Explaining Cortical Responses with Efficient Coding

Efficiency Hypothesis

Information-theoretic formulation: maximize *mutual information* between stimuli and responses of a population of neurons.

Ingredients:

- ensemble (or model) of stimuli
- choice of neural population
- definition of “response”
- limitations on responses

Conceptual problem: Purely bottom-up (ignores task).

Practical problem: information is hard to estimate.

Simplistic Efficiency Hypothesis

Ignoring noise allows us reduce the hypothesis to two simple components:

- individual responses to stimulus ensemble should have maximum entropy (subject to response limitations)
- responses to stimulus ensemble should be statistically independent

[Attneave, ’54; Barlow, ’61]

Intuition (on board...)

Maximum Entropy

The distribution over response r with maximum entropy subject to a constraint of the form:

$$\mathcal{E}(f(r)) = c$$

is

$$\mathcal{P}(r) \propto \exp(-\lambda(c) f(r))$$

Examples:

- $f(r) = r^2$ gives a (half-)Gaussian
- $f(r) = |r|$ gives an exponential
- $f(r) = \sqrt{|r|}$ gives something even sparser

Current Biology, Vol. 13, 493–497, March 18, 2003, ©2003 Elsevier Science Ltd. All rights reserved. DOI 10.1016/S0960-9822(03)00135-0

The Cost of Cortical Computation

Peter Lennie*
Center for Neural Science
New York University
4 Washington Place
New York, New York 10003

Summary

Electrophysiological recordings show that individual neurons in cortex are strongly activated when engaged in appropriate tasks, but they tell us little about how many neurons might be engaged by a task, which is important to know if we are to understand how cortex encodes information. For human cortex, I estimate the cost of individual spikes, then, from the known energy consumption of cortex, I establish how many neurons can be active concurrently. The cost of a single spike is high, and this severely limits, possibly to fewer than 1%, the number of neurons that can be substantially active concurrently. The high cost of spikes requires the brain not only to use representational codes that rely on very few active neurons, but also to allocate its energy resources flexibly among cortical regions according to task demand. The latter constraint explains the investment in local control of hemodynamics, exploited by functional magnetic resonance imaging, and the need for mechanisms of selective attention.

Introduction

rat neocortex. Neurons in human neocortex are larger than those in rat and receive and make more synapses, but they are not otherwise known to differ in their basic structure or organization [5]. Thus, with appropriate scaling of parameters for the larger neurons, Attwell and Laughlin's analysis can be used to estimate the energy consumed by a pyramidal neuron in human neocortex.

In different mammals, the number of neurons under a unit area of cortical surface is relatively constant (~100,000/mm^2), except in primate striate cortex, where it may be twice as high [6]. Increasing brain size brings an increase in cortical thickness and a proportionately lower density of neurons [5, 6] without an increase in cell body size, which remains approximately constant at 15 μm diameter [7]. The volume of axons and dendrites increases with cortical thickness. This reflects an increase in the lengths of dendrites and axons without an increase in diameter [5]. Table 1 summarizes relevant statistics for human cortex.

Postsynaptic Potentials

Individual synapses are assumed to be the same in rat and human neurons, so the energy costs associated with transmitter uptake and release will be the same, as will the current flow through receptor channels. Given (from Table 1) 7×10^8 synapses per mm^3 of cortex, and 40,000 neurons/mm^3, the average neuron will make 17,500 synaptic contacts. If we use this number, and assume a 50% failure rate [8, 9], the cost of EPSPs arising from a single spike will be 1.2×10^9 ATP molecules [4].

Maximum Entropy [Baddeley etal, 97]

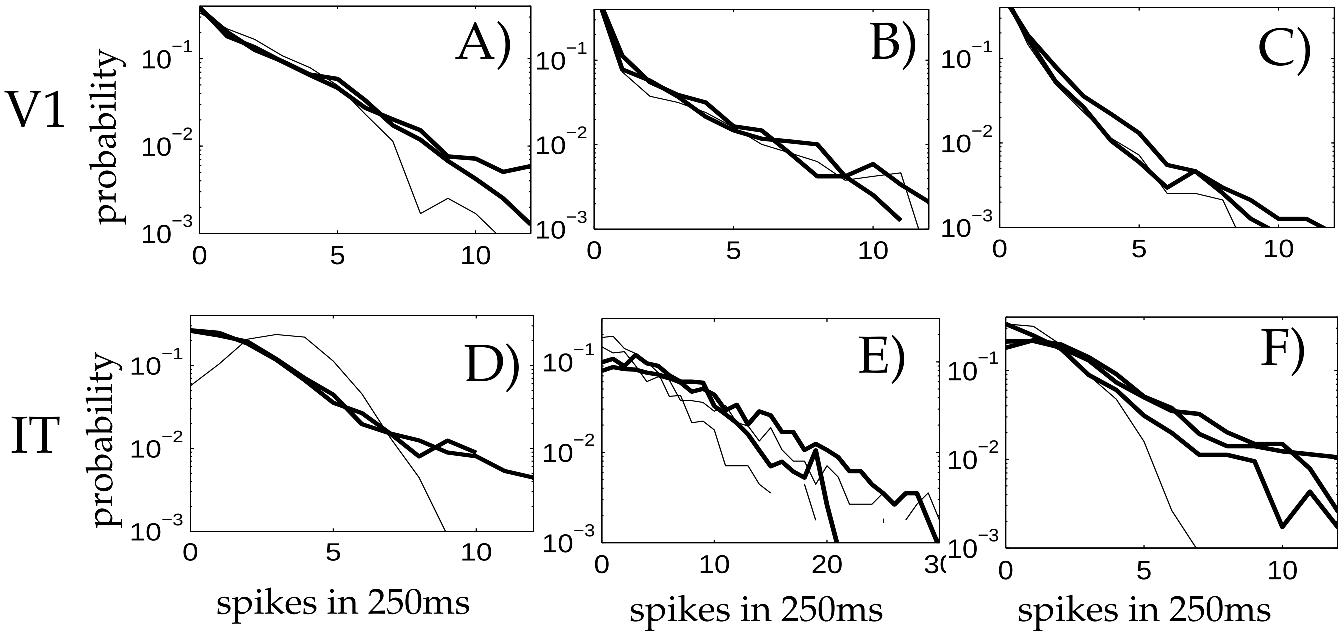

Distributions appear exponential, optimally efficient under a constraint on the mean firing rate.

Caveats/concerns:

- Firing rate constraint not always the right one. [Laughlin '81]
- Why are responses to white noise also somewhat exponential?
- Noise probably matters here

Testing the Independence Hypothesis

- Experimental: Examine dependency between neuronal responses under natural stimulation conditions [eg: Vinje/Gallant ’00; Nirenberg/Latham ’01]

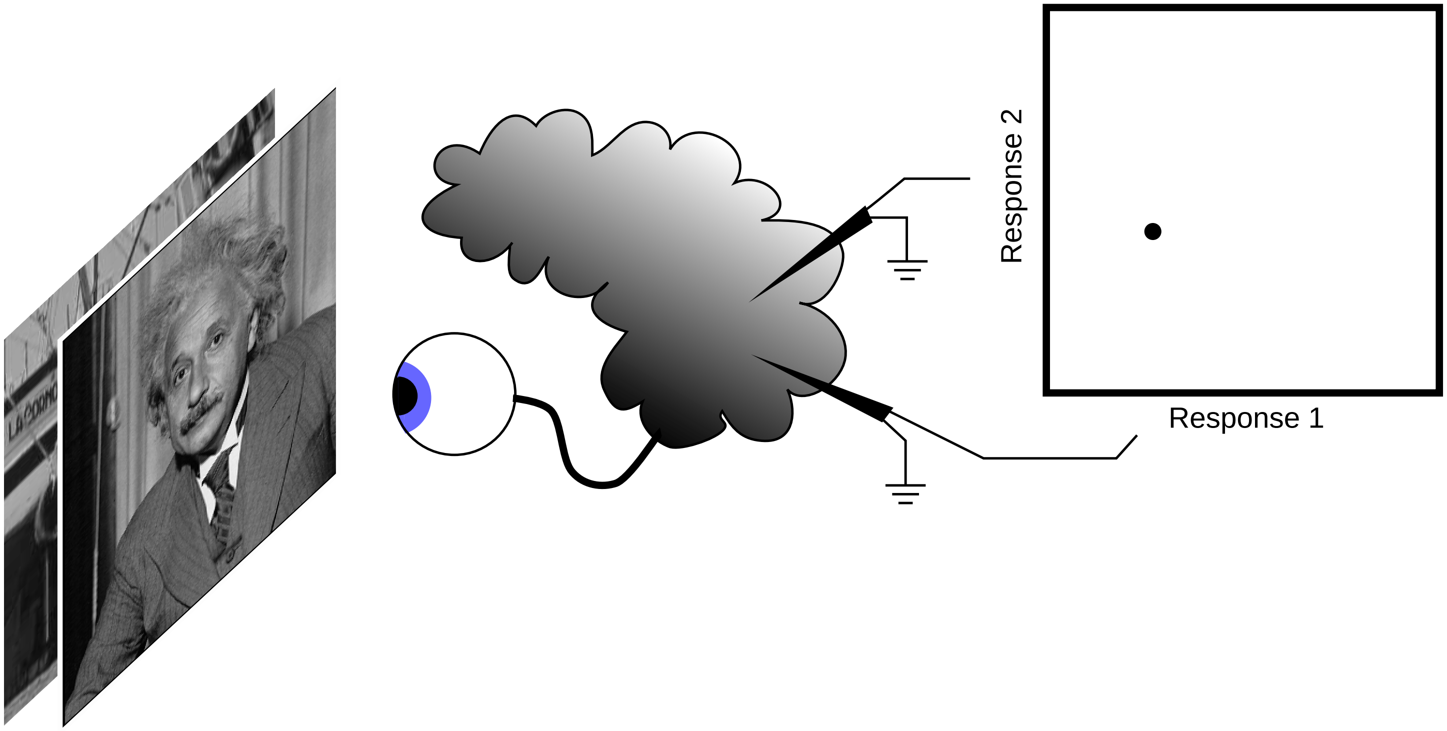

Testing the Independence Hypothesis

- Experimental: Examine dependency between neuronal responses under natural stimulation conditions [eg: Vinje/Gallant ′00; Nirenberg/Latham ′01]

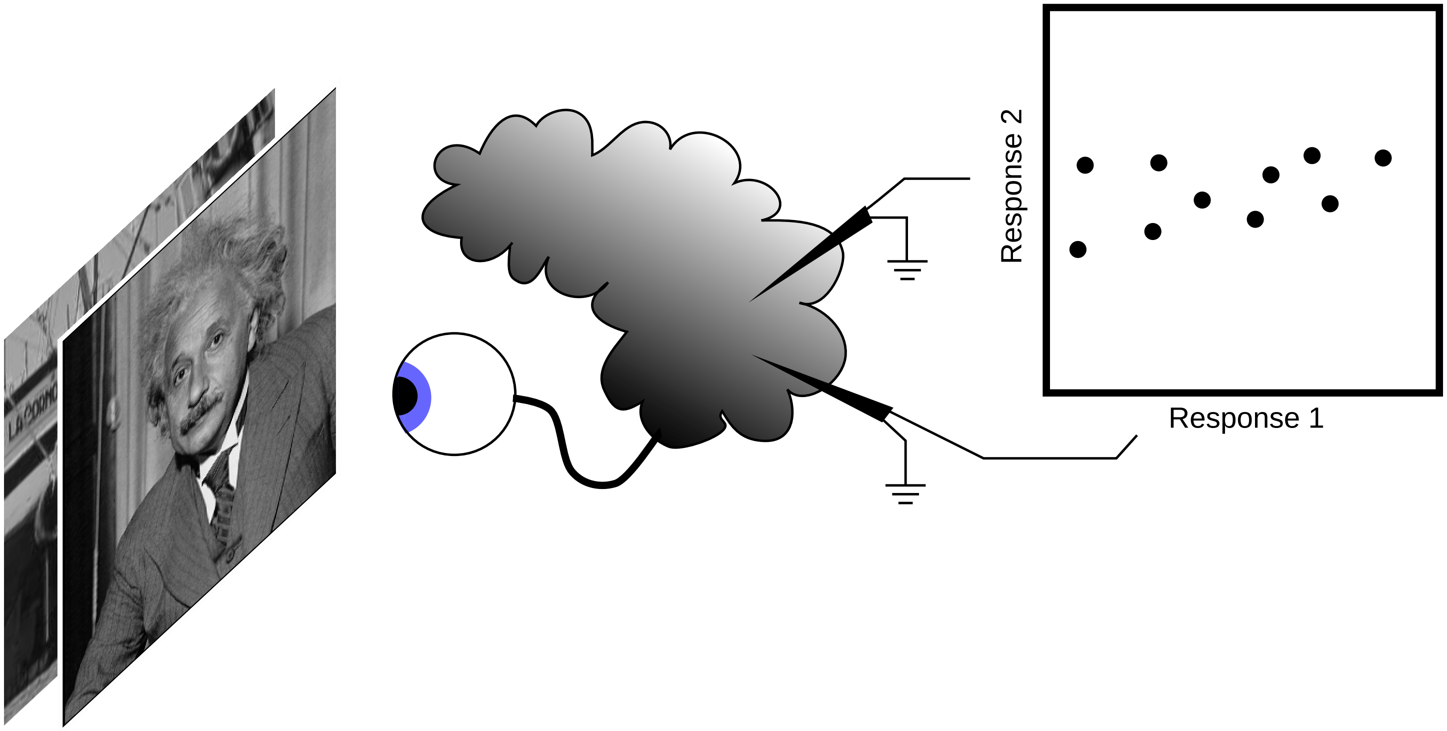

Testing the Independence Hypothesis

- Experimental: Examine dependency between neuronal responses under natural stimulation conditions [eg: Vinje/Gallant ′00; Nirenberg/Latham ′01]

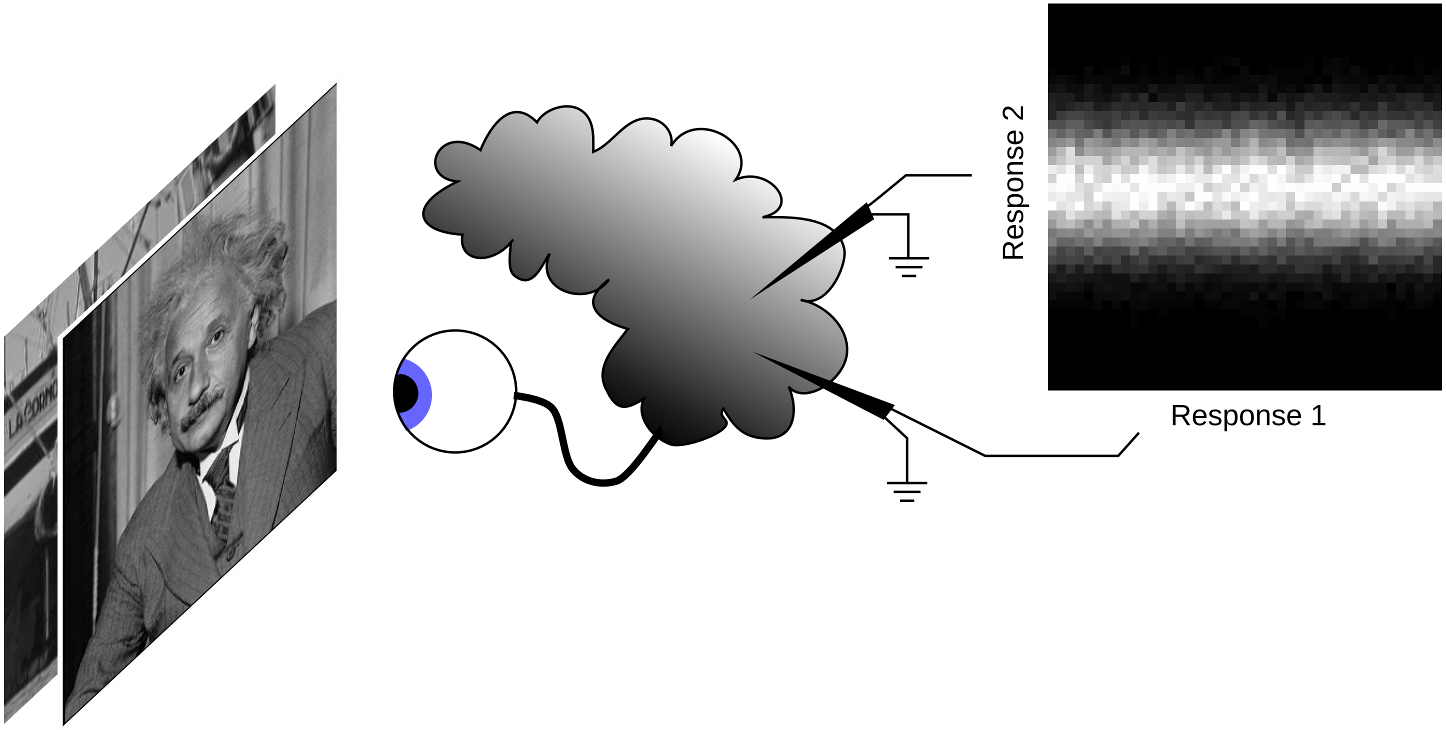

Testing the Independence Hypothesis

- Experimental: Examine dependency between neuronal responses under natural stimulation conditions [eg: Vinje/Gallant ′00; Nirenberg/Latham ′01]
- Theoretical: Derive a model that can produce independent responses; Compare with physiology

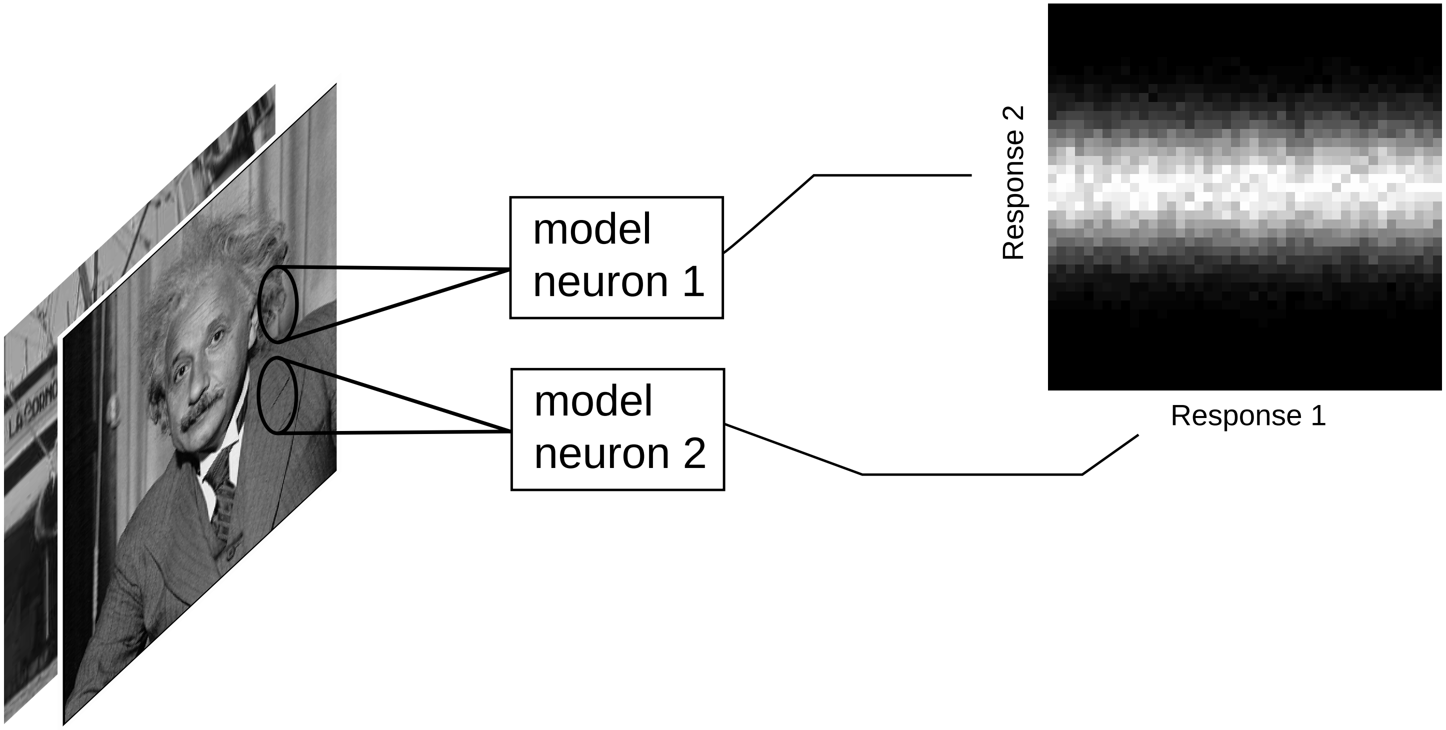

Toward Independent Image Components...

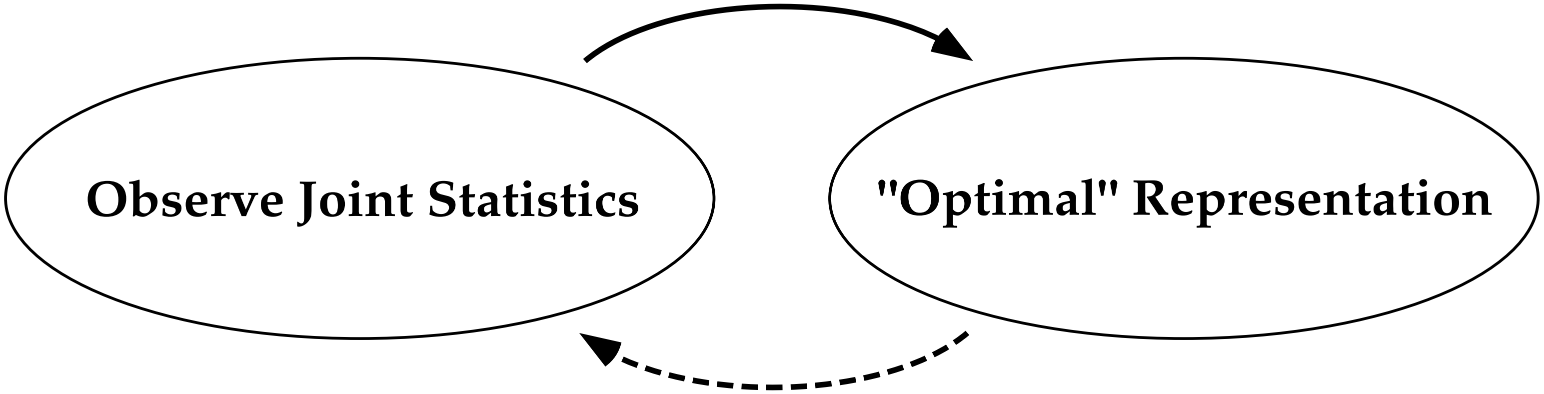

A bottom-up methodology for constructing image models, or image processing systems.

Extensive history in engineering and theoretical neuroscience.

Principal Component Analysis (PCA)

Find linear transform (specifically, rotation and axis re-scaling) that transforms the covariance matrix to the identity.

Geometrically, convert the ellipse that describes the variance in all directions into a circle!

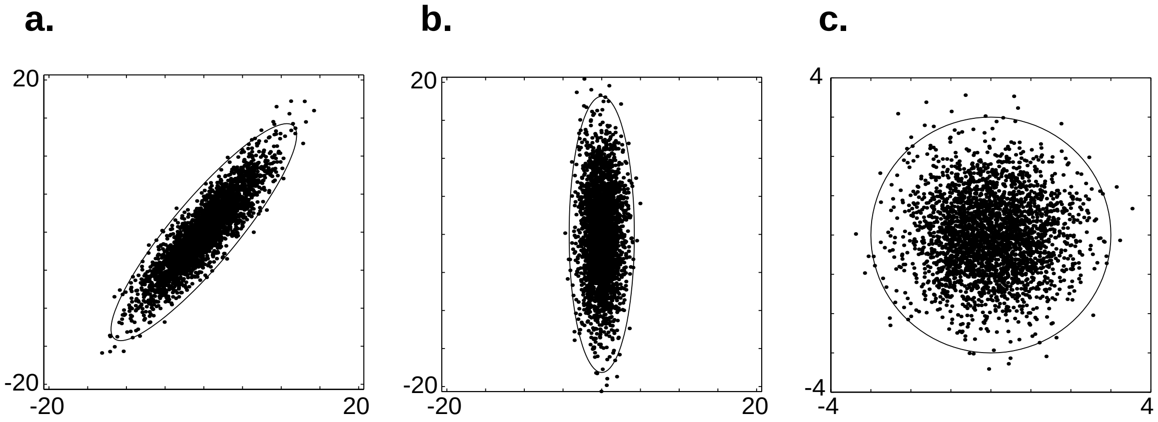

PCA on images: Spectral image model

Reminder of previous lecture:

1. Characterize image statistics by covariance matrix
2. Translation-invariance $\Longrightarrow$ Fourier diagonalizes
3. Variances in Fourier domain (power spectrum) falls like $1/f^p$:
   - verified empirically for "typical" photographic images
   - guaranteed if we assume scale-invariance
4. To complete model, assume Gaussianity (which is maximum entropy under variance constraint)

PCA is Insufficient

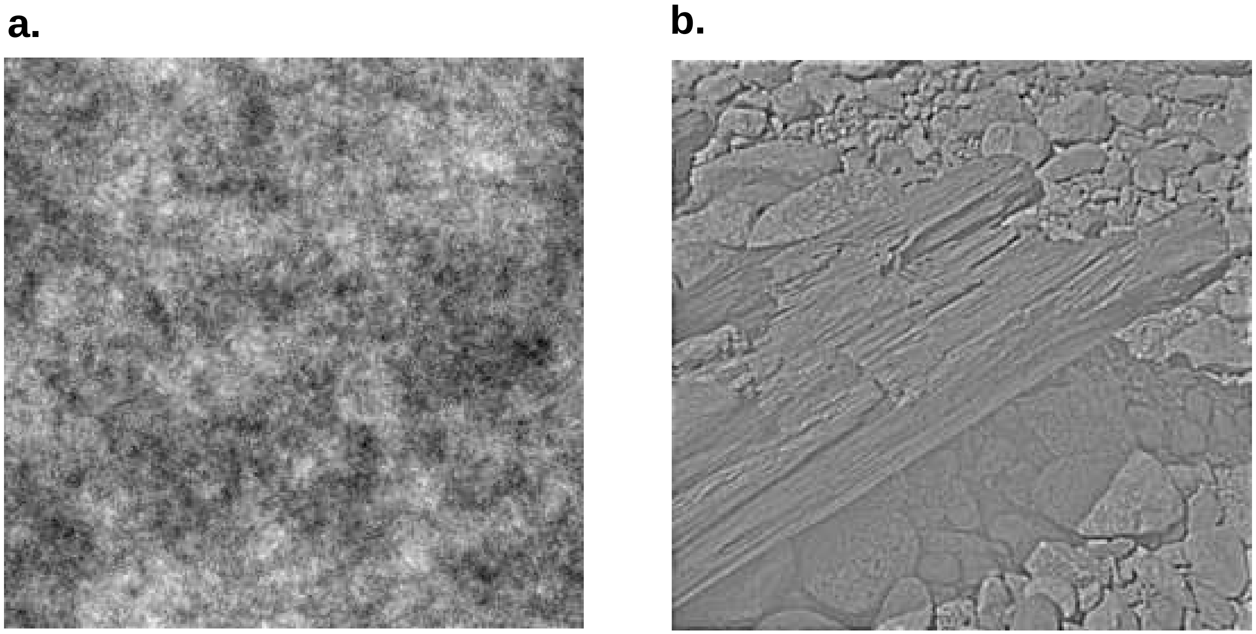

left: $1/f$ Gaussian noise. **right**: whitened natural image.

Bandpass Filters Reveal non-Gaussian Behaviors

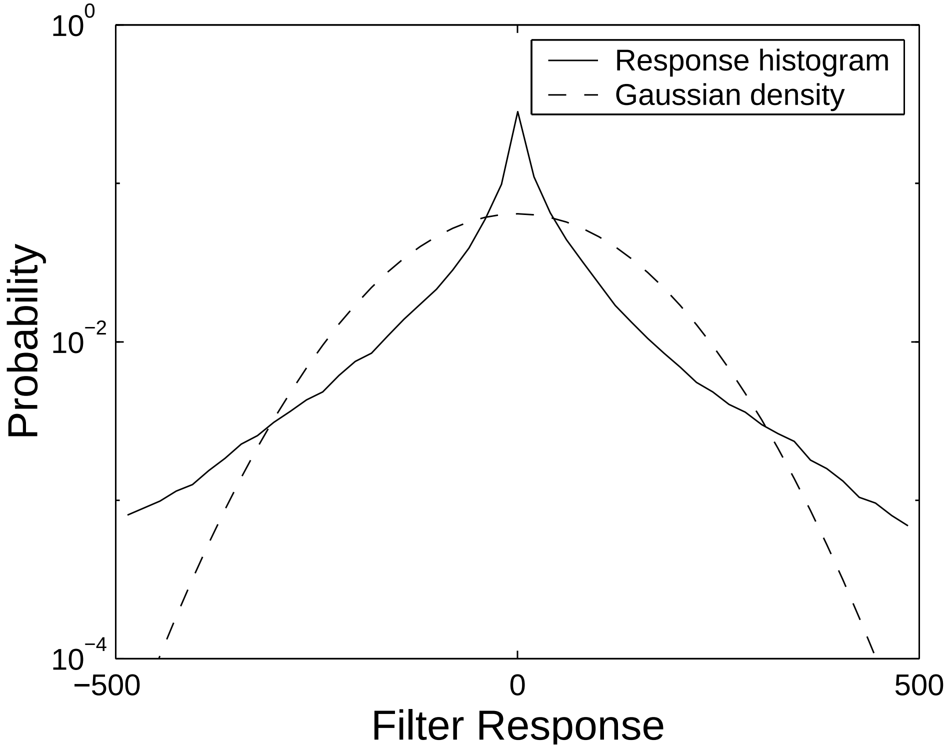

Marginal densities of bandpass filtered images are non-Gaussian.

[Field ’87; Mallat ’89; Zetzsche ’90]

Multi-dimensional Gaussians

- Characterized by mean (vector) and covariance (matrix)
- Remain Gaussian under linear transformation
- Specifically, conditionals (slices) and marginals (projections) are Gaussian
- Separable product of i.i.d. Gaussians is spherically symmetric
- Central limit theorem: sums of i.i.d. random variables become Gaussian
- Maximizes entropy for a given mean and covariance
- Heisenberg: Fourier transform of a Gaussian is Gaussian, and minimizes variance product

PCA on Linear Combination of non-Gaussian Sources

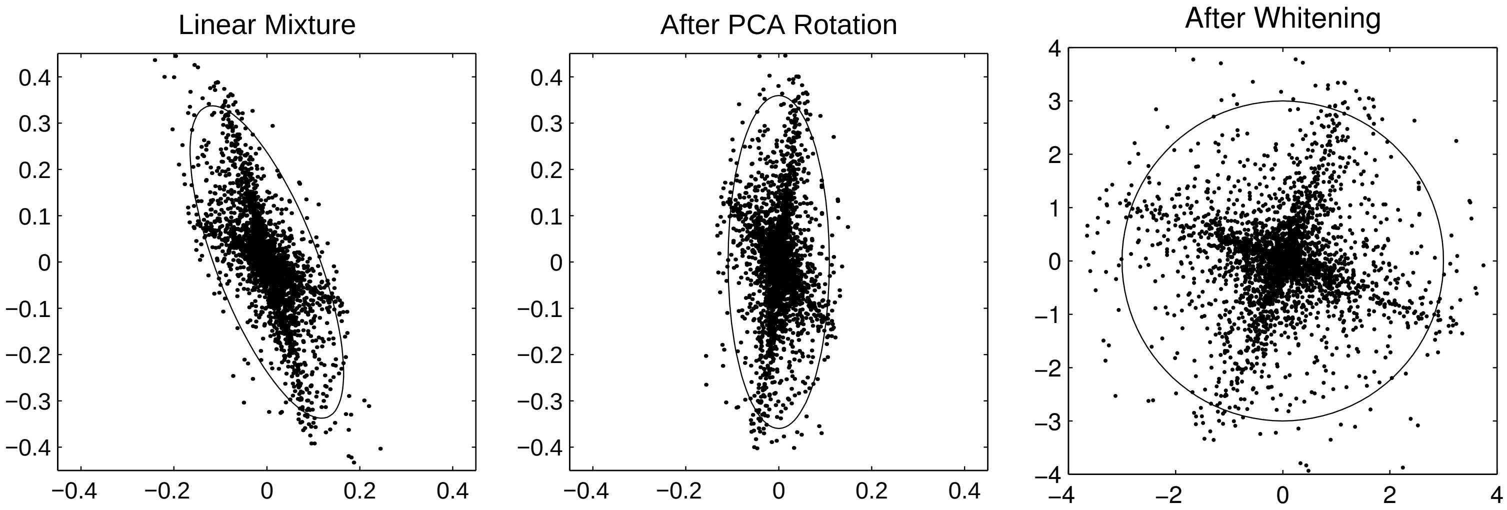

2nd-order whitening does not necessarily recover independent sources!

Need an additional rotation matrix...

Independent Component Analysis (ICA)

Seek linear transform that maximizes statistical independence of transform coefficients.

A variety of algorithms exist.

ICA on Images

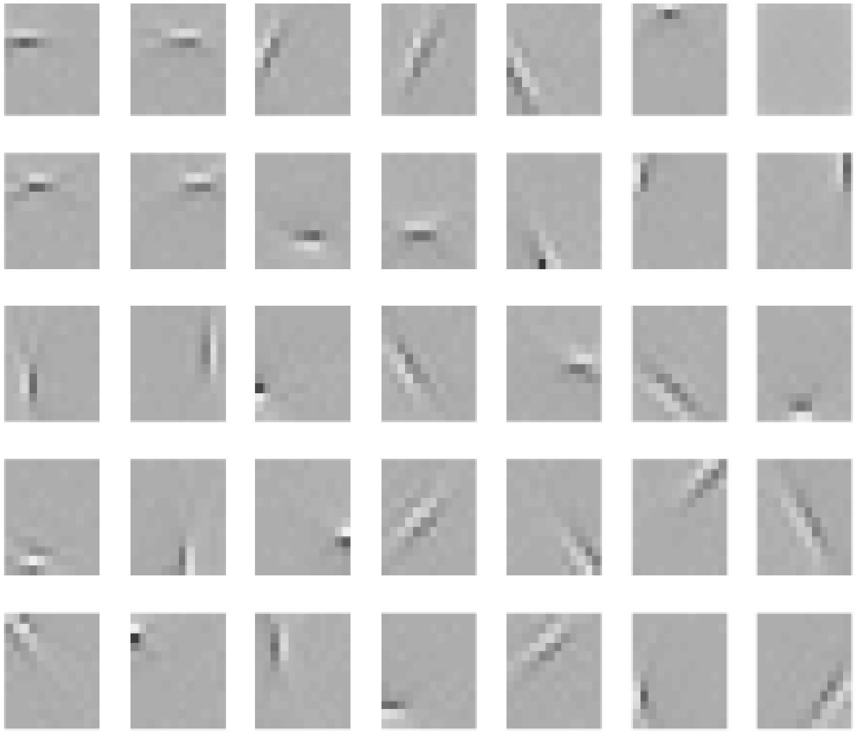

Linear operators with maximally independent (or maximally non-Gaussian) responses are oriented bandpass filters (loosely, wavelets)

[Bell/Sejnowski ’97; Olshausen/Field ’96]

Sample Kurtosis vs. Filter Bandwidth

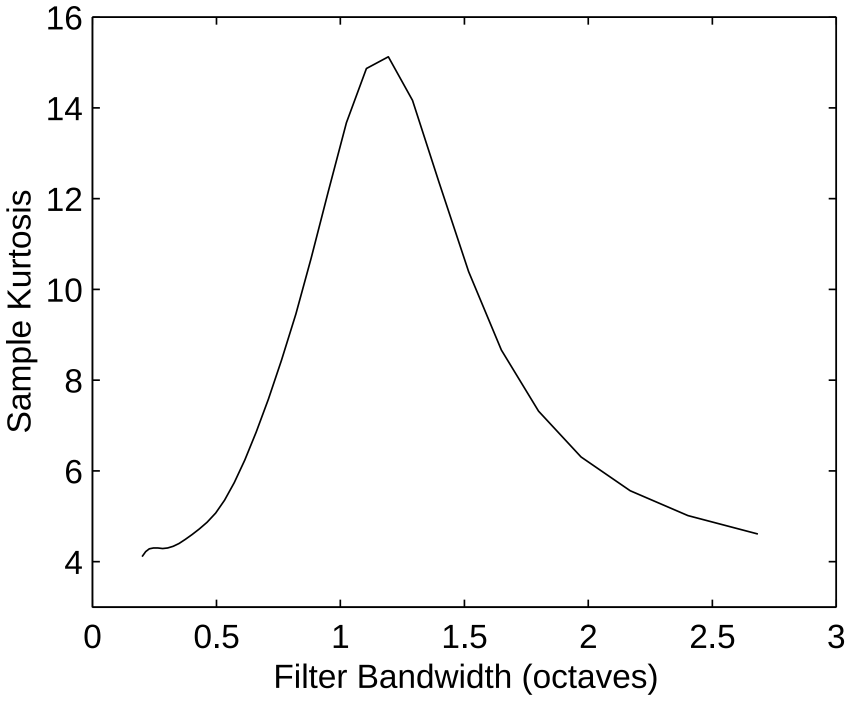

For most images, maximum is near one octave [after Field, 1987]

.

Model II: Wavelet+generalizedGaussian Marginals

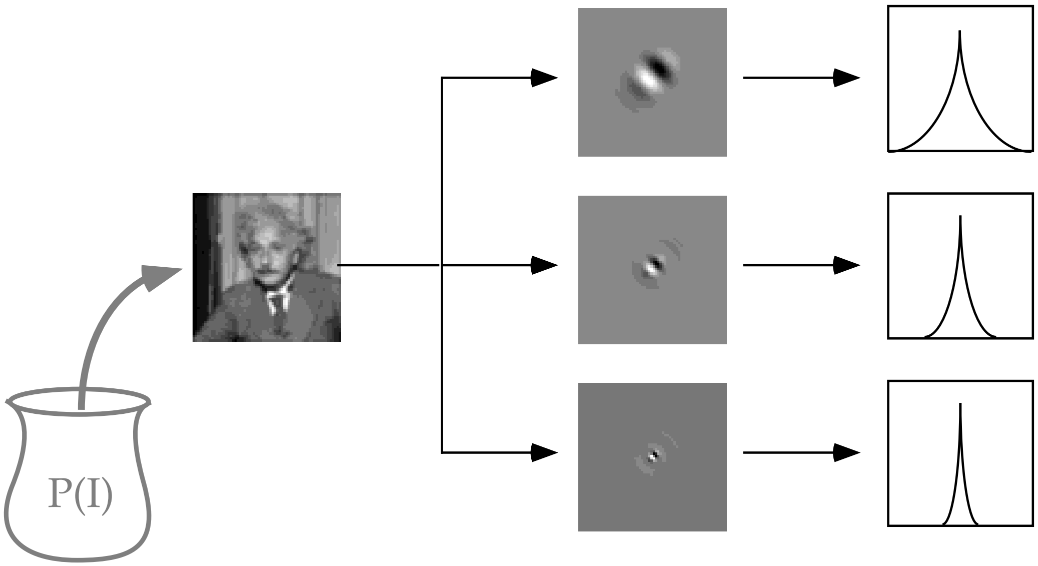

Image:

Project onto:

Coefficient
density:

Coefficient densities are (approximately) uncorrelated *generalized* Gaussians (Mallat89, Simoncelli/Adelson96):

$$f(c) \propto e^{-|c/s|^p}, \quad p \in [0.5, 0.8].$$

Coefficient Dependency

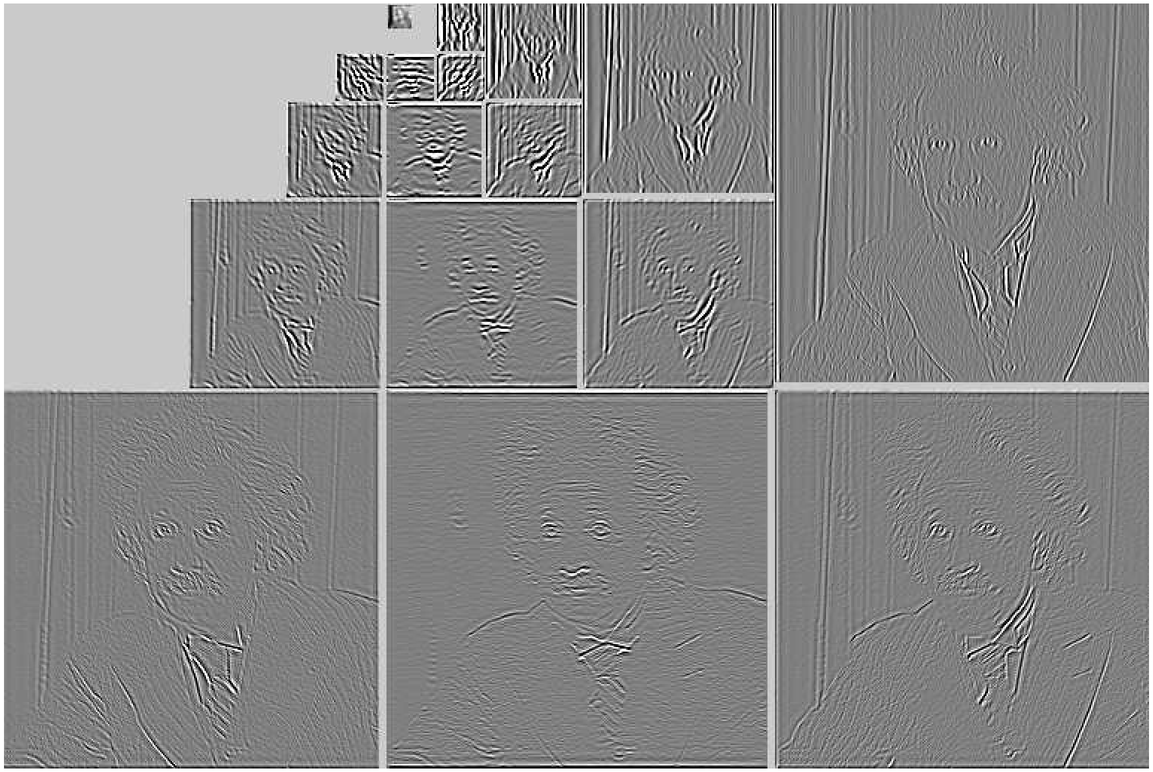

Large-magnitude subband coefficients are found at neighboring positions, orientations, and scales.

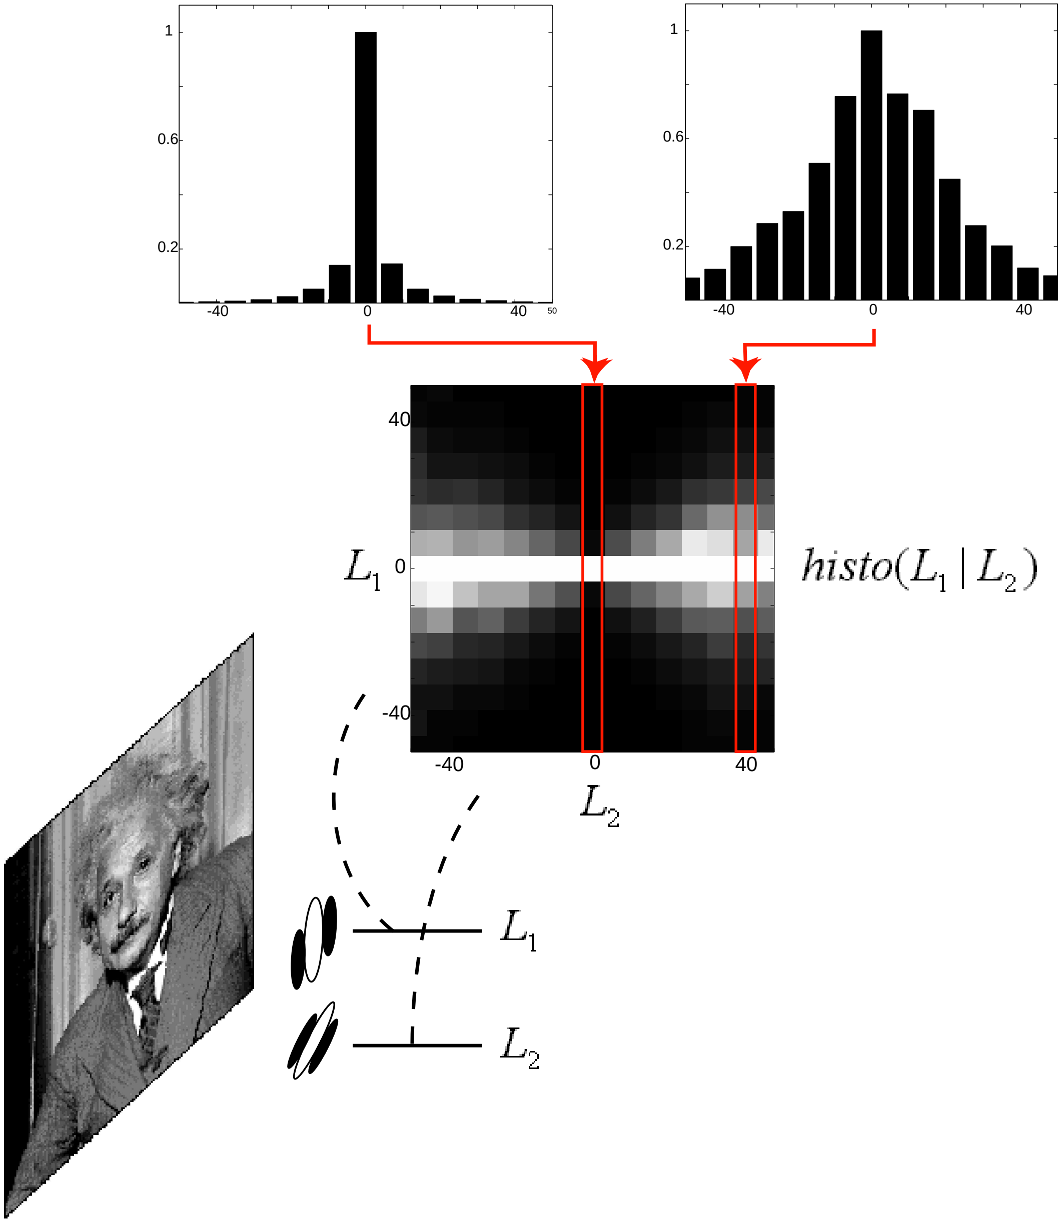

Linear responses are not independent, even for optimized filters!

Different Filter Pairs:

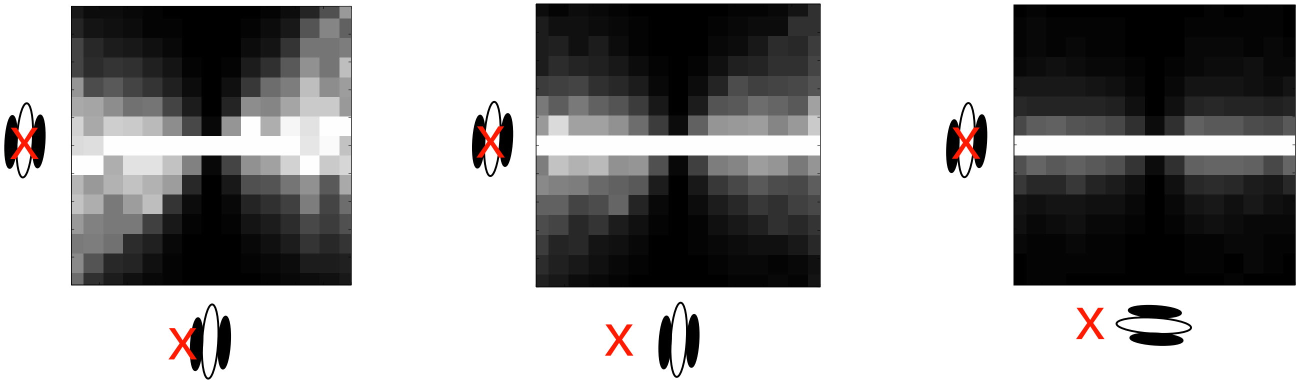

Different Images:

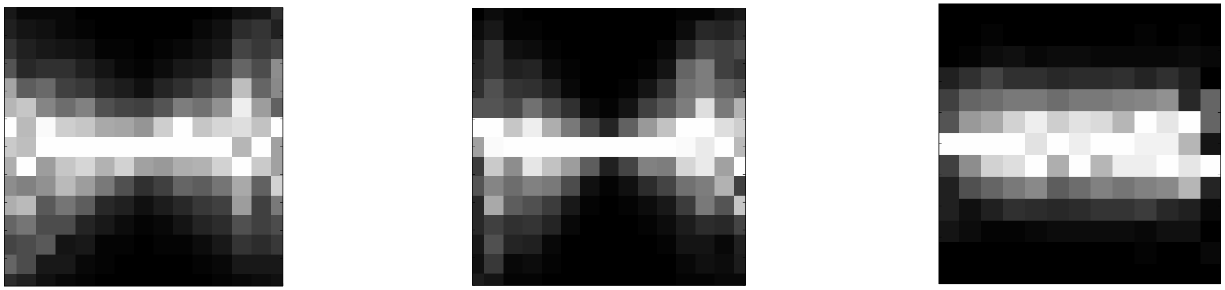

Baboon

Flowers

White noise

Modeling the Dependency

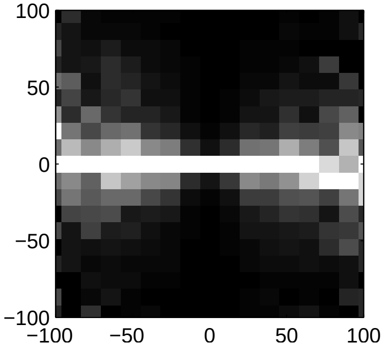

One filter:

$$\mathrm{var}(L_1|L_2) = w\ L_2^2 + \sigma^2$$

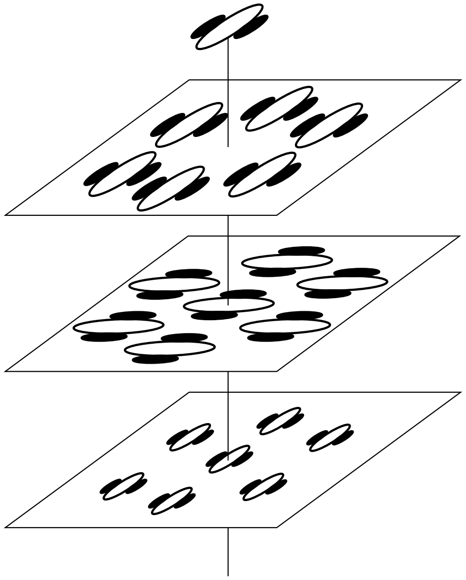

Generalized neighborhood:

$$\mathrm{var}(L_1|\{L_n\}) = \Sigma_n w_n\ {L_n}^2 + \sigma^2$$

Reducing Dependency

$$R_n = \frac{{L_n}^2}{\Sigma_k \;\, w_{nk} \, {L_k}^2 + \sigma^2}$$

Original Responses

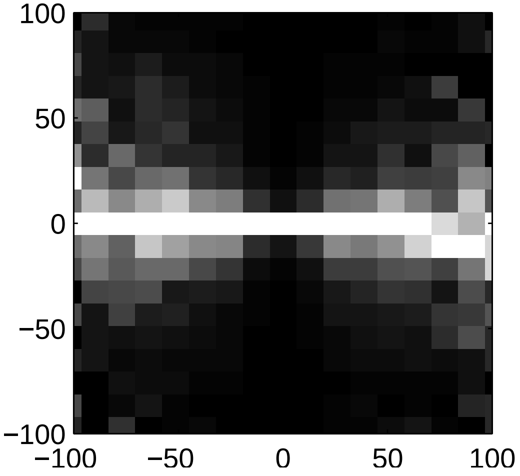

After Division

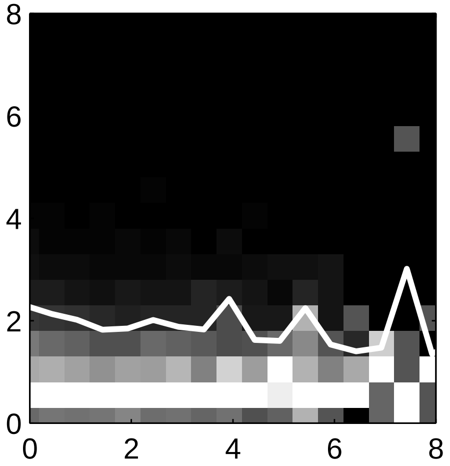

Functional Physiological Model

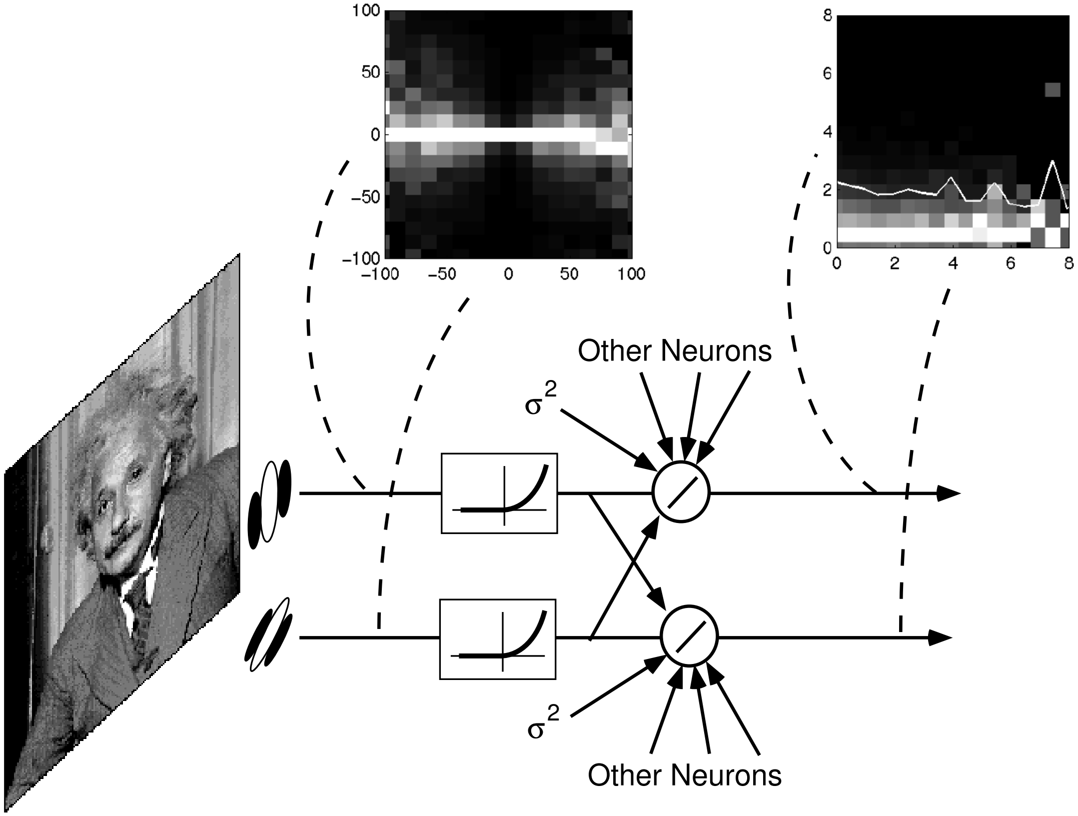

- Divisive suppression by *weighted* sum of neighbors increases independence
- Optimal weights determined from image statistics

Divisive Normalization: Physiological Evidence

Steady-state neural responses = linear projection, rectification, and division by the summed responses of other neurons [Heeger '92; Carandini/Heeger/Movshon '97]

Such models can account for some nonlinear striate cortical behaviors. Examples [Carandini et al. 1997]:

- Tuning curves independent of contrast
- Contrast saturation level depends on stimulus parameters
- Cross-orientation suppression
- Increasing phase lag at lower contrast

Methods

- Statistically-determined model:
  1. Linear basis: multi-scale, oriented 3rd derivative operators
  2. “Neuron”: vertical, optimal spatial frequency 0.125 cycle/pixel
  3. Neighborhood: 2 scales, 4 orientations, 3×3 array
  4. Weights: optimized (ML) for statistics of 10 images (faces, landscapes, and animals).

- Physiological simulations:
  1. Compute linear responses of full neighborhood
  2. Square
  3. Divide chosen neuron response by weighted combination of squared neighbor responses.

[Schwartz & Simoncelli ’01]

Parameter Optimization

Assume a Gaussian form for the conditional distribution:

$$\mathcal{P}\left(L_n \mid \{L_k\}\right) \;\sim\; \mathcal{N}\left(0;\ \sum_k\ w_{nk}\ |L_k|^2 + \sigma^2\right)$$

Maximize the likelihood over the image data:

$$\hat{w}_{nk}, \hat{\sigma} = \arg\max_{w_{nk},\sigma} \prod_i \frac{1}{\sqrt{2\pi\, \Sigma_k\ \ w_{nk}\ |L_k|^2 + \sigma^2}} \exp\left[\frac{-{L^2}_n}{2\, \Sigma_k\ \ w_{nk}\ |L_k|^2 + \sigma^2}\right]$$

Cross-orientation Suppression

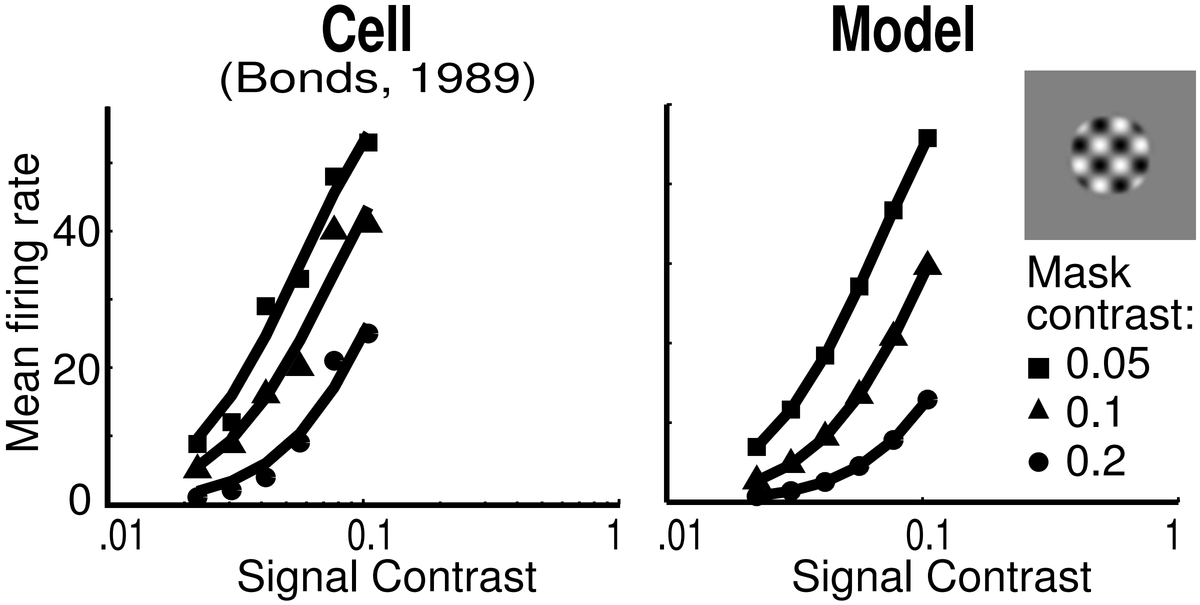

Tuning Curves Independent of Contrast

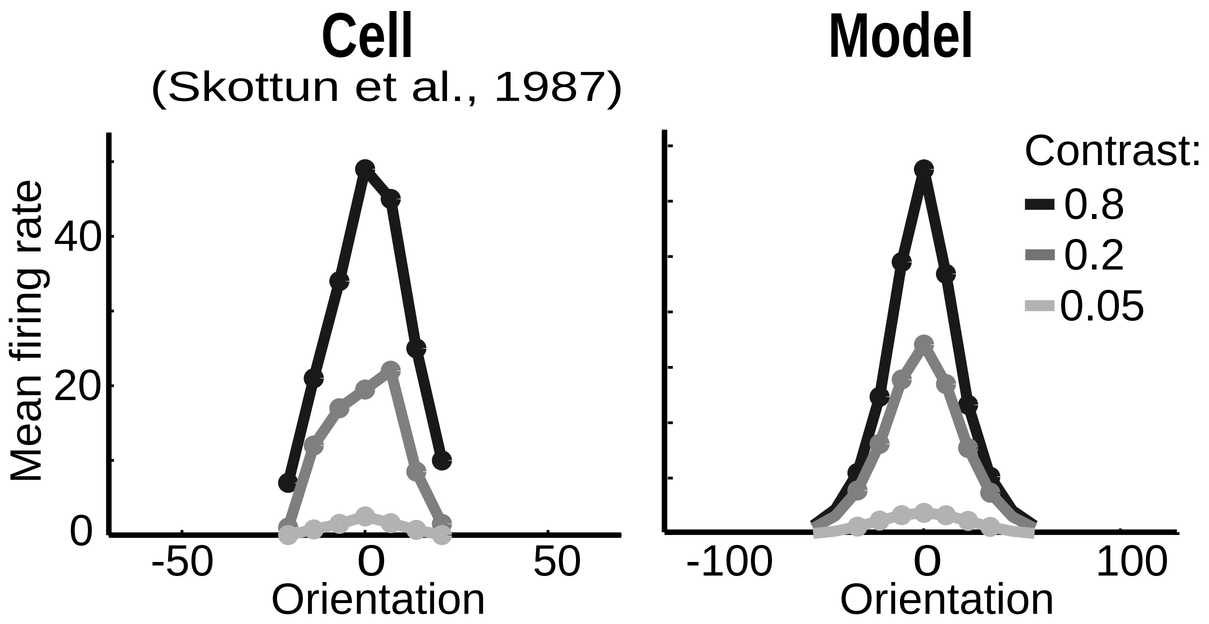

Surround Suppression

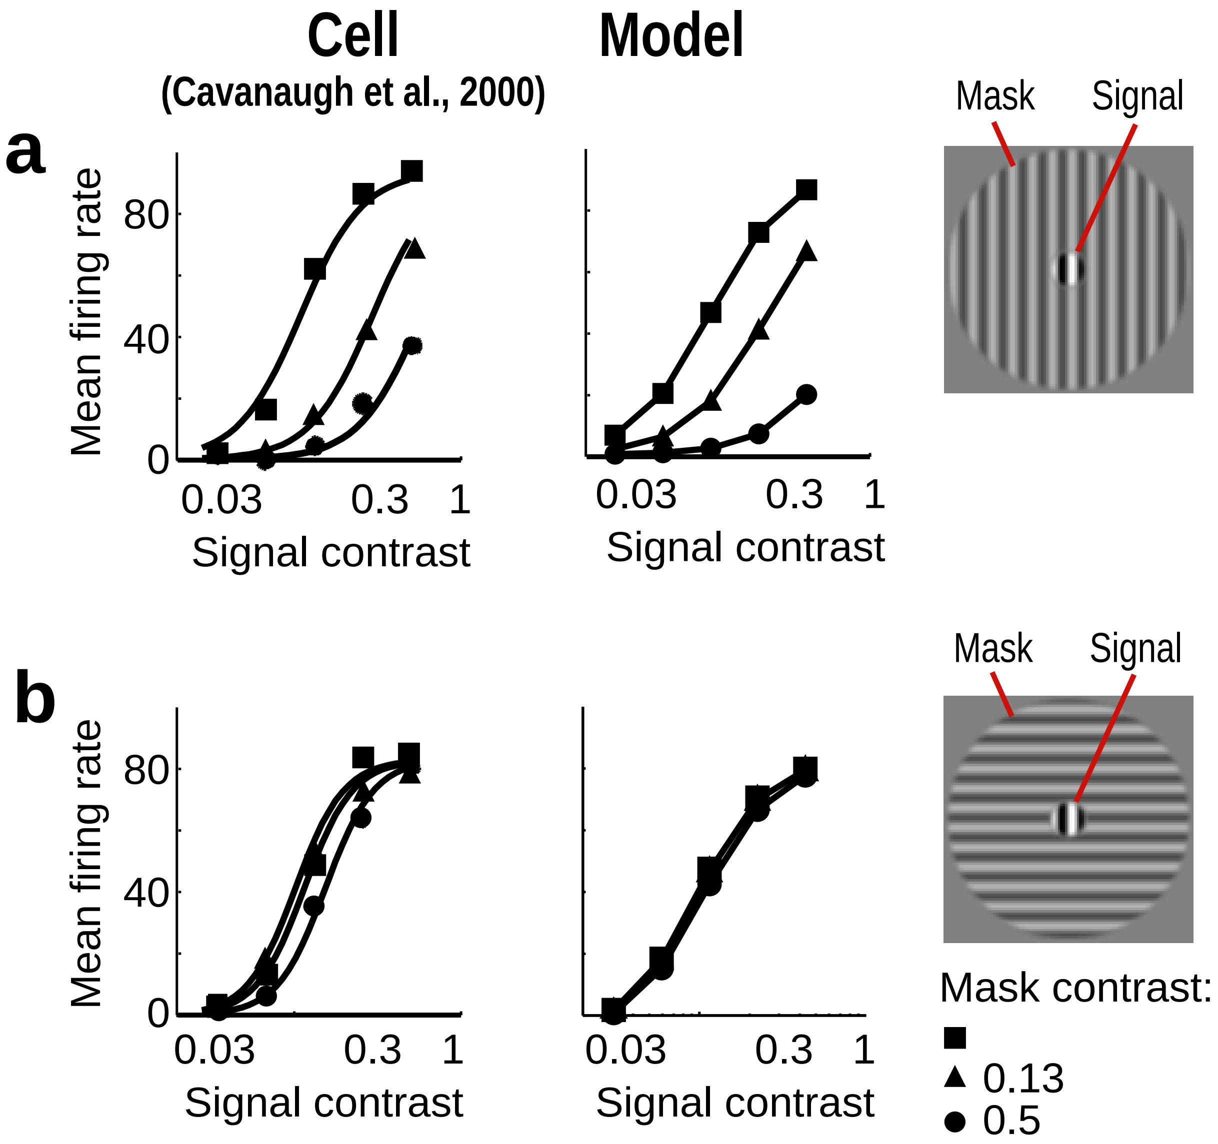

Surround Spatial Frequency

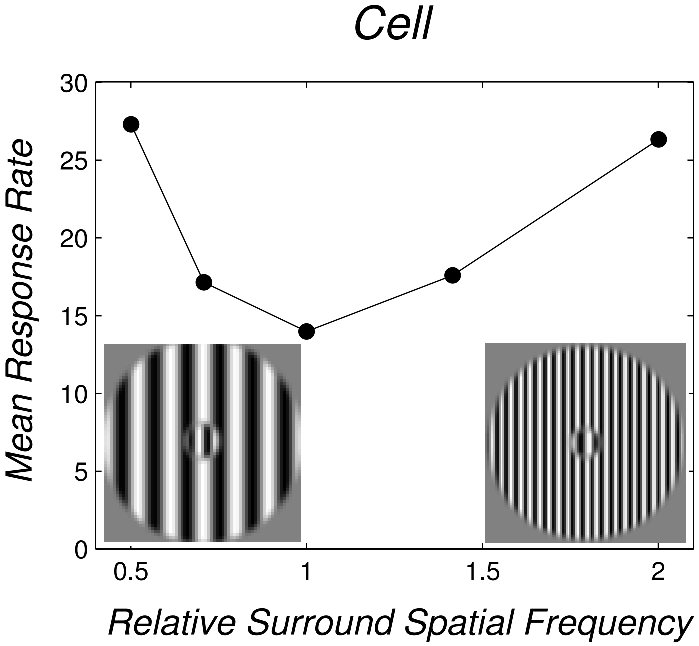

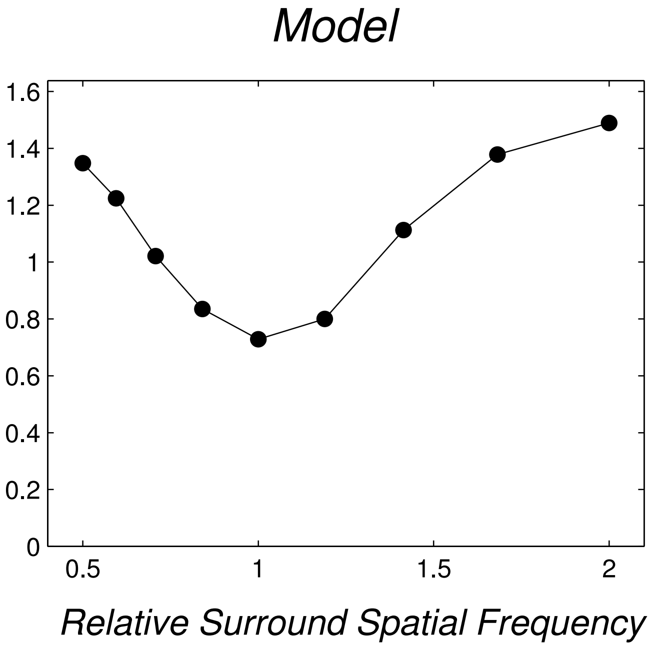

[Data: Müller, Krauskopf, & Lennie.]

Stimulus Diameter

Cell

(Cavanaugh et al., 2000)

Model

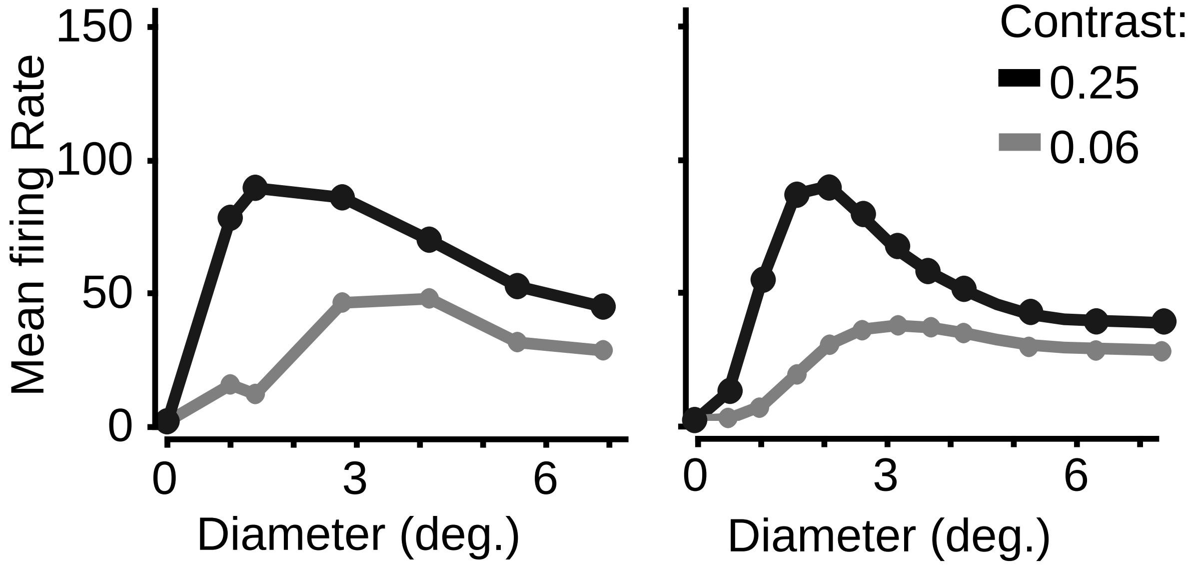

Natural Sounds

Different Sounds:

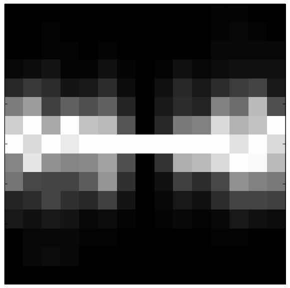

Cat

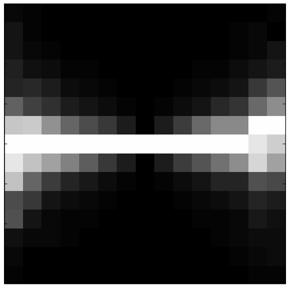

Speech

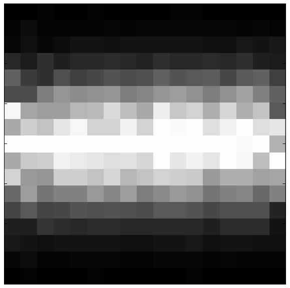

White noise

Different Filter Pairs:

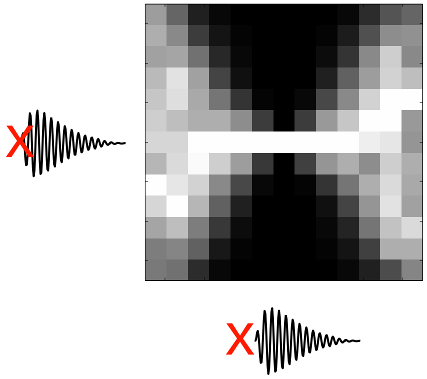

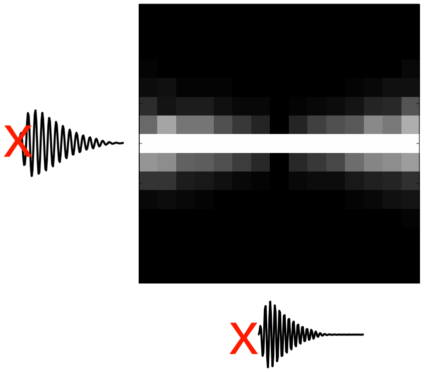

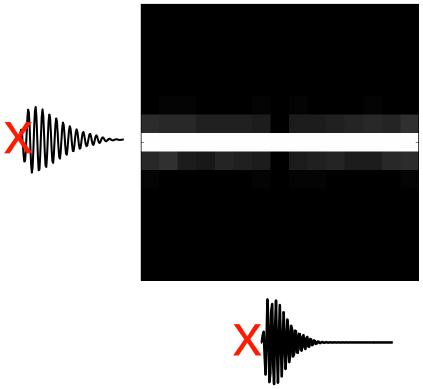

Auditory Nerve Fiber (cat)

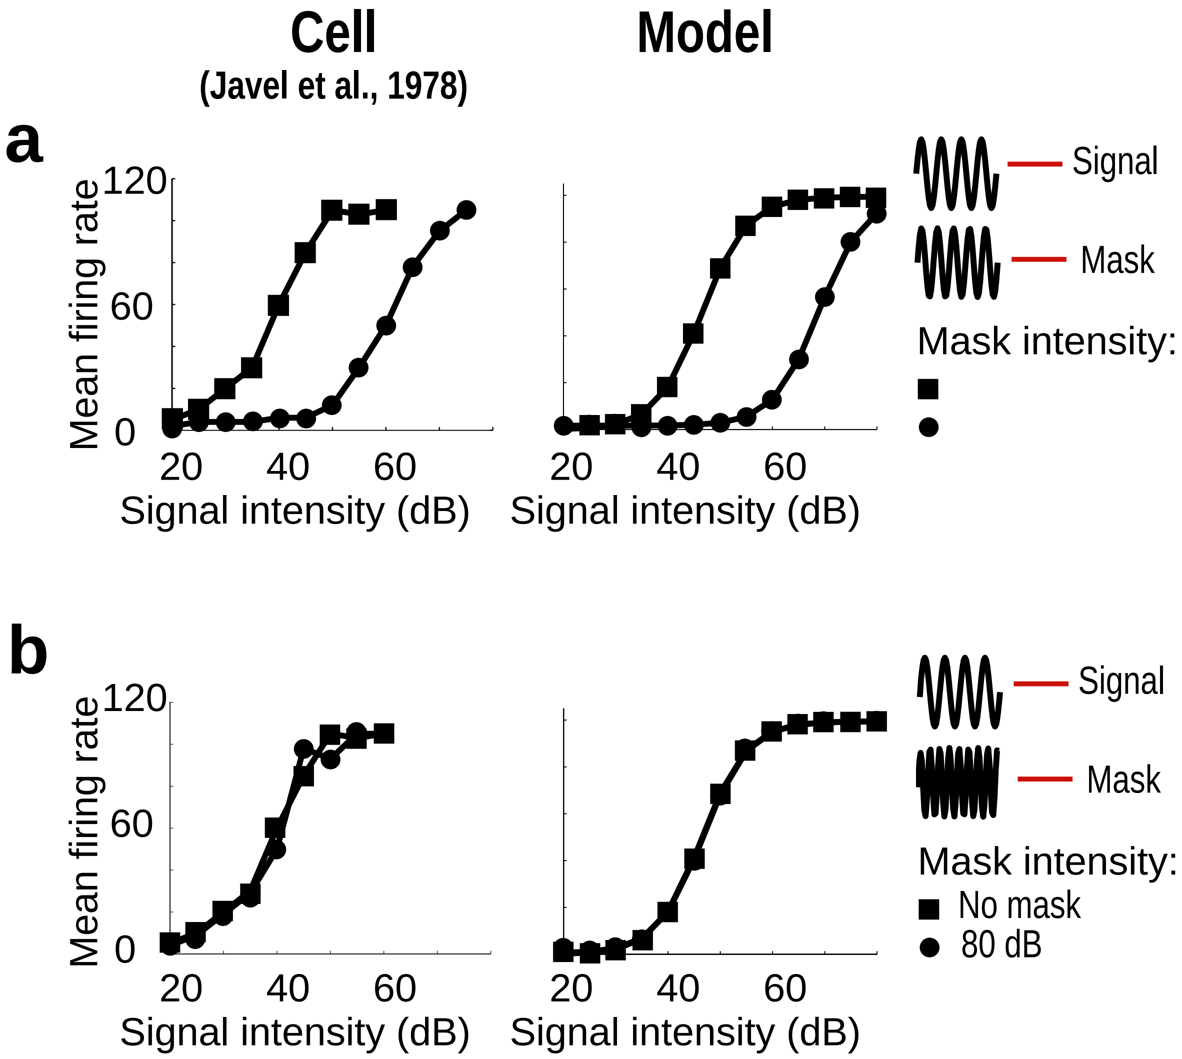

Auditory Nerve Fiber (cat)

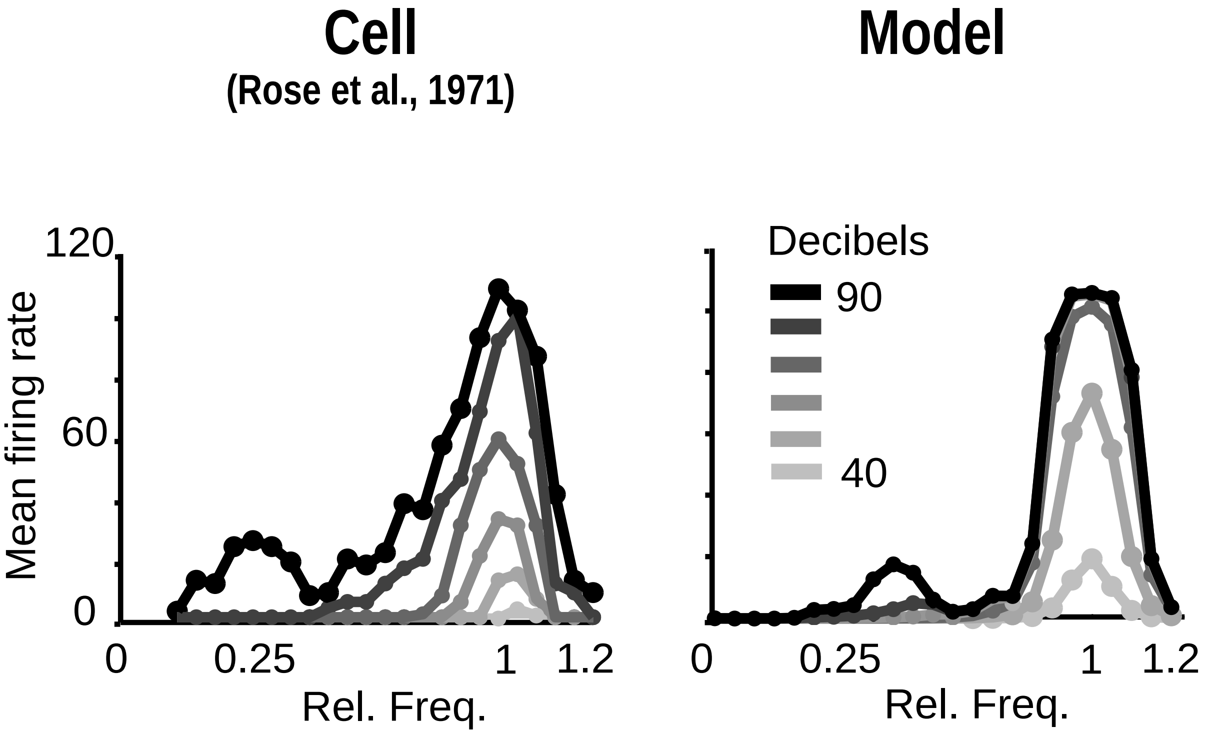

Cell

(Rose et al., 1971)

Model

Summary

Divisive gain control can get us closer to a statistically independent decomposition for images (or sounds)

- Ecological justification for physiological models
- Quantitative predictions of sensory nonlinearities:
  - suppression by non-optimal stimuli
  - intensity-independence of some tuning curves (eg orientation)
  - intensity-independence of other tuning curves (eg diameter)
- Canonical computation?
- Numerous applications in signal/image processing

General Summary

Efficiency hypothesis provides a principle linking environmental properties to response properties of sensory neurons.

There are many ways to study/test this.

How far can we go, given that it's a bottom-up hypothesis?

Some interesting possibilities

Adaptation [Laughlin ′81; Barlow & Foldiak; Wainwright ′99; Fairhall etal ′01; Wainwright etal, ′02]

New classes of “naturalistic” stochastic stimuli (e.g., textures)

Synthetic texture images (aperiodic)

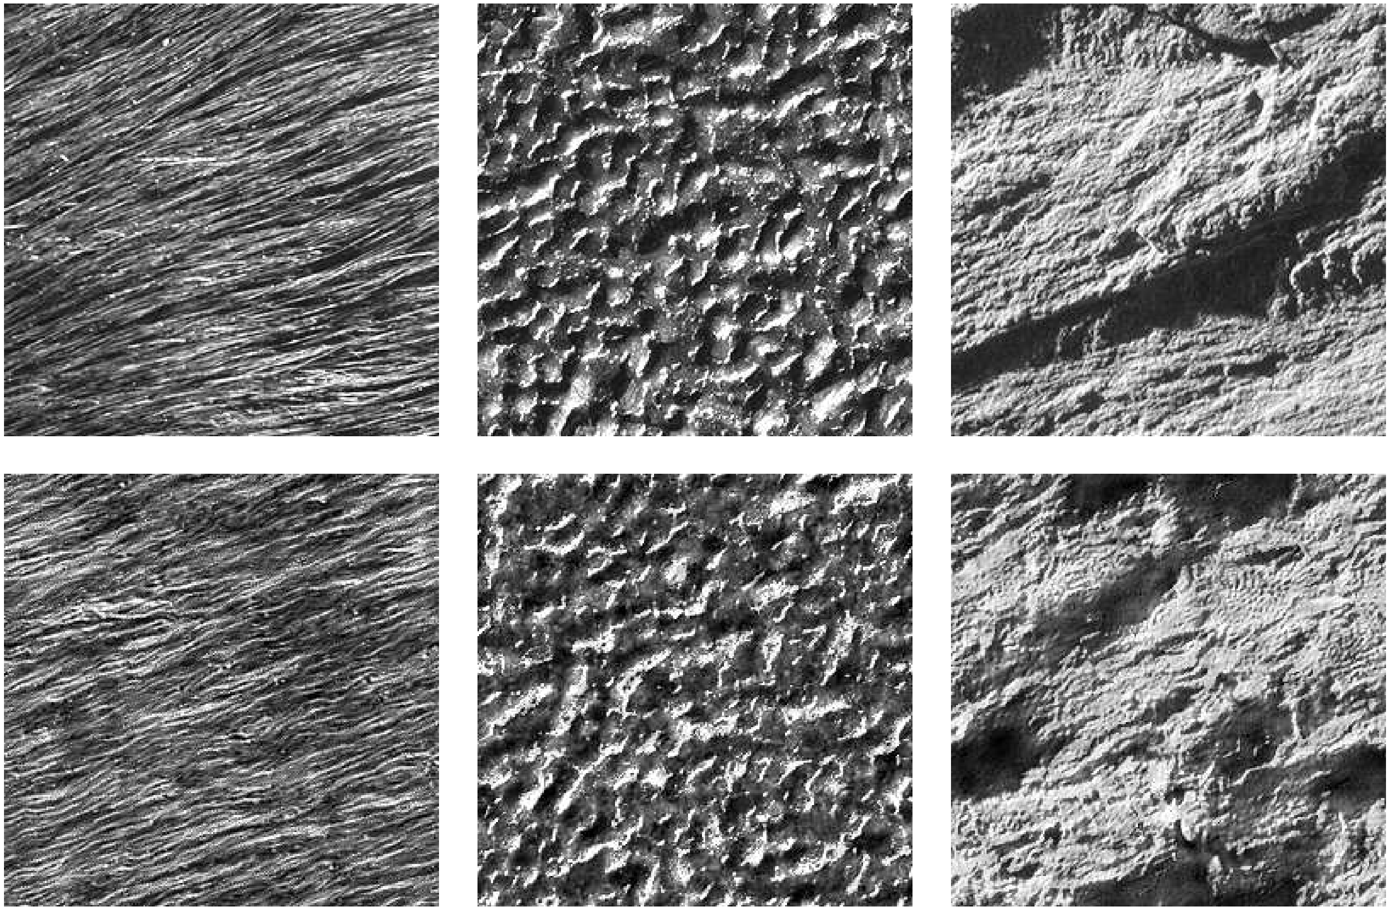

[Portilla & Simoncelli, 2000]