CSE/NEUBEH 528

Lecture 12: Unsupervised Learning

(Chapters 8 & 10)

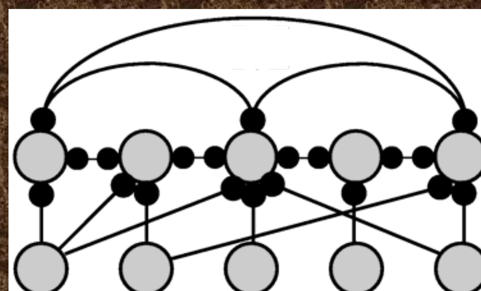

Image from http://clasdean.la.asu.edu/news/images/ubep2001/neuron3.jpg
Lecture figures are from Dayan & Abbott's book
http://people.brandeis.edu/~abbott/book/index.html

Gameplan for Today

(Copyright, Warner Brothers)

- Unsupervised (Representational) Learning
  - Hebb rule and Principal Component Analysis (PCA)
  - Causal Models
  - Generative versus Recognition Models
  - Density Estimation and EM
  - Sparse Coding & Independent Component Analysis (ICA)

Flashback: Hebb Rule

- Consider a linear neuron:

  $v = \mathbf{w}^T \mathbf{u} = \mathbf{u}^T \mathbf{w}$

- Basic Hebb Rule: $\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u} v$
- Pure Hebb only increases synaptic weights (LTP)
  - What about LTD?
- Covariance rules: $\tau_w \frac{d\mathbf{w}}{dt} = (\mathbf{u} - \boldsymbol{\theta}_{\mathrm{u}}) v$

  $$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u}(v - \theta_v)$$

Are these learning rules stable?

On Board Analysis, leading up to Oja's rule

Variants of the Hebb Rule

- Pure Hebb $\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u}v$
- Covariance rules: $\tau_w \frac{d\mathbf{w}}{dt} = (\mathbf{u} - \boldsymbol{\theta}_u)v$

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u}(v - \theta_v)$$

- Oja's Rule: $\tau_w \frac{d\mathbf{w}}{dt} = (\mathbf{u} - \alpha \mathbf{w} v)v$ (stable, $\|w\|^2 \rightarrow 1/\alpha$)

Hebbian Learning in the Brain

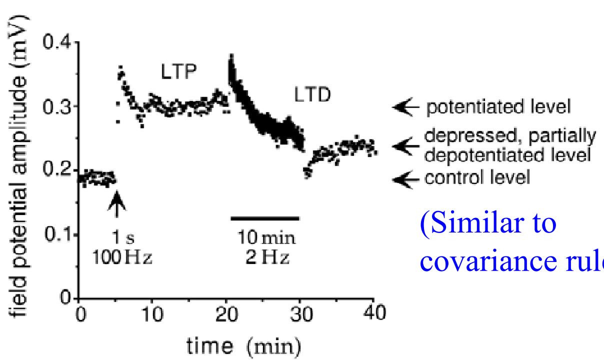

(Similar to covariance rule)

LTP = Long Term Potentiation
LTD = Long Term Depression

What does Hebbian Learning do?

- Consider a linear neuron:

  $v = \mathbf{w}^T \mathbf{u} = \mathbf{u}^T \mathbf{w}$
- Basic Hebb Rule: $\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u} v$
- What is the average effect of this rule over many inputs?

  $$\tau_w \frac{d\mathbf{w}}{dt} = \langle \mathbf{u} v \rangle = Q\mathbf{w}$$
- Q is the input correlation matrix: $Q = \langle \mathbf{u}\mathbf{u}^T \rangle$

What does Hebbian Learning do?

Eigenvector analysis of Hebb rule…

Hebb Rule implements PCA!

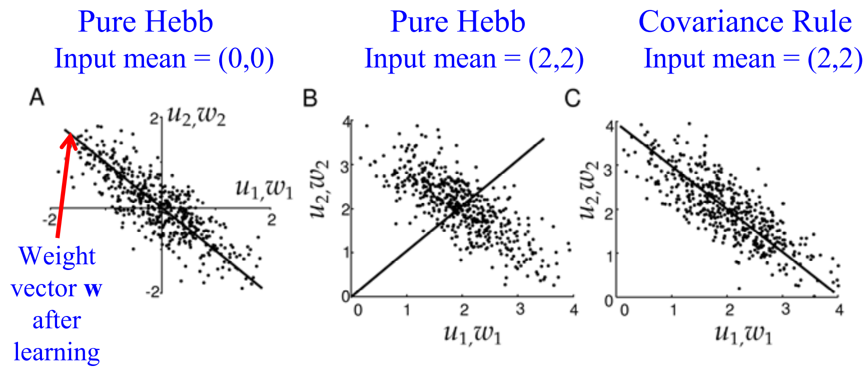

Hebb rule *rotates* weight vector to align with principal eigenvector of input correlation/covariance matrix (i.e. direction of maximum variance)

What about this data?

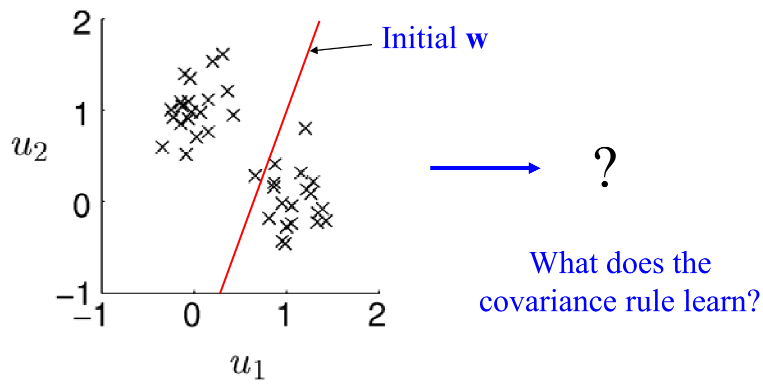

PCA does not correctly describe the data

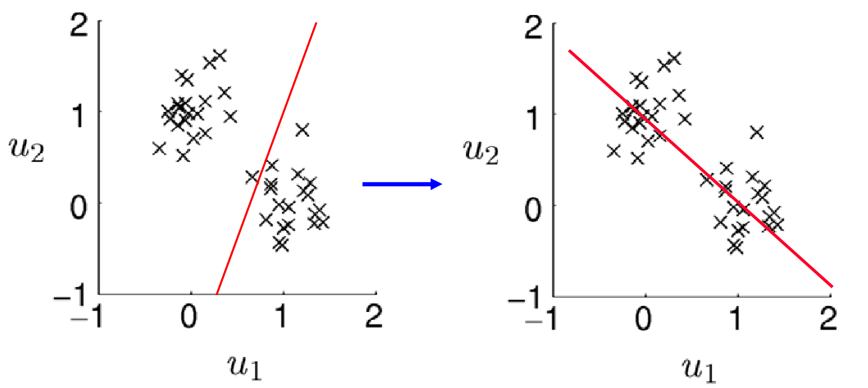

Input data is made up of two clusters (Gaussians) → two "causes"

Causal Models

- Main goal of unsupervised learning: Learn the "Causes" underlying the input data
- Example: Learn the means and variances of the two Gaussians A and B that generated this data
- Want: Two neurons A and B that learn the means and variances based solely on input data (samples from distribution)

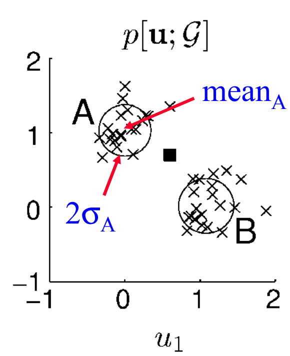

Generative versus Recognition Models

Causes $\mathbf{v}$ $p[\mathbf{v} \mid \mathbf{u}; G]$ (posterior)

Generative model

Recognition model

Data $\mathbf{u}$ $p[\mathbf{u} \mid \mathbf{v}; G]$ (data likelihood)

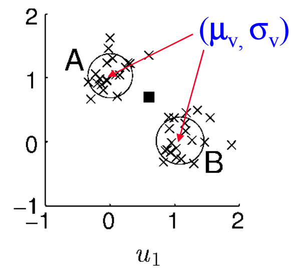

$$p[\mathbf{u}; G] = \sum_{\mathbf{v}} p[\mathbf{u} \mid \mathbf{v}; G] p[\mathbf{v}; G]$$

Prior $= \gamma_v$

Parameters $G = (\mu_v, \sigma_v, \gamma_v)$

EM algorithm for Learning Data Clusters

- Stands for Expectation-Maximization algorithm
- Repeat the following two steps until convergence:
  - E step: Compute recognition distribution (v = A or B) for each $\mathbf{u}$:

$$p[v \mid \mathbf{u}; G] = \frac{p[\mathbf{u} \mid v; G] p[v; G]}{p[\mathbf{u}; G]} = \frac{N(\mathbf{u}; \boldsymbol{\mu}_v, \sigma_v I) \cdot \gamma_v}{\sum_v N(\mathbf{u}; \mu_v, \sigma_v I) \cdot \gamma_v} \quad \text{(Bayes rule)}$$

  - M step: Change parameters G using results from E step

$$\gamma_v = \sum_{\mathbf{u}} p[v \mid \mathbf{u}; G] / N_u, \quad \boldsymbol{\mu}_v = \frac{\sum_{\mathbf{u}} p[v \mid \mathbf{u}; G] \cdot \mathbf{u}}{\sum_{\mathbf{u}} p[v \mid \mathbf{u}; G]},$$

(Learn parameters)

$$\sigma_v^2 = \frac{\sum_{\mathbf{u}} p[v \mid \mathbf{u}; G] \, |\mathbf{u} - \boldsymbol{\mu}_v|^2}{\sum_{\mathbf{u}} p[v \mid \mathbf{u}; G]}$$

Results from the EM algorithm

Input data:

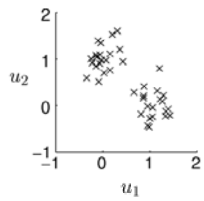

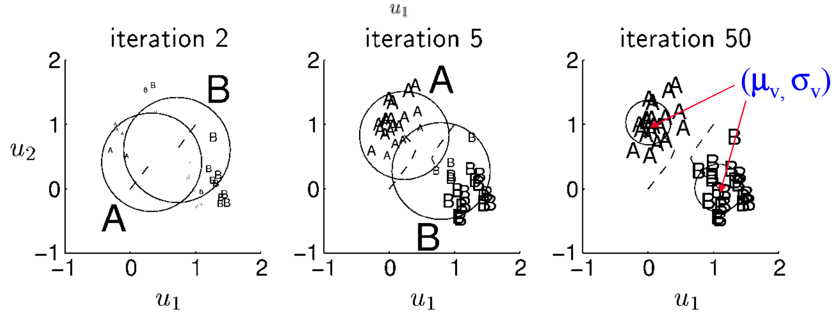

Next Class: Supervised Learning

- Things to do:
  - Finish reading Chapters 8 and 10
  - Do Homework #4 (last homework!)
  - Start mini-project