Lateral organization & computation cont'd

lateral organization

Why the organization? The level of abstraction?

- Keep similar features together for feedforward integration.
- Lateral computations to group features of similar type—segmentation
- Efficiency constraints
  - Minimum wiring constraint

    to keep similar features near..
    but V1 is ~ 2D, and many features!

  - Efficient representation of sensory input & cost of neural activity
  - Efficient representations for learning

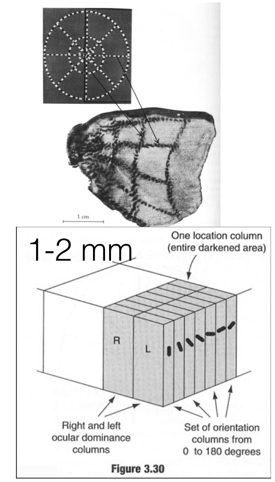

how can layout be learned?

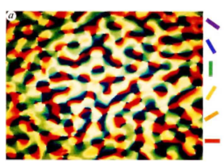

Ts'o, D. Y., Frostig, R. D., Lieke, E. E., & Grinvald, A. (1990, 27 July 1990). Functional Organization of Primate Visual Cortex Revealed by High Resolution Optical Imaging. Science, 249, 417-420.

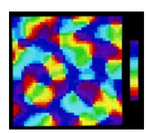

Durbin, R., & Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. Nature, 343, 644-647.

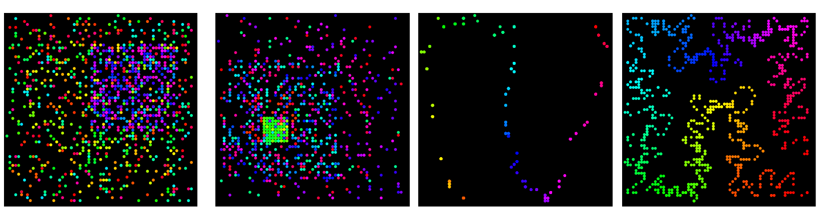

Kohonen map demo: Mapping 2D to 1D

Just V1?

Tanaka, K. (2003). Columns for complex visual object features in the inferotemporal cortex: clustering of cells with similar but slightly different stimulus selectivities. Cereb Cortex, 13(1), 90-99.

lateral organization: "maps"

Why the organization? The level of abstraction?

- Keep similar features together for feedforward integration.
- Lateral computations to group features of similar type—segmentation
- Efficiency constraints
  - Minimum wiring constraint
  - Efficient representation of sensory input & cost of neural activity
  - Efficient representations for learning

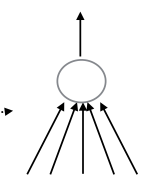

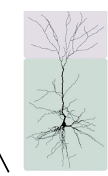

limited dendritic spread

lateral computations?

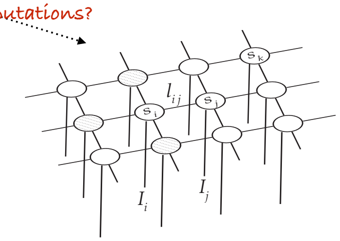

Markov Random Field models

Grouping

link contours with similar orientations

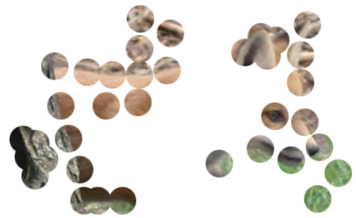

link regions with similar colors, textures

What should the local features be? How many different types?

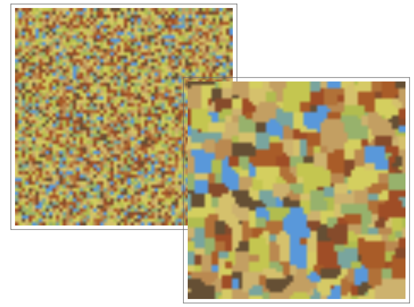

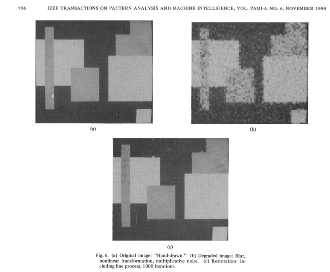

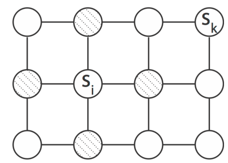

prior

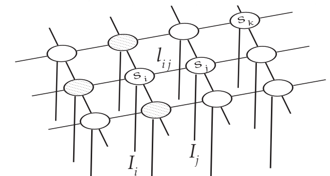

prior + likelihood

..but would the visual system need to "denoise"?

what is noise anyway?

Human fMRI evidence for lateral computations?

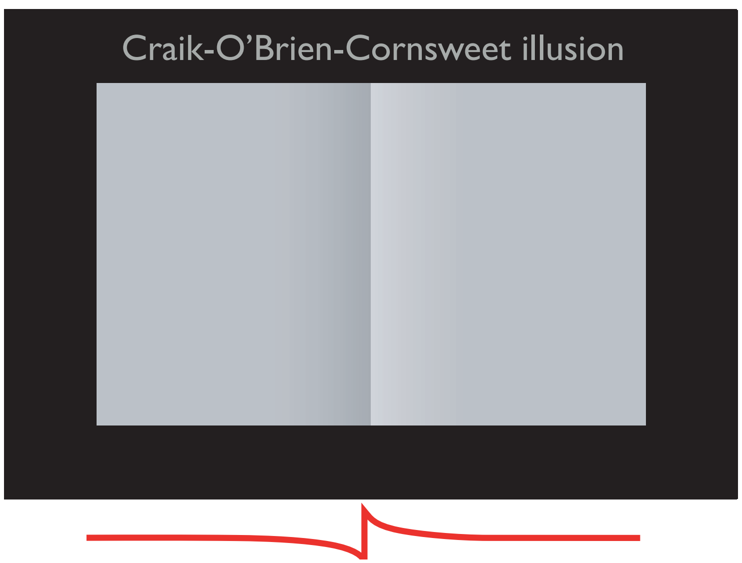

What are the features that are being linked?

image = f(pigment, illumination) ~ r(x,y) × e(x,y)

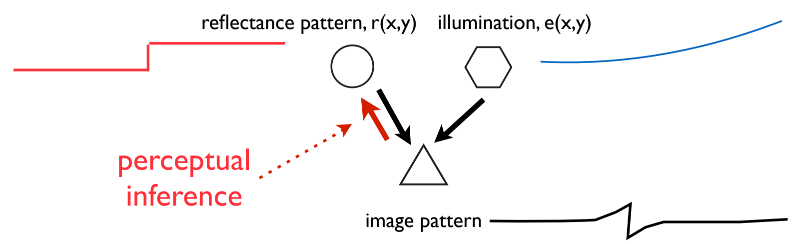

estimate pigment property--the *reflectance*, and discount illumination

prior probabilistic structure:
- illumination spatially smooth
- reflectance is piece-wise constant.
- E.g. gibbs sampler texture demo

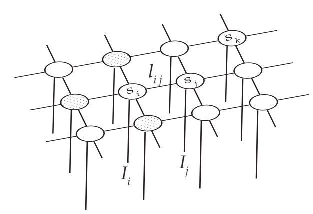

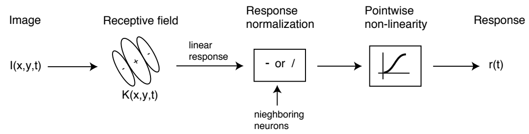

Convolution

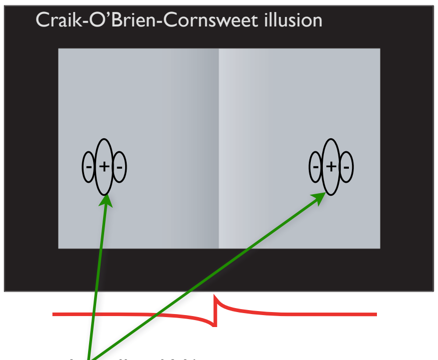

V1 response follows perceived lightness, not physical intensity

Localized V1 responses here should be the same with standard feedforward model

Purely lateral? Don't know. But neuroimaging effect persists with when attention is diverted.

Boyaci, H., Fang, F., Murray, S. O., & Kersten, D. (2007). Current Biology, 17(11), 989–993.

lateral organization

Why the organization? The level of abstraction?

- Keep similar features together for feedforward integration.
- Lateral computations to group features of similar type—segmentation
- Efficiency constraints
  - Minimum wiring constraint
  - Efficient representation of sensory input & cost of neural activity
  - Efficient representations for learning

how can receptive field weights be learned?

both unsupervised, and supervised learning methods

Unsupervised learning of receptive fields

- Unsupervised learning assumes there is statistical structure to be discovered in the sensory input
- Exploit regularities in natural image input to either reduce redundancy or dimensionality, or reduce #active neurons with minimal loss of information.

"efficient coding theories"

Types of structure

1rst order

What to do with the structure?

Recode to eliminate it

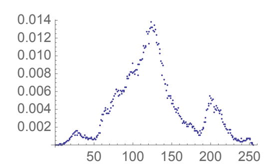

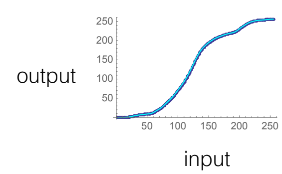

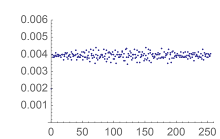

Can no longer predict, *a priori*, that the response to any intensity is probably localized near a mode.

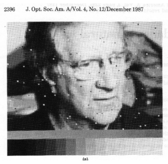

Kersten, D. (1987). Predictability and redundancy of natural images. J Opt Soc Am A, 4(12), 2395-2400.

Types of structure

2nd order

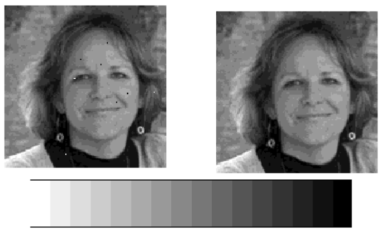

Pixel colors can predict the colors of their neighbors

Gives rise to neural network models that are closely related to principles of image compression developed in signal processing theory, as in "difference coding"

R(x) = L(x) - L(x-1)

which exploits the observation that L(x) is often ~ L(x-1)

this looks like lateral inhibition!

$$R(x) = L(x) - \sum_{x' \neq x} w(x - x') L(x')$$

Types of structure

2nd order

Dimensionality reduction via Principal Components Analysis (PCA)

2 pixel example:

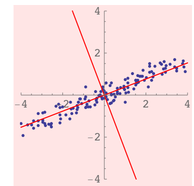

decorrelates the input

and provides the basis for throwing out dimensions

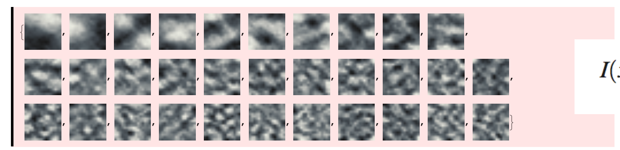

$$I(x,y) = \sum_{i=1}^{n} A_i(x,y) s_i$$

Principal Components Analysis (PCA) with neural networks

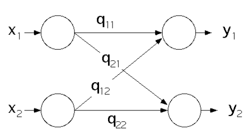

Hebbian learning + Oja's rule to normalize weights:

$$\Delta q_{ij} = \alpha \left(x_j y_i - q_{ij} y_i^2 \right)$$

Oja's rule automatically normalizes:

$$\sum_{i,j} q_{ij}^2 = 1$$

…but because of symmetry, this network will only pull out the first principal component, and does it twice (in this case)

A solution?

$$\Delta q_{ij} = \alpha \left(x_j y_i - y_i \sum_{k=1}^{i} q_{kj} y_k \right)$$

Sanger, T. (1989). Optimal unsupervised learning in a single-layer linear feedforward neural network. Neural Networks, 2, 459-473.

…but this still seems dissatisfying because one neuron would do lots of work, the next less so, and the next even less, etc..

A solution?

“autoencoder networks”

use backprop to find weights that encourage L to predict its own values: input L close to the output L’ :

training pairs: $\{L_i, L_i\}$

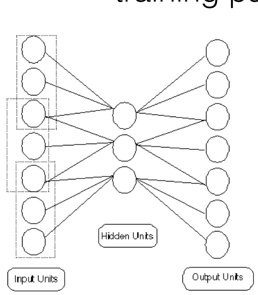

$L \sim L'$

finds subspace that captures larger fraction of the variance

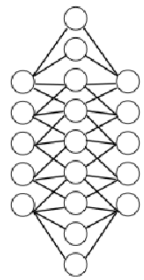

reduce or expand dimensionality

In PCA, the number of basis functions or vectors is less than or equal to the dimensionality of the input

But what if “efficiency” has another meaning, e.g. represent the input with as few features as possible?

…and we allow for over-complete representations where the number of feature detectors could be more than the dimensionality of the input

$$I(x,y) = \sum_{i=1}^{n} A_i(x,y) s_i \qquad s_i = \sum_{x,y} W_i(x,y) I(x,y).$$

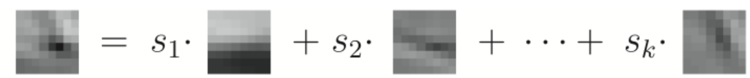

(see Lecture 5)

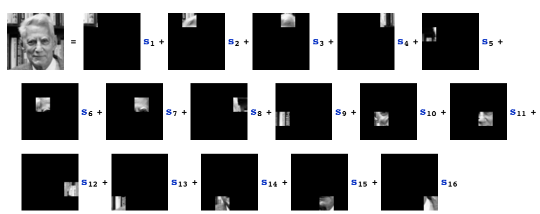

only a few features required for one image…but what if we wanted to have a set of features, or “dictionary” that was in “good” for all natural images?

Good, efficient representation is interpreted as finding the receptive field weights that minimize the sum of squared errors AND # active neurons

so given L(x,y) in a set of images find the $A_i(x,y)$’s that minimize:

$$[L(x,y) - \sum_i s_i A_i(x,y)]^2 + \sum_i B(s_i)$$

the $A_i(x,y)$’s

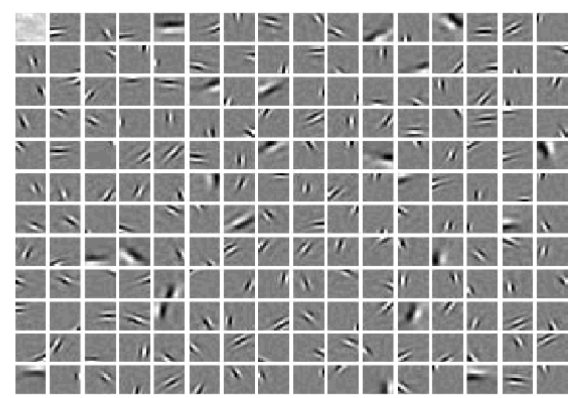

Olshausen & Field’s model of V1 receptive fields

captures localized sensitivities to orientation and spatial frequency

Higher-order structure?

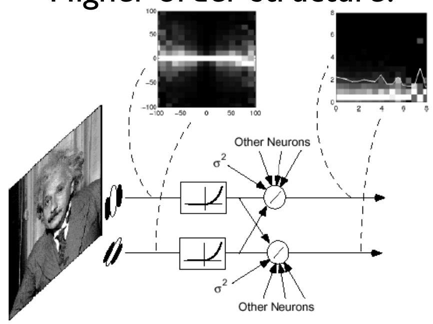

Figure 1: Illustration of image statistics as seen through two neighboring receptive fields. Left image: Joint conditional histogram of two linear coefficients. Pixel intensity corresponds to frequency of occurrence of a given pair of values, except that each column has been independently rescaled to fill the full intensity range. Right image: Joint histogram of divisively normalized coefficients (see text).

responses of linear model neurons with receptive fields that are close in space, preferred orientation or spatial frequency are not statistically independent

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. Nature Neuroscience, 4(8), 819–825.

Higher-order structure?

Accounts for neurophysiological responses of neurons in V1.

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. Nature Neuroscience, 4(8), 819–825.

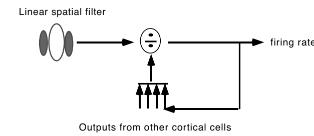

$$R_i = \sigma\left(\sum_{j=1}^{n} w_{ij}\, L_j\right) \Big/ \sum_{k \in N_i} R_k^2$$

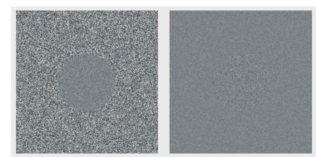

From Heeger

The middle disks have the same physical luminance variance, but the one on the right appears more "contrasty", i.e. to have higher variance.

This may be a behavioral consequence of an underlying non-linearity in the spatial filtering properties of V1 neurons involving "divisive normalization" derived from measures of the activity of other nearby neurons.

More on decorrelation:

contingent adaptation

Contingent Adaptation: McCollough effect

Out[215]=

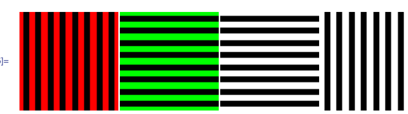

McCollough, C. (1965, 3 September 1965). Color Adaptation of Edge-Detectors in the Human Visual System. Science, 149, 1115-1116.

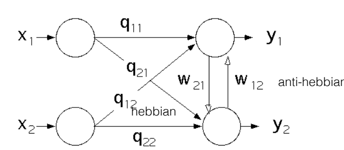

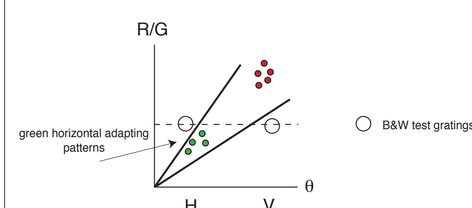

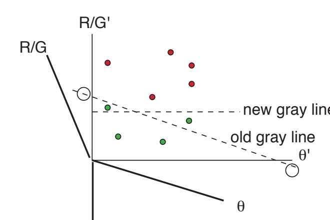

Decorrelation due to adaptation

Barlow, H. B., & Foldiak, P. (1989). Adaptation and decorrelation in the cortex. In C. Miall, R. M. Durban, & G. J. Mitchison (Ed.), The Computing Neuron Addison-Wesley.

Lateral organization

How do neural populations represent information?

Assumption: lateral organization involves features at the same level of abstraction

Mathematica notebook