Some background for next week:

Learning network parameters

Models of neural interaction

- Modeling neural interactions
  - Markov Random Fields -> neural interactions
- Applications to visual surface perception
  - surface interpolation, texture modeling

Models of neural interactions

- Theory
- The role of lateral, local interactions in perception
  - filling-in, texture analysis, normalization..

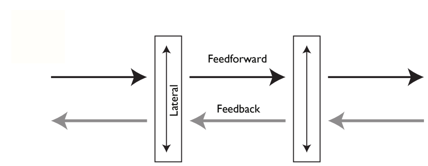

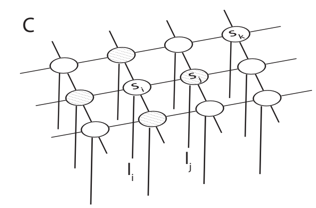

$$p(S_i|S_1, S_2, ...S_n) = p(S_i|S_j, j \in N_i)$$

Applications in vision

- S represents a surface property
  - e.g. "intrinsic images": depth, shape, lightness, …
- S represents image intensity: texture models
- S represents neural activity V that in turn represents an inferred surface property

$$p\left(V_i \mid V_j,\ j \in N_i\right) = \kappa\, e^{-\sum_{j \in N(i)} f\left(V_i - V_j\right)}$$

Examples of 2D interpolation

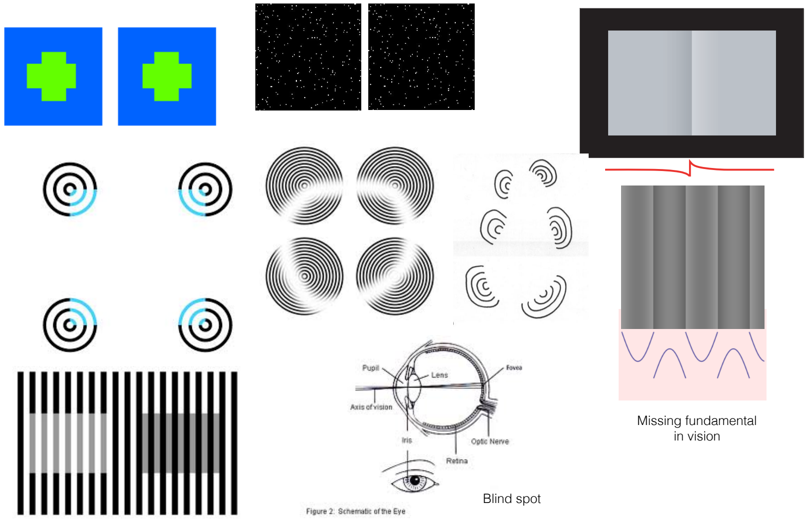

Applications in texture learning and synthesis

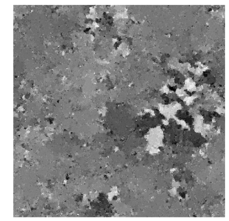

Zhu & Mumford, 1997

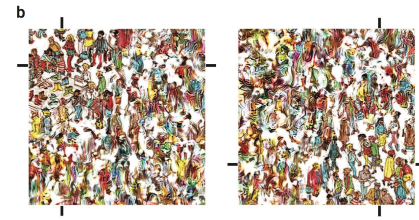

Freeman & Simoncelli (2011)

What good are probabilistic models of “unit” or neural interactions for studies of biological vision?

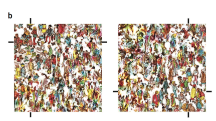

- Stimulus generation based on physical interactions ✓
- Explaining perceptual grouping in terms of priors on natural surface properties ✓
- Models of neural interactions
  - theoretical framework ✓
  - experimental predictions ?

Learning network parameters

- Unsupervised
  - data is a collection of inputs, e.g. images
- Supervised
  - data is a collection of input/output pairs

    e.g. an image and its depth map
- Learning depends on an underlying an inference algorithm

Inference algorithm review

- Boltzmann machine

$$P_{BM}(\vec{x},\vec{I}) = \frac{1}{Z}\exp\{\sum_{ij} T_{ij} x_i I_j + (1/2)\sum_{ij}\theta_{ij} x_i x_j\}$$

- Restricted Boltzmann machine
  - no interactions between hidden units
  - independent factors

$$P_{RBM}(\vec{x}|\vec{I}) = \frac{1}{Z_T}\exp\{\sum_{ij} T_{ij} x_i I_j\} = \prod_{i=1}^{n} P(x_i|\vec{I}),$$

Inference review

Inference

- can be done by drawing samples, e.g. Gibbs sampling
- estimating the mode or mean
  - annealing (slow)
  - mean field theory algorithms

Unsupervised learning

- Values of hidden units determined by the visual inputs
- Hidden units represent internal"explanations" of the inputs
- Original Boltzmann algorithm:
  - The weights get adjusted through experience to move P(x) close to P'(x) where
    - P(V) - probability of visible units taking on certain values determined by visual input from the environment
    - P'(V) - probability that the visible units take on certain values while the network is running without visual input

Make weights explicit

$$\{I^\mu : \mu = 1, ..., N\}$$

$$P(I|\lambda) = \sum_x P(I,x|\lambda)$$

$$\lambda^* = \arg\max_\lambda \prod_{\mu=1}^{N} \sum_{x^\mu} P(x^\mu, I^\mu|\lambda)$$

Notation: x includes visible V, and invisible units.

λ includes the weights from visible units and the ones between invisible units

Supervised learning

- Regression
  - Binary regression & the artificial neuron
  - Linear regression
  - Non-linear regression
    - polynomials, RBFs
    - perceptron and "back-prop"

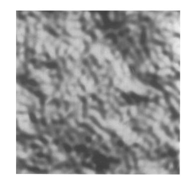

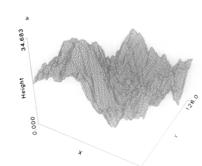

Knill, D. C., & Kersten, D. (1990). Learning a near-optimal estimator for surface shape from shading. Computer Vision, Graphics, and Image Processing, 50(1), 75–100.

Next week

- Deep-belief networks
  - Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. Neural Computation, 18(7), 1527–1554.
- Experimental support for neural networks that learn the statistics of their input
  - Berkes, P., Orban, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. Science, 331(6013), 83–87.