. N. (1992).
What does the retina know about natural scenes.
Neural Comp., 4:196–210.

Hyvärinen, A., Hurri, J., and Hoyer, P. (2009).
Natural Image Statistics.
Spinger.

Pitkow, X. and Meister, M. (2012).
Decorrelation and efficient coding by retinal ganglion cells.
Nat Neurosci, 15(4):628–635.

Van Nes, F. and Bouman, M. (1967).
Spatial modulation transfer in the human eye.
J Opt Soc Am, 57:401–406.