

Machine Learning

Unsupervised Methods

Part 3

Sepp Hochreiter

Institute of Bioinformatics

Johannes Kepler University, Linz, Austria

Outline



- 1 Introduction
- 2 Basic Terms and Concepts
- 3 Principal Component Analysis
- 4 Independent Component Analysis
- 5 Factor Analysis
- 6 Scaling and Projection Methods
- 7 Clustering
- 8 Biclustering
- 9 Hidden Markov Models
- 10 Boltzmann Machines

7 Clustering

- 7.1 Mixture Models
- 7.2 k-Means Clustering
- 7.3 Hierarchical Clustering
- 7.4 Similarity-Based Clustering

8 Biclustering

8.1 Types of Biclusters

8.2 Overview of Biclustering Methods

8.3 FABIA Biclustering

8.4 Examples

9 Hidden Markov Models

9.1 Hidden Markov Models in Bioinformatics

9.2 Hidden Markov Model Basics

9.3 Expectation Maximization for HMM: Baum-Welch Algorithm

9.4 Viterby Algorithm

9.5 Input Output Hidden Markov Models

9.6 Factorial Hidden Markov Models

9.7 Memory Input Output Factorial Hidden Markov Models

9.8 Tricks of the Trade

9.9 Profile Hidden Markov Models

10 Boltzmann Machines

10.1 The Boltzmann Machine

10.2 Learning in the Boltzmann Machine

10.3 The Restricted Boltzmann Machine

Chapter 7

Clustering

7 Clustering

- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

8 Biclustering

- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

9 Hidden Markov Models

- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

10 Boltzmann Machines

- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

Clustering is one of the most popular unsupervised techniques

Clusters in the data are regions where observations group together
→ regions of high data density

clusters may correspond to a **prototype** from which observations are obtained via noise perturbations

Clustering extracts structures and can identify new data classes

important application of clustering: **data visualization**

observations are represented by prototypes: **vector quantization**

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Mixture models locally assign in the feature space a component which represents a cluster.

Component j out of l components has parameters like location μ_j and width or shape Σ_j . It has in every case a weight w_j that gives the local probability mass.

generative framework: w_j is the probability $p(j)$ of choosing component j , which has density $p(\mathbf{x} | j, \theta_j)$, where θ_j summarizes the parameters of component j

θ summarize all parameters, which gives the generative model

$$p(\mathbf{x} | \theta) = \sum_{j=1}^l p(j) p(\mathbf{x} | j, \theta_j)$$

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

For clustering, Bayes' formula can be used:

$$p(j | \mathbf{x}, \boldsymbol{\theta}) = \frac{p(\mathbf{x} | j, \boldsymbol{\theta}_j) p(j)}{p(\mathbf{x} | \boldsymbol{\theta})}$$

Observation \mathbf{x} is assigned to the component j with largest posterior $p(j | \mathbf{x}, \boldsymbol{\theta})$

Before an observation was seen, each component or cluster has the prior probability; after observing data some clusters may be more or less probable of having produced the data, therefore the prior probability changes to the posterior.

Mixture components can be Poissons, or negative Binomials as used at our institute for analyzing sequencing data.

Optimizing Mixture Models



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

$$\text{log-likelihood } \ln \mathcal{L} = \sum_{i=1}^n \ln p(\mathbf{x}_i | \boldsymbol{\theta})$$

derivative with respect to parameters $\boldsymbol{\theta}_j$ of component j is

$$\frac{\partial}{\partial \boldsymbol{\theta}_j} \ln \mathcal{L} = \sum_{i=1}^n \frac{1}{p(\mathbf{x}_i | \boldsymbol{\theta})} \sum_{k=1}^l p(k) \frac{\partial}{\partial \boldsymbol{\theta}_j} p(\mathbf{x}_i | k, \boldsymbol{\theta}_k) = \sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\theta}_j) \frac{\partial}{\partial \boldsymbol{\theta}_j} \ln p(\mathbf{x}_i | j, \boldsymbol{\theta}_j)$$

we used Bayes' formula

$$p(j | \mathbf{x}_i, \boldsymbol{\theta}_j) = \frac{p(\mathbf{x}_i | j, \boldsymbol{\theta}_j) p(j)}{p(\mathbf{x}_i | \boldsymbol{\theta})}$$

The derivative of the log-likelihood of the model with respect to the parameters of the j -th component:

posterior expectation of the derivative of the log-likelihood of component j

Mixture of Gaussians



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

We will now consider mixture of Gaussian (MoG): $\theta_j = (\mu_j, \Sigma_j)$

$$p(\mathbf{x}_i | j, \theta_j) = p(\mathbf{x}_i | j, \mu_j, \Sigma_j) \sim \mathcal{N}(\mu_j, \Sigma_j)$$

$$p(\mathbf{x} | \theta) = \sum_{j=1}^l w_j \mathcal{N}(\mu_j, \Sigma_j)$$

$$\sum_{j=1}^l w_j = 1, \quad w_j \geq 0$$

$$\mathcal{N}(\mu_j, \Sigma_j)(\mathbf{x}) = (2\pi)^{-m/2} |\Sigma_j|^{-1/2} \exp\left(-\frac{1}{2} (\mathbf{x} - \mu_j)^T \Sigma_j^{-1} (\mathbf{x} - \mu_j)\right)$$

Exponential distributions like the Gaussians are convenient because the logarithm inverts the exponential function:

$$\ln p(\mathbf{x} | j, \mu_j, \Sigma_j) = -\frac{m}{2} \ln(2\pi) - \frac{1}{2} \ln |\Sigma_j| - \frac{1}{2} (\mathbf{x} - \mu_j)^T \Sigma_j^{-1} (\mathbf{x} - \mu_j)$$

derivatives

$$\frac{\partial}{\partial \mu_j} \ln p(\mathbf{x} | j, \mu_j, \Sigma_j) = \Sigma_j^{-1} (\mathbf{x} - \mu_j)$$

$$\frac{\partial}{\partial \Sigma_j} \ln p(\mathbf{x} | j, \mu_j, \Sigma_j) = \frac{1}{2} (\Sigma_j^T)^{-1} + \frac{1}{2} \Sigma_j^{-T} (\mathbf{x} - \mu_j) (\mathbf{x} - \mu_j)^T \Sigma_j^{-T}$$

Mixture of Gaussians



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

EM-algorithm:

E-step:

$$p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) = \frac{w_j \mathcal{N}(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)(\mathbf{x}_i)}{\sum_{t=1}^l w_t \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)(\mathbf{x}_i)}$$

M-step:

$$w_j^{\text{new}} = \frac{1}{n} \sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)$$

$$\boldsymbol{\mu}_j^{\text{new}} = \frac{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) \mathbf{x}_i}{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)}$$

$$\boldsymbol{\Sigma}_j^{\text{new}} = \frac{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) (\mathbf{x}_i - \boldsymbol{\mu}_j) (\mathbf{x}_i - \boldsymbol{\mu}_j)^T}{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)}$$

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Maximum a posteriori MoG:

- proper prior for Σ_j : Wishart density $\mathcal{W}(\Sigma^{-1} \mid \alpha, \Psi)$
- proper prior for weights w_j : Dirichlet density $\mathcal{D}(\mathbf{w} \mid \gamma)$
- proper prior for means μ_j : Gaussian density $\mathcal{N}(\mu \mid \nu, \eta^{-1} \Sigma)$

$$\mathcal{W}(\Sigma^{-1} \mid \alpha, \Psi) = c(\alpha, \Psi) |\Sigma^{-1}|^{\alpha-(m+1)/2} \exp(-\text{tr}(\Psi \Sigma^{-1}))$$

$$\mathcal{D}(\mathbf{w} \mid \gamma) = c(\gamma) \prod_{j=1}^l w_j^{\gamma_j-1}$$

$$\mathcal{N}(\mu \mid \nu, \eta^{-1} \Sigma) = (2\pi)^{-m/2} |\eta^{-1} \Sigma_j|^{-1/2} \exp\left(-\frac{\eta}{2} (\mu - \nu)^T \Sigma_j^{-1} (\mu - \nu)\right)$$

normalizing constants: $\alpha > (m - 1)/2, c(\gamma), c(\alpha, \Psi)$

“tr” is the trace operator

Mixture of Gaussians: MAP and EM



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

expectation-maximization:

E-step:

$$p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) = \frac{w_j \mathcal{N}(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)(\mathbf{x}_i)}{\sum_{t=1}^l w_t \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)(\mathbf{x}_i)}$$

M-step:

$$w_j^{\text{new}} = \frac{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) + \gamma - 1}{n + l (\gamma - 1)}$$

$$\boldsymbol{\mu}_j^{\text{new}} = \frac{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) \mathbf{x}_i + \eta \boldsymbol{\nu}_j}{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) + \eta}$$

$$\boldsymbol{\Sigma}_j^{\text{new}} = \left(\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) (\mathbf{x}_i - \boldsymbol{\mu}_j) (\mathbf{x}_i - \boldsymbol{\mu}_j)^T + \right.$$

$$\left. \eta (\boldsymbol{\nu}_j - \boldsymbol{\mu}_j) (\boldsymbol{\nu}_j - \boldsymbol{\mu}_j)^T + 2 \boldsymbol{\Psi} \right)$$

$$\left(\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) + 2 \alpha - m \right)^{-1}$$

Mixture of Gaussians: Derivative of MAP EM



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Constraints:
$$\sum_{j=1}^l w_j^{\text{new}} = 1$$

Lagrangian for the constrained optimization problem for the w_j :

$$L = \sum_{i=1}^n \sum_{j=1}^l p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) \ln w_j^{\text{new}} + \ln \mathcal{D}(\mathbf{w} | \gamma) - \lambda \left(\sum_{j=1}^l w_j^{\text{new}} - 1 \right)$$

Setting the derivative to zero:

$$\frac{\partial L}{\partial w_j} = \sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) (w_j^{\text{new}})^{-1} + (\gamma - 1) (w_j^{\text{new}})^{-1} - \lambda = 0$$

$$\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) + (\gamma - 1) = \lambda w_j^{\text{new}}$$

Summing over j :
$$\sum_{i=1}^n \sum_{j=1}^l p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) + l(\gamma - 1) = \lambda \sum_{j=1}^l w_j^{\text{new}}$$

$$n + l(\gamma - 1) = \lambda$$

We obtain

$$w_j^{\text{new}} = \frac{\sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) + (\gamma - 1)}{n + l(\gamma - 1)}$$

Mixture of Gaussians: Derivative of MAP EM



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Derivatives of the posterior with respect to parameters of mixture components are set to zero:

$$\frac{\partial L}{\partial \mu_j} = \sum_{i=1}^n p(j | \mathbf{x}_i, \mu_j, \Sigma_j) \Sigma_j^{-1} (\mathbf{x}_i - \mu_j) + \eta \Sigma_j^{-1} (\nu_j - \mu_j) = 0$$

gradient with respect to Σ_j^{-1} must be zero, too:

$$\frac{\partial L}{\partial \Sigma_j} = \frac{\partial L}{\partial \Sigma_j^{-1}} \frac{\partial \Sigma_j^{-1}}{\partial \Sigma_j} = -\Sigma_j^{-2} \frac{\partial L}{\partial \Sigma_j^{-1}}$$

We obtain

$$\begin{aligned} \frac{\partial L}{\partial \Sigma_j^{-1}} &= \frac{1}{2} \sum_{i=1}^n p(j | \mathbf{x}_i, \mu_j, \Sigma_j) \left(\Sigma_j - (\mathbf{x}_i - \mu_j) (\mathbf{x}_i - \mu_j)^T \right) + \\ &\frac{1}{2} \left(\Sigma_j - \eta (\nu_j - \mu_j) (\nu_j - \mu_j)^T \right) + \Sigma_j (\alpha - (m + 1)/2) - \Psi = 0 \end{aligned}$$

where we used $\frac{\partial \ln |U|}{\partial U} = U^{-1}$ for $U = \Sigma_j^{-1}$

Mixture of Gaussians: MAP Hyperparameters



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Default hyperparameters:

$$\alpha = \frac{m}{2}$$

$$\Psi = \frac{1}{2}I \quad \text{OR} \quad \Psi = \frac{1}{2}\text{covar}(\mathbf{x})$$

$$\gamma = 1$$

$$\eta = 0$$

$$\nu_j = \text{mean}(\mathbf{x}_j)$$

A prior on the mean is in most cases not useful except a preferred region is known.

The posterior $p(j | \mathbf{x}_i, \mu_j, \Sigma_j)$ can be used for clustering: data belongs to the cluster for which the posterior is largest.

Mixture of Poissons



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

Mixture of Poissons with count data x :
$$p(x) = \sum_{j=1}^l w_j P(x; \lambda_j)$$

P is the **probability mass function** of the Poisson distribution:

$$P(x; \lambda) = \frac{1}{x!} e^{-\lambda} \lambda^x$$

Bayes framework for model selection:
$$\mathbf{w} = (w_1, \dots, w_l)$$
$$\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_l)$$

posterior:
$$p(\mathbf{w}, \boldsymbol{\lambda} | x) = \frac{p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w}) p(\boldsymbol{\lambda})}{\int p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w}) p(\boldsymbol{\lambda}) d\mathbf{w} d\boldsymbol{\lambda}}$$

Mixture of Poissons



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Dirichlet prior $\mathbf{w}^l = (w_2, \dots, w_l)$ $w_1 = 1 - \sum_{j=2}^l w_j$

$$p(\mathbf{w}) = D(\mathbf{w}^l; \boldsymbol{\gamma}) = b(\boldsymbol{\gamma}) \prod_{j=1}^l w_j^{\gamma_j - 1}$$

Each component w_j is distributed according to a beta distribution:

$$\text{mean}(w_j) = \frac{\gamma_j}{\gamma_s}$$

$$\text{mode}(w_j) = \frac{\gamma_j - 1}{\gamma_s - l}$$

$$\text{var}(w_j) = \frac{\gamma_j (\gamma_s - \gamma_j)}{\gamma_s^2 (\gamma_s + 1)}$$

$$\gamma_s = \sum_{j=1}^l \gamma_i$$

Mixture of Poissons



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Prior on λ_j is a uniform distribution on $(0, 1/t]$: $p(\lambda_j) = t$

posterior of the model parameters:

$$\begin{aligned} p(\mathbf{w}, \boldsymbol{\lambda} | x) &= \frac{p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w}) p(\boldsymbol{\lambda})}{\int p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w}) p(\boldsymbol{\lambda}) d\mathbf{w} d\boldsymbol{\lambda}} = \frac{p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w})}{\int p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w}) d\mathbf{w} d\boldsymbol{\lambda}} \\ &= \frac{1}{c(x)} p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w}) \end{aligned}$$

where $c(x)$ is independent of the model parameters

Mixture of Poissons



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

upper bound on log posterior with new variables and **one** data point

$$\hat{w}_j, \sum_{j=1}^l \hat{w}_j = 1$$

$$\begin{aligned} -\log p(\mathbf{w}, \boldsymbol{\lambda} | x) &= -\log(p(x | \mathbf{w}, \boldsymbol{\lambda}) p(\mathbf{w}) / c(x)) = -\log \sum_{j=1}^l w_j P(x; \lambda_j) - \log p(\mathbf{w}) + \log(c(x)) \\ &= -\log \sum_{j=1}^l \frac{\hat{w}_j}{\hat{w}_j} w_j P(x; \lambda_j) - \log p(\mathbf{w}) + \log(c(x)) \leq -\sum_{j=1}^l \hat{w}_j \log \frac{w_j P(x; \lambda_j)}{\hat{w}_j} - \log p(\mathbf{w}) + \log(c(x)) \\ &= -\sum_{j=1}^l \hat{w}_j \log(w_j P(x; \lambda_j)) - \log p(\mathbf{w}) + \sum_{j=1}^l \hat{w}_j \log \hat{w}_j + \log(c(x)) \end{aligned}$$

where we applied Jensen's inequality and for

$$\hat{w}_j = p(j | x, \mathbf{w}, \boldsymbol{\lambda}) = \frac{w_j P(x; \lambda_j)}{p(x | \mathbf{w}, \boldsymbol{\lambda})} \quad \text{we obtain}$$

$\log \frac{w_j P(x; \lambda_j)}{\hat{w}_j} = \log p(x | \mathbf{w}, \boldsymbol{\lambda})$ and the inequality becomes an equality

Mixture of Poissons



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Count data set: $\{x_1, \dots, x_n\}$

Component posterior for component j and count x_i :

$$w_{ji} = p(j | x_i, \mathbf{w}, \boldsymbol{\lambda}) = \frac{p(j) p(x_i | j, \mathbf{w}, \boldsymbol{\lambda})}{p(x_i | \mathbf{w}, \boldsymbol{\lambda})} = \frac{w_j P(x_i; \lambda_j)}{p(x_i | \mathbf{w}, \boldsymbol{\lambda})}$$

We introduce for each x_i variables \hat{w}_{ji} , $\sum_{j=1}^l \hat{w}_{ji} = 1$ which approximate $p(j | \mathbf{w}, x_i, \boldsymbol{\lambda})$ and are independent of parameters

$$\hat{w}_{ji} = \frac{w_j^{\text{old}} P(x_i; \lambda_j^{\text{old}})}{p(x_i; \mathbf{w}^{\text{old}}, \boldsymbol{\lambda}^{\text{old}})}$$

for the estimation the actual parameters “old” are used instead of the optimal parameters

Mixture of Poissons



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

upper bound B on the $1/n$ scaled negative log-posterior for **all** data:

$$B = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^l \hat{w}_{ji} \log(w_j P(x; \lambda_j)) - \frac{1}{n} \log p(\mathbf{w}) + \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^l \hat{w}_{ji} \log \hat{w}_{ji} + \frac{1}{n} \sum_{i=1}^n \log c(x_i)$$

Mixture of Poissons



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Optimizing \mathbf{w} :
$$\min_{\mathbf{w}} -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^l \hat{w}_{ji} \log w_j - \frac{1}{n} \log p(\mathbf{w})$$

s.t.
$$\sum_{j=1}^l w_j = 1$$

Lagrangian:

$$\begin{aligned} L &= -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^l \hat{w}_{ji} \log w_j - \frac{1}{n} \log p(\mathbf{w}) + \rho \left(\sum_{j=1}^l w_j - 1 \right) \\ &= -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^l \hat{w}_{ji} \log w_j - \frac{1}{n} \sum_{j=1}^l (\gamma_j - 1) \log w_j + \rho \left(\sum_{j=1}^l w_j - 1 \right) \end{aligned}$$

Derivatives zero:

$$\frac{\partial L}{\partial w_j} = -\frac{1}{n} \sum_{i=1}^n \hat{w}_{ji} \frac{1}{w_j} - \frac{1}{n} \frac{1}{w_j} (\gamma_j - 1) + \rho = 0 \quad \text{multiplying by } w_j$$

$$-\frac{1}{n} \sum_{i=1}^n \hat{w}_{ji} - \frac{1}{n} (\gamma_j - 1) + \rho w_j = 0$$

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Summation over j : $1 + \frac{1}{n} (\gamma_s - l) = \rho$

Inserting this expression:

$$-\frac{1}{n} \sum_{i=1}^n \hat{w}_{ji} - \frac{1}{n} (\gamma_j - 1) + \left(1 + \frac{1}{n} (\gamma_s - l)\right) w_j = 0$$

$$w_j^{\text{new}} = \frac{\hat{w}_j + \frac{1}{n} (\gamma_j - 1)}{1 + \frac{1}{n} (\gamma_s - l)} \quad \hat{w}_j = \frac{1}{n} \sum_{i=1}^n \hat{w}_{ji}$$

Note that:

$$\begin{aligned} w_j &= p(j) = p(j | \mathbf{w}, \boldsymbol{\lambda}) = \int p(j, x | \mathbf{w}, \boldsymbol{\lambda}) dx = \int p(j | x, \mathbf{w}, \boldsymbol{\lambda}) p(x | \mathbf{w}, \boldsymbol{\lambda}) dx \\ &= \mathbb{E}_{p(x|\mathbf{w},\boldsymbol{\lambda})}(p(j | x, \mathbf{w}, \boldsymbol{\lambda})) \approx \frac{1}{n} \sum_{i=1}^n p(j | x_i, \mathbf{w}, \boldsymbol{\lambda}) = \frac{1}{n} \sum_{i=1}^n w_{ji} = \hat{w}_j \end{aligned}$$

Mixture of Poissons



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Optimizing λ_j :

$$\min_{\lambda_j} \left(-\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^l \hat{w}_{ji} \log P(x_i; \lambda_j) \right)$$

$$\log P(x_i; \lambda_j) = -\log(x_i!) - \lambda_j + x_i \log(\lambda_j)$$

the derivative of the objective with respect to λ_j is

$$-\frac{1}{n} \sum_{i=1}^n \left(-1 + \frac{x_i}{\lambda_j} \right) \hat{w}_{ji}$$

Multiplying by λ_j and solving for it:

$$\lambda_j^{\text{new}} = \frac{\sum_{i=1}^n x_i \hat{w}_{ji}}{\sum_{i=1}^n \hat{w}_{ji}}$$

Mixture of Poissons



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

update rules:

$$\hat{w}_{ji} = \frac{w_j^{\text{old}} P(x_i; \lambda_j^{\text{old}})}{p(x_i | \mathbf{w}^{\text{old}}, \boldsymbol{\lambda}^{\text{old}})}$$

$$w_j^{\text{new}} = \frac{\frac{1}{n} \sum_{i=1}^n \hat{w}_{ji} + \frac{1}{n} (\gamma_j - 1)}{1 + \frac{1}{n} (\gamma_s - l)}$$

$$\lambda_j^{\text{new}} = \frac{\frac{1}{n} \sum_{i=1}^n \hat{w}_{ji} x_i}{\frac{1}{n} \sum_{i=1}^n \hat{w}_{ji}}$$

Mixture of Gaussians: Initialization



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

initialization of the MoG:

- “em”: first several low tolerance fast runs then a precise slow run

- “rnd”: random initializations and pick the best

- “svd”: singular value decomposition to find a good initialization

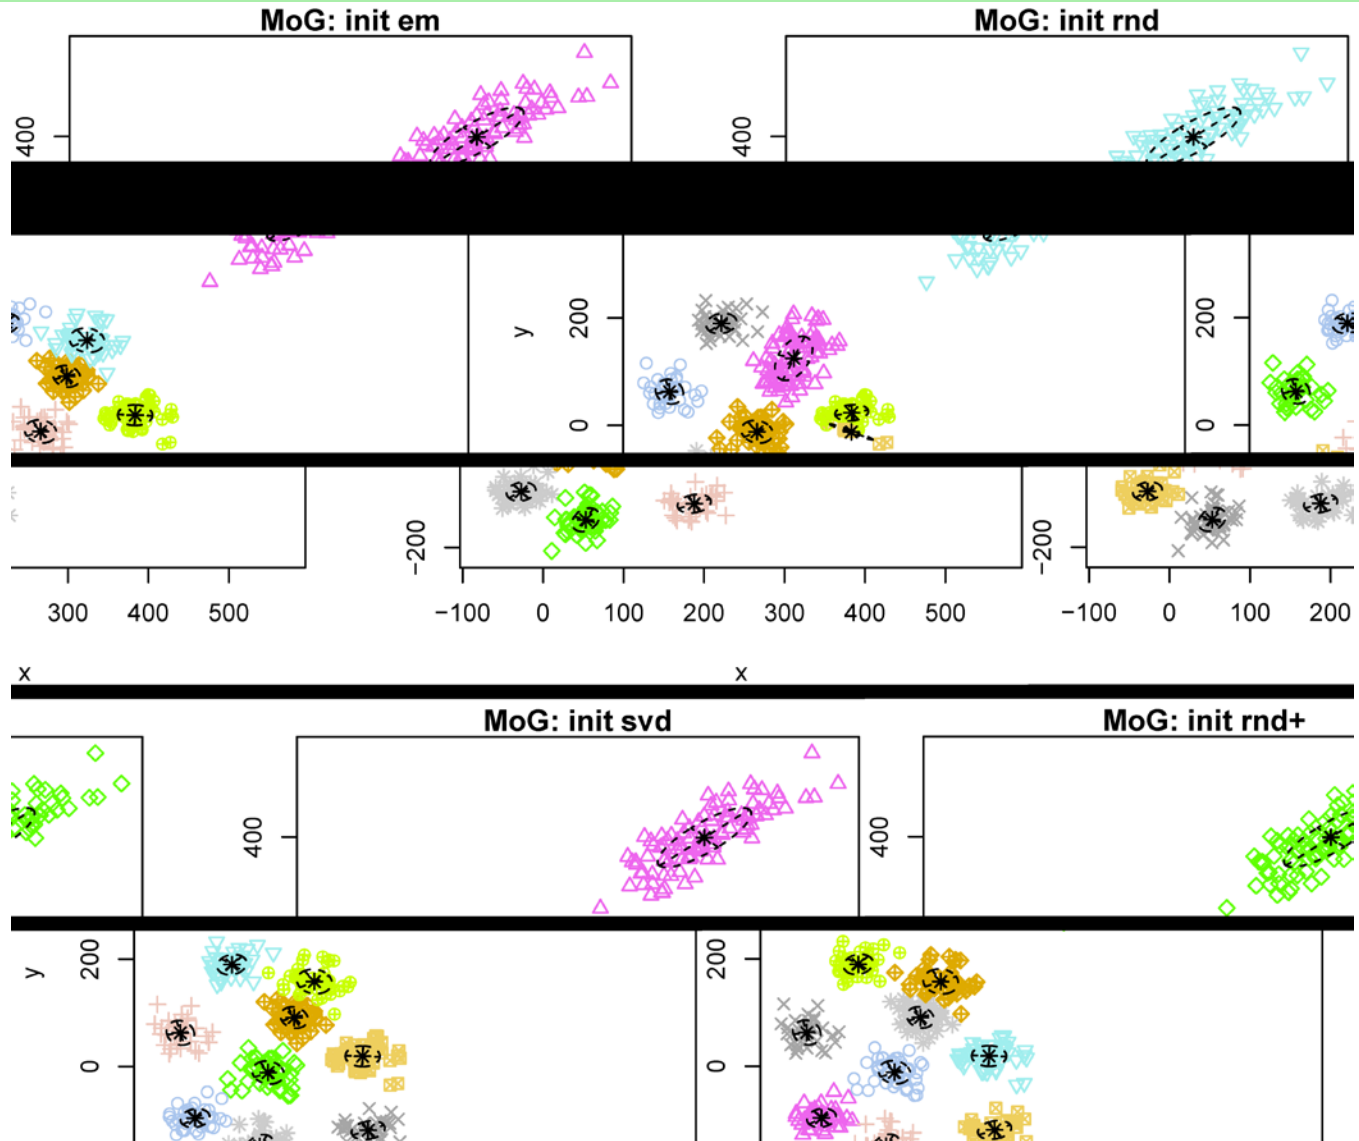
Mixture of Gaussians: Example

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine



Mixture of Gaussians: Constraints



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

MoG for the iris data set

constraints on the parameters:

- “spherical”: covariance matrix is a multiple of the identity
- “diagonal”: diagonal covariance matrix (clusters along axis)
- “volume”: weighting factor or prior for the components

univariate mixture:

"E" = equal variance (one-dimensional)
"V" = variable variance (one-dimensional)

multivariate mixture:

"EII" = spherical, equal volume
"VII" = spherical, unequal volume
"EEI" = diagonal, equal volume and shape
"VEI" = diagonal, varying volume, equal shape
"EVI" = diagonal, equal volume, varying shape
"VVI" = diagonal, varying volume and shape
"EEE" = ellipsoidal, equal volume, shape, and orientation
"EEV" = ellipsoidal, equal volume and equal shape
"VEV" = ellipsoidal, equal shape
"VVV" = ellipsoidal, varying volume, shape, and orientation

single component:

"X" = univariate normal
"XII" = spherical multivariate normal
"XXI" = diagonal multivariate normal
"XXX" = ellipsoidal multivariate normal

Mixture of Gaussians: Example

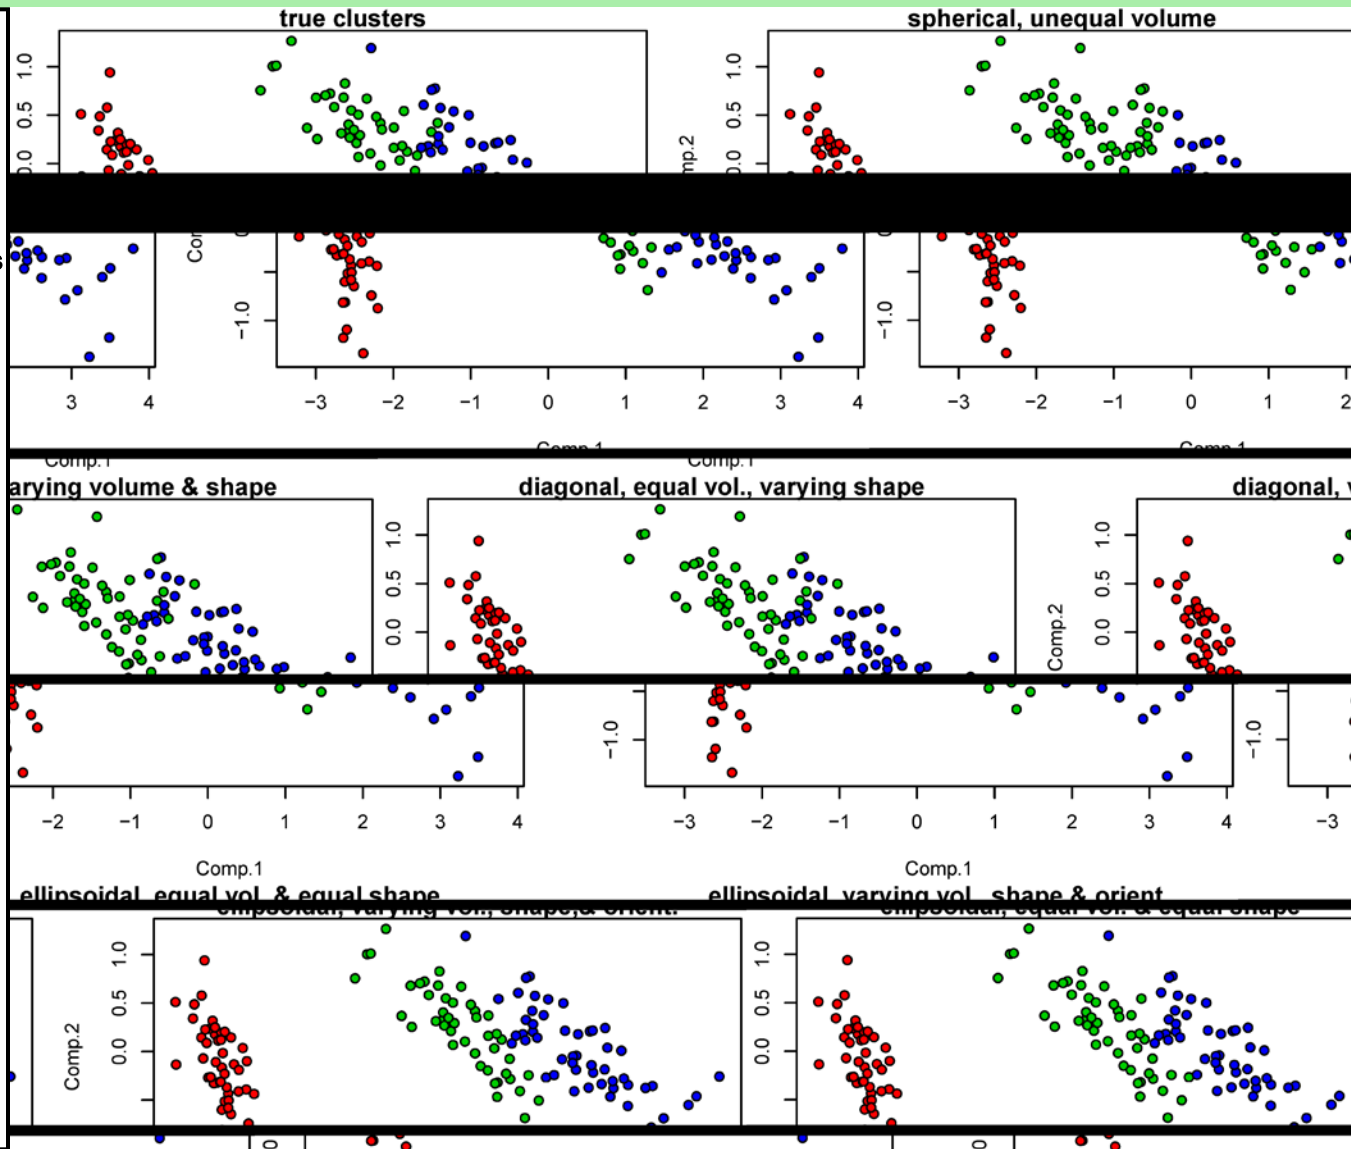


7 Clustering
 7.1 Mixture Models
 7.2 k-Means
 7.3 Hierarchical
 7.4 Similarity-Based

8 Biclustering
 8.1 Types of Biclusters
 8.2 Overview Methods
 8.3 FABIA
 8.4 Examples

9 Hidden Markov Models
 9.1 HMMs in Bioinf.
 9.2 HMM Basics
 9.3 EM / Baum-Welch
 9.4 Viterby
 9.5 Input Output HMM
 9.6 Factorial HMM
 9.7 Memory IOHMM
 9.8 Tricks of the Trade
 9.9 Profile HMM

10 Boltzmann Machines
 10.1 The Boltzmann Machine
 10.2 Learning
 10.3 Restricted Boltzmann Machine



MoG with 3 comp.

spherical, unequal vol.

diagonal, equal volume, varying shape

diagonal, varying volume & shape

ellipsoidal, equal vol. and equal shape

ellipsoidal, varying vol., shape, and orientation

Mixture of Gaussians: Example

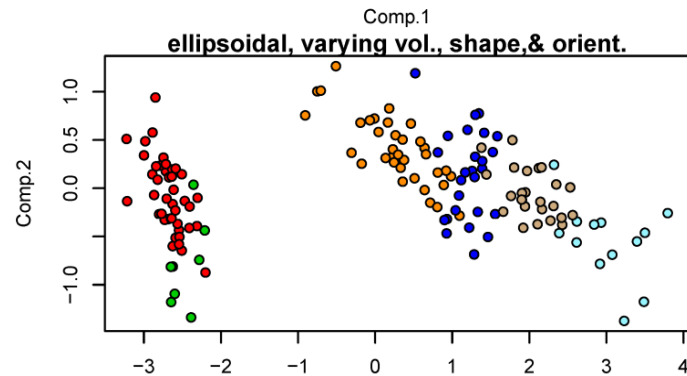


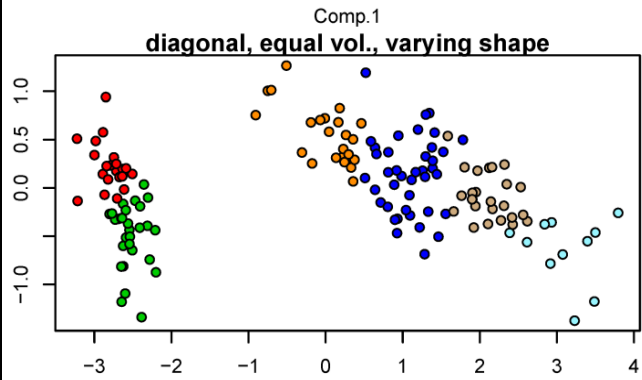

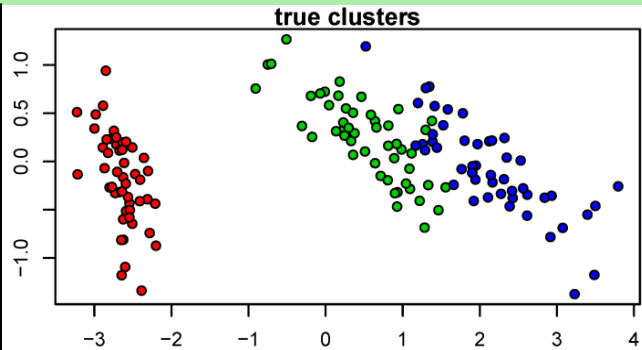


7 Clustering
 7.1 Mixture Models
 7.2 k-Means
 7.3 Hierarchical
 7.4 Similarity-Based

8 Biclustering
 8.1 Types of Biclusters
 8.2 Overview Methods
 8.3 FABIA
 8.4 Examples

9 Hidden Markov Models
 9.1 HMMs in Bioinf.
 9.2 HMM Basics
 9.3 EM / Baum-Welch
 9.4 Viterby
 9.5 Input Output HMM
 9.6 Factorial HMM
 9.7 Memory IOHMM
 9.8 Tricks of the Trade
 9.9 Profile HMM

10 Boltzmann Machines
 10.1 The Boltzmann Machine
 10.2 Learning
 10.3 Restricted Boltzmann Machine



MoG with 6 comp.

spherical, unequal vol.

diagonal, equal volume, varying shape

diagonal, varying volume & shape

ellipsoidal, equal vol. and equal shape

ellipsoidal, varying vol., shape, and orientation

Mixture of Gaussians: Example

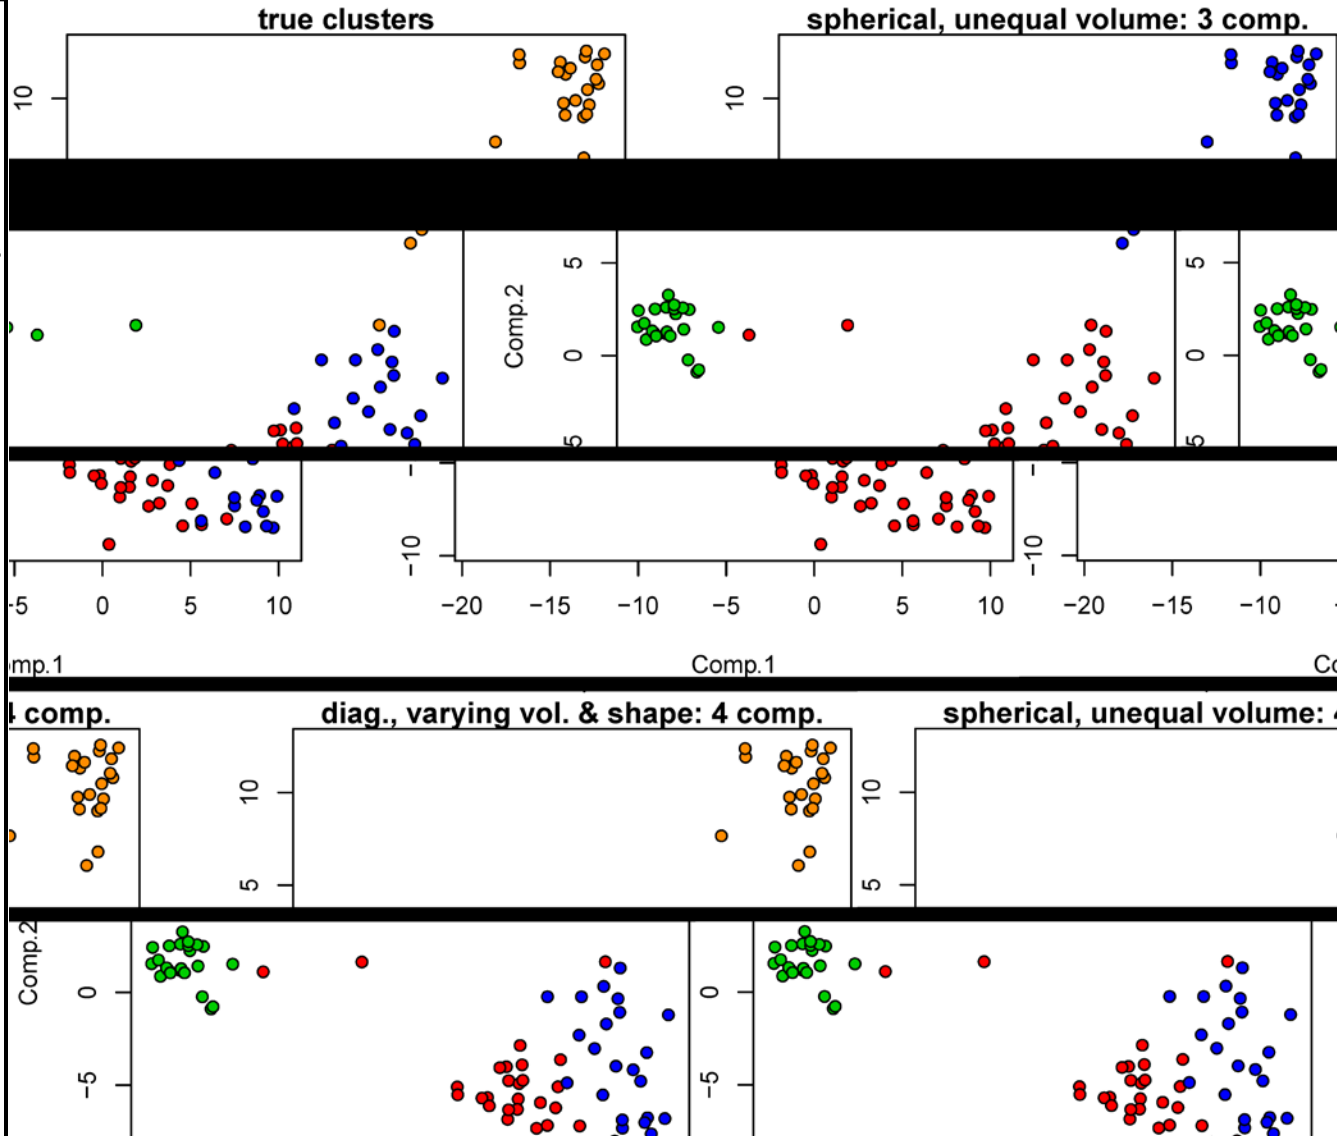


7 Clustering
 7.1 Mixture Models
 7.2 k-Means
 7.3 Hierarchical
 7.4 Similarity-Based

8 Biclustering
 8.1 Types of Biclusters
 8.2 Overview Methods
 8.3 FABIA
 8.4 Examples

9 Hidden Markov Models
 9.1 HMMs in Bioinf.
 9.2 HMM Basics
 9.3 EM / Baum-Welch
 9.4 Viterby
 9.5 Input Output HMM
 9.6 Factorial HMM
 9.7 Memory IOHMM
 9.8 Tricks of the Trade
 9.9 Profile HMM

10 Boltzmann Machines
 10.1 The Boltzmann Machine
 10.2 Learning
 10.3 Restricted Boltzmann Machine



MoG applied to multiple tissues: 76 genes with largest variance

“spherical, unequal volume”

“diagonal, equal volume, varying shape”

Mixture of Gaussians: Example

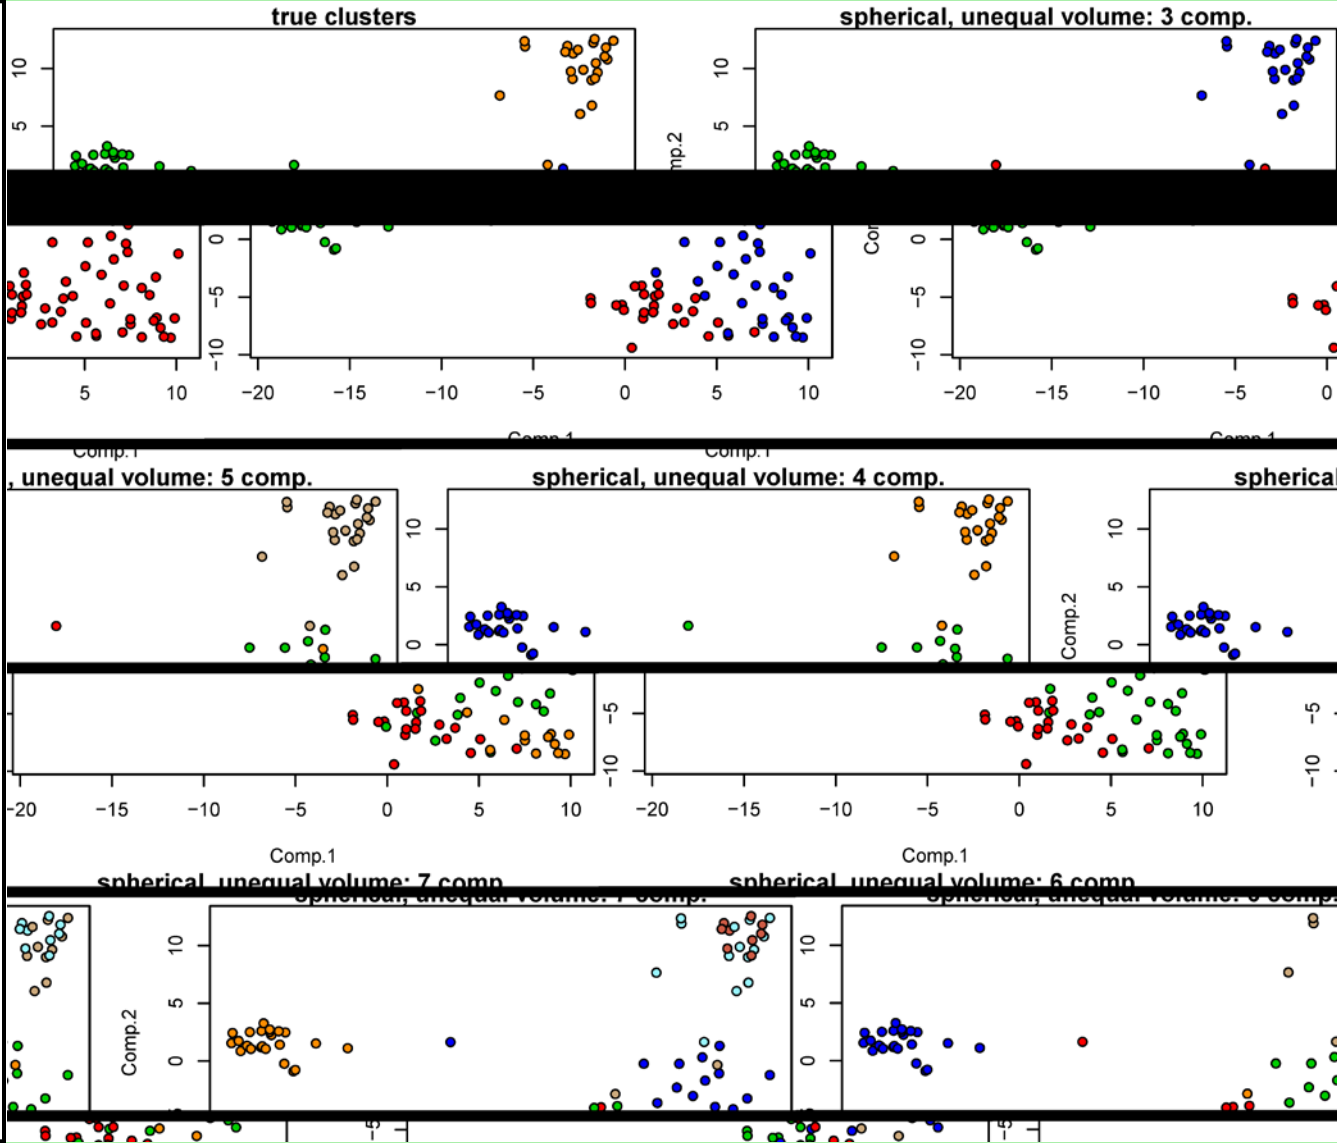


- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine



MoG applied to multiple tissues

spherical compo. with unequal volume and 3 to 7 comp.

k-Means Clustering



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

k-means clustering is probably the best known clustering algorithm

k-means clustering is obtained from mixture clustering if the model assumptions are simplified:

- equal weight (equal volume) for each component: $w_j = \frac{1}{l}$
- spherical and equal (between components) covariance $\Sigma_j^{-1} = I$
- hard (discrete) cluster membership (a sample belongs to a cluster or not)

The only remaining parameters are the cluster centers.

A sample belongs to the cluster with the closest center:

$$p(j | \mathbf{x}_i, \boldsymbol{\mu}_j) = \begin{cases} 1 & \text{if } j = c_{\mathbf{x}_i} = \arg \min_k \|\mathbf{x}_i - \boldsymbol{\mu}_k\| \\ 0 & \text{otherwise} \end{cases}$$

Center updates (according to the mixture model: mean of members):

$$\boldsymbol{\mu}_j^{\text{new}} = \frac{1}{n_j} \sum_{i=1, j=c_{\mathbf{x}_i}}^n \mathbf{x}_i, \quad n_j = \sum_{i=1}^n p(j | \mathbf{x}_i, \boldsymbol{\mu}_j, \Sigma_j) = \sum_{i=1, j=c_{\mathbf{x}_i}}^n 1$$

k-Means Clustering



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Given: data $\{\mathbf{x}\} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, number of clusters l

BEGIN initialization

initialize the cluster centers $\mu_j, 1 \leq j \leq l$

END initialization

BEGIN Iteration

Stop=false

while Stop=false **do**

for ($i = 1 ; i \leq n ; i ++$) **do**

 assign \mathbf{x}_i to the nearest μ_j

end for

for ($j = 1 ; j \leq l ; j ++$) **do**

$$\mu_j^{\text{new}} = \frac{1}{n_j} \sum_{i=1, j=c_{\mathbf{x}_i}}^n \mathbf{x}_i$$

end for

if stop criterion fulfilled **then**

 Stop=true

end if

end while

END Iteration

k-means algorithm

k -Means Clustering



7 Clustering
7.1 Mixture Models
7.2 k -Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

k -means clustering

- fast
- robust (outliers)
- simple (advantage or disadvantage)
- prone to initialization

center near some outliers → center will stay on the outliers even if some cluster are not modeled

k-Means Clustering



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

membership continuous: $p^b(j | \mathbf{x}_i, \boldsymbol{\mu}_j) = \frac{\|\mathbf{x}_i - \boldsymbol{\mu}_j\|^{-2/(b-1)}}{\sum_{k=1}^l \|\mathbf{x}_i - \boldsymbol{\mu}_k\|^{-2/(b-1)}}$
(softmax)

update rule

$$\boldsymbol{\mu}_j^{\text{new}} = \frac{\sum_{i=1}^n p^b(j | \mathbf{x}_i, \boldsymbol{\mu}_j) \mathbf{x}_i}{\sum_{i=1}^n p^b(j | \mathbf{x}_i, \boldsymbol{\mu}_j)}$$

The following objective is minimized:

$$\sum_{j=1}^l \sum_{i=1}^n p^b(j | \mathbf{x}_i, \boldsymbol{\mu}_j) \mathbf{x}_i \|\mathbf{x}_i - \boldsymbol{\mu}_j\|^2$$

This algorithm is called **fuzzy k-means clustering**

k-Means Clustering



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Given: data $\{\mathbf{x}\} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, number of clusters l , parameter b

BEGIN initialization

initialize the cluster centers μ_j , $1 \leq j \leq l$, and $w_j(\mathbf{x}_i) = p(j | \mathbf{x}_i, \mu_j)$ so that $\sum_{j=1}^l w_j(\mathbf{x}_i) = 1$, $w_j(\mathbf{x}_i) \geq 0$.

END initialization

BEGIN Iteration

```
Stop=false
while Stop=false do
```

$$\mu_j^{\text{new}} = \frac{\sum_{i=1}^n w_j(\mathbf{x}_i) \mathbf{x}_i}{\sum_{i=1}^n w_j(\mathbf{x}_i)}$$

$$w_j(\mathbf{x}_i) = \frac{\|\mathbf{x}_i - \mu_j\|^{-2/(b-1)}}{\sum_{k=1}^l \|\mathbf{x}_i - \mu_k\|^{-2/(b-1)}}$$

```
    if stop criterion fulfilled then
        Stop=true
    end if
end while
END Iteration
```

fuzzy
k-means
algorithm

k-Means Clustering

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based


- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

artificial data set in two dimensions with five clusters

- color indicate cluster membership
- filled circles mark the cluster centers






k-Means Clustering

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine



Local minima are shown:

- top row one cluster explains two true clusters while one true cluster is divided into two model clusters
- lower row: three model clusters share one true cluster

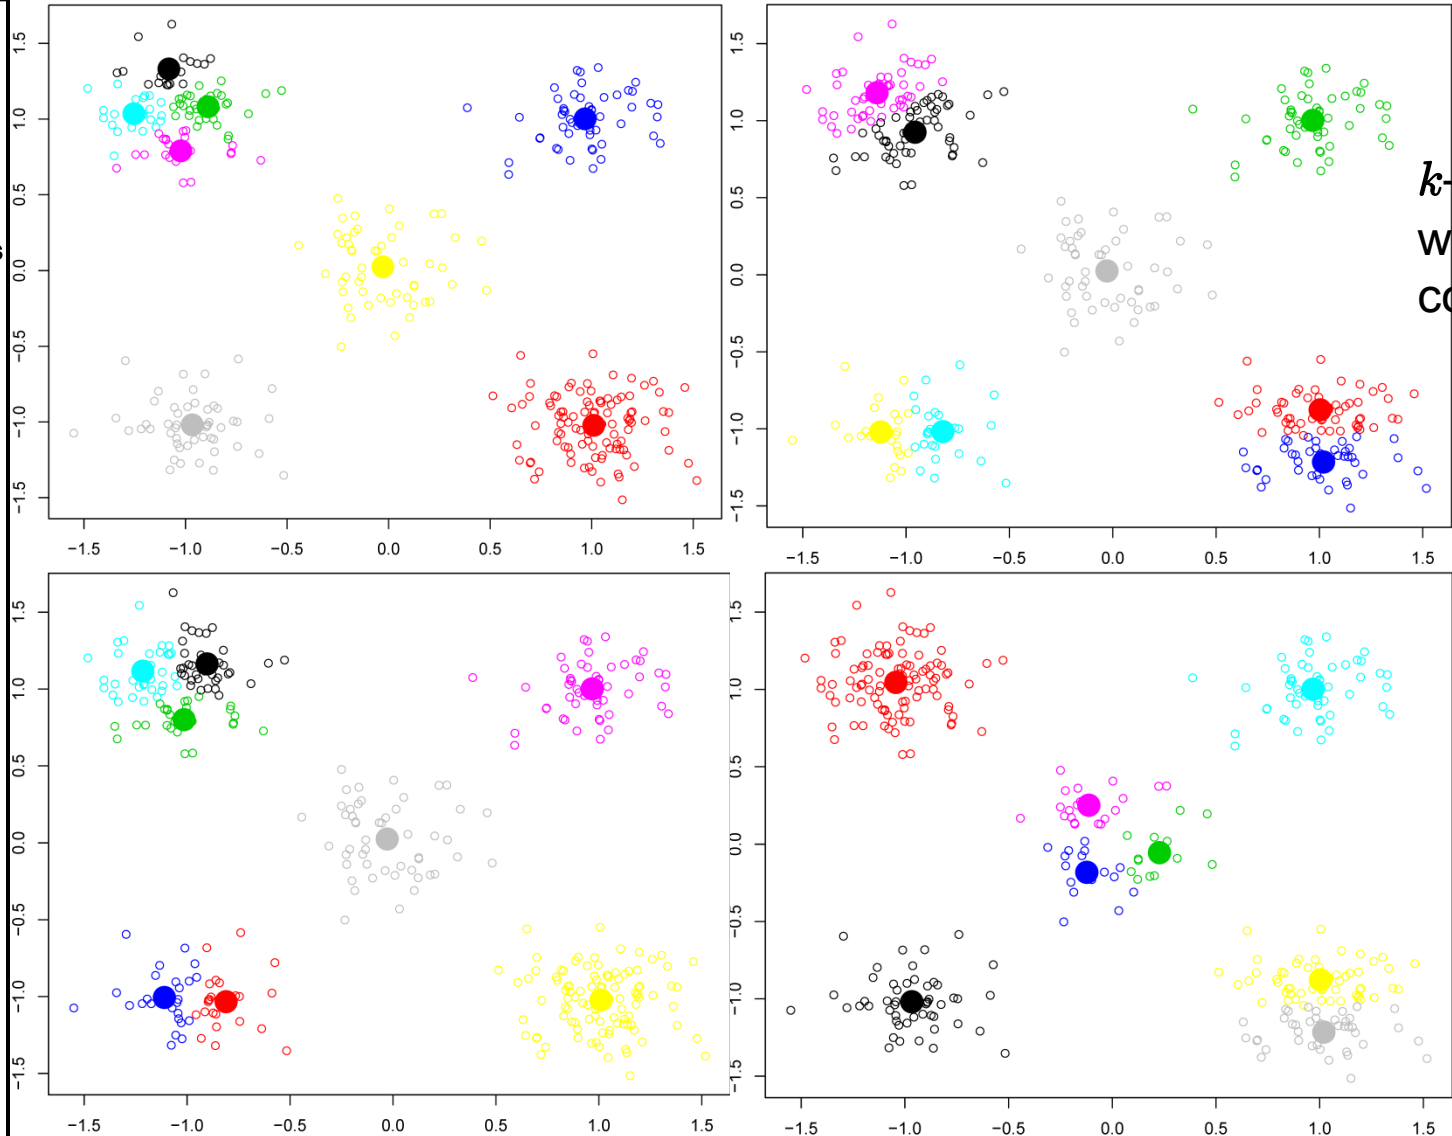
k-Means Clustering

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine



k-means
with 8
comp.




k-Means Clustering

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine



k-means applied to the Iris data

Upper right: typical quite good solution
only errors at class borders


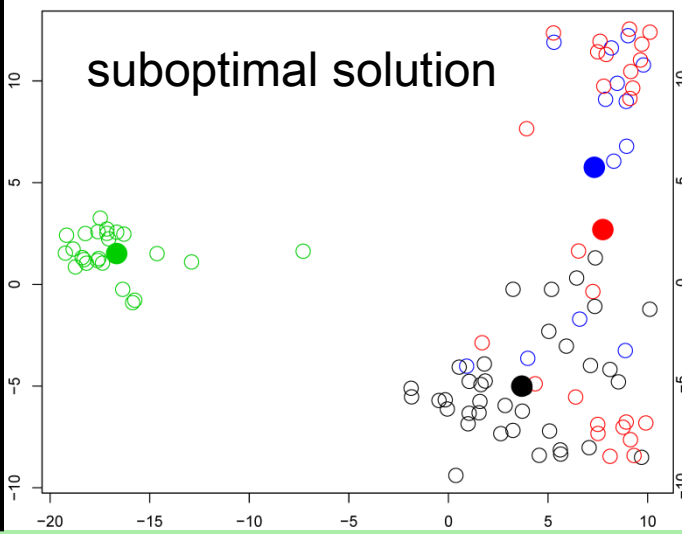
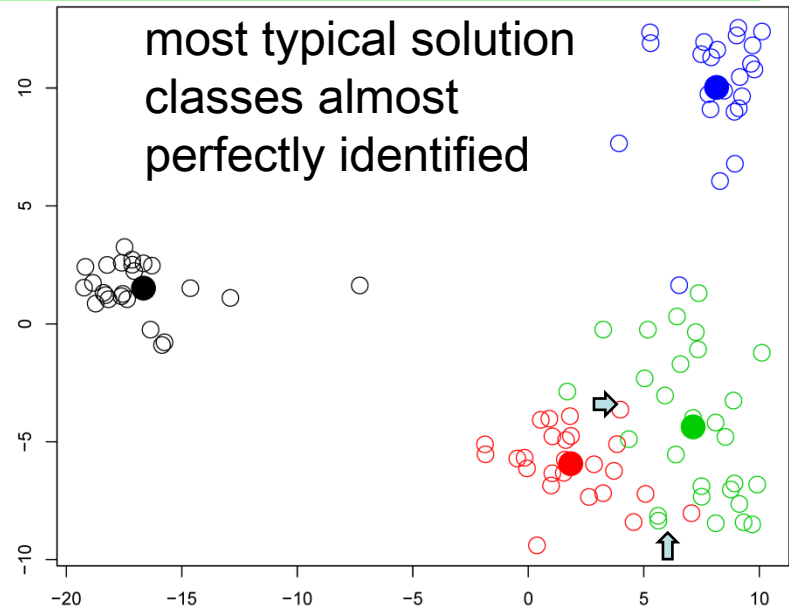
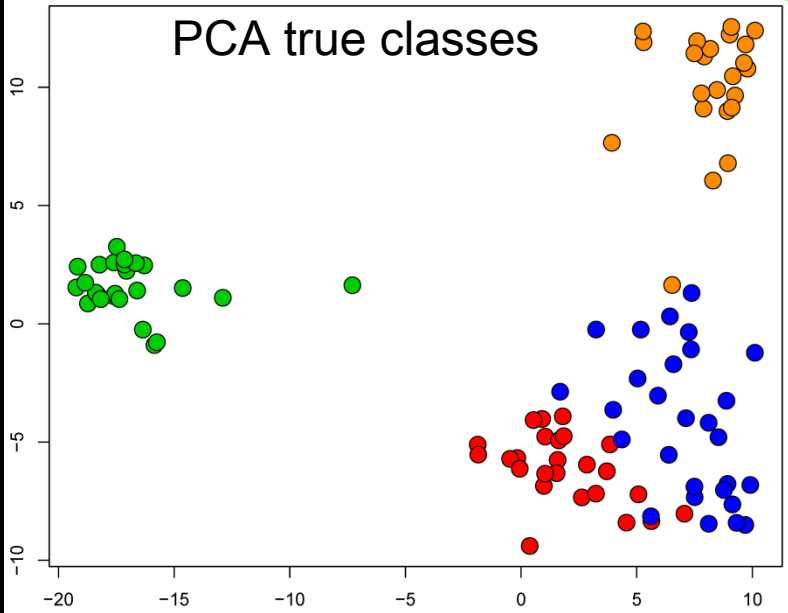
Lower left: another typical solution
which is not as good

Can these solutions be distinguished?

k-Means Clustering



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

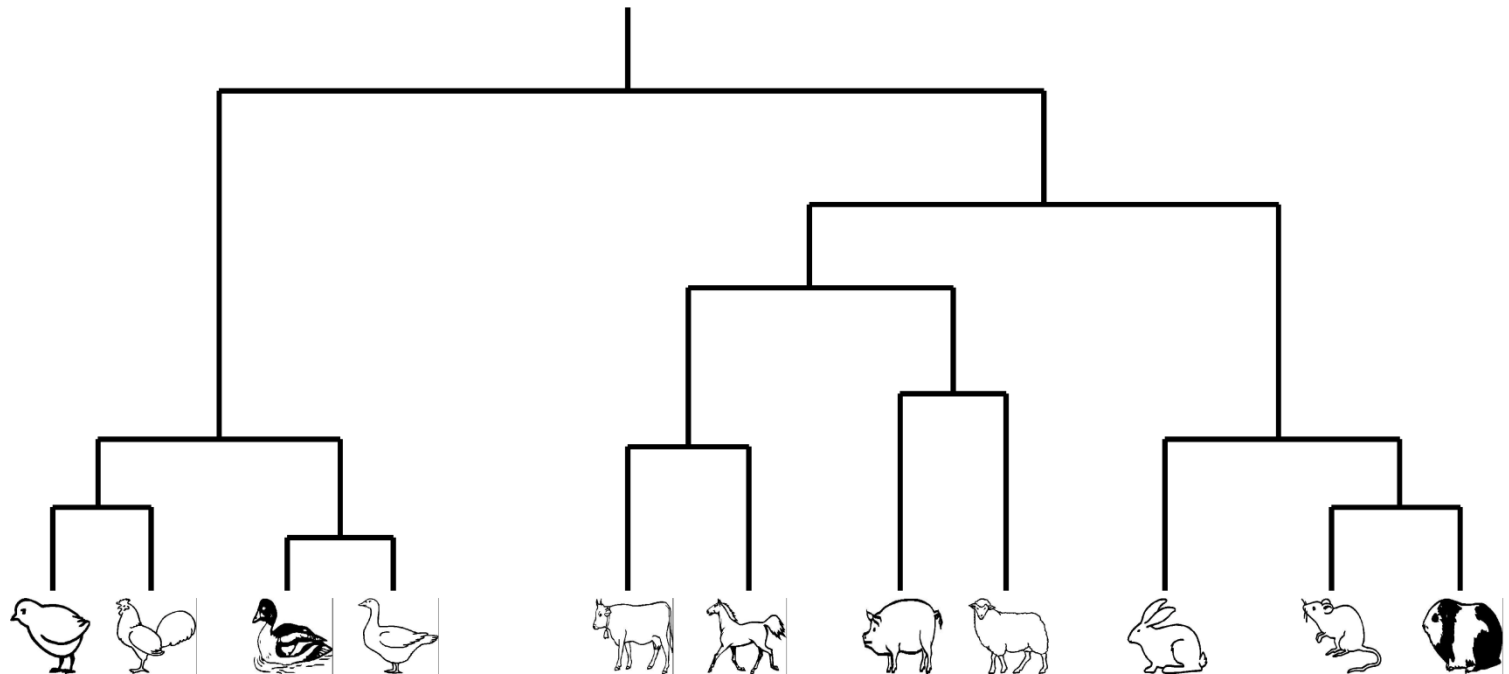


k-means applied to multiple tissues

Hierarchical Clustering

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Hierarchical clustering supplies distances between clusters which are captured in a dendrogram. These distances allow to merge or cut clusters. Clustering is done agglomerative (bottom up) or divisive (top down)



Hierarchical Clustering: Cluster Distance = Linkage



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

agglomerative hierarchical clustering (bottom up) merges the closest clusters to new clusters.

Starts with clusters single observations and iteratively merges clusters.

different distance measures between clusters A and B are used:

$$\begin{aligned}d_{\min}(A, B) &= \min_{\mathbf{a} \in A, \mathbf{b} \in B} \|\mathbf{a} - \mathbf{b}\| && \text{(single linkage)} \\d_{\max}(A, B) &= \max_{\mathbf{a} \in A, \mathbf{b} \in B} \|\mathbf{a} - \mathbf{b}\| && \text{(complete linkage)} \\d_{\text{avg}}(A, B) &= \frac{1}{n_A n_B} \sum_{\mathbf{a} \in A} \sum_{\mathbf{b} \in B} \|\mathbf{a} - \mathbf{b}\| && \text{(average linkage)} \\d_{\text{mean}}(A, B) &= \|\bar{\mathbf{a}} - \bar{\mathbf{b}}\| && \text{(average linkage)}\end{aligned}$$

where n_A (n_B) is the number of elements in A (B) and $\bar{\mathbf{a}}$ ($\bar{\mathbf{b}}$) is the mean of cluster A (B).

For the element distance $\|\cdot\|$ any distance measure is possible like the Euclidean distance, the Manhattan distance, or the Mahalanobis distance.

Hierarchical Clustering: Linkage



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based
8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples
9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM
10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

single element clusters: distance measures are equivalent

For more elements in cluster:

- **complete linkage** d_{\max} avoids that clusters are elongated in some direction (smallest distance between points remains small).
→ cluster may not be well separated.
- **single linkage** d_{\min} ensures that each pair of elements from different clusters has a minimal distance. Single linkage clustering is relevant for **leave-one-cluster-out** cross-validation, which assumes that a whole new group of objects is unknown and left out.
- **average linkage** d_{avg} is “Unweighted Pair Group Method using arithmetic Averages” (UPGMA)

Divisive or top down clustering is often based on graph theoretic considerations. First the **minimal spanning tree** is built.

Then the **largest edge is removed** which gives two clusters.

Now the second largest edge can be removed and so on.

It might be more appropriate to compute the average edge length within a cluster and find the edge which is considerably larger than other edges in the cluster.

Hierarchical Clustering: Examples



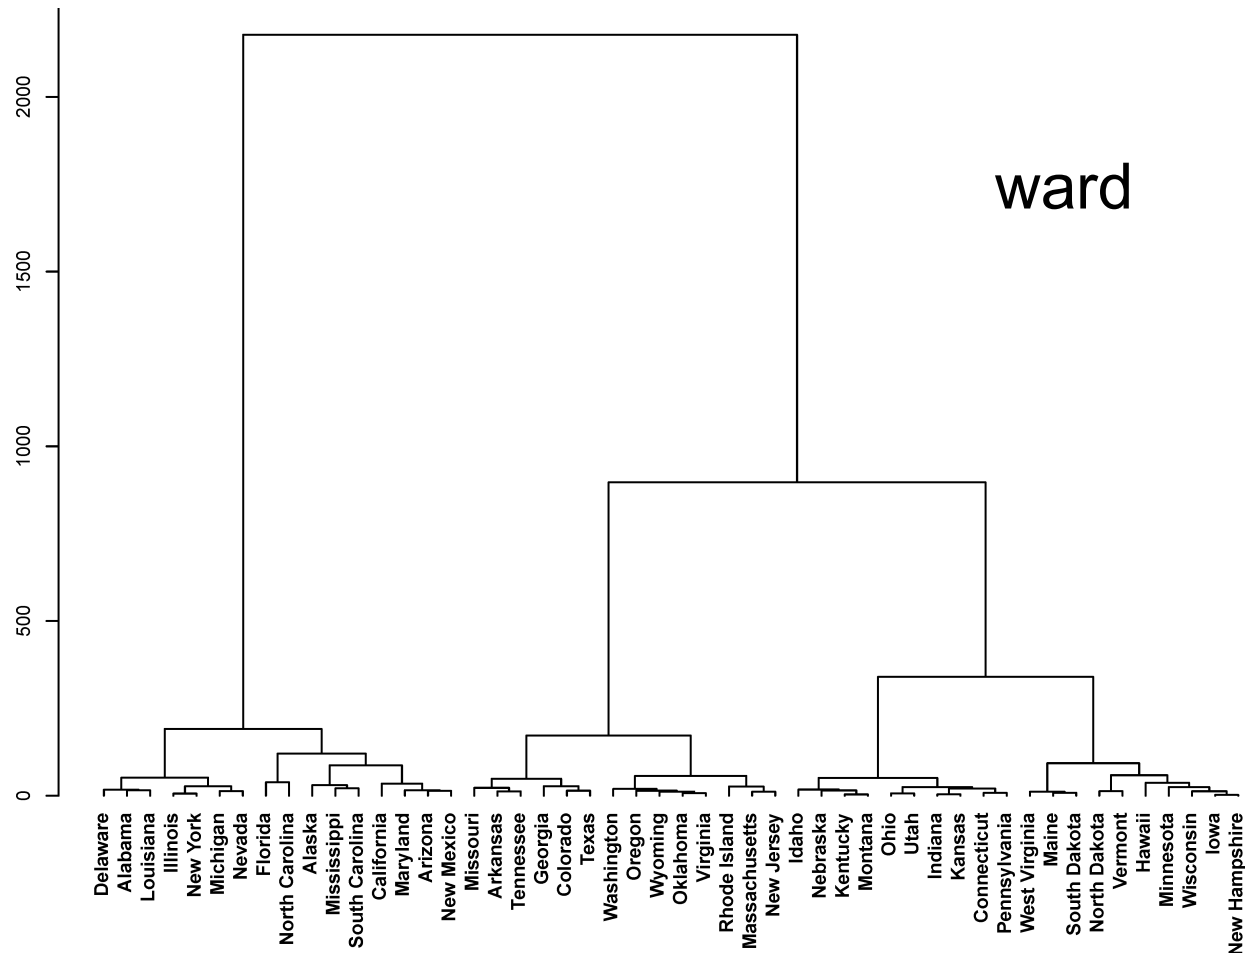
- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

hierarchical clustering of the US Arrest data distance measures “ward”, “single”, “complete”, “average”, “mcquitty”, “median”, and “centroid”



Hierarchical Clustering: Examples


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Hierarchical Clustering US Arrests: single



Hierarchical Clustering: Examples


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Hierarchical Clustering US Arrests: average



Hierarchical Clustering: Examples


- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Hierarchical Clustering US Arrests: mcquitty



Hierarchical Clustering: Examples


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Hierarchical Clustering US Arrests: median



Hierarchical Clustering: Examples


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine


Hierarchical Clustering US Arrests: centroid



Hierarchical Clustering: Examples

hierarchical clustering of the five cluster data set: Ward's distance is perfect. To determine the clusters, the dendrogram was cut

Hierarchical Clustering Five Cluster: ward



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Hierarchical Clustering: Examples

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based


- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM


- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

hierarchical clustering of the five cluster data set


Hierarchical Clustering Five Cluster: complete




Hierarchical Clustering Five Cluster: average



Hierarchical Clustering Five Cluster: single



Hierarchical Clustering Five Cluster: median



Hierarchical Clustering: Examples

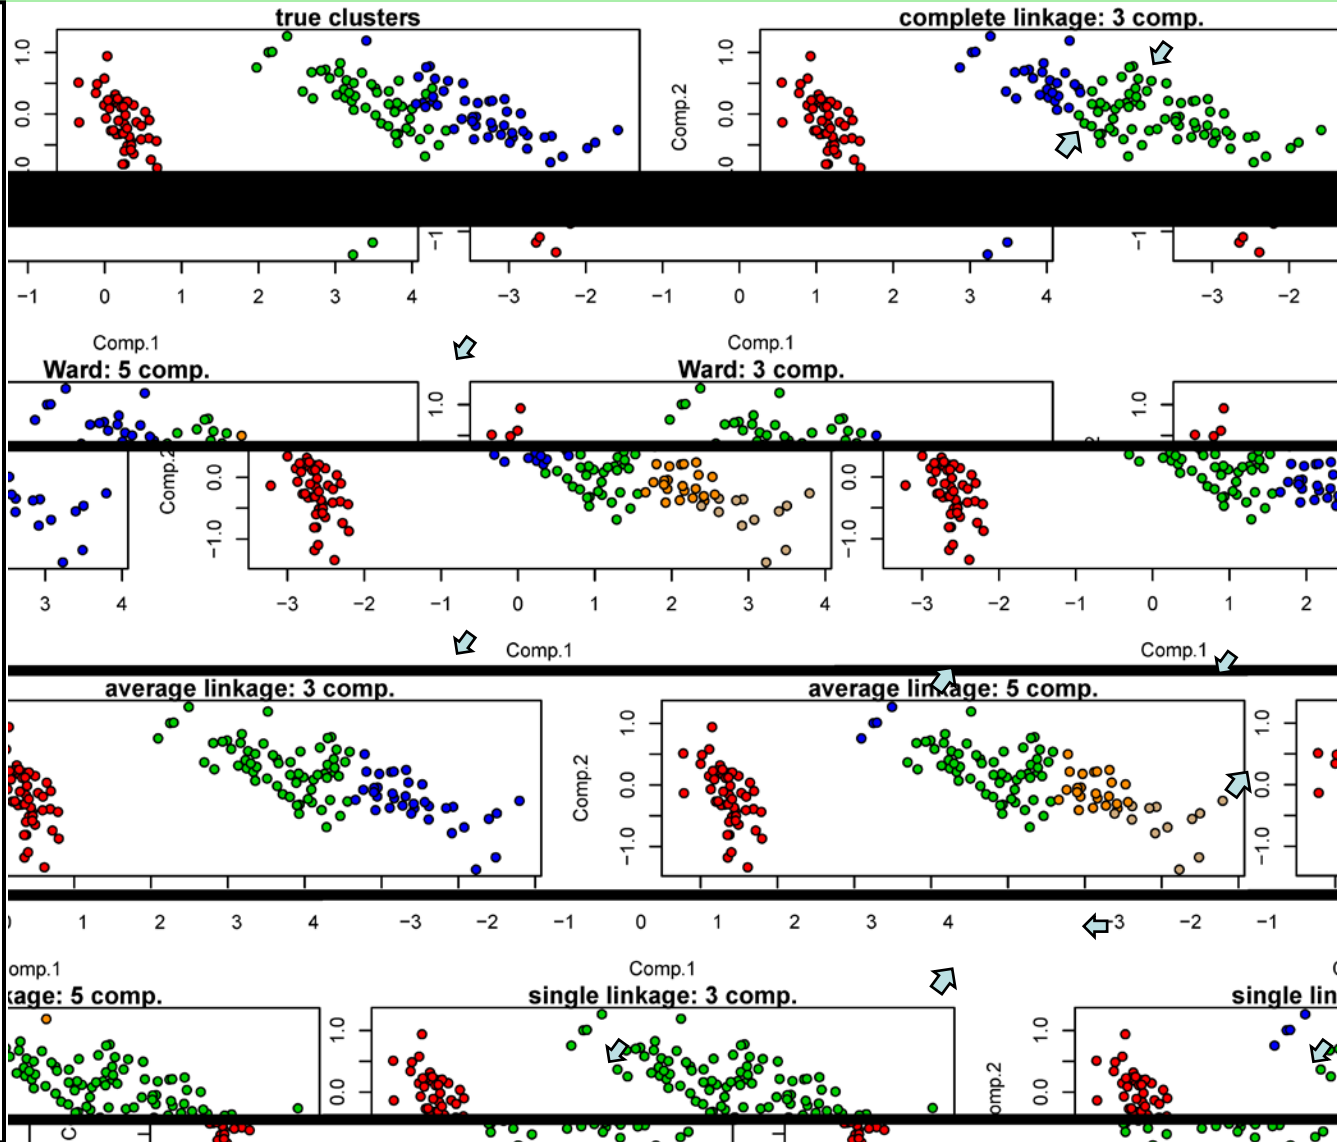


7 Clustering
 7.1 Mixture Models
 7.2 k-Means
 7.3 Hierarchical
 7.4 Similarity-Based

8 Biclustering
 8.1 Types of Biclusters
 8.2 Overview Methods
 8.3 FABIA
 8.4 Examples

9 Hidden Markov Models
 9.1 HMMs in Bioinf.
 9.2 HMM Basics
 9.3 EM / Baum-Welch
 9.4 Viterby
 9.5 Input Output HMM
 9.6 Factorial HMM
 9.7 Memory IOHMM
 9.8 Tricks of the Trade
 9.9 Profile HMM

10 Boltzmann Machines
 10.1 The Boltzmann Machine
 10.2 Learning
 10.3 Restricted Boltzmann Machine



hierarchical clustering of the iris data set

Hierarchical Clustering: Examples



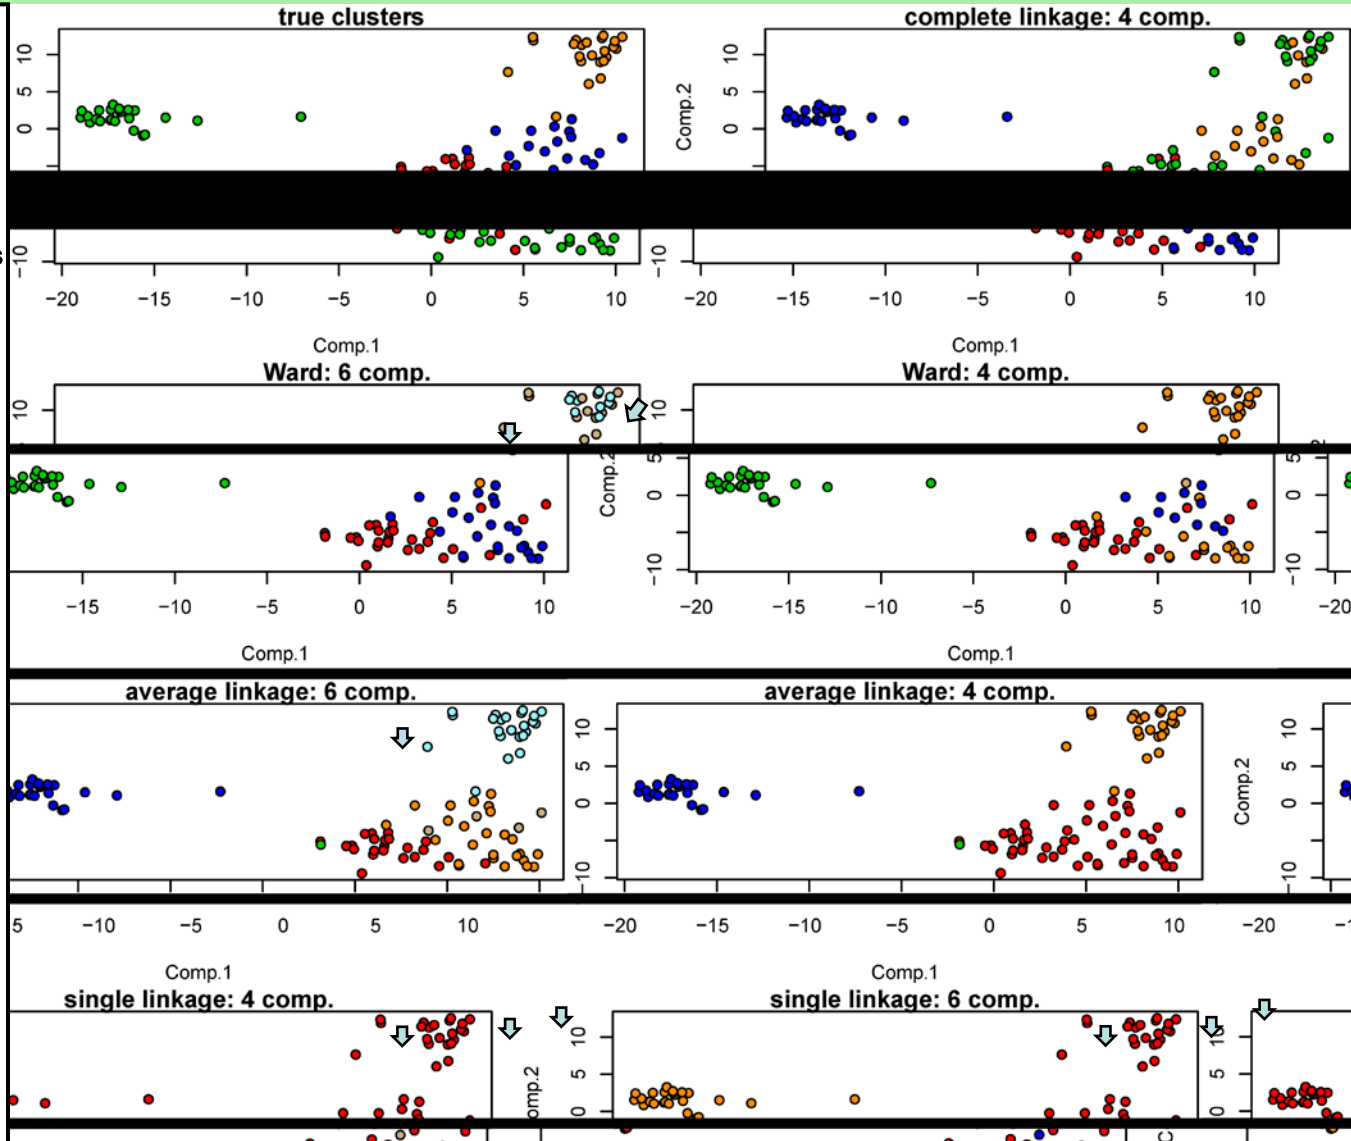
hierarchical clustering of the multiple tissue data

7 Clustering
 7.1 Mixture Models
 7.2 k-Means
 7.3 Hierarchical
 7.4 Similarity-Based

8 Biclustering
 8.1 Types of Biclusters
 8.2 Overview Methods
 8.3 FABIA
 8.4 Examples

9 Hidden Markov Models
 9.1 HMMs in Bioinf.
 9.2 HMM Basics
 9.3 EM / Baum-Welch
 9.4 Viterby
 9.5 Input Output HMM
 9.6 Factorial HMM
 9.7 Memory IOHMM
 9.8 Tricks of the Trade
 9.9 Profile HMM

10 Boltzmann Machines
 10.1 The Boltzmann Machine
 10.2 Learning
 10.3 Restricted Boltzmann Machine



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Similarity-based clustering uses similarities between objects but does **not** require to represent the objects via **feature vectors**

Similarities:

- links in the web domain
- interactions of humans (facebook)
- co-occurrences of objects (co-expression of genes or co-citations)
- spacial distances (cities on a map or atoms in a molecule)
- co-processing (compressing two documents or sorting two sets)
- alignment of two sequences
- alignment of two structures

Similarity-Based Clustering: Aspect Model



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

The **aspect model** considers discrete data of observations that are pairs (x, y) , which are counted.

Example: “person x buys product y ” or “person x participates in y ”
Applications are document-word or sample-gene relations.

the model is
$$p(x, y) = \sum_z p(z) p(x | z) p(y | z)$$

class variable: $z \in \{z_1, \dots, z_l\}$

probability of the observation: $p(x, y)$

Model assumption: x and y are independent conditioned on z

$$p(x, y) = \sum_z p(x, y, z) = \sum_z p(z) p(x, y | z) = \sum_z p(z) p(x | z) p(y | z)$$

class conditional probabilities: $p(x | z)$ and $p(y | z)$

hidden factor z has an effect on the occurrence of both x and y

Similarity-Based Clustering: Aspect Model



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

John Paulos' example in ABCNews.com:

“Consumption of hot chocolate is correlated with low crime rate, but both are responses to cold weather.”

- x = consumption of hot chocolate
- y = crime rate
- z = cold weather

The maximum likelihood model parameters $p(x | z)$ and $p(y | z)$ can be estimated by an EM algorithm:

E-step

$$p(z | x, y) = \frac{p(z) p(x | z) p(y | z)}{\sum_{z'} p(z') p(x | z') p(y | z')}$$

M-step

$$p(z) = \frac{\sum_{x, y} n(x, y) p(z | x, y)}{\sum_x n(x, y)}$$
$$p(y | z) = \frac{\sum_x n(x, y) p(z | x, y)}{p(z)}$$
$$p(x | z) = \frac{\sum_y n(x, y) p(z | x, y)}{p(z)}$$

$n(x, y)$ is the count of observations (x, y) , that is, the row of x and the column of y in the data matrix

Similarity-Based Clustering: Aspect Model



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

For the aspect model, clustering of the x can be based on

$$p(z | x) = \frac{p(x | z) p(z)}{p(x)}$$

$$p(x) = \sum_{z,y} p(z) p(x | z) p(y | z)$$

z indicates the cluster, that is, each z represents one cluster

Analog formulas are obtained for clustering y or pairs (x, y)

Similarity-Based Clustering: Affinity Propagation



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Affinity propagation is a **similarity-based** clustering method that is also **exemplar-based** clustering

exemplar-based clustering enforces cluster centers to be data points, the “**prototypes**” or “**exemplars**”

exemplar-based clustering is the **k -centers clustering** which starts with an initial set of randomly selected exemplars and iteratively refines this set so as to decrease the sum of squared errors.

- only for small number of clusters
- good initialization required

Affinity propagation overcomes the problems of **k -centers clustering**

Similarity-Based Clustering: Affinity Propagation



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

- **similarities** between object i and object k : $s(i, k)$
- **preferences**: $s(k, k)$ how likely object k becomes an exemplar
- **responsibilities**: $r(i, k)$ messages sent from object i to candidate exemplar k . responsibility reflects the evidence that k serves as an exemplar for i : how well can k represent the object i .
- **availabilities**: $a(i, k)$ messages sent from candidate exemplar k to object i . Availability reflects the evidence that k is indeed an exemplar.

initialization: $a(i, k) = 0$

updates:

$$r(i, k) = s(i, k) - \max_{k', k' \neq k} \{a(i, k') + s(i, k')\}$$
$$a(i, k) = \min \left\{ 0, r(k, k) + \sum_{i', i' \notin \{i, k\}} \max\{0, r(i', k)\} \right\}$$
$$a(k, k) = \sum_{i', i' \neq k} \max\{0, r(i', k)\}$$

Similarity-Based Clustering: Affinity Propagation

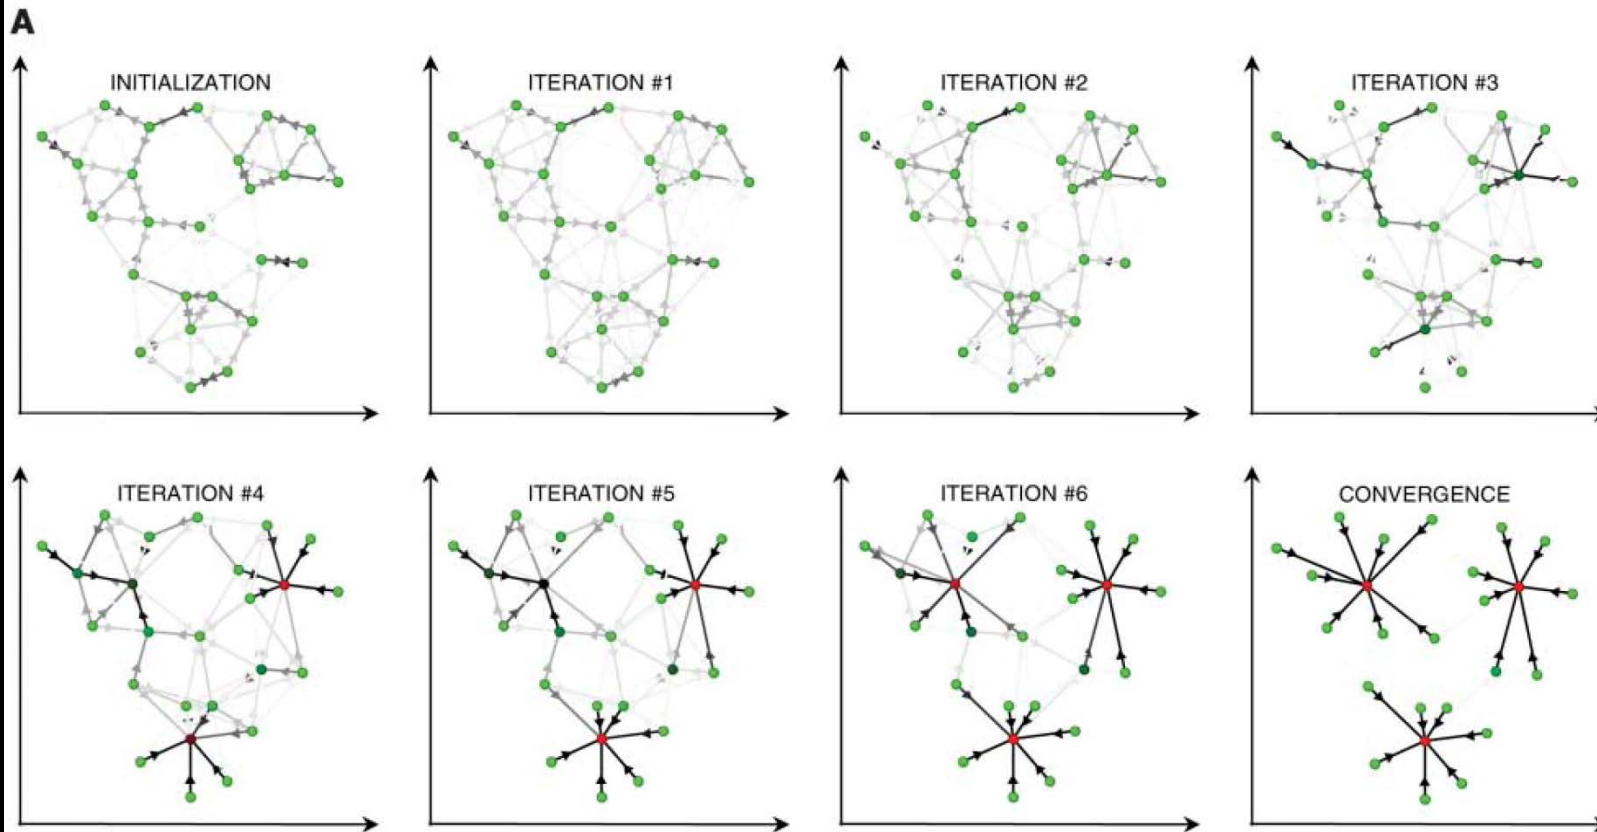
7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

different iterations of affinity propagation and messages sent



Similarity-Based Clustering: Affinity Propagation

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based


- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM


- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

specific messages passing in the algorithm of affinity propagation

Sending responsibilities



Sending availabilities



Similarity-Based Clustering: Affinity Propagation

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Example of affinity propagation with images of faces

The 15 images with highest squared error under either affinity propagation or k -centers clustering are shown in the top row. The middle and bottom rows show the exemplars assigned by the two methods, and the boxes show which of the two methods performed better for that image, in terms of squared error. Affinity propagation found higher-quality exemplars.



Similarity-Based Clustering: Affinity Propagation

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

(A) Similarities between pairs of sentences in the AP manuscript were constructed by matching words. Four identified exemplars are shown.

(B) AP was applied to similarities derived from air-travel efficiency between the 456 busiest commercial airports in Canada and the United States.

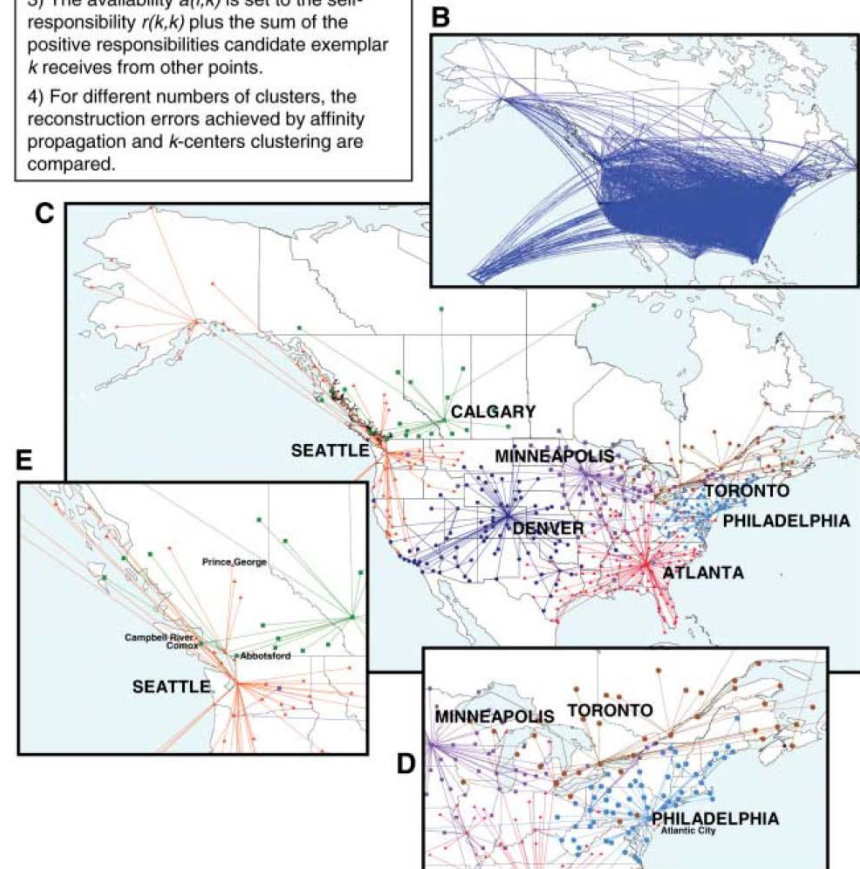
(C) Seven exemplars identified by AP are color-coded, and the assignments of other cities to these exemplars is shown.

(D) The inset shows that the Canada-USA border roughly divides the Toronto and Philadelphia clusters, due to a larger availability of domestic flights vs. international flights.

(E) The west coast has extraordinarily frequent airline service between Vancouver and Seattle connects Canadian cities in the northwest to Seattle.

A Exemplar sentences from draft of this paper:

- 1) Affinity propagation identifies exemplars by recursively sending real-valued messages between pairs of data points.
- 2) The number of identified exemplars (number of clusters) is influenced by the values of the input preferences, but also emerges from the message-passing procedure.
- 3) The availability $a(i,k)$ is set to the self-responsibility $r(k,k)$ plus the sum of the positive responsibilities candidate exemplar k receives from other points.
- 4) For different numbers of clusters, the reconstruction errors achieved by affinity propagation and k -centers clustering are compared.

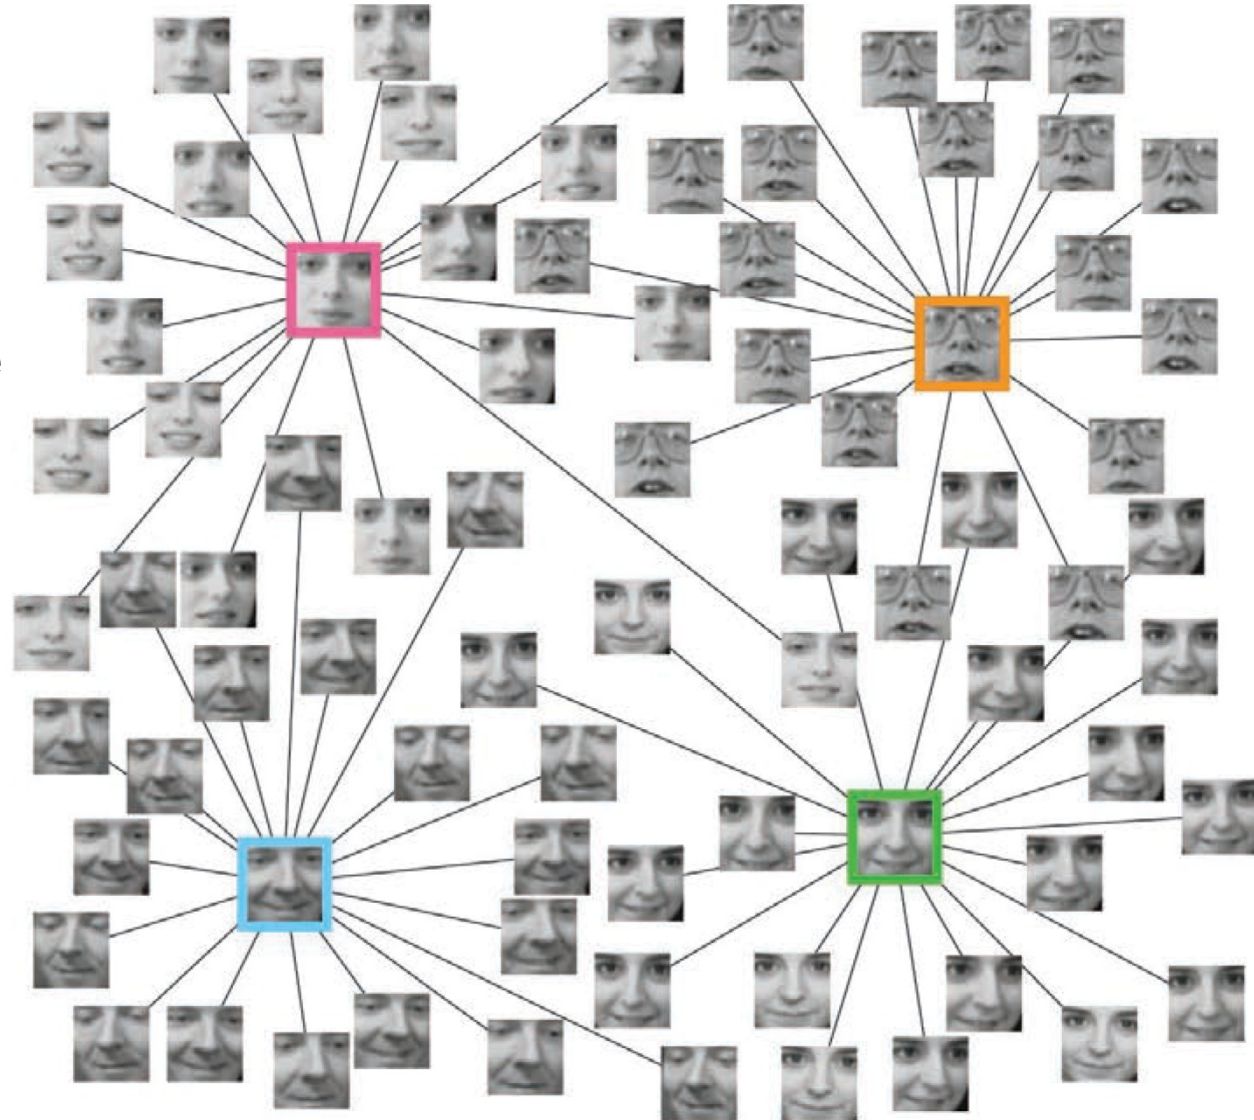


Similarity-Based Clustering: Affinity Propagation

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

affinity propagation applied to face images

exemplars are highlighted by colored boxes



Affinity Propagation: Examples

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based


8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

data set with 6 clusters in a 2-D space, where three clusters are smaller (smaller variance) than the other three clusters

- Circles indicate the centers of the clusters
- Blue are the large clusters and brown the small clusters



Affinity Propagation: Examples


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

- AP exemplars: black rectangles
- The two large clusters are identified
- AP prefers spherical clusters which helps here



Affinity Propagation: Examples

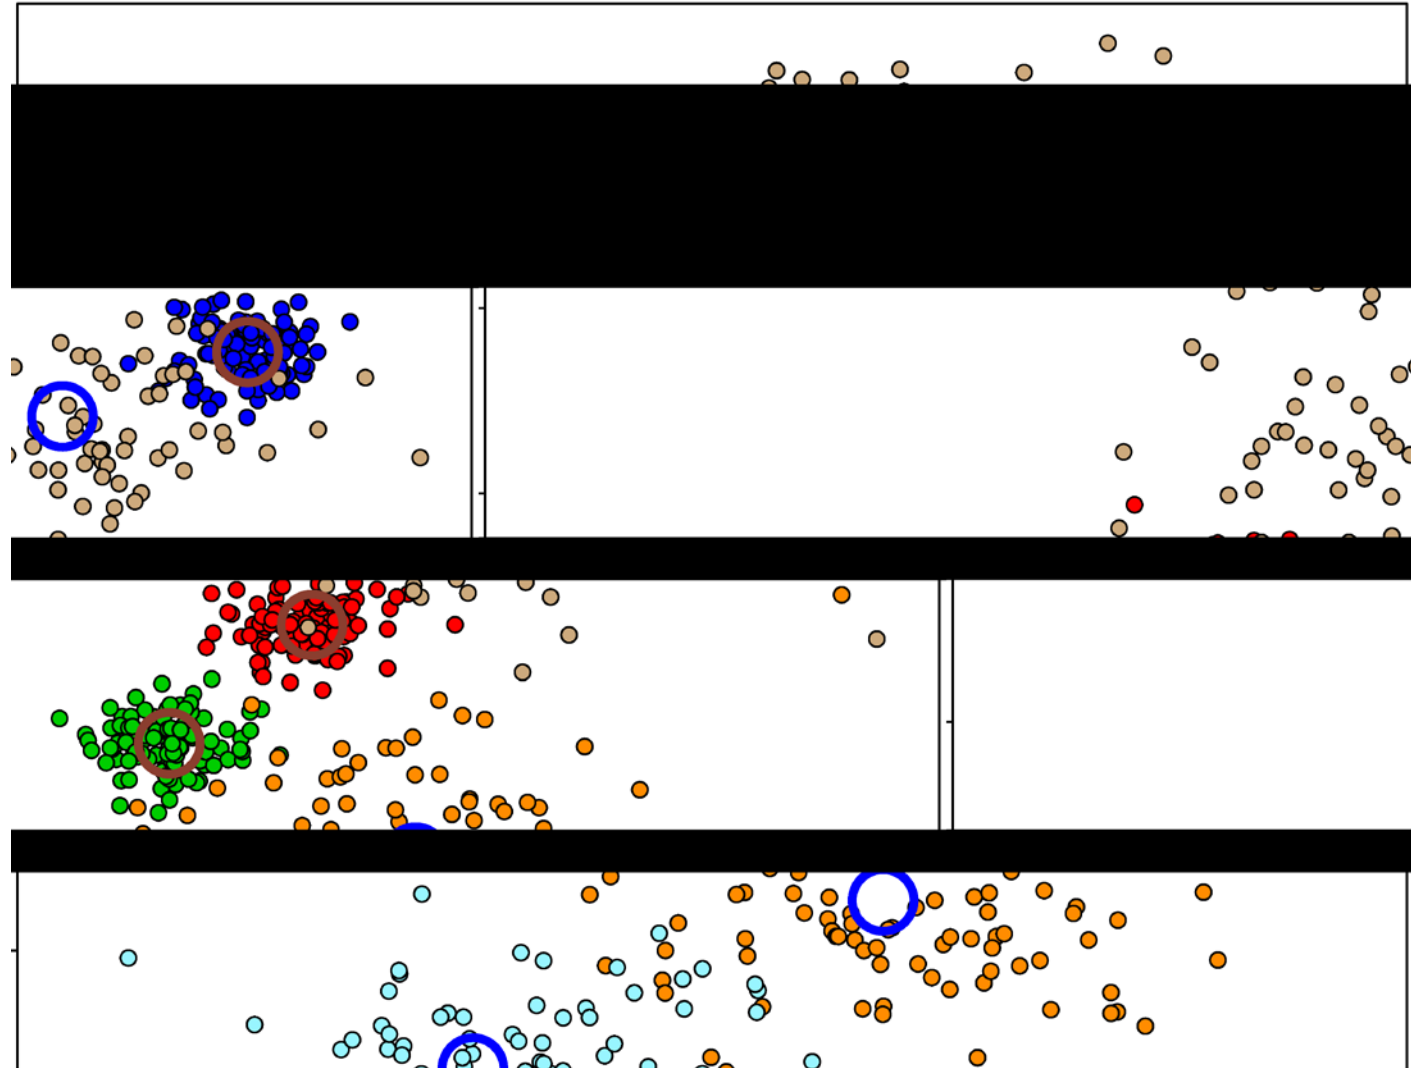
- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

The x7A data set: another version of the x7 data generation



Affinity Propagation: Examples

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based


- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

upper right: a large

equal large clusters,
therefore cannot detect
the small cluster within
the large cluster



Affinity Propagation: Examples

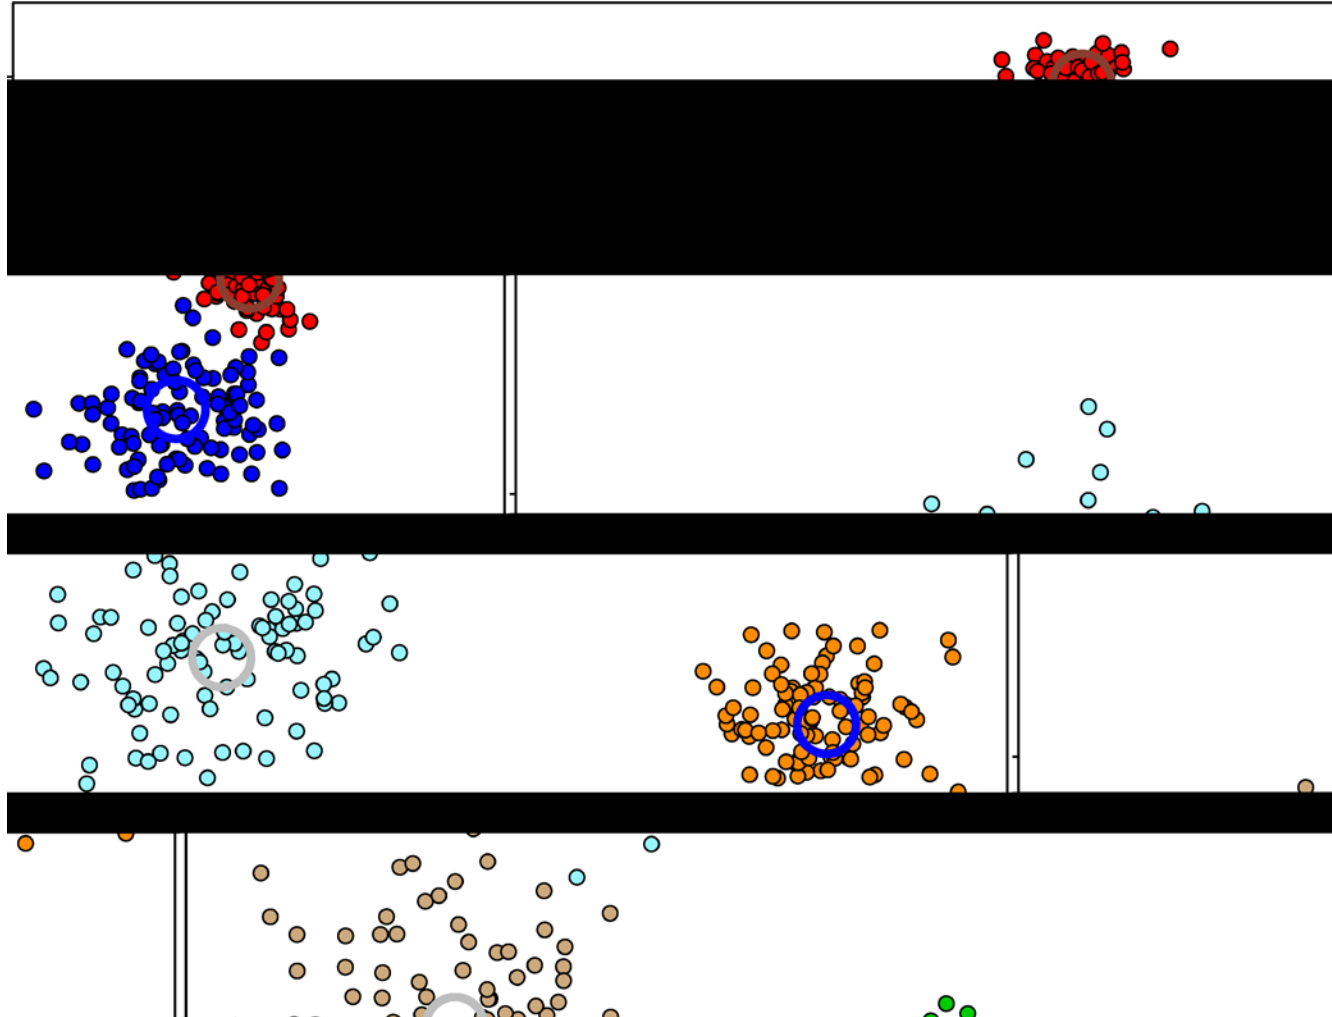
x9 data set: 6 clusters of which 2 are small, 2 medium sized, 2 large with respect to the variance

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine



Affinity Propagation: Examples


- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

AP applied to x9: one cluster is not separated into the two true clusters



Affinity Propagation: Examples


- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

6-D data set with different cluster sizes (variation of their elements)
one large and three smaller clusters: 100 samples; PCA down-projection



Affinity Propagation: Examples

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

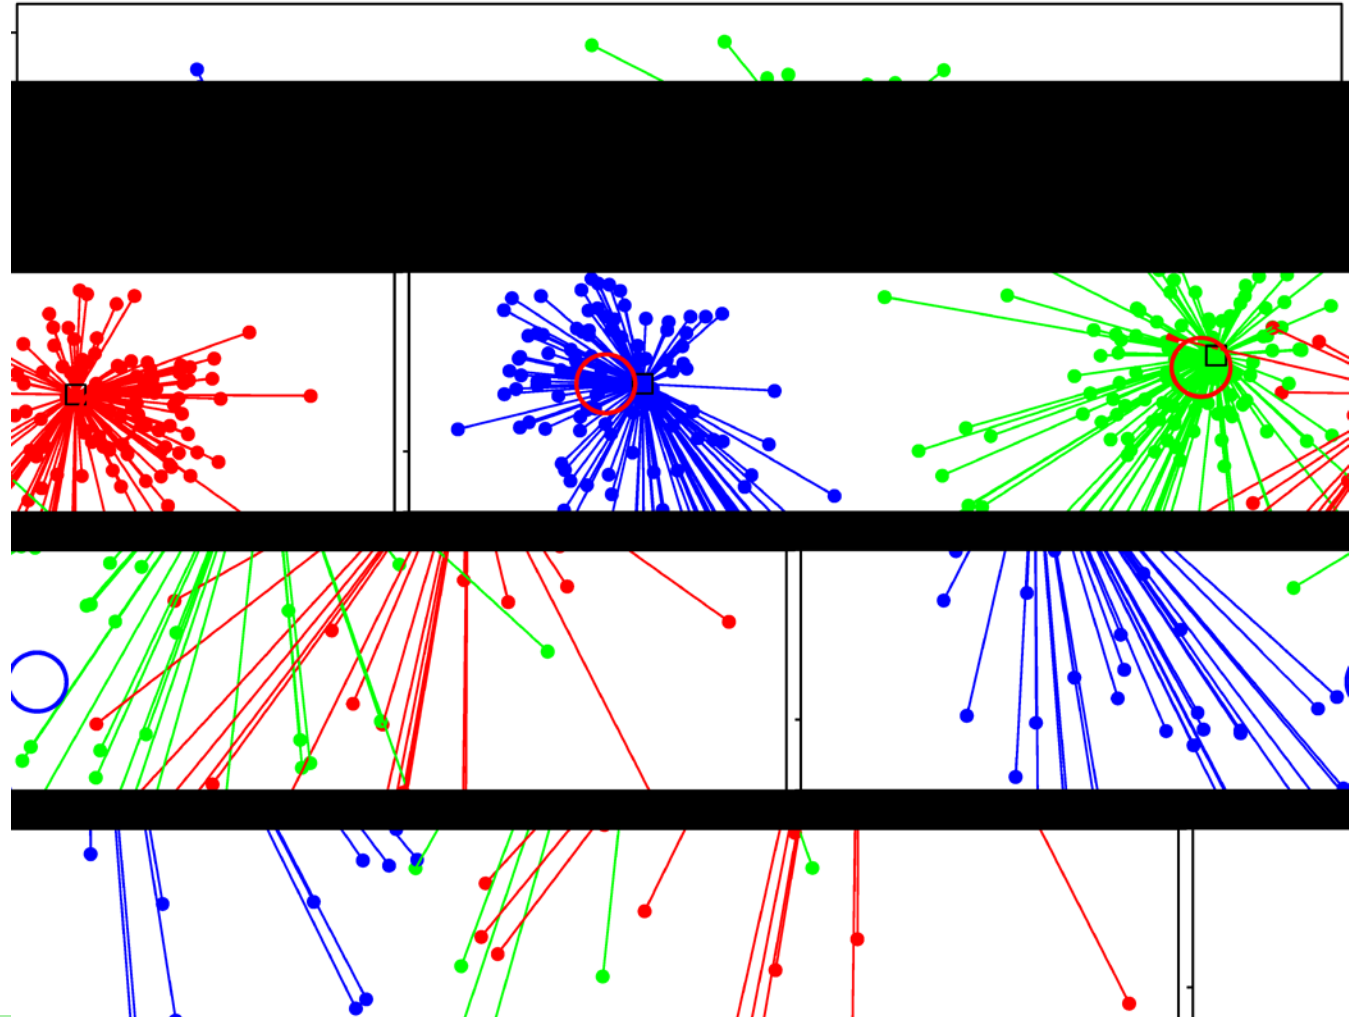
- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

affinity propagation

AP does not detect the large cluster

Elements of the large cluster are assigned to the smaller cluster

AP has problems with the different cluster sizes because it cannot adjust the variance



Affinity Propagation: Examples


Another data set of d6: d6A down-projected by PCA

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine



Affinity Propagation: Examples

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

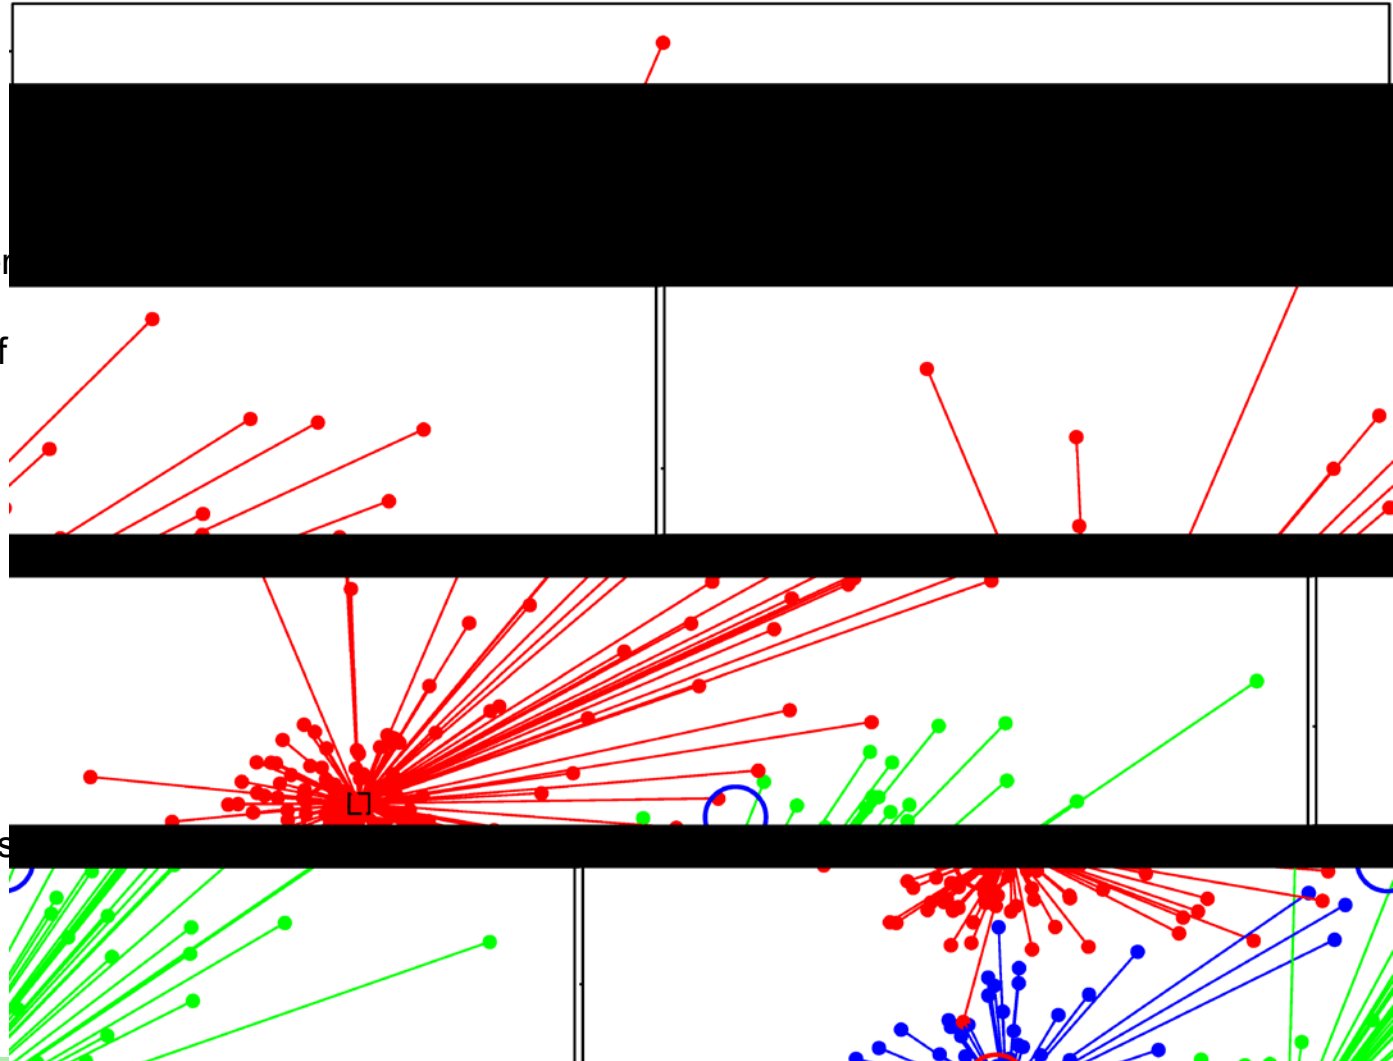
10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

affinity propagation for d6A

Again AP does not detect the large cluster

Elements of the large cluster are assigned to the smaller cluster

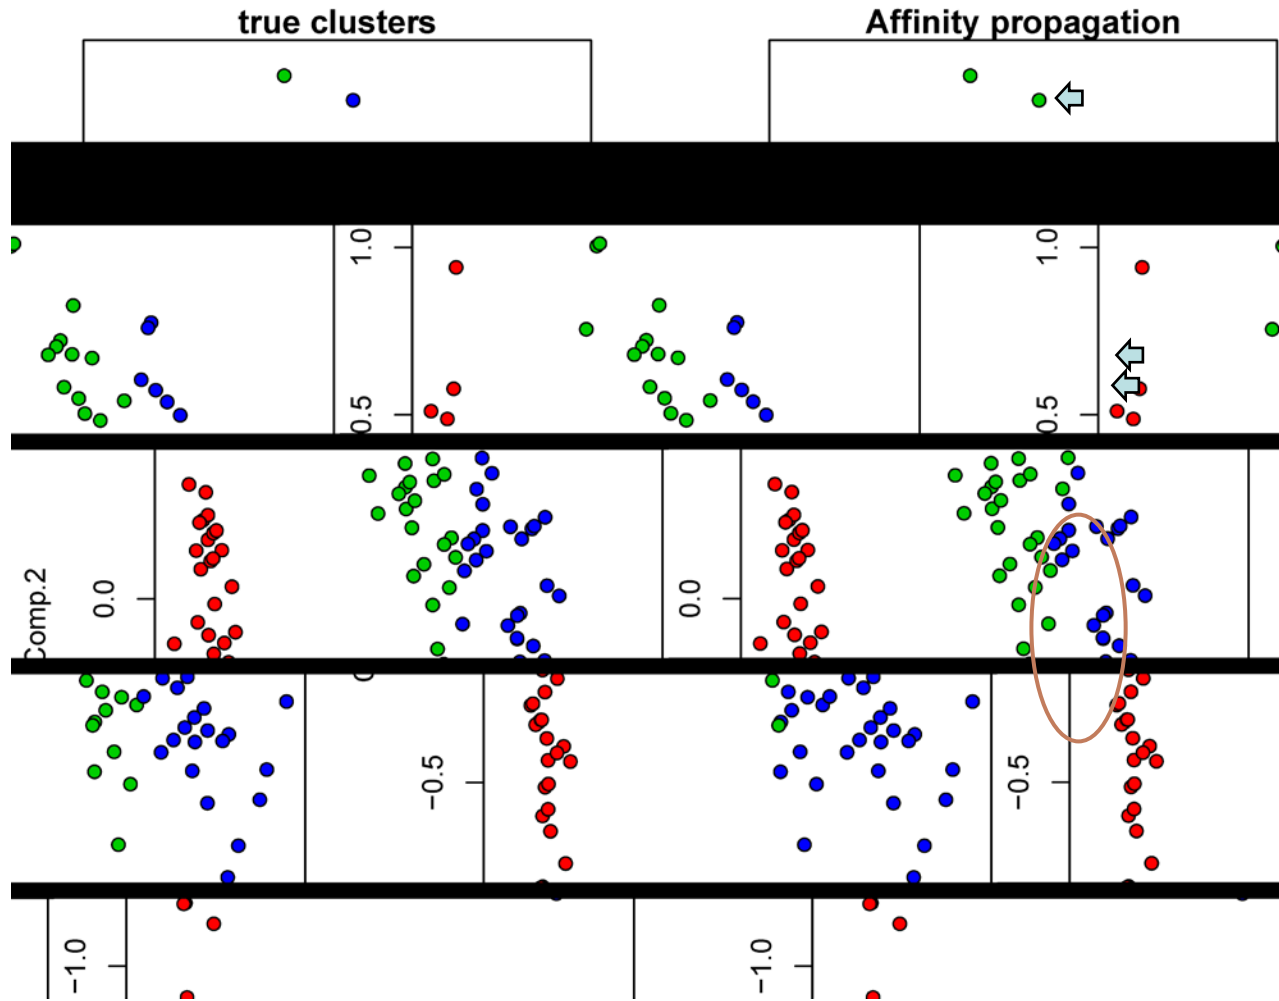
AP has problems with the different cluster sizes because it cannot adjust the variance



Affinity Propagation: Examples

affinity propagation applied to the iris data


- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine



Affinity Propagation: Examples

affinity propagation applied to the multiple tissues: $p=0$


- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine



Affinity Propagation: Examples

affinity propagation applied to the multiple tissues: $p=1$

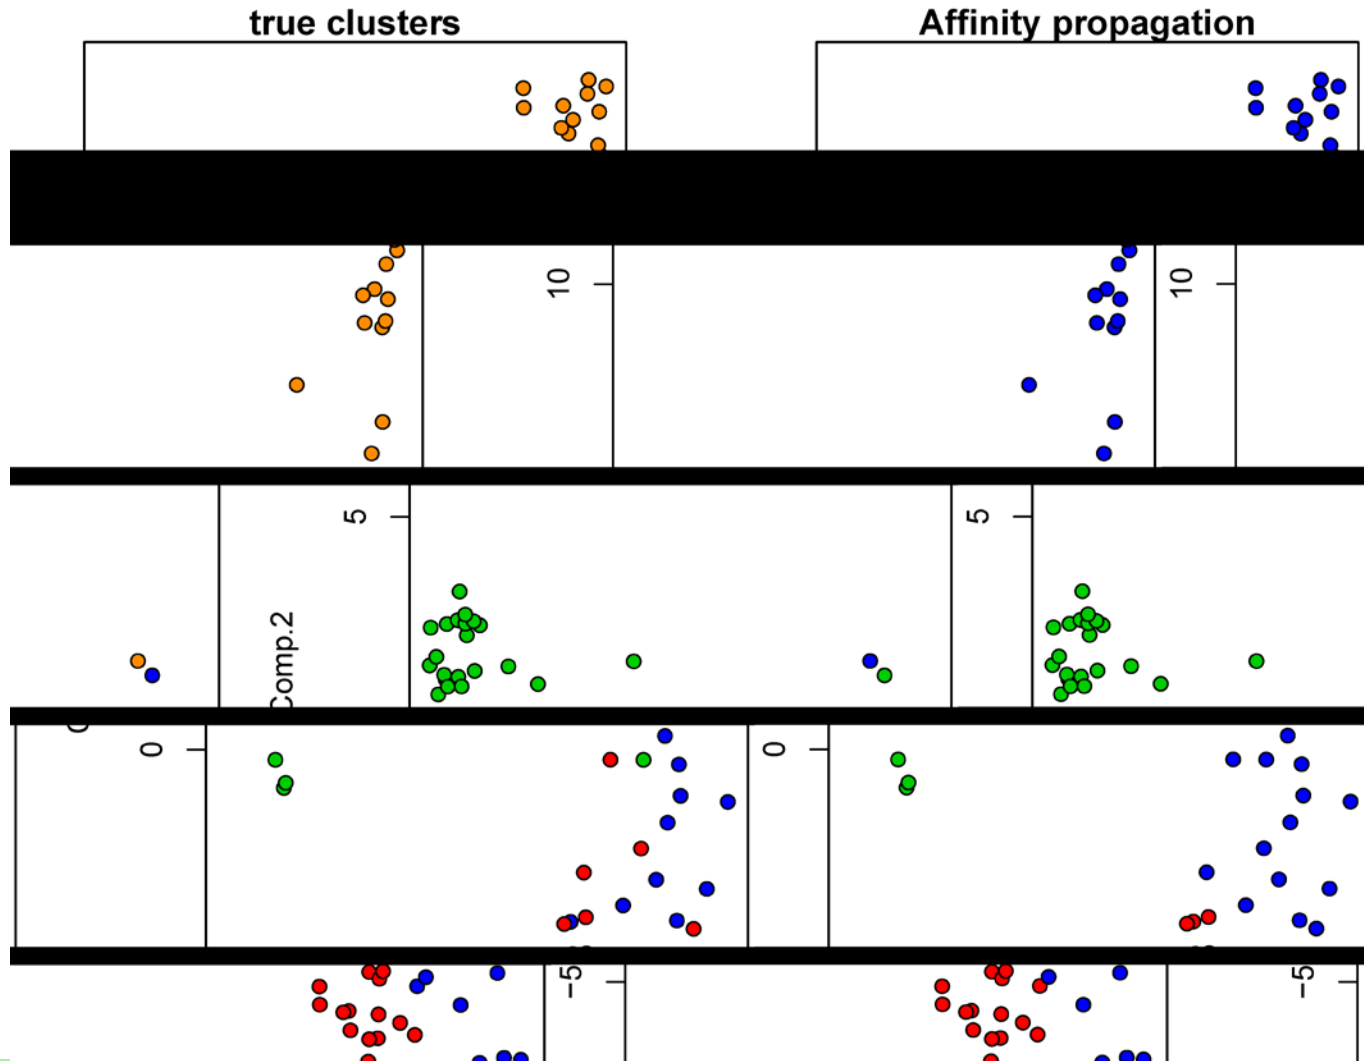
- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine



Affinity Propagation: Examples

affinity propagation applied to the multiple tissues: $p=-1$ (three clusters)

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine



Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Similarity-based mixture models are mixture models that use only similarities between objects but not feature vectors

- similarity between object i and object j : $p(\mathbf{x}_j | \mathbf{x}_i)$ (given)
- stochastic neighbor embedding: probabilities are computed

Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Similarity-based mixture of Gaussians

For \mathbf{x} the similarities $k(\mathbf{x}, \mathbf{x}_i)$ to each element of $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ are given, in particular all $k(\mathbf{x}_j, \mathbf{x}_i)$ for $1 \leq i, j \leq n$ are given.

all update rules and parameters are expressed by these similarities

$$\text{Gaussian mixture model } p(\mathbf{x}) = \sum_{i=1}^K p(i) p(\mathbf{x} | i) = \sum_{i=1}^K p(i) k(\mathbf{x}; \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

where $\sum_{i=1}^K p(i) = 1$ and $k(\mathbf{x}; \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ is the Gaussian density with mean

$\boldsymbol{\mu}_i$ and variance $\boldsymbol{\Sigma}_i$ evaluated at \mathbf{x}

We introduce a variance σ_i for component i by $\frac{1}{\sigma_i^m} k(\mathbf{x}; \boldsymbol{\mu}_i)^{1/\sigma_i^2}$ and obtain for the likelihood of \mathbf{x} :

$$p(\mathbf{x}) = \sum_{j=1}^K p(j) \frac{1}{\sigma_j^m} k(\mathbf{x}, \boldsymbol{\mu}_j)^{1/\sigma_j^2}$$

Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

We set $k(\mathbf{x}; \boldsymbol{\mu}_i, \sigma_i^2) = \frac{1}{\sigma_i^m} k(\mathbf{x}, \boldsymbol{\mu}_i)^{1/\sigma_i^2}$

The posterior gives the probability that \mathbf{x} belongs to cluster i :

$$p(i | \mathbf{x}) = \frac{p(i) p(\mathbf{x} | i)}{\sum_{j=1}^K p(j) p(\mathbf{x} | j)} = \frac{p(i) k(\mathbf{x}; \boldsymbol{\mu}_i, \sigma_i^2)}{\sum_{j=1}^K p(j) k(\mathbf{x}; \boldsymbol{\mu}_j, \sigma_j^2)} = \frac{p(i) k(\mathbf{x}; \boldsymbol{\mu}_i, \sigma_i^2)}{p(\mathbf{x})}$$

With $\alpha_i = p(i)$ and $\alpha_{ik} = p(i | \mathbf{x}_k, \boldsymbol{\alpha})$ we obtain

$$\begin{aligned} \alpha_i &= p(i) = p(i | \boldsymbol{\alpha}) = \int p(i, \mathbf{x} | \boldsymbol{\alpha}) d\mathbf{x} = \int p(i | \mathbf{x}, \boldsymbol{\alpha}) p(\mathbf{x} | \boldsymbol{\alpha}) d\mathbf{x} = E_{p(\mathbf{x} | \boldsymbol{\alpha})}(p(i | \mathbf{x}, \boldsymbol{\alpha})) \\ &\approx \frac{1}{n} \sum_{k=1}^n p(i | \mathbf{x}_k, \boldsymbol{\alpha}) = \frac{1}{n} \sum_{k=1}^n \alpha_{ik} \end{aligned}$$

Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

The likelihood for one data point x is

$$p(\mathbf{x}) = \sum_{i=1}^K \alpha_i \frac{1}{\sigma_i^m} k(\mathbf{x}_k; \boldsymbol{\mu}_i)^{1/\sigma_i^2} = \sum_{i=1}^K \alpha_i k(\mathbf{x}_k; \boldsymbol{\mu}_i, \sigma_i^2)$$

The objective is the negative log posterior:

$$\begin{aligned} B = & -\frac{1}{n} \sum_{k=1}^n \sum_{i=1}^K \hat{\alpha}_{ik} \log(k(\mathbf{x}_k; \boldsymbol{\mu}_i, \sigma_i^2)) - \frac{1}{n} \sum_{k=1}^n \sum_{i=1}^K \hat{\alpha}_{ik} \log(\alpha_i) \\ & - \frac{1}{n} \log p(\boldsymbol{\alpha}) - \frac{1}{n} \log p(\boldsymbol{\sigma}^2) + \frac{1}{n} \sum_{k=1}^n \sum_{i=1}^K \hat{\alpha}_{ik} \log \hat{\alpha}_{ik} + \frac{1}{n} \sum_{k=1}^n \log c(\mathbf{x}_k) \end{aligned}$$

The objective contains the Dirichlet prior $p(\boldsymbol{\alpha})$ and Wishart prior $p(\boldsymbol{\sigma}^2)$

$$\boldsymbol{\sigma}^2 = (\sigma_1^2, \sigma_2^2, \dots, \sigma_K^2)$$

Similarity-Based Clustering: Mixture Models



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Using $\hat{\alpha}_{ik} = \frac{\alpha_i^{\text{old}} k(\mathbf{x}_k; \boldsymbol{\mu}_i, \sigma_i^2)}{p(\mathbf{x}_k)}$ and $\hat{\alpha}_i = \frac{1}{n} \sum_{k=1}^n \hat{\alpha}_{ik}$ and $\gamma_s = \sum_{i=1}^n \gamma_i$

the update rules are:

$$\alpha_i^{\text{new}} = \frac{n \hat{\alpha}_i + \gamma_i - 1}{\gamma_s}$$

$$(\sigma_i^2)^{\text{new}} = \frac{\sum_{k=1}^n \hat{\alpha}_{ik} (-2 \log k(\mathbf{x}_k, \boldsymbol{\mu}_i)) + w S}{n \hat{\alpha}_i + w}$$

hyper-parameters γ_i come from the Dirichlet prior

hyper-parameters w and S come from the Wishart prior

Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Representing the center

EM algorithm for MoG has center update:

$$\boldsymbol{\mu}_i = \frac{\sum_{k=1}^n \hat{\alpha}_{ik} \mathbf{x}_k}{\sum_{k=1}^n \hat{\alpha}_{ik}}$$

For the optimal center: $\log k(\mathbf{x}_k; \boldsymbol{\mu}_i, \sigma_i^2) = -n \log(\sigma_i) - \sigma_i^{-2} (-\log k(\mathbf{x}_k, \boldsymbol{\mu}_i))$

$$\frac{\partial B}{\partial \boldsymbol{\mu}_i} = -\frac{1}{n} \sum_{k=1}^n \hat{\alpha}_{ik} \frac{1}{\sigma_i^2} \frac{\partial \log k(\mathbf{x}_k, \boldsymbol{\mu}_i)}{\partial \boldsymbol{\mu}_i} = 0$$

We assume:

$$\log k(\mathbf{x}_k, \boldsymbol{\mu}_i) = -\frac{1}{2} \|\mathbf{x}_k - \boldsymbol{\mu}_i\|^2 = -\frac{1}{2} (\mathbf{x}_k - \boldsymbol{\mu}_i)^T (\mathbf{x}_k - \boldsymbol{\mu}_i)$$

$$\frac{\partial \log k(\mathbf{x}_k, \boldsymbol{\mu}_i)}{\partial \boldsymbol{\mu}_i} = -\mathbf{x}_k + \boldsymbol{\mu}_i$$

Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

For the update rules we only require $\|\mathbf{x}_l - \boldsymbol{\mu}_i\| = -2 \log k(\mathbf{x}_l, \boldsymbol{\mu}_i)$

We assumed for the similarities $k(\mathbf{x}_k, \mathbf{x}_j) = \frac{1}{(2\pi)^{m/2}} \exp\left(-\frac{1}{2} \|\mathbf{x}_k - \mathbf{x}_j\|^2\right)$

from which follows that

$$-2 \log k(\mathbf{x}_k, \mathbf{x}_j) = n \log(2\pi) + \|\mathbf{x}_k - \mathbf{x}_j\|^2 = n \log(2\pi) + \mathbf{x}_k^T \mathbf{x}_k - 2 \mathbf{x}_k^T \mathbf{x}_j + \mathbf{x}_j^T \mathbf{x}_j$$

and, therefore, we have

$$\mathbf{x}_k^T \mathbf{x}_j = \frac{1}{2} n \log(2\pi) + \frac{1}{2} \mathbf{x}_k^T \mathbf{x}_k + \frac{1}{2} \mathbf{x}_j^T \mathbf{x}_j + \log k(\mathbf{x}_k, \mathbf{x}_j)$$

$$\begin{aligned} \|\mathbf{x}_l - \boldsymbol{\mu}_i\|^2 &= (\mathbf{x}_l - \boldsymbol{\mu}_i)^T (\mathbf{x}_l - \boldsymbol{\mu}_i) = \mathbf{x}_l^T \mathbf{x}_l - 2 \mathbf{x}_l^T \boldsymbol{\mu}_i + \boldsymbol{\mu}_i^T \boldsymbol{\mu}_i = \mathbf{x}_l^T \mathbf{x}_l - 2 \frac{\sum_{k=1}^n \hat{\alpha}_{ik} \mathbf{x}_k^T \mathbf{x}_l}{\sum_{k=1}^n \hat{\alpha}_{ik}} + \frac{\sum_{k,j=1}^n \hat{\alpha}_{ik} \hat{\alpha}_{ij} \mathbf{x}_k^T \mathbf{x}_j}{(\sum_{k=1}^n \hat{\alpha}_{ik})^2} \\ &= \mathbf{x}_l^T \mathbf{x}_l - \frac{\sum_{k=1}^n \hat{\alpha}_{ik} (\mathbf{x}_k^T \mathbf{x}_k + 2 \log k(\mathbf{x}_k, \mathbf{x}_l))}{\sum_{k=1}^n \hat{\alpha}_{ik}} - n \log(2\pi) - \mathbf{x}_l^T \mathbf{x}_l \\ &\quad + \frac{\sum_{k,j=1}^n \hat{\alpha}_{ik} \hat{\alpha}_{ij} (1/2 n \log(2\pi) + 1/2 \mathbf{x}_k^T \mathbf{x}_k + 1/2 \mathbf{x}_j^T \mathbf{x}_j + \log k(\mathbf{x}_k, \mathbf{x}_j))}{(\sum_{k=1}^n \hat{\alpha}_{ik})^2} \\ &= - \frac{\sum_{k=1}^n \hat{\alpha}_{ik} \mathbf{x}_k^T \mathbf{x}_k}{\sum_{k=1}^n \hat{\alpha}_{ik}} - 2 \frac{\sum_{k=1}^n \hat{\alpha}_{ik} \log k(\mathbf{x}_k, \mathbf{x}_l)}{\sum_{k=1}^n \hat{\alpha}_{ik}} - n \log(2\pi) \\ &\quad + 1/2 n \log(2\pi) + \frac{1}{2} \frac{\sum_{k=1}^n \hat{\alpha}_{ik} \mathbf{x}_k^T \mathbf{x}_k}{\sum_{k=1}^n \hat{\alpha}_{ik}} + \frac{1}{2} \frac{\sum_{j=1}^n \hat{\alpha}_{ij} \mathbf{x}_j^T \mathbf{x}_j}{\sum_{k=1}^n \hat{\alpha}_{ik}} + \frac{\sum_{k,j=1}^n \hat{\alpha}_{ik} \hat{\alpha}_{ij} \log k(\mathbf{x}_k, \mathbf{x}_j)}{(\sum_{k=1}^n \hat{\alpha}_{ik})^2} \\ &= - 1/2 n \log(2\pi) - 2 \frac{\sum_{k=1}^n \hat{\alpha}_{ik} \log k(\mathbf{x}_k, \mathbf{x}_l)}{\sum_{k=1}^n \hat{\alpha}_{ik}} + \frac{\sum_{k,j=1}^n \hat{\alpha}_{ik} \hat{\alpha}_{ij} \log k(\mathbf{x}_k, \mathbf{x}_j)}{(\sum_{k=1}^n \hat{\alpha}_{ik})^2} \end{aligned}$$

Similarity-Based Clustering: Mixture Models



- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

we can compute $k(\mathbf{x}_k, \boldsymbol{\mu}_i) = \frac{1}{(2\pi)^{m/2}} \exp\left(-\frac{1}{2} \|\mathbf{x}_k - \boldsymbol{\mu}_i\|^2\right)$

If we ignore $1/(2\pi)^{m/2}$ and define $k(\mathbf{x}, \mathbf{x}_i) = \exp\left(-\frac{1}{2} \|\mathbf{x} - \mathbf{x}_i\|^2\right)$

then we would obtain

$$\|\mathbf{x}_l - \boldsymbol{\mu}_i\|^2 = -2 \frac{\sum_{k=1}^n \hat{\alpha}_{ik} \log k(\mathbf{x}_k, \mathbf{x}_l)}{\sum_{k=1}^n \hat{\alpha}_{ik}} + \frac{\sum_{k,j=1}^n \hat{\alpha}_{ik} \hat{\alpha}_{ij} \log k(\mathbf{x}_k, \mathbf{x}_j)}{\left(\sum_{k=1}^n \hat{\alpha}_{ik}\right)^2}$$

and

$$k(\mathbf{x}_k, \boldsymbol{\mu}_i) = \exp\left(-\frac{1}{2} \|\mathbf{x}_k - \boldsymbol{\mu}_i\|^2\right)$$

Scaling of k does not change the updates.

Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Update Rules

$$k(\mathbf{x}_k; \boldsymbol{\mu}_i, \sigma_i^2) = \frac{1}{\sigma_i^m} k(\mathbf{x}_k, \boldsymbol{\mu}_i)^{1/\sigma_i^2}$$

$$p(\mathbf{x}_k) = \sum_{i=1}^l \alpha_i^{\text{old}} k(\mathbf{x}_k; \boldsymbol{\mu}_i, \sigma_i^2)$$

$$\hat{\alpha}_{ik} = \frac{\alpha_i^{\text{old}} k(\mathbf{x}_k; \boldsymbol{\mu}_i, \sigma_i^2)}{p(\mathbf{x}_k)}$$

$$\hat{\alpha}_i = \frac{1}{n} \sum_{k=1}^n \hat{\alpha}_{ik}$$

$$\alpha_i^{\text{new}} = \frac{n \hat{\alpha}_i + \gamma_i - 1}{\gamma_s}$$

$$(\sigma_i^2)^{\text{new}} = \frac{\sum_{k=1}^n \hat{\alpha}_{ik} (-2 \log k(\mathbf{x}_k, \boldsymbol{\mu}_i)) + w S}{n \hat{\alpha}_i + w}$$

$$-2 \log k(\mathbf{x}_l, \boldsymbol{\mu}_i) = \|\mathbf{x}_l - \boldsymbol{\mu}_i\|^2 = \frac{\sum_{k=1}^n \hat{\alpha}_{ik} (-2 \log k(\mathbf{x}_k, \mathbf{x}_l))}{n \hat{\alpha}_i} - \frac{1}{2} \frac{\sum_{k,j=1}^n \hat{\alpha}_{ik} \hat{\alpha}_{ij} (-2 \log k(\mathbf{x}_k, \mathbf{x}_j))}{n^2 \hat{\alpha}_i^2}$$

$$k(\mathbf{x}_l, \boldsymbol{\mu}_i)^{\text{new}} = \exp\left(-\frac{1}{2} \|\mathbf{x}_l - \boldsymbol{\mu}_i\|^2\right)$$

Similarity-Based Clustering: Mixture Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

The covariance matrix can also be represented by the similarities. However, we would require a matrix inversion where the dimension of this matrix is the number of observations.

This is computationally very expensive and the algorithm is no longer stable.

Similarity-Based Mixture Models: Examples

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based


- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM


- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

x7 data set clustered by similarity-based mixture clustering

x7 data



result affinity propagation



Similarity-Based Mixture Models: Examples

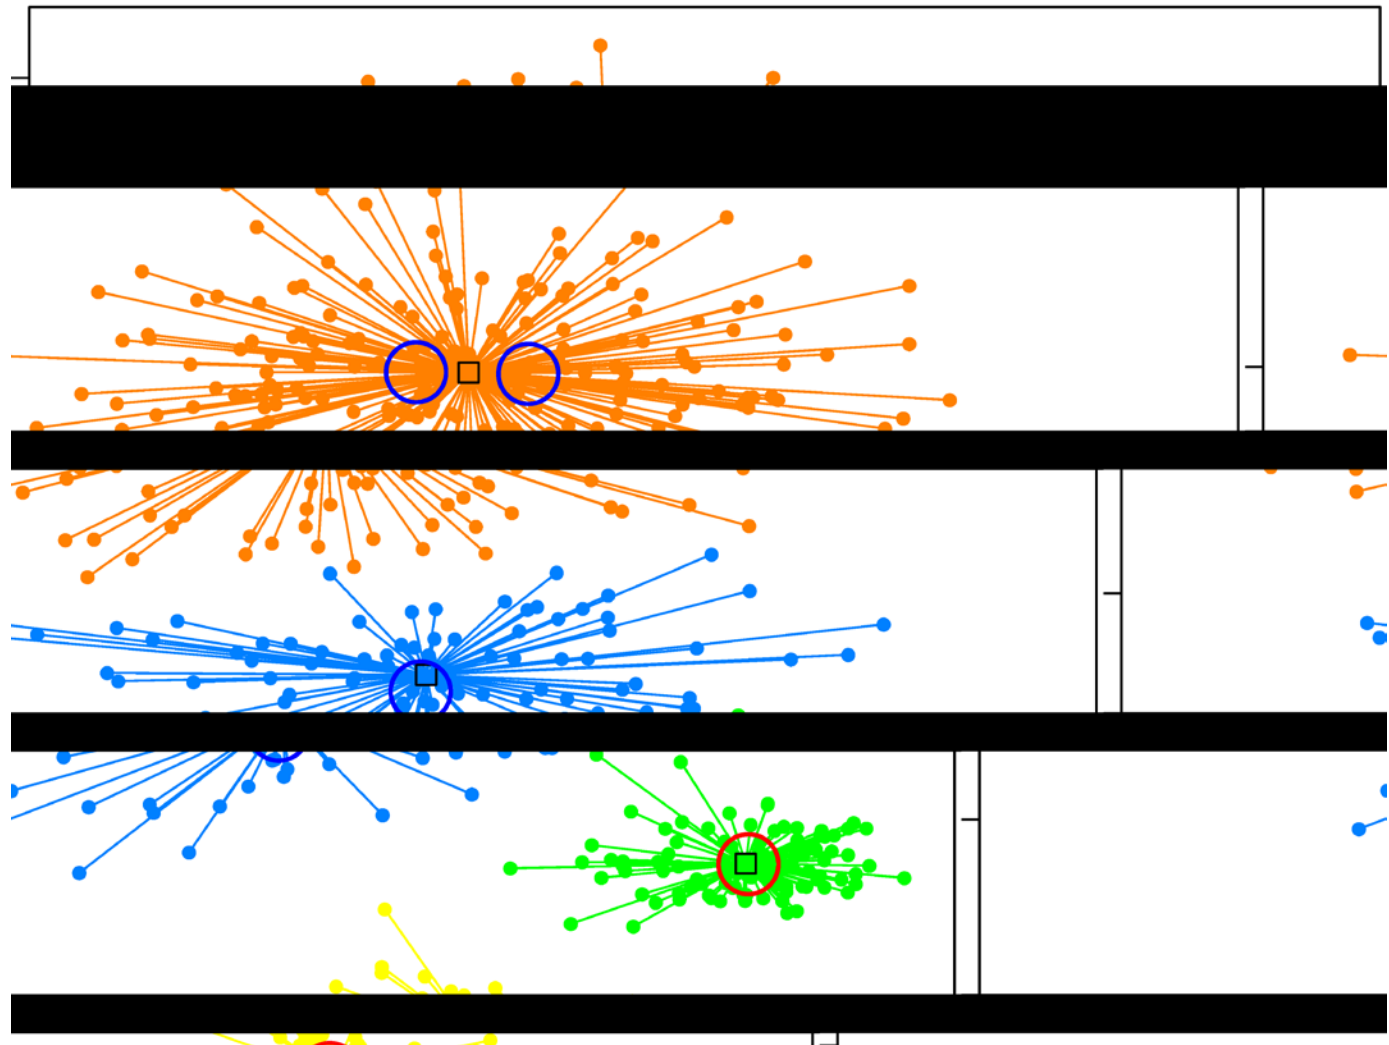
- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

EMcluster



Similarity-Based Mixture Models: Examples




7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples


9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Alpha parameters during EMCluster learning



Sigma parameters (variance) during EMCluster learning



Similarity-Based Mixture Models: Examples

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based


- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM


- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

x7A

x7A data



result affinity propagation



Similarity-Based Mixture Models: Examples


- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

EMcluster



Similarity-Based Mixture Models: Examples

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based


- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM


- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

x9

x9 data



result affinity propagation



Similarity-Based Mixture Models: Examples

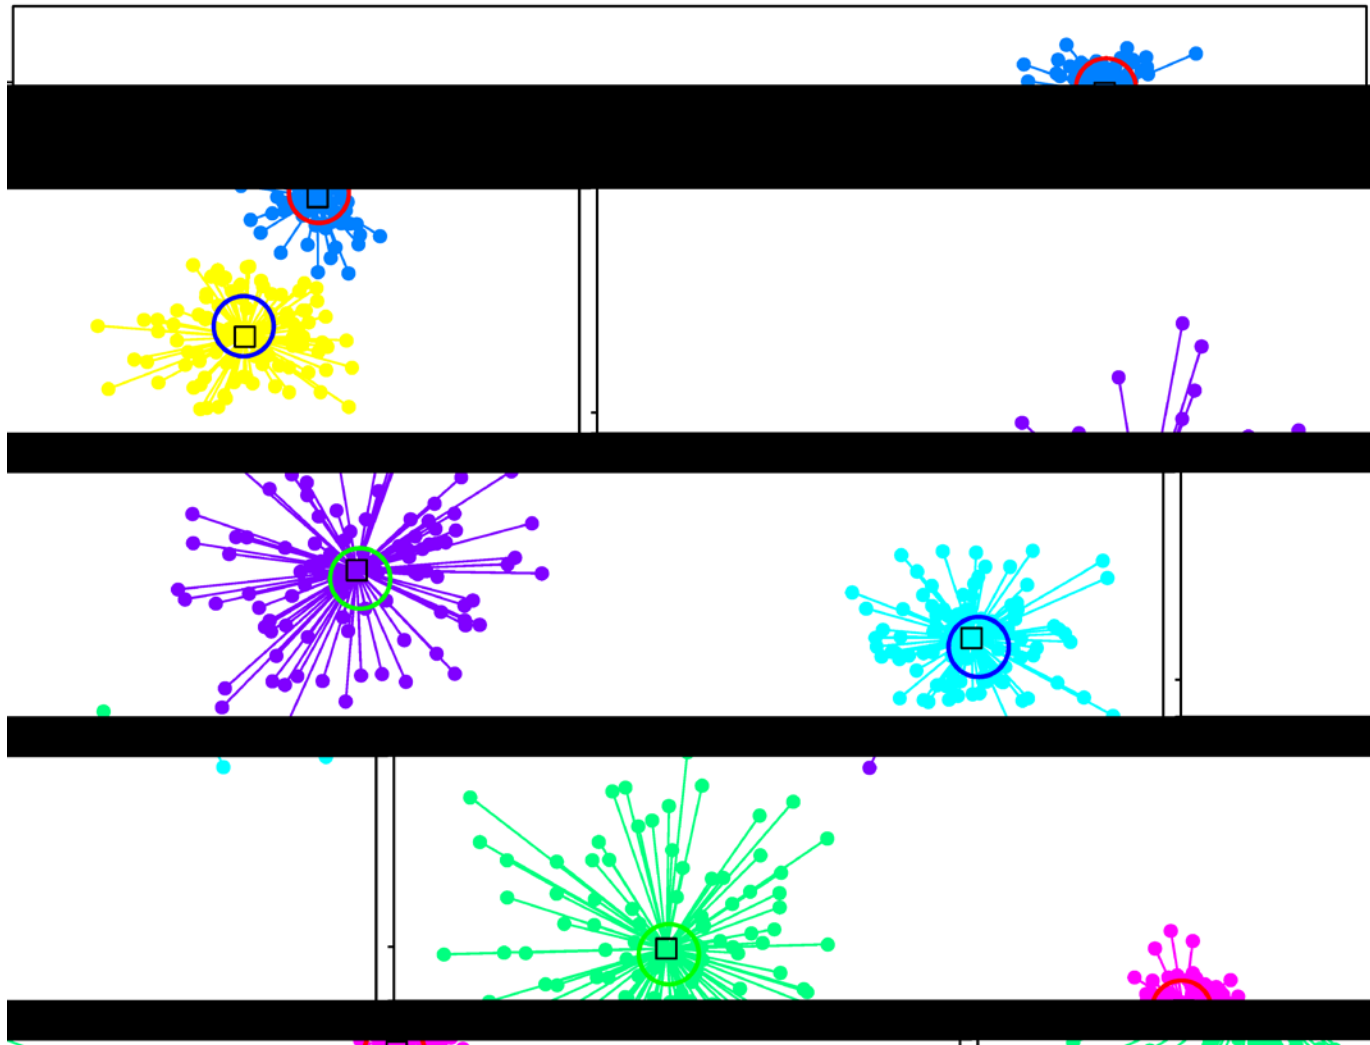
- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

EMcluster



Similarity-Based Mixture Models: Examples

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based


- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM


- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

d6

d6 data



result affinity propagation



Similarity-Based Mixture Models: Examples

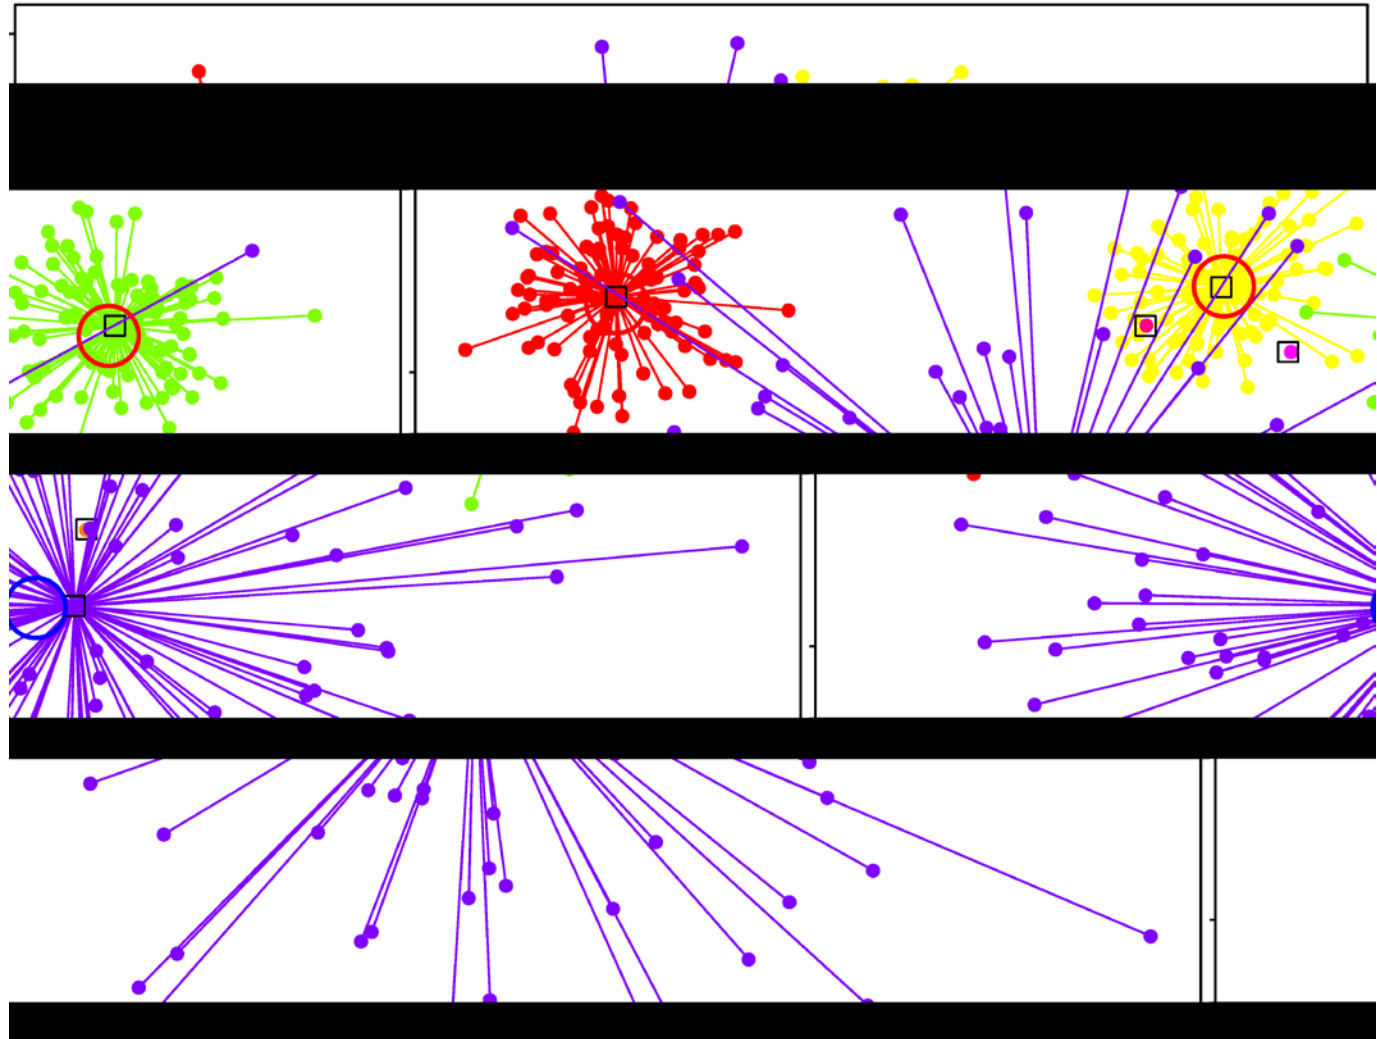
- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

EMcluster



Similarity-Based Mixture Models: Examples

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based


- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM


- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

d6A

d6A data



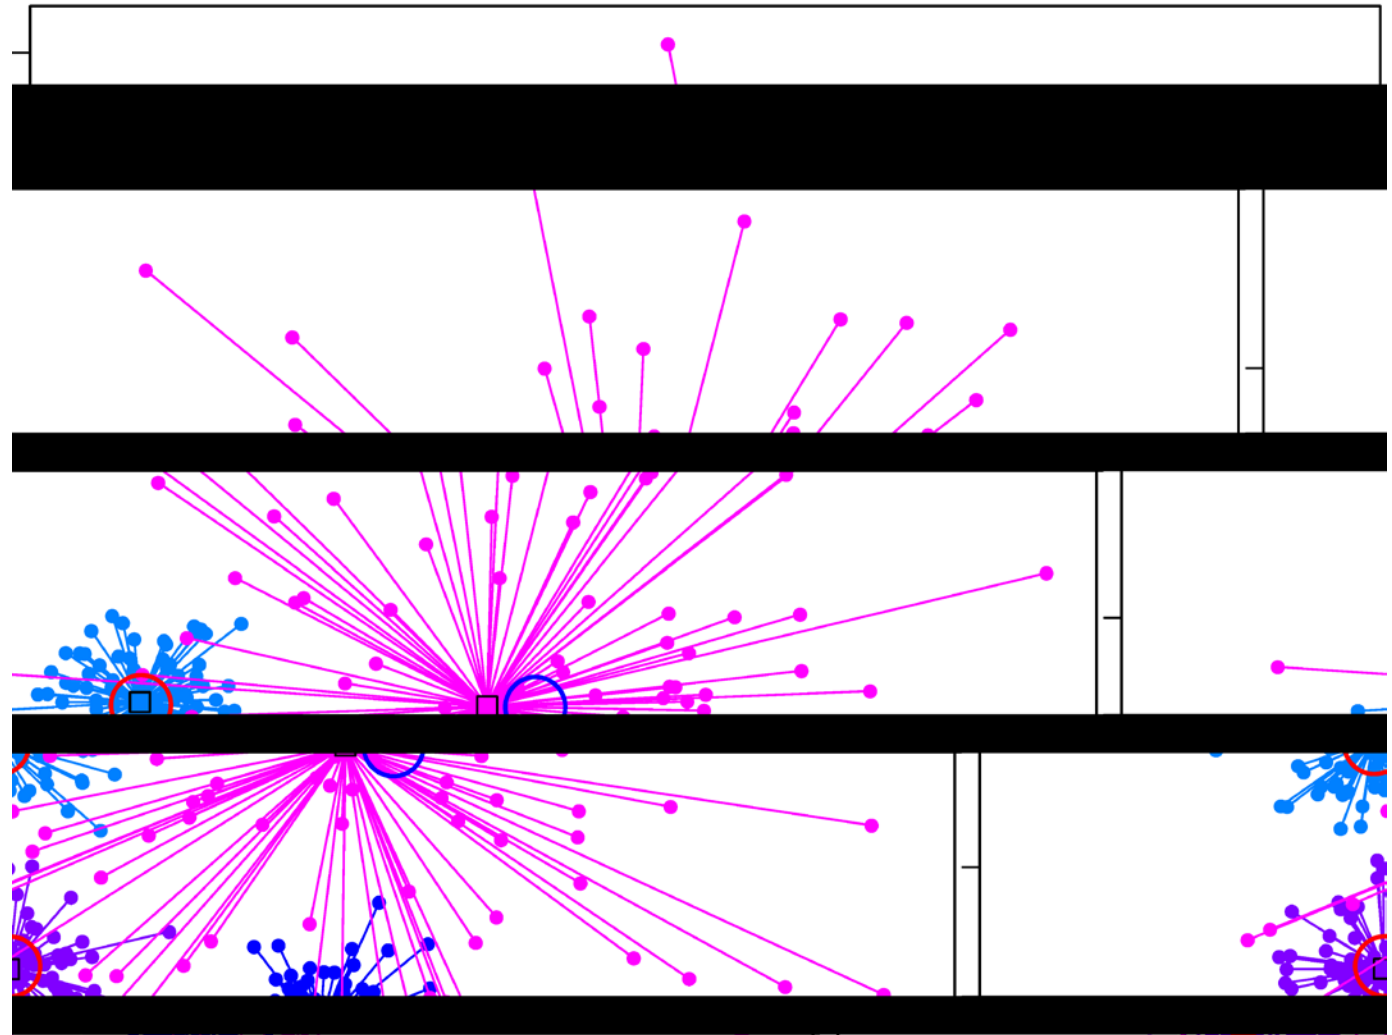
result affinity propagation



Similarity-Based Mixture Models: Examples

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

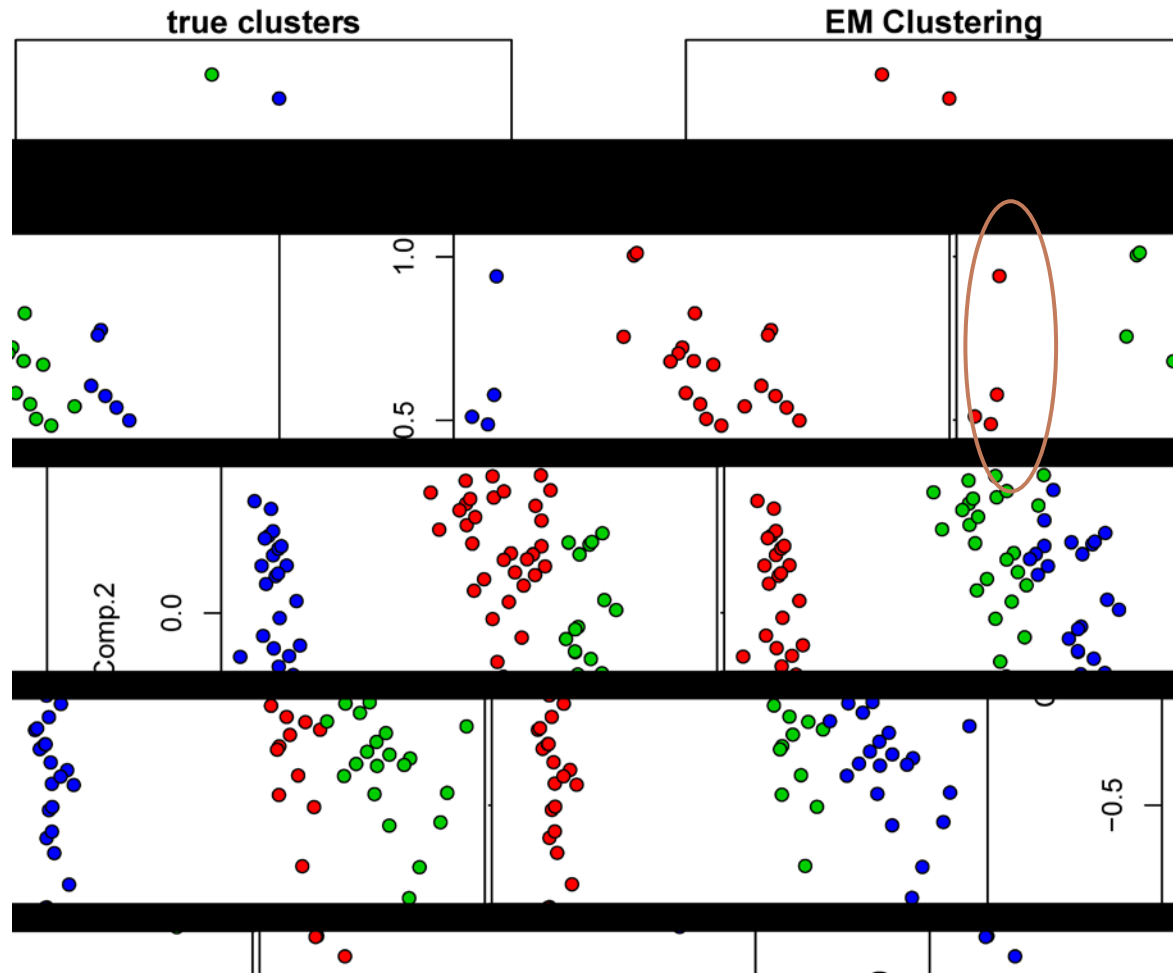
EMcluster



Similarity-Based Mixture Models: Examples

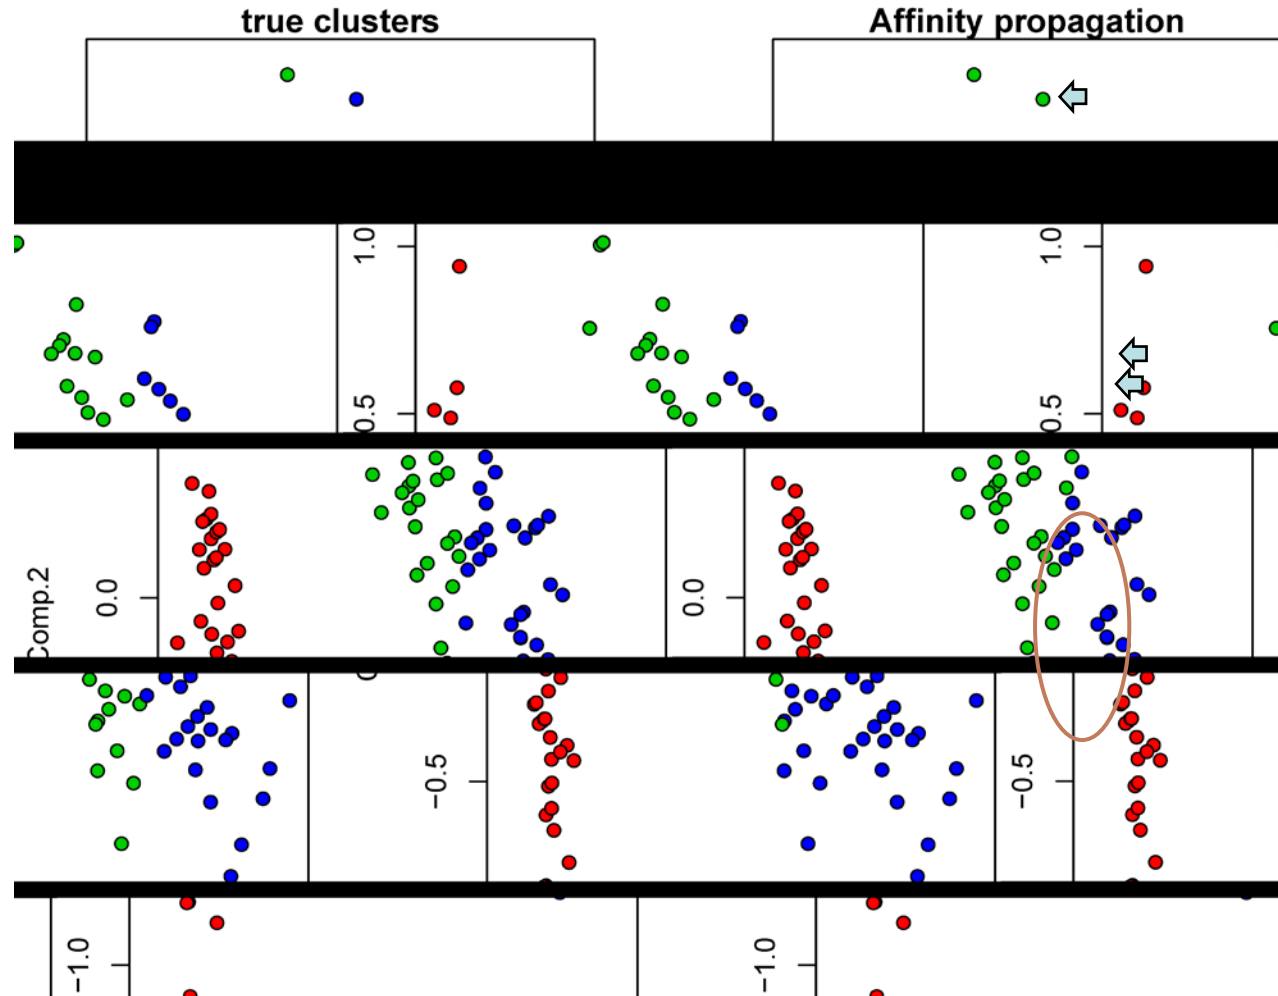
- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

similarity-based mixture clustering of the iris data



Similarity-Based Mixture Models: Examples


- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine



Similarity-Based Mixture Models: Examples


- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

similarity-based mixture clustering of the multiple tissues



Similarity-Based Mixture Models: Examples

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering

8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models

9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines

10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Biclustering simultaneously clusters the rows and the columns of a matrix

Bicluster:

- pair of a set of row elements and a set of column elements
- column elements are similar to each other on row elements and vice versa

row and column elements can belong to **multiple** or to **no** bicluster

contrast to standard clustering

- row elements clustered only on a subgroup of column elements
- multiple or not cluster membership

Biclustering

- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

8 Biclustering

- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

9 Hidden Markov Models


- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

10 Boltzmann Machines

- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

examples of
biclusters

for visualization
the biclusters
consist of
adjacent row and
column elements



Biclustering: Types of Biclusters

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Types of bicluster:

- constant values
- constant rows values
- constant column values
- additive coherent values
- multiplicative coherent values
- general coherent values

1.0	4.0	5.0	0.0	1.5
4.0	7.0	8.0	3.0	4.5
3.0	6.0	7.0	2.0	3.5
5.0	8.0	9.0	4.0	5.5
2.0	5.0	6.0	1.0	2.5

coherent values: additive

1.0	0.5	2.0	0.2	0.8
2.0	1.0	4.0	0.4	1.6
3.0	1.5	6.0	0.6	2.4
4.0	2.0	8.0	0.8	3.2
5.0	2.5	10.0	1.0	4.0

coherent values: multiplicative

2.0	2.0	2.0	2.0	2.0
2.0	2.0	2.0	2.0	2.0
2.0	2.0	2.0	2.0	2.0
2.0	2.0	2.0	2.0	2.0
2.0	2.0	2.0	2.0	2.0

constant bicluster

1.0	1.0	1.0	1.0	1.0
2.0	2.0	2.0	2.0	2.0
3.0	3.0	3.0	3.0	3.0
4.0	4.0	4.0	4.0	4.0
4.0	4.0	4.0	4.0	4.0

constant row bicluster

1.0	2.0	3.0	4.0	5.0
1.0	2.0	3.0	4.0	5.0
1.0	2.0	3.0	4.0	5.0
1.0	2.0	3.0	4.0	5.0
1.0	2.0	3.0	4.0	5.0

constant column bicluster

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

Biclustering methods:

- variance minimization methods
- two-way clustering methods
- motif and pattern recognition methods
- probabilistic and generative approaches

variance minimization methods define clusters as blocks in the matrix with minimal deviation of their elements: Cheng-Church δ -biclusters; δ -cluster methods search for blocks of elements having a deviation below δ ; δ -pClusters search sub-matrices with pairwise edge differences less than δ

Two-way clustering methods apply conventional clustering to the columns and rows and (iteratively) combine the results: Coupled Two-Way Clustering (CTWC), Interrelated Two-Way Clustering (ITWC), Double Conjugated Clustering (DCC)

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Motif and pattern recognition methods define a bicluster as samples sharing a common pattern or motif: xMOTIF, Order-Preserving Sub-Matrices (OPSM), Spectral clustering (SPEC), Iterative Signature Algorithm (ISA)

Probabilistic and generative methods use model-based techniques to define biclusters: Statistical-Algorithmic Method for Bicluster Analysis (SAMBA), Probabilistic Relational Models (PRMs), ProBic, cMonkey, plaid models, Bayesian BiClustering model (BBC), Factor Analysis for Bicluster Acquisition (FABIA)

BBC and FABIA are generative models:

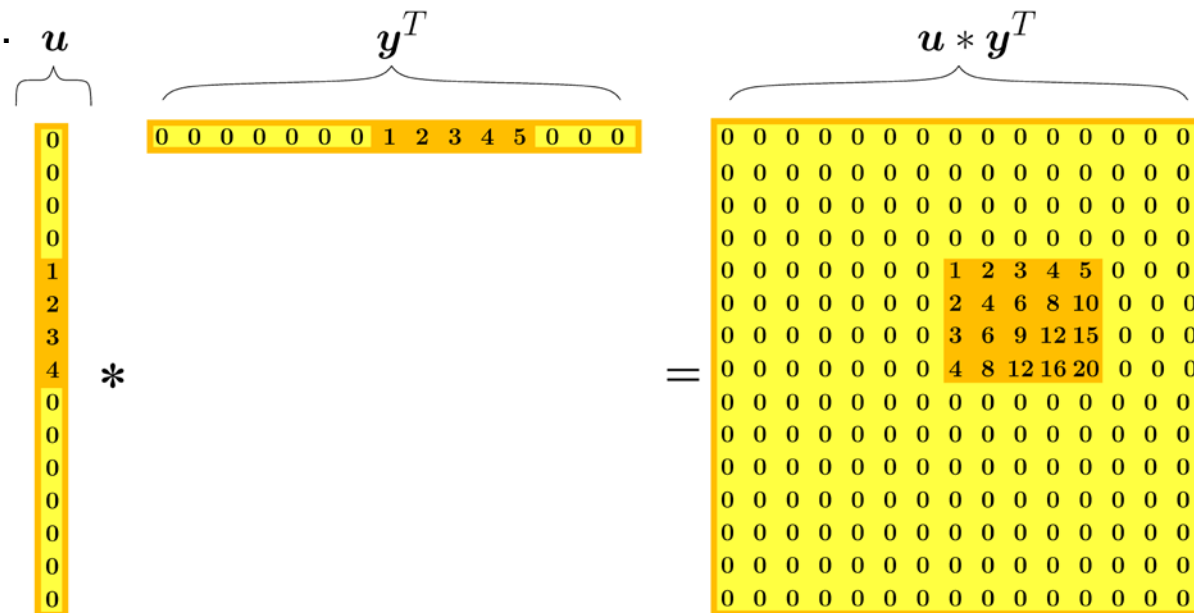
1. well-understood model selection techniques (maximum likelihood)
2. hyperparameter selection within the Bayesian framework
3. signal-to-noise ratios
4. model comparisons via the likelihood or posterior
5. tests like the likelihood ratio test
6. global model to explain all data

Biclustering: FABIA

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

FABIA biclustering defines a bicluster as an outer product $u y^T$. The vector u corresponds to a **prototype column vector** that contains zeros for features not participating in the bicluster. The vector y is a vector of **factors** with which the prototype column vector is scaled for each sample and contains zeros for samples not participating in the bicluster.

Vectors containing many zeros or values close to zero are called **sparse vectors**.



u y^T $u * y^T$

u : 0, 0, 0, 0, 1, 2, 3, 4, 5, 0, 0, 0

y^T : 0, 0, 0, 0, 0, 0, 1, 2, 3, 4, 5, 0, 0, 0

$u * y^T$ matrix:

0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0

Biclustering: FABIA



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

model for l biclusters and additive noise:

$$\mathbf{X} = \sum_{j=1}^l \mathbf{u}_j \mathbf{y}_j^T + \mathbf{\Upsilon} = \mathbf{Y} \mathbf{U}^T + \mathbf{\Upsilon}$$

data matrix: $\mathbf{X} \in \mathbb{R}^{n \times m}$

additive noise: $\mathbf{\Upsilon} \in \mathbb{R}^{n \times m}$

sparse prototype vector: $\mathbf{u}_i \in \mathbb{R}^m$

sparse vector of factors: $\mathbf{y}_j \in \mathbb{R}^n$ ($\mathbf{y}_j^T = (y_{1j}, \dots, y_{nj})$ values for samples)

sparse prototype matrix: $\mathbf{U} \in \mathbb{R}^{m \times l}$

sparse factor matrix: $\mathbf{Y} \in \mathbb{R}^{n \times l}$

$$\mathbf{x}_i = \sum_{j=1}^l \mathbf{u}_j y_{ij} + \boldsymbol{\epsilon}_i = \mathbf{U} \tilde{\mathbf{y}}_i + \boldsymbol{\epsilon}_i$$

column of the noise matrix: $\boldsymbol{\epsilon}_i$

i -th row of factor matrix: $\tilde{\mathbf{y}}_i = (y_{i1}, \dots, y_{il})^T$ (one value per bicluster)

Biclustering: FABIA



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

factor analysis model with l factors:

$$\mathbf{x} = \sum_{j=1}^l \mathbf{u}_j \tilde{y}_j + \boldsymbol{\epsilon} = \mathbf{U} \tilde{\mathbf{y}} + \boldsymbol{\epsilon}$$

observations: \mathbf{x}

loading matrix: \mathbf{U}

j -th factor: \tilde{y}_j

vector of factors: $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_l)^T$

additive noise: $\boldsymbol{\epsilon} \in \mathbb{R}^m$

Assumptions:

noise is independent of factors

$$\tilde{\mathbf{y}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

$$\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Psi})$$

$\boldsymbol{\Psi} \in \mathbb{R}^{l \times l}$ is diagonal (independent Gaussian noise)

parameter \mathbf{U} explains the dependent (common) and $\boldsymbol{\Psi}$ the independent variance in the observations

Biclustering: FABIA



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

FABIA formulation requires sparse factors and loadings to represent biclusters. Sparseness is obtained by a component-wise independent Laplace distribution:

$$p(\tilde{\mathbf{y}}) = \left(\frac{1}{\sqrt{2}}\right)^l \prod_{j=1}^l e^{-\sqrt{2} |\tilde{y}_j|}$$

$$p(\mathbf{u}_j) = \left(\frac{1}{\sqrt{2}}\right)^m \prod_{k=1}^m e^{-\sqrt{2} |u_{kj}|}$$

Unfortunately, for FABIA, the likelihood is analytically intractable because of the Laplacian priors.

Expectation maximization (EM) has been extended to variational expectation maximization for model selection and finding the optimal parameters. The variational EM maximizes the posterior.

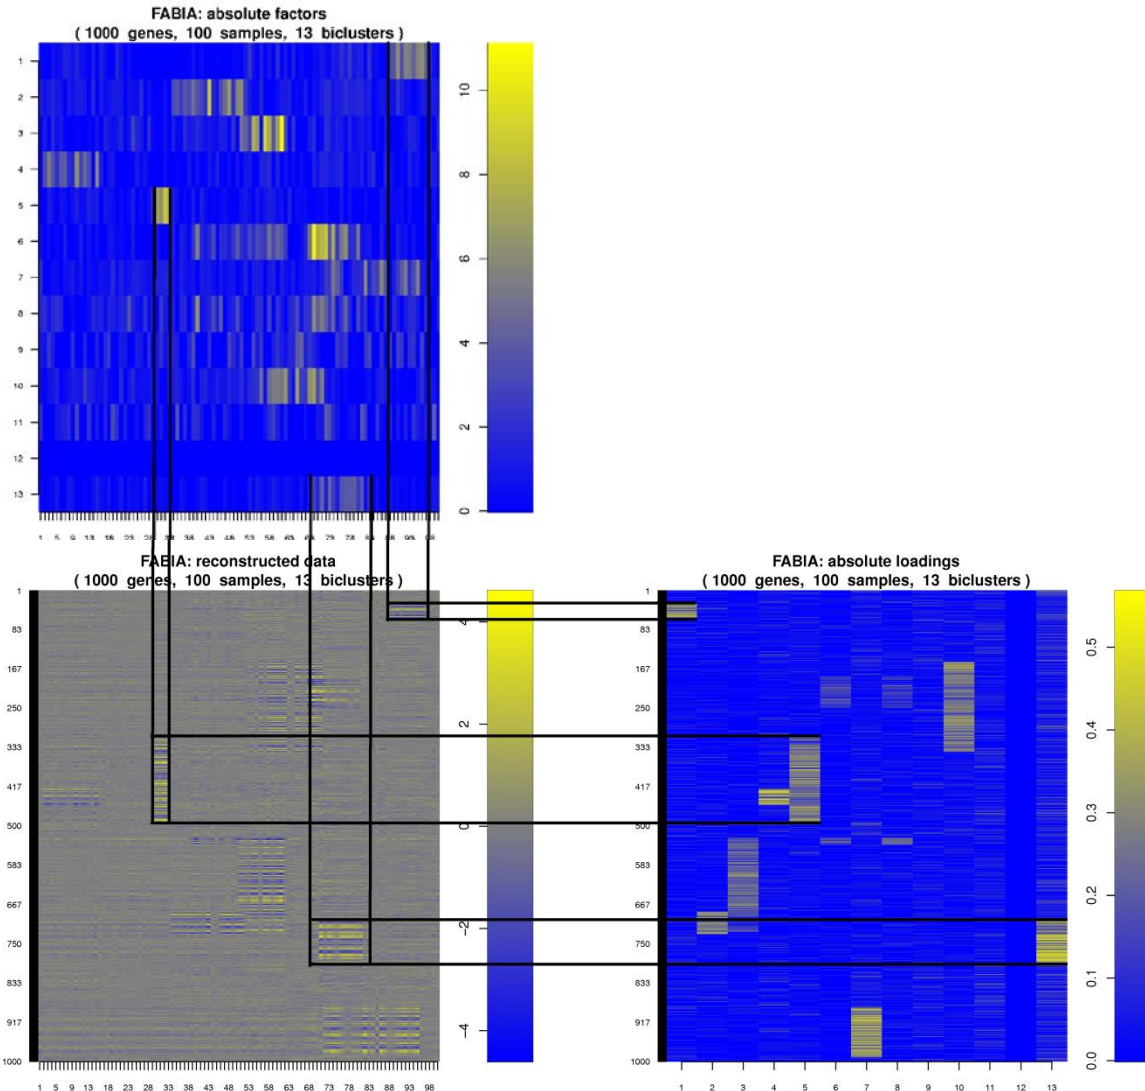
Biclustering: FABIA

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine



Biclustering: Examples


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

We test FABIA on a 50-dimensional data set with linearly mixed super-Gaussians



- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

FABIA biclustering of the iris data

The loadings of FABIA are

	bicluster2	bicluster3	bicluster1
Sepal.Length	0.6490253	-0.2155847	0
Sepal.Width	-0.1828042	-0.3511878	0
Petal.Length	1.7465614	0.0000000	0
Petal.Width	0.7285857	-0.0453688	0

Only two biclusters have been found.

- first bicluster “bicluster2” focuses on “Petal.Length” which is correlated to “Petal.Width” and “Sepal.Length”. **This bicluster is related to petal** but includes also sepal length.
- second bicluster “bicluster3” removed “Petal.Length” and “Petal.Width” due to sparseness and focused on “Sepal.Width” including also “Sepal.Length”. **This bicluster is related to sepal.**

Biclustering: Examples


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

biplot that plots both the features and the samples into one plot



First bicluster on the x-axis is driven by Petal.Length and separates the species

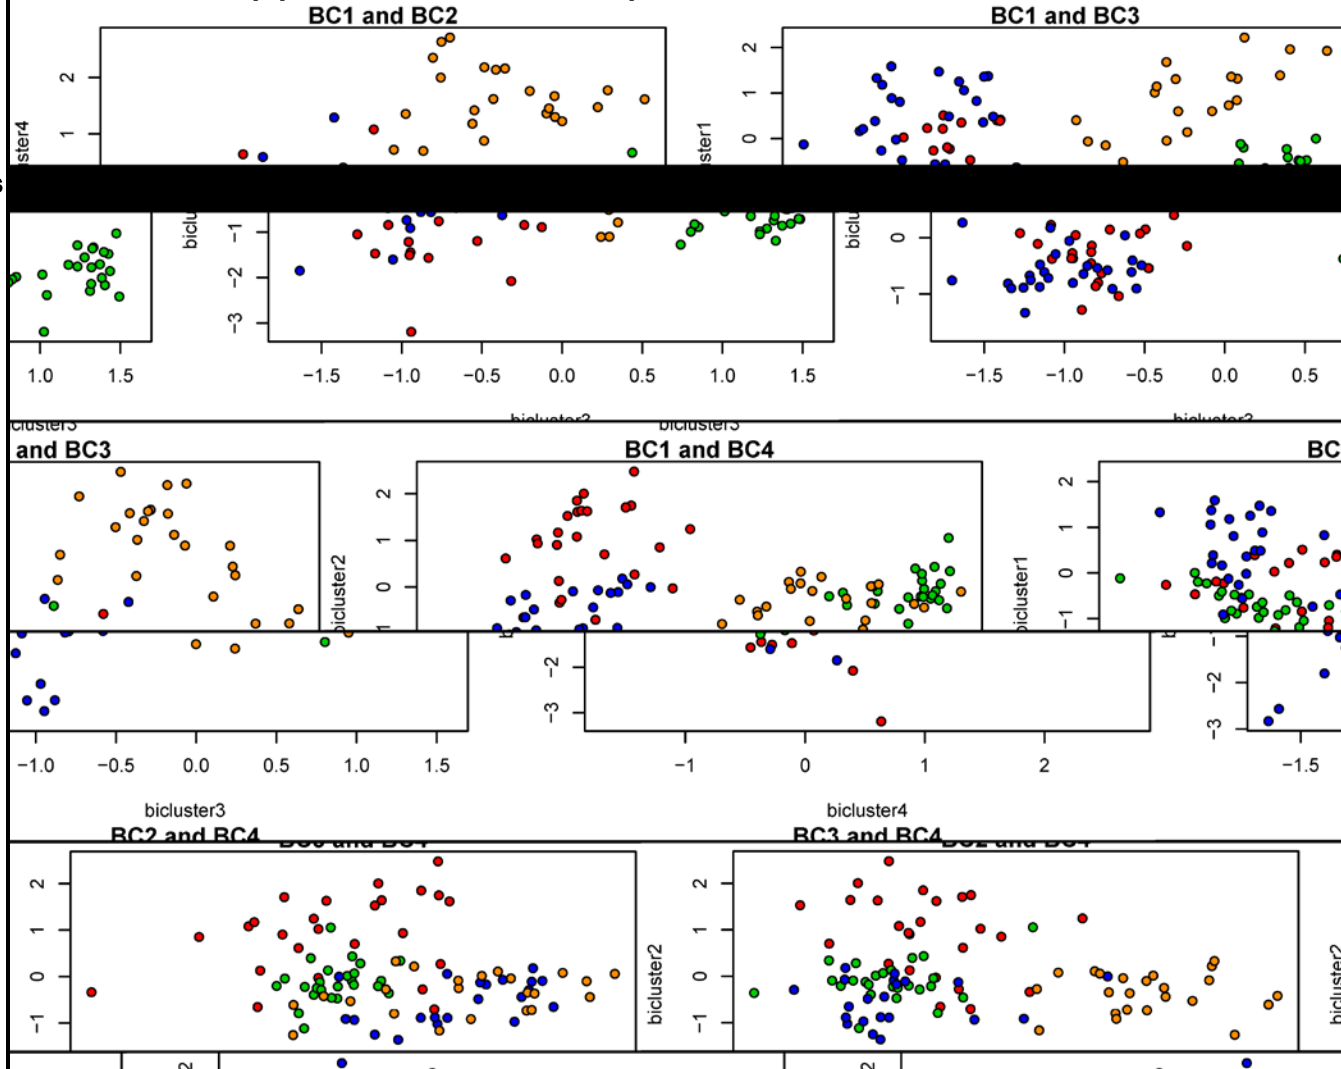
Second bicluster is driven by Sepal.Width

Sepal.Length belongs to both clusters

Biclustering: Examples

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

FABIA applied to the multiple tissue data: 200 iterations



first bicluster (“bicluster3”) separates prostate (green) and colon (orange)

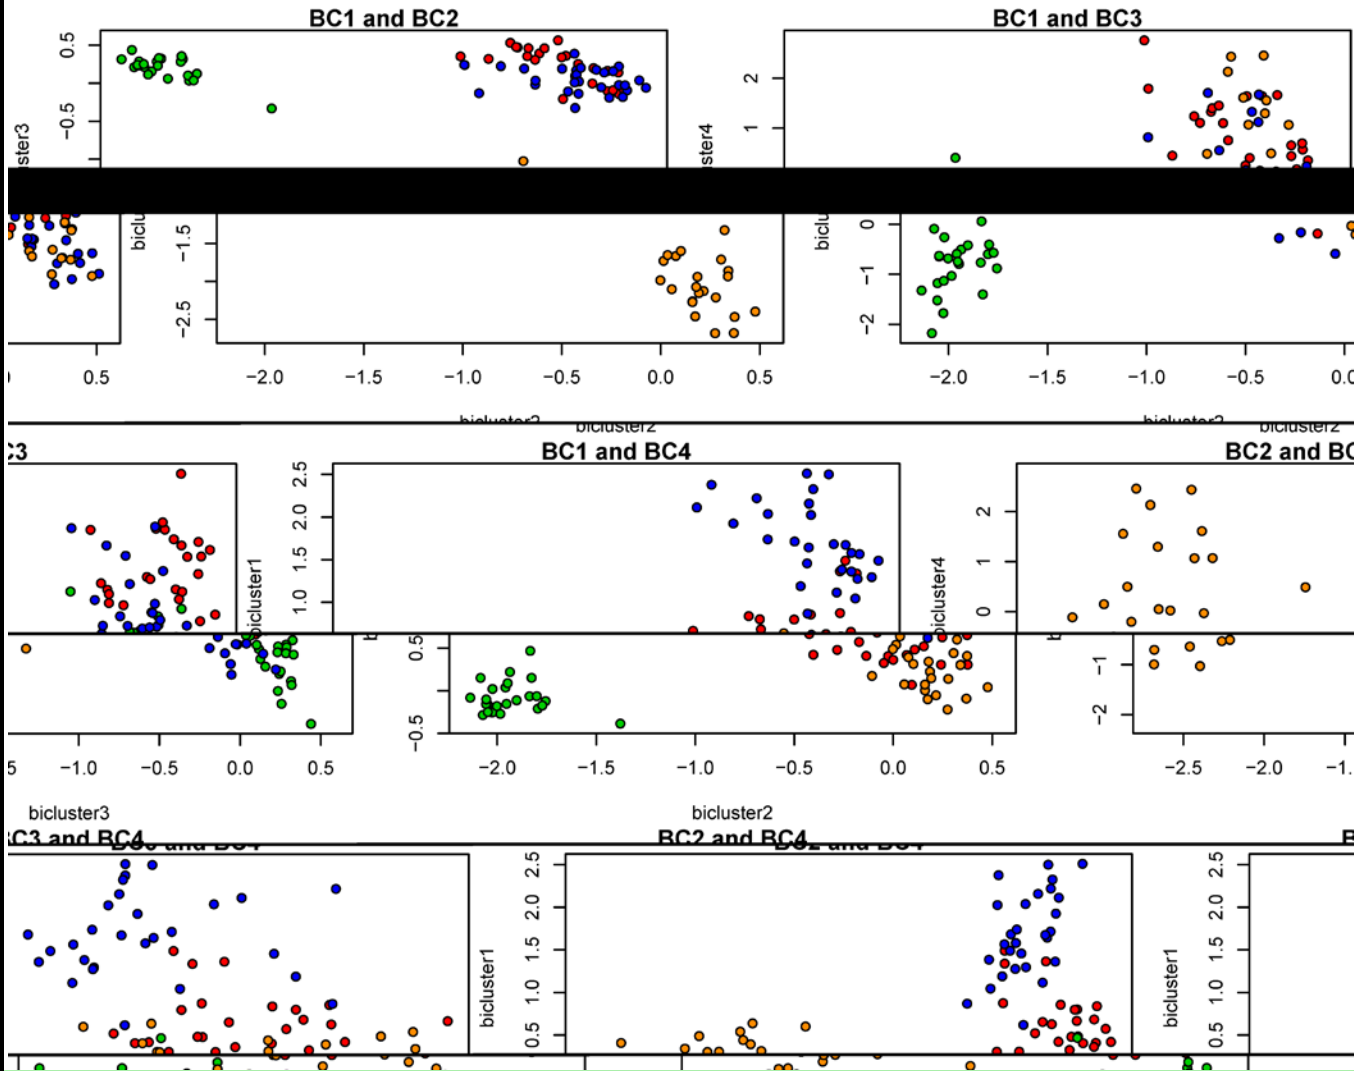
second bicluster (“bicluster4”) separates colon (orange)

fourth bicluster (“bicluster2”) separates breast (red), though not perfectly.

Biclustering: Examples

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

FABIA applied to the multiple tissue data: 500 iterations



first bicluster
("biclust2")
separates
prostate
(green)

second
bicluster
("biclust3")
separates
colon
(orange)

fourth
bicluster
("biclust2")
separates
lung (blue)
from the rest

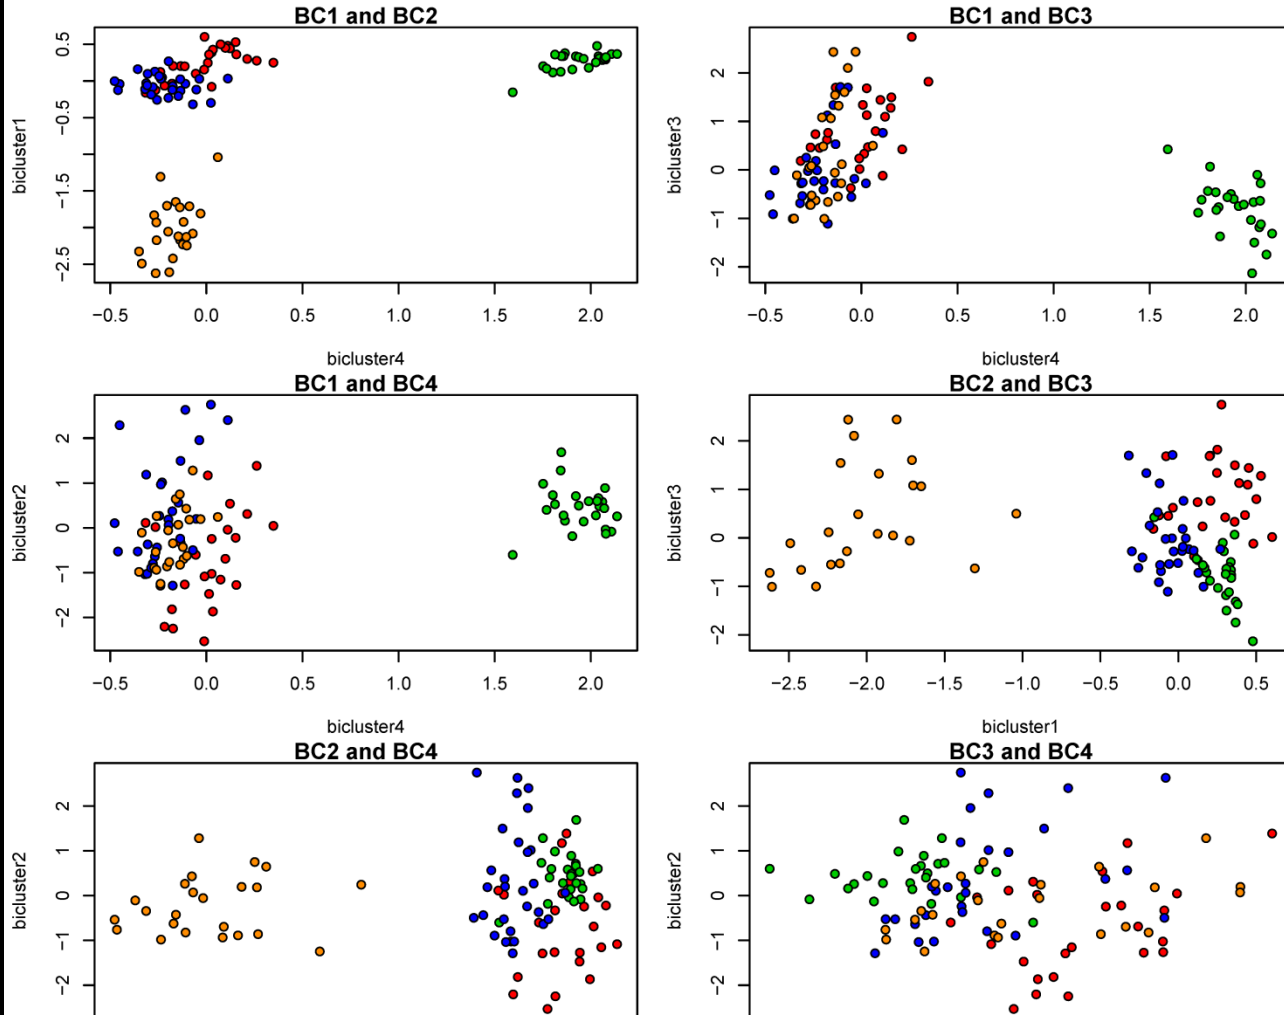
Separation is
better than
with 200
iterations

Biclustering: Examples



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

FABIA applied to the multiple tissue data: 2,000 iterations



first bicluster (“biclustern4”) clearly separates prostate (green)

The second bicluster (“biclustern1”) separates colon (orange)

third and fourth bicluster (“biclustern3/2”) help to separate lung (blue) and breast (red)

more iterations: the solutions become more sparse and more clear separation

Biclustering: Examples



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

Since the FABIA solution is sparse, it allows for an interpretation.
most relevant genes of the first bicluster:

"KLK3" "ACPP" "KLK2" "CUTL2" "RNF41"

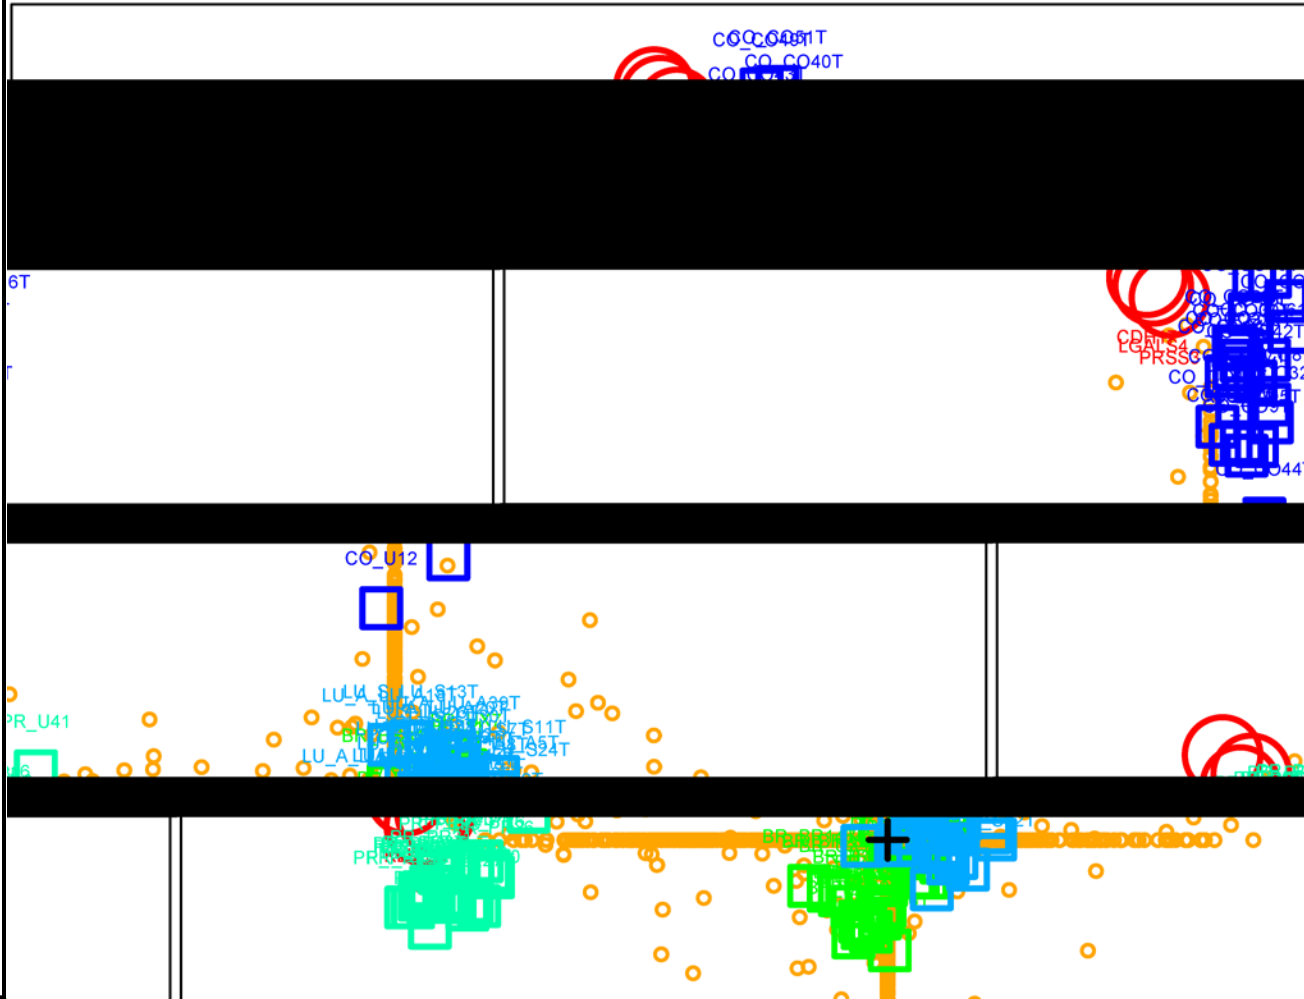
- KLK3 is known as the prostate specific antigen. “Serum level of this protein, called PSA in the clinical setting, is useful in the diagnosis and monitoring of prostatic carcinoma.”
- ACPP “is synthesized under androgen regulation and is secreted by the epithelial cells of the prostate gland.”
- KLK2 “is primarily expressed in prostatic tissue and is responsible for cleaving pro-prostate-specific antigen into its enzymatically active form.”

This means that the most relevant genes of the first bicluster are all strongly associated with prostate tissue.

Biclustering: Examples



biplot of the first and second bicluster of FABIA (2000 iterations)



The large red circles are the features (genes) driving the bicluster while the golden small circles are features with minor influence. Sparseness pushes most features to zero in one direction (the axes are golden due to the features pushed to zero). The turquoise prostate samples at the x-axis are separated from the other samples.

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Biclustering: Examples

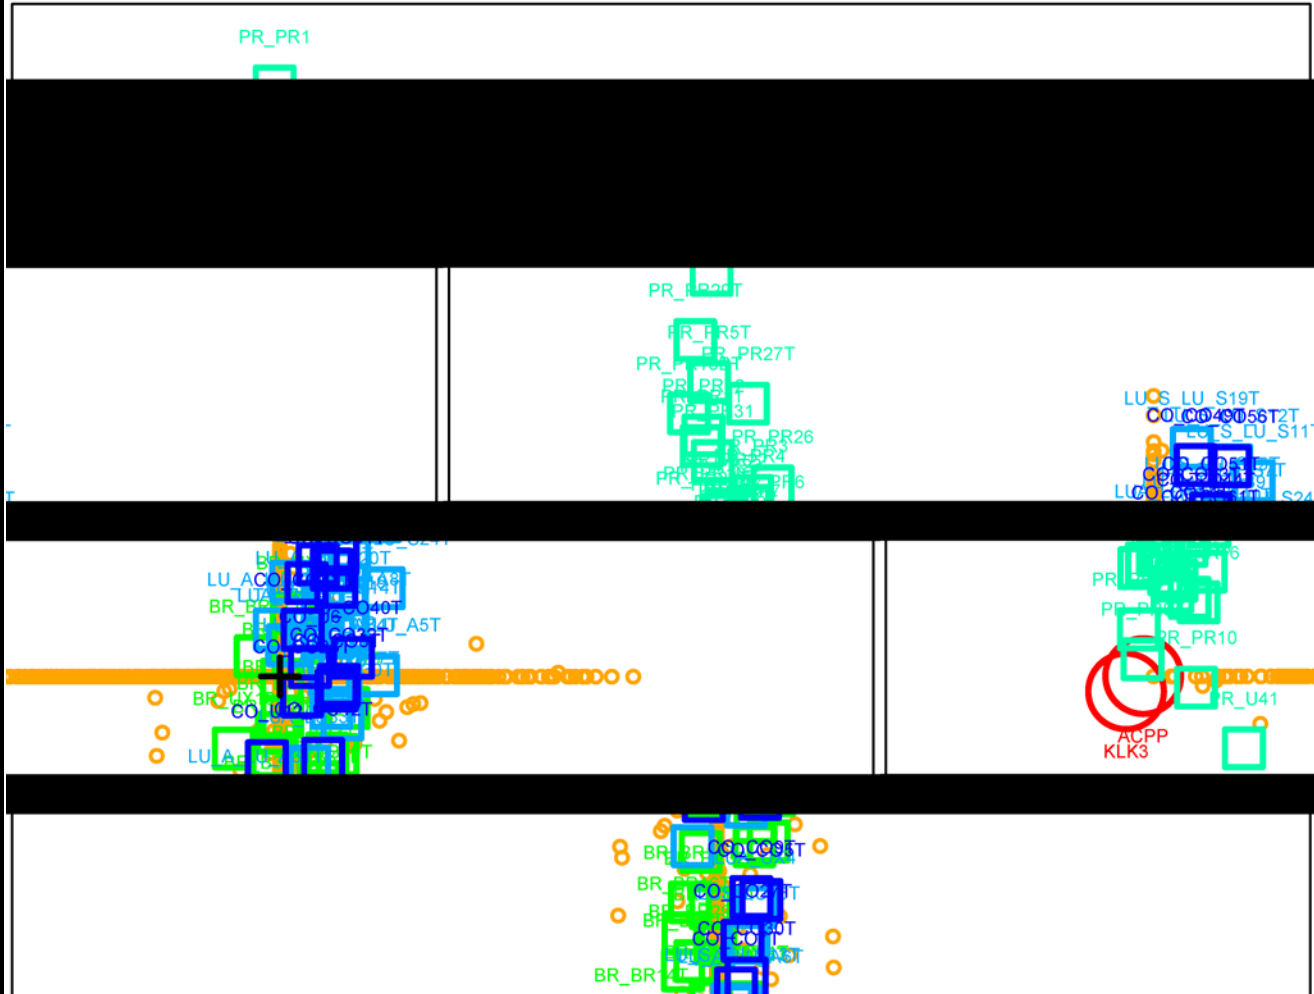
7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Biplot for the first and third bicluster: 2,000 iterations



For the first bicluster at the x-axis the genes KLK3 (prostate specific antigen) and ACPP (secreted by the epithelial cells of the prostate gland) can be recognized. These genes separate the turquoise prostate samples at the x-axis from the other samples.

Biclustering: Examples

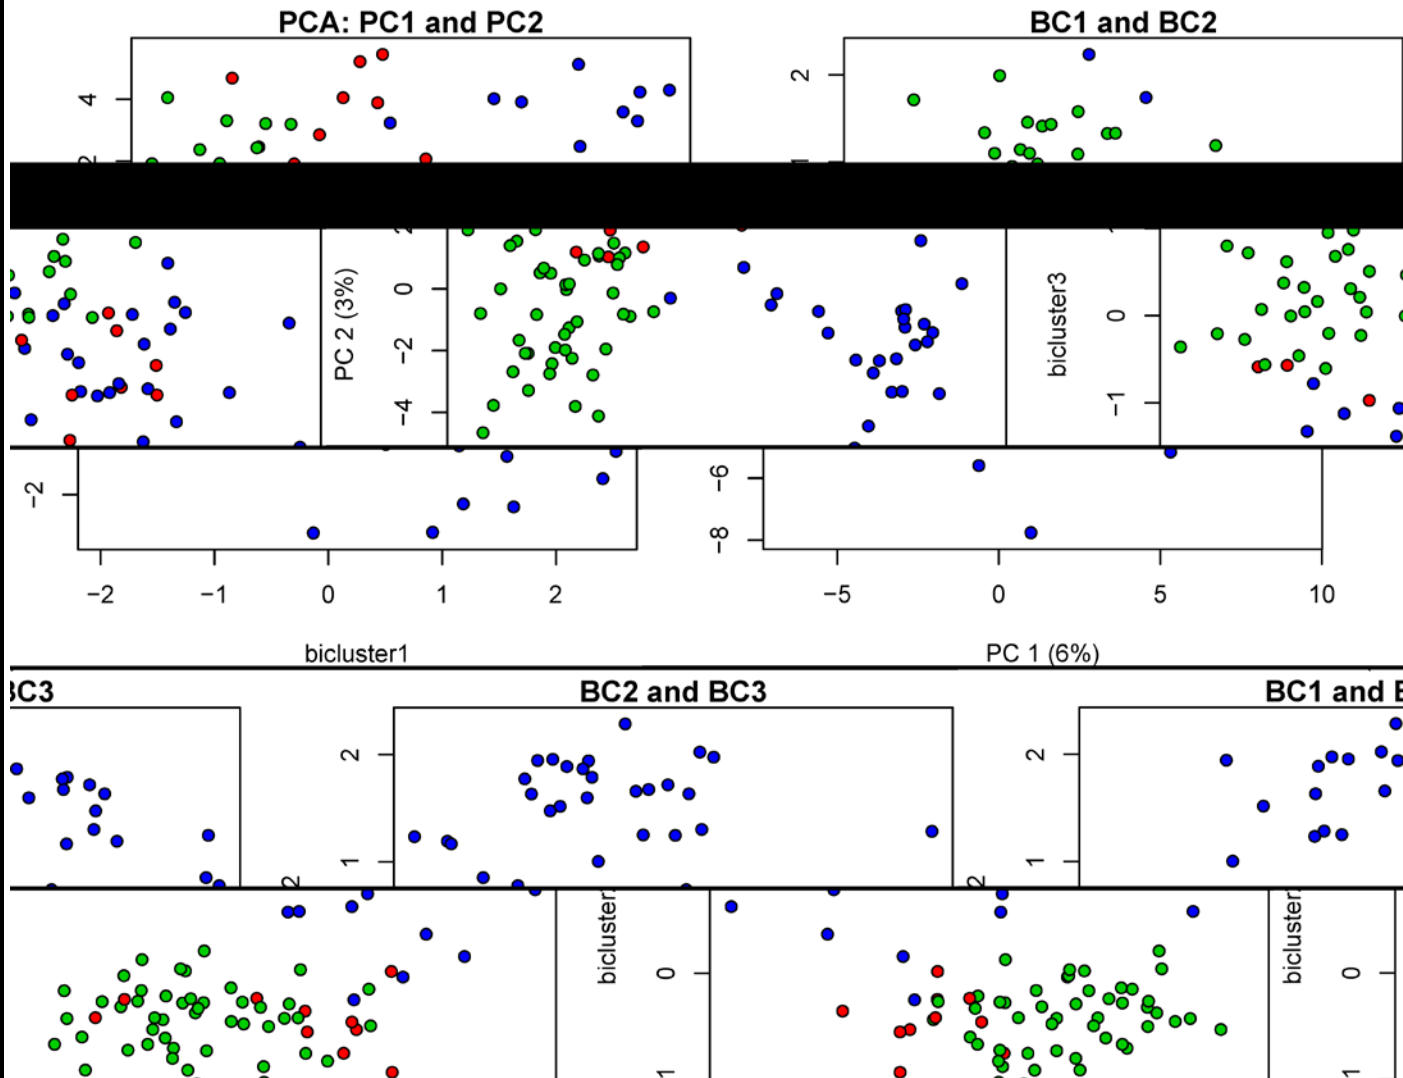
7 Clustering
 7.1 Mixture Models
 7.2 k-Means
 7.3 Hierarchical
 7.4 Similarity-Based

8 Biclustering
 8.1 Types of Biclusters
 8.2 Overview Methods
 8.3 FABIA
 8.4 Examples

9 Hidden Markov Models
 9.1 HMMs in Bioinf.
 9.2 HMM Basics
 9.3 EM / Baum-Welch
 9.4 Viterby
 9.5 Input Output HMM
 9.6 Factorial HMM
 9.7 Memory IOHMM
 9.8 Tricks of the Trade
 9.9 Profile HMM

10 Boltzmann Machines
 10.1 The Boltzmann Machine
 10.2 Learning
 10.3 Restricted Boltzmann Machine

FABIA applied to the breast cancer data:



Both PCA and biclustering separate the blue class but have problems to separate the red class.

Biclustering: Examples


- 7 Clustering
- 7.1 Mixture Models
- 7.2 k-Means
- 7.3 Hierarchical
- 7.4 Similarity-Based

- 8 Biclustering
- 8.1 Types of Biclusters
- 8.2 Overview Methods
- 8.3 FABIA
- 8.4 Examples

- 9 Hidden Markov Models
- 9.1 HMMs in Bioinf.
- 9.2 HMM Basics
- 9.3 EM / Baum-Welch
- 9.4 Viterby
- 9.5 Input Output HMM
- 9.6 Factorial HMM
- 9.7 Memory IOHMM
- 9.8 Tricks of the Trade
- 9.9 Profile HMM

- 10 Boltzmann Machines
- 10.1 The Boltzmann Machine
- 10.2 Learning
- 10.3 Restricted Boltzmann Machine

Biplot breast cancer



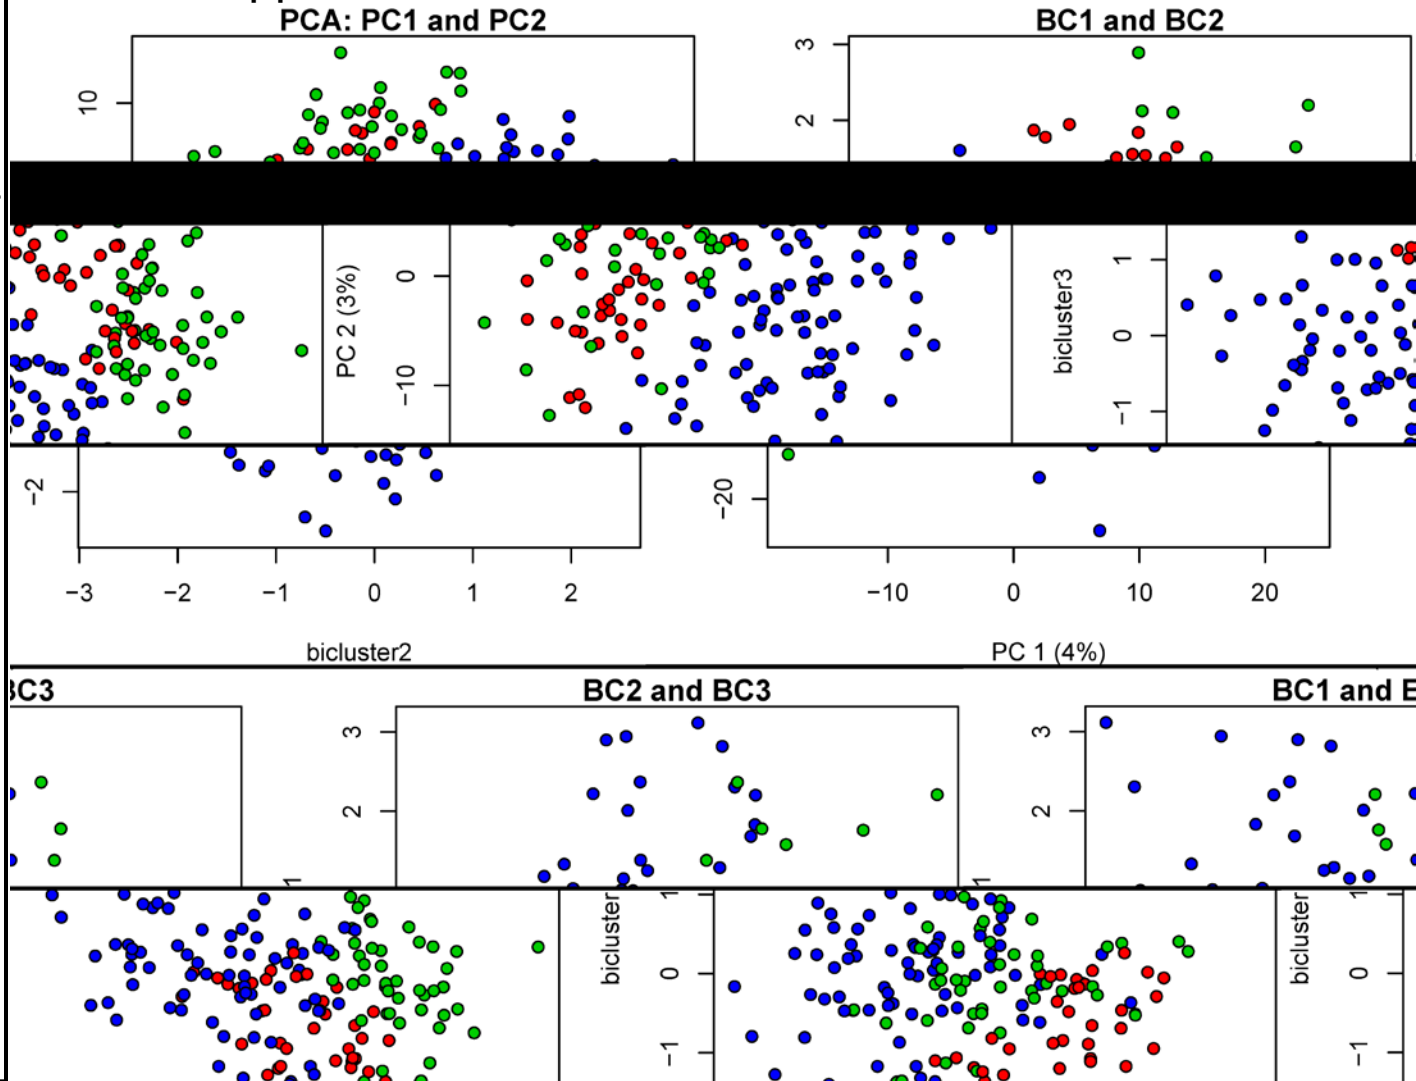
The biplot for bicluster 2 and 3.

The subclasses are quite good separated.

Biclustering: Examples

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

FABIA applied to the DLBCL data:



PCA separates the blue subclass.

Biclustering separates subclasses better because red cluster is better separated.

Biclustering: Examples




7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine



FABIA applied to DLBCL data.

The biplot for bicluster 1 and 2.

The subclasses are quite good separated.

Chapter 9

Hidden Markov Models

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

first chapter devoted to unsupervised learning

→ generative models

most prominent generative model in bioinformatics
hidden Markov model

well suited for analyzing protein or DNA sequences: discrete

Hidden Markov Models in Bioinformatics



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models

9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

generative approach: generating output sequences

New sequences: similar to the model-building sequences

Similar: → certain patterns in a sequences

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

HMMs for gene prediction

exons and introns are identified e.g. GENSCAN

base-pair specificity: 50% - 80%

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Profile HMMs

store a multiple alignment in a hidden Markov model

new sequences: evaluated by likelihood

HMMs are build from unaligned sequences: local maxima

avoid: → deterministic annealing (HMMER)

software packages: HMMER
SAM

- cannot discover long-range dependencies
- cannot deal with real-valued inputs
- cannot detect higher order correlations
- cannot use a negative set

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Other HMMs Applications

- remote homology detection
- scoring
- combined with trees
- data base is build on HMMs: PFAM (protein family database) \

Hidden Markov Model Basics

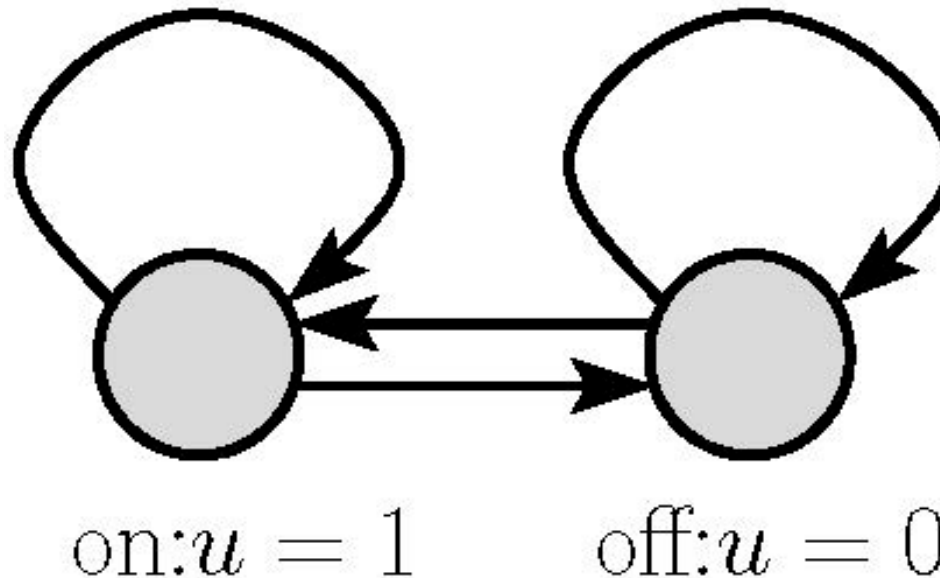
7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

graph of connected hidden states $u \in \{1, \dots, S\}$
each state produces a probabilistic output




Hidden Markov Model Basics

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

model **evolves over time** t (sequence position)

At each step it jumps into another state or remains in current

$$\text{At } t: u_t \in \{1, \dots, S\}$$



Hidden Markov Model Basics


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

all possible sequences of values of the hidden state:



Each path from left to right is a possible sequence

Hidden Markov Model Basics: Transition

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples


9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

transition probabilities: $p(a | b)$ $a, b \in \{1, \dots, S\}$
 b is the current state, a next state

Markov assumption: next state only depends on current state

Higher order hidden Markov models: eg second order $p(a | b, c)$



Hidden Markov Model Basics: Transition

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM


10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

transition probability: $p(u_t | u_{t-1})$

start state probability: $p_S(u_1)$

probability of observing sequence $u^T = (u_1, u_2, u_3, \dots, u_T)$ length T

$$p(u^T) = p_S(u_1) \prod_{t=2}^T p(u_t | u_{t-1})$$



Hidden Markov Model Basics: Emission




- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine


emission probability $p_E(x_t | u_t)$

x_t : element of the output alphabet of size P

$x_t \in \{A, T, C, G\}$



Shine-Dalgarno pattern for ribosome binding regions



Hidden Markov Model Basics: Probability



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

joint probability of the output sequence $x^T = (x_1, x_2, x_3, \dots, x_T)$ of length T and the hidden state sequence $u^T = (u_1, u_2, u_3, \dots, u_T)$

$$p(u^T, x^T) = p_S(u_1) \prod_{t=2}^T p(u_t | u_{t-1}) \prod_{t=1}^T p_E(x_t | u_t)$$

probability of the output sequence \rightarrow marginalization:

$$p(x^T) = \sum_{u^T} p(u^T, x^T) = \sum_{u^T} p_S(u_1) \prod_{t=2}^T p(u_t | u_{t-1}) \prod_{t=1}^T p_E(x_t | u_t)$$

\sum_{u^T} sum over all possible sequences of length T of values $\{1, \dots, S\}$

(first order) Markov assumption \rightarrow recursively compute this sum

Hidden Markov Model Basics: Forward Pass



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

$x^t = (x_1, x_2, x_3, \dots, x_t)$ prefix sequence of x^T

$p(x^t, u_t)$ probability of the model producing x^t and being in u_t at end

$$p(x^t, u_t) = p(x_t | x^{t-1}, u_t) p(x^{t-1}, u_t) =$$

$$p_E(x_t | u_t) \sum_{u_{t-1}} p(x^{t-1}, u_t, u_{t-1}) =$$

$$p_E(x_t | u_t) \sum_{u_{t-1}} p(u_t | x^{t-1}, u_{t-1}) p(x^{t-1}, u_{t-1}) =$$

$$p_E(x_t | u_t) \sum_{u_{t-1}} p(u_t | u_{t-1}) p(x^{t-1}, u_{t-1})$$

Markov assumptions $p(x_t | x^{t-1}, u_t) = p_E(x_t | u_t)$
 $p(u_t | x^{t-1}, u_{t-1}) = p(u_t | u_{t-1})$

marginalization $p(x^{t-1}, u_t) = \sum_{u_{t-1}} p(x^{t-1}, u_t, u_{t-1})$

Cond. prob. $p(x^{t-1}, u_t, u_{t-1}) = p(u_t | x^{t-1}, u_{t-1}) p(x^{t-1}, u_{t-1})$

Hidden Markov Model Basics: Forward Pass



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

each recursion step: sum over all S u_{t-1}

for each value u_t

for each time step t

Complexity: $O(T S^2)$

recursion starts with $p(x^1, u_1) = p_S(u_1) p_E(x_1 | u_1)$

and the probability is $p(x^T) = \sum_{u_T} p(x^T, u_T)$

“forward pass” or “forward phase” to compute the probability of x^T

Hidden Markov Model Basics: Forward Pass



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

HMM Forward Pass

Given: sequence $x^T = (x_1, x_2, x_3, \dots, x_T)$, state values $u \in \{1, \dots, S\}$, start probabilities $p_S(u_1)$, transition probabilities $p(u_t | u_{t-1})$, and emission probabilities $p_E(x_t | u_t)$; Output: likelihood $p(x^T)$ and $p(x^t, u_t)$

BEGIN initialization

$$p(x^1, u_1) = p_S(u_1) p_E(x_1 | u_1)$$

END initialization

BEGIN Recursion

```
for (t = 2 ; t ≤ T ; t++) do
  for (a = 1 ; a ≤ S ; a++) do
```

$$p(x^t, u_t = a) = p_E(x_t | u_t = a) \sum_{u_{t-1}=1}^S p(u_t = a | u_{t-1}) p(x^{t-1}, u_{t-1})$$

```
  end for
```

```
end for
```

END Recursion

BEGIN Compute Likelihood

$$p(x^T) = \sum_{a=1}^S p(x^T, u_T = a)$$

END Compute Likelihood

Expectation Maximization for HMM: Baum-Welch Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Learning / parameter selection

Parameters: S start probabilities $p_S(u_1)$
 S^2 transitions probabilities $p(u_t | u_{t-1})$
 SP emission probabilities $p_E(x_t | u_t)$

training sequences: $\{x^i\}$, $1 \leq i \leq l$ -- not subsequence!!

parameter optimization: maximize likelihood

Expectation Maximization algorithm

EM algorithm defines a lower bound on the likelihood:

$$\mathcal{F}(Q, w) = \int_U Q(u | x) \ln p(x, u; w) du - \int_U Q(u | x) \ln Q(u | x) du$$

Expectation Maximization for HMM: Baum-Welch Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

$Q(\mathbf{u} \mid \mathbf{x})$: estimation for $p(\mathbf{u} \mid \mathbf{x}; \mathbf{w})$

discrete HMMs: \mathbf{u} is the sequence of hidden states

\mathbf{x} the sequence of output states

\mathbf{w} all probabilities (parameters) in the model

discrete: integral \rightarrow sum

estimation for the state sequence:

$$Q(\mathbf{u} \mid \mathbf{x}) = p(u_1 = a_1, u_2 = a_2, \dots, u_T = a_T \mid \mathbf{x}^T; \mathbf{w}^{\text{old}})$$

Expectation Maximization for HMM: Baum-Welch Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

$$\mathcal{F}(Q, \mathbf{w}) = \sum_{a_1=1}^S \dots \sum_{a_T=1}^S$$

$$p(u_1 = a_1, u_2 = a_2, \dots, u_T = a_T \mid x^T; \mathbf{w}^{\text{old}}) \ln p(x^T, u^T; \mathbf{w}) -$$

$$\sum_{a_1=1}^S \dots \sum_{a_T=1}^S p(u_1 = a_1, u_2 = a_2, \dots, u_T = a_T \mid x^T; \mathbf{w}^{\text{old}})$$

$$\ln p(u_1 = a_1, u_2 = a_2, \dots, u_T = a_T \mid x^T; \mathbf{w}^{\text{old}}) =$$

$$\sum_{a_1=1}^S \dots \sum_{a_T=1}^S p(u_1 = a_1, u_2 = a_2, \dots, u_T = a_T \mid x^T; \mathbf{w}^{\text{old}})$$

$$\ln p(x^T, u^T; \mathbf{w}) - c$$

c independent of \mathbf{w}

Expectation Maximization for HMM: Baum-Welch Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

$$\ln p(u^T, x^T; \mathbf{w}) =$$

$$\ln p_S(u_1) + \sum_{t=2}^T \ln p(u_t | u_{t-1}) + \sum_{t=1}^T \ln p_E(x_t | u_t)$$

a_t summed out

$$\mathcal{F}(Q, \mathbf{w}) = \sum_{a=1}^S p(u_1 = a | x^T; \mathbf{w}^{\text{old}}) \ln p_S(u_1 = a) +$$

$$\sum_{t=1}^T \sum_{a=1}^S p(u_t = a | x^T; \mathbf{w}^{\text{old}}) \ln p_E(x_t | u_t = a) +$$

$$\sum_{t=2}^T \sum_{a=1}^S \sum_{b=1}^S p(u_t = a, u_{t-1} = b | x^T; \mathbf{w}^{\text{old}}) \ln p(u_t = a | u_{t-1} = b) - c$$

Expectation Maximization for HMM: Baum-Welch Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

parameters w :

start probabilities $p_S(a)$

emission probabilities $p_E(x | a)$

transition probabilities $p(a | b)$

Constraints:

$$\sum_a p_S(a) = 1$$

$$\sum_x p_E(x | a) = 1$$

$$\sum_a p(a | b) = 1$$

Now: $w = w^{\text{old}}$

Expectation Maximization for HMM: Baum-Welch Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

M-step optimization problem

$$\min_{\mathbf{w}} \quad \sum_t \sum_a c_{ta} \ln w_a$$
$$\text{s.t.} \quad \sum_a w_a = 1$$

Lagrangian

$$L = \sum_t \sum_a c_{ta} \ln w_a - \lambda \left(\sum_a w_a - 1 \right)$$

$$\frac{\partial L}{\partial w_a} = \sum_t c_{ta} \frac{1}{w_a} - \lambda = 0 \quad \Rightarrow \quad \sum_t c_{ta} - \lambda w_a = 0$$

summing over a gives

$$\sum_a \sum_t c_{ta} = \lambda$$

Expectation Maximization for HMM: Baum-Welch Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

We obtain $w_a = \frac{\sum_t c_{ta}}{\sum_a \sum_t c_{ta}}$

(constraint) **maximization step (M-step)** for the different probabilities:

$$p_S(a) = \frac{p(u_1 = a \mid x^T; \mathbf{w})}{\sum_{a'} p(u_1 = a' \mid x^T; \mathbf{w})}$$

$$p_E(x \mid a) = \frac{\sum_{t=1}^T \delta_{x_t=x} p(u_t = a \mid x^T; \mathbf{w})}{\sum_y \sum_{t=1}^T \delta_{x_t=y} p(u_t = a \mid x^T; \mathbf{w})}$$

$$p(a \mid b) = \frac{\sum_{t=2}^T p(u_t = a, u_{t-1} = b \mid x^T; \mathbf{w})}{\sum_{a'} \sum_{t=2}^T p(u_t = a', u_{t-1} = b \mid x^T; \mathbf{w})}$$

Expectation Maximization for HMM: Baum-Welch Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

$$p_S(a) = p(u_1 = a \mid x^T; \mathbf{w})$$

$$p_E(x \mid a) = \frac{\sum_{t=1}^T \delta_{x_t=x} p(u_t = a \mid x^T; \mathbf{w})}{\sum_{t=1}^T p(u_t = a \mid x^T; \mathbf{w})}$$

$$p(a \mid b) = \frac{\sum_{t=2}^T p(u_t = a, u_{t-1} = b \mid x^T; \mathbf{w})}{\sum_{t=2}^T p(u_{t-1} = b \mid x^T; \mathbf{w})}$$

estimation step (E-step):

Estimate $p(u_t = a \mid x^T; \mathbf{w})$ and $p(u_t = a, u_{t-1} = b \mid x^T; \mathbf{w})$

suffix sequence $x^{t \leftarrow T} = (x_t, x_{t+1}, \dots, x_T)$ of length $T - t + 1$

probability $p(x^{t+1 \leftarrow T} \mid u_t = a)$ of suffix sequence

$x^{t+1 \leftarrow T} = (x_{t+1}, \dots, x_T)$ if starting from $u_t = a$

Expectation Maximization for HMM: Baum-Welch Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

$$p(u_t = a \mid x^T; \mathbf{w}) = \frac{p(u_t = a, x^T; \mathbf{w})}{p(x^T)} =$$

$$\frac{p(x^t, u_t = a; \mathbf{w}) p(x^{t+1 \leftarrow T} \mid u_t = a)}{p(x^T)}$$

$$p(u_t = a, u_{t-1} = b \mid x^T; \mathbf{w}) = \frac{p(u_t = a, u_{t-1} = b, x^T; \mathbf{w})}{p(x^T)} =$$

$$p(x^{t-1}, u_{t-1} = b; \mathbf{w}) p(u_t = a \mid u_{t-1} = b) p_E(x_t \mid u_t = a) \\ p(x^{t+1 \leftarrow T} \mid u_t = a) / p(x^T)$$

$$p(u_t = a \mid x^T; \mathbf{w}) = \sum_{b=1}^S p(u_t = a, u_{t-1} = b \mid x^T; \mathbf{w})$$

Expectation Maximization for HMM: Baum-Welch Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

backward recursion for computing $p(x^{t+1\leftarrow T} \mid u_t = a)$

$$p(x^{t+1\leftarrow T} \mid u_t = a) = \sum_{b=1}^S p_E(x_{t+1} \mid u_{t+1} = b) p(u_{t+1} = b \mid u_t = a) p(x^{t+2\leftarrow T} \mid u_{t+1} = b)$$

starting conditions

$$p(x^{T\leftarrow T} \mid u_{T-1} = a) = \sum_{b=1}^S p_E(x_T \mid u_T = b) p(u_T = b \mid u_{T-1} = a)$$

or, alternatively

$$\forall_a : p(x^{T+1\leftarrow T} \mid u_T = a) = 1$$

Expectation Maximization for HMM: Baum-Welch Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

HMM Backward Pass

Given: sequence $x^T = (x_1, x_2, x_3, \dots, x_T)$, state values $u \in \{1, \dots, S\}$, start probabilities $p_S(u_1)$, transition probabilities $p(u_t | u_{t-1})$, and emission probabilities $p_E(x_t | u_t)$; Output: likelihood $p(x^T)$ and $p(x^{t+1 \leftarrow T} | u_t = a)$

BEGIN initialization

$$\forall_a : p(x^{T+1 \leftarrow T} | u_T = a) = 1$$

END initialization

BEGIN Recursion

```
for (t = T - 1 ; t ≥ 1 ; t --) do
  for (a = 1 ; a ≤ S ; a ++ ) do
```

$$p(x^{t+1 \leftarrow T} | u_t = a) = \sum_{b=1}^S p_E(x_{t+1} | u_{t+1} = b) p(u_{t+1} = b | u_t = a) p(x^{t+2 \leftarrow T} | u_{t+1} = b) .$$

```
  end for
```

```
end for
```

END Recursion

BEGIN Compute Likelihood

$$p(x^T) = \sum_{a=1}^S p_S(u_1 = a) (p(x^{1 \leftarrow T} | u_1 = a))$$

END Compute Likelihood

Expectation Maximization for HMM: Baum-Welch Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

HMM EM Algorithm

Given: l training sequences $(x^T)^i = (x_1^i, x_2^i, x_3^i, \dots, x_T^i)$
for $1 \leq i \leq l$, state values $u \in \{1, \dots, S\}$,
start probabilities $p_S(u_1)$,
transition probabilities $p(u_t | u_{t-1})$, and
emission probabilities $p_E(x | u)$;
Output: updated values of $p_S(u)$, $p_E(x | u)$,
and $p(u_t | u_{t-1})$

BEGIN initialization

initialize start probabilities $p_S(u_1)$,
transition probabilities $p(u_t | u_{t-1})$, and
emission probabilities $p_E(x | u)$;
Output: updated values of $p_S(u)$, $p_E(x | u)$,
and $p(u_t | u_{t-1})$

END initialization

```
Stop=false  
while Stop=false do  
  for ( $i = 1$  ;  $i \leq l$  ;  $i++$ ) do
```

Forward Pass

```
    forward pass for  $(x^T)^i$ 
```

Backward Pass

```
    backward pass for  $(x^T)^i$ 
```

EM-STEP

```
  end for  
  if stop criterion fulfilled then  
    Stop=true  
  end if  
end while
```

Expectation Maximization for HMM: Baum-Welch Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

HMM EM Algorithm

E-Step

for ($a = 1 ; a \leq S ; a ++$) do
 for ($b = 1 ; b \leq S ; b ++$) do

$$p(u_t = a, u_{t-1} = b \mid (x^T)^i; \mathbf{w}) = p((x^{t-1})^i, u_{t-1} = b; \mathbf{w}) p(u_t = a \mid u_{t-1} = b) p_E(x_t^i \mid u_t = a) p((x^{t+1 \leftarrow T})^i \mid u_t = a) / p((x^T)^i)$$

 end for
end for
for ($a = 1 ; a \leq S ; a ++$) do

$$p(u_t = a \mid (x^T)^i; \mathbf{w}) = \sum_{b=1}^S p(u_t = a, u_{t-1} = b \mid (x^T)^i; \mathbf{w})$$

end for

HMM EM Algorithm

M-Step

for ($a = 1 ; a \leq S ; a ++$) do

$$p_S(a) = p(u_1 = a \mid (x^T)^i; \mathbf{w})$$

end for
for ($a = 1 ; a \leq S ; a ++$) do
 for ($x = 1 ; x \leq P ; x ++$) do

$$p_E(x \mid a) = \frac{\sum_{t=1}^T \delta_{x_t^i=x} p(u_t = a \mid (x^T)^i; \mathbf{w})}{\sum_{t=1}^T p(u_t = a \mid (x^T)^i; \mathbf{w})}$$

 end for
end for
for ($a = 1 ; a \leq S ; a ++$) do
 for ($b = 1 ; b \leq S ; b ++$) do

$$p(a \mid b) = \frac{\sum_{t=2}^T p(u_t = a, u_{t-1} = b \mid (x^T)^i; \mathbf{w})}{\sum_{t=2}^T p(u_{t-1} = b \mid (x^T)^i; \mathbf{w})}$$

 end for
end for

Viterby Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

likelihood of x^T is an integral - more exactly a sum - over all probabilities of possible sequences of hidden states multiplied by the probability that the hidden sequence emits x^T

Often a specific hidden sequence dominates sum:

$$(u^T)^* = \arg \max_{u^T} p(u^T | x^T) = \arg \max_{u^T} p(u^T, x^T)$$


semantic meaning: hidden states are interpretable what is $(u^T)^*$

Bioinformatics: $(u^T)^*$ needed for alignment sequence \leftrightarrow HMM

Alignment \rightarrow obtained through dynamic programming
path in from left to right in figure

Viterby Algorithm

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine



circles: represented by a matrix V

$V_{t,a}$: maximal probability of a sequence of length t ending in state a

$$V_{t,a} = \max_{u^{t-1}} p(x^t, u^{t-1}, u_t = a)$$

Viterby Algorithm



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

Markov conditions \rightarrow recursive formulae

$$V_{t,a} = p_E(x_t | u_t = a) \max_b p(u_t = a | u_{t-1} = b) V_{t-1,b}$$

Initialization

$$V_{1,a} = p_S(a) p_E(x_1 | u_1 = a)$$

Result

$$\max_{u^T} p(u^T, x^T) = \max_a V_{T,a} .$$

sequence by backtracing

$$b(t, a) = \arg \max_b p(u_t = a | u_{t-1} = b) V_{t-1,b}$$

complexity of **Viterby algorithm**: $O(T S^2)$

(maximum over S terms)

Viterby Algorithm



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

HMM Viterby

Given: sequence $x^T = (x_1, x_2, x_3, \dots, x_T)$, state values $u \in \{1, \dots, S\}$, start probabilities $p_S(u_1)$, transition probabilities $p(u_t | u_{t-1})$, and emission probabilities $p_E(x_t | u_t)$; Output: most likely sequence of hidden state values $(u^T)^*$ and its probability $p(x^T, (u^T)^*)$

BEGIN initialization

$$V_{1,a} = p_S(a)p_E(x_1 | u_1 = a)$$

END initialization

BEGIN Recursion

```
for (t = 2 ; t ≤ T ; t++) do
  for (a = 1 ; a ≤ S ; a++) do
```

$$V_{t,a} = p_E(x_t | u_t = a) \max_b p(u_t = a | u_{t-1} = b) V_{t-1,b}$$

$$b(t, a) = \arg \max_b p(u_t = a | u_{t-1} = b) V_{t-1,b}$$

```
  end for
```

```
end for
```

END Recursion

BEGIN Compute Probability

$$p(x^T, (u^T)^*) = \max_{a=1}^S V(T, a)$$

END Compute Probability

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

HMM Viterby-Backtracing

BEGIN Backtracing

$$s = \arg \max_{a=1}^S V(T, a)$$

print s

for ($t = T$; $t \geq 2$; $t --$) **do**

$$s = b(t, s)$$

print s

end for

END Backtracing

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Viterby algorithm to improve a multiple alignment:


1. initialize HMM
2. align all sequences to the HMM via the Viterby algorithm
3. make frequency counts per column and compute the transition probabilities to update the HMM
4. if not converged go to step 2

Input Output Hidden Markov Models

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Input Output Hidden Markov Models (IOHMMs)

Output sequence $x^T = (x_1, x_2, x_3, \dots, x_T)$ conditioned on input sequence $y^T = (y_1, y_2, y_3, \dots, y_T)$



Input Output Hidden Markov Models



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

probabilities are conditioned on the input:

$$p_S(u_1 | y_1) \quad p(u_t | y_t, u_{t-1}) \quad p_E(x_t | y_t, u_t)$$

IOHM: classification including negative examples

$$y_T = 1 \quad y_T = -1$$

Subclasses of the negative class which is very similar to the positive class can be better discriminated


parameters increase proportional to input symbols

Learning analog to original method

Factorial Hidden Markov Models

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

hidden state is divided into more components



Factorial Hidden Markov Models



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

The transition probability of hidden state variable u_i is conditioned on all u_k with $k \leq i$

emission probability depends on all hidden states

Idea: u_1 evolves very slowly, u_2 evolves faster, etc
state variable depends on all slower ones

factorial HMMs is computationally expensive to learn

approximative methods have been developed to speed up learning

Memory Input Output Factorial Hidden Markov Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Sean Eddy

“HMMs are reasonable models of linear sequence problems, but they don't deal well with correlations between residues or states, especially long-range correlations.”

“The state path of an HMM has no way of remembering what a distant state generated when a second state generates its residue.”

only way an HMM can store information:

→go into a certain state value and don't change it

BUT: event which led the HMM enter the fixed state is memorized

Memory Input Output Factorial Hidden Markov Models



- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

$$p(u_t = a \mid u_{t-1} = a) = 1$$

if the state takes on the value a then the state is fixed forever

memory is enforced by fixing $p(u_t = a \mid u_{t-1} = a) = 1$

However: after the storage process systems dynamics nor other events to memorize can be dealt with

→ factorial HMM

Storing events → input output HMMs


Memory-based Input-Output Factorial HMM

Initially: all state variables have “uncommitted” values
then various inputs can trigger the memory state

Memory Input Output Factorial Hidden Markov Models

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

number of updates required to train three models:
input output HMM (IOHMM),
input output factorial HMM (IOFHMM),
and “Memory-based Input-Output Factorial HMM” (MIOFHMM)




- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

1. Bioinformatics: delete states which do not emit symbols
2. variable length of the sequences
carefully compare likelihoods → length dependent
3. small likelihood → log-space
4. kept probabilities above a threshold ϵ to allow all solutions
5. EM-algorithm: probabilities > 0 → set to zero after learning
→ helps to generalize
6. prone to local minima → global optimization strategies
(deterministic annealing)

Profile Hidden Markov Models

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Profile Hidden Markov Models code a multiple sequence alignment position-specific scoring system → search databases for homologous



1. The top row with states indicated with circles are a pattern.
2. The diamond states are inserted strings.
3. The bottom row with states indicated as squares are deletions, where a letter from the pattern is skipped.

Profile Hidden Markov Models


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

HMMER package by Sean Eddy



1. squares "Mx": consensus string
2. circles "Dx": are deletion states (non-emitting states)
3. diamonds "Ix": insertion states

Profile Hidden Markov Models



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

package SAM

Sequence Alignment and Modeling system (SAM -
<http://www.cse.ucsc.edu/research/compbio/sam.html>)

creating, refining, and using HMMs

SAM models: multiple alignment as HMMER

Databases: Protein FAMily database (Pfam)

67% proteins contain at least one Pfam profile HMM

Profile Hidden Markov Models


7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

HMM for splice site detection



Chapter 10

The Boltzmann Machine

Boltzmann Machines

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine


Boltzmann machine: stochastic recurrent neural network proposed by Geoffrey Hinton and Terry Sejnowski in 1985.

Boltzmann machines: not useful for practical problems

restricted Boltzmann machines: gained high popularity in the context of **deep learning**

name from statistical mechanics (normalizing distribution and used for sampling)

A graphical representation of an example Boltzmann machine. Each undirected edge represents dependency. In this example there are 3 **hidden units** and 4 **visible units**



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

Boltzmann machine: network of units with an “energy” of its current state; **binary stochastic units, global energy**

$$E = - \left(\sum_{i,j;i < j} w_{ij} s_i s_j + \sum_i \theta_i s_i \right)$$

- w_{ij} is the **connection strength** between unit j and unit i
- s_i is the **state** of unit i which is either 1 or 0
- θ_i is the **bias** or the activation threshold of unit i
- $w_{ii} = 0$ for all i , that is, units **do not have self-connections**
- $w_{ij} = w_{ji}$ for all i, j , that is, connections are **symmetric**

Weights are written by a symmetric matrix \mathbf{W} with zero diagonal

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

energy difference of a single unit:
$$\Delta E_i = \sum_j w_{ij} s_j + \theta_i$$

Boltzmann distribution: energy of a state is proportional to the negative log probability of that state.

Boltzmann Factor: energy difference if one unit is flipped

$$\Delta E_i = -k_B T \ln(p_{i=\text{off}}) - (-k_B T \ln(p_{i=\text{on}}))$$

k_B : Boltzmann's constant is absorbed into the temperature T

energy difference at temperature T :

$$\begin{aligned} \frac{\Delta E_i}{T} &= \ln(p_{i=\text{on}}) - \ln(p_{i=\text{off}}) \\ \Leftrightarrow \frac{\Delta E_i}{T} &= \ln(p_{i=\text{on}}) - \ln(1 - p_{i=\text{on}}) \\ \Leftrightarrow \frac{\Delta E_i}{T} &= \ln\left(\frac{p_{i=\text{on}}}{1 - p_{i=\text{on}}}\right) \\ \Leftrightarrow -\frac{\Delta E_i}{T} &= \ln\left(\frac{1 - p_{i=\text{on}}}{p_{i=\text{on}}}\right) \\ \Leftrightarrow -\frac{\Delta E_i}{T} &= \ln\left(\frac{1}{p_{i=\text{on}}} - 1\right) \\ \Leftrightarrow \exp\left(-\frac{\Delta E_i}{T}\right) &= \frac{1}{p_{i=\text{on}}} - 1 \end{aligned}$$

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

→ probability that the i -th unit is on:

$$p_{i=\text{on}} = \frac{1}{1 + \exp\left(-\frac{\Delta E_i}{T}\right)}$$

→ logistic function is used as activation probability

The network **runs iteratively** choosing a unit and setting its state according to its probability from above.

When the machine is “**at thermal equilibrium**” the probability distribution of global states has converged.

starts with a high temperature which gradually decreases

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Learning in the Boltzmann Machine

Training: weights so that the global states with the **highest probabilities** will get the **lowest energies**

- “**visible**” units V : receive binary training input from the “environment”
- “**hidden**” units H

distribution over the training set: $P^+(V)$

distribution over converged global states and after marginalization over the hidden units: $P^-(V)$

Boltzmann Machine learning: $P^-(V) \approx P^+(V)$

Objective is the Kullback-Leibler divergence:

$$D_{\text{KL}}(P^+ \parallel P^-) = \sum_v P^+(v) \ln \left(\frac{P^+(v)}{P^-(v)} \right)$$

sum is over all the possible states of V

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Analog to the EM algorithm, Boltzmann machine training has two phases that are performed alternating:

- “**positive**” phase: visible units' states are clamped to the training data
- “**negative**” phase the network runs freely

gradient of the objective with respect to a given weight:

$$\frac{\partial D}{\partial w_{ij}} = -\frac{1}{\eta} (p_{ij}^+ - p_{ij}^-)$$

- p_{ij}^+ is the probability of units i and j both being on when the machine is at equilibrium on the **positive phase**
- p_{ij}^- is the probability of units i and j both being on when the machine is at equilibrium on the **negative phase**
- η is the learning rate

Learning by using only “local information” → biological plausible

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov
Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann
Machines
10.1 The Boltzmann
Machine
10.2 Learning
10.3 Restricted
Boltzmann Machine

training the bias weight:
$$\frac{\partial D}{\partial \theta_i} = -\frac{1}{\eta} (p_i^+ - p_i^-)$$

Problems:

- the time to equilibrium grows exponentially with the size and with the magnitude of the connection strengths
- probabilities between zero and one \rightarrow more plastic, variance trap net effect \rightarrow noise causes the connection strengths to random walk until the activities saturate

7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based


8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

Restricted Boltzmann Machine (RBM)

- no intralayer connections between hidden or visible units
- activities of its hidden units are a representation of the visible units
- **deep learning** idea: hidden units of one RBM are the visible units of a higher-level RBM → stacked RBMs



Graphical representation of a restricted Boltzmann machine. The four blue units represent hidden units, and the three red units represent visible states. In restricted Boltzmann machines there are only connections (dependencies) between hidden and visible units, and none between units of the same type.

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based

- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples

- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM

- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

$w_{i,j}$: connection between hidden unit h_j and visible unit v_i

bias weight visible unit: a_i

bias weight hidden unit: b_j

energy:

$$E(\mathbf{v}, \mathbf{h}) = - \sum_i a_i v_i - \sum_j b_j h_j - \sum_{i,j} h_j v_i w_{i,j}$$

$$E(\mathbf{v}, \mathbf{h}) = - \mathbf{a}^T \mathbf{v} - \mathbf{b}^T \mathbf{h} - \mathbf{h}^T \mathbf{W} \mathbf{v}$$

Probability distributions over hidden and visible vectors:

$p(\mathbf{v}, \mathbf{h}) = \frac{1}{Z} e^{-E(\mathbf{v}, \mathbf{h})}$, where Z is a normalizing constant or the **partition function** defined as the sum of $e^{-E(\mathbf{v}, \mathbf{h})}$ over all possible configurations.

Marginalizing \rightarrow probability of a visible (input) vector:

$$p(\mathbf{v}) = \frac{1}{Z} \sum_{\mathbf{h}} e^{-E(\mathbf{v}, \mathbf{h})}$$

- 7 Clustering
 - 7.1 Mixture Models
 - 7.2 k-Means
 - 7.3 Hierarchical
 - 7.4 Similarity-Based
- 8 Biclustering
 - 8.1 Types of Biclusters
 - 8.2 Overview Methods
 - 8.3 FABIA
 - 8.4 Examples
- 9 Hidden Markov Models
 - 9.1 HMMs in Bioinf.
 - 9.2 HMM Basics
 - 9.3 EM / Baum-Welch
 - 9.4 Viterby
 - 9.5 Input Output HMM
 - 9.6 Factorial HMM
 - 9.7 Memory IOHMM
 - 9.8 Tricks of the Trade
 - 9.9 Profile HMM
- 10 Boltzmann Machines
 - 10.1 The Boltzmann Machine
 - 10.2 Learning
 - 10.3 Restricted Boltzmann Machine

Mutual conditional independence:

$$p(\mathbf{v} | \mathbf{h}) = \prod_{i=1}^m p(v_i | \mathbf{h}) \quad p(h_j = 1 | \mathbf{v}) = \sigma(b_j + \sum_{i=1}^m w_{i,j} v_i)$$
$$p(\mathbf{h} | \mathbf{v}) = \prod_{j=1}^n p(h_j | \mathbf{v}) \quad p(v_i = 1 | \mathbf{h}) = \sigma(a_i + \sum_{j=1}^n w_{i,j} h_j)$$

where σ is the sigmoid function

Training $\arg \max_{\mathbf{w}} \prod_{\mathbf{v} \in V} p(\mathbf{v})$ $\arg \max_{\mathbf{w}} E \left(\sum_{\mathbf{v} \in V} \log p(\mathbf{v}) \right)$

Boltzmann Machines: RBM



7 Clustering
7.1 Mixture Models
7.2 k-Means
7.3 Hierarchical
7.4 Similarity-Based

8 Biclustering
8.1 Types of Biclusters
8.2 Overview Methods
8.3 FABIA
8.4 Examples

9 Hidden Markov Models
9.1 HMMs in Bioinf.
9.2 HMM Basics
9.3 EM / Baum-Welch
9.4 Viterby
9.5 Input Output HMM
9.6 Factorial HMM
9.7 Memory IOHMM
9.8 Tricks of the Trade
9.9 Profile HMM

10 Boltzmann Machines
10.1 The Boltzmann Machine
10.2 Learning
10.3 Restricted Boltzmann Machine

train RBMs by **contrastive divergence** (CD):

1. Take a training sample \mathbf{v} , compute the probabilities of the hidden units and sample a hidden activation vector \mathbf{h} from this probability distribution
2. Compute the outer product of \mathbf{v} and \mathbf{h} and call this the “positive gradient”
3. From \mathbf{h} , sample a reconstruction \mathbf{v}' of the visible units, then resample the hidden activations \mathbf{h}' from this
4. Compute the outer product of \mathbf{v}' and \mathbf{h}' and call this the “negative gradient”
5. Let the weight update to $w_{i,j}$ be the “positive gradient” minus the “negative gradient”, times some learning rate η