

Computer vision: models, learning and inference

Chapter 7

Modeling Complex Densities

Structure

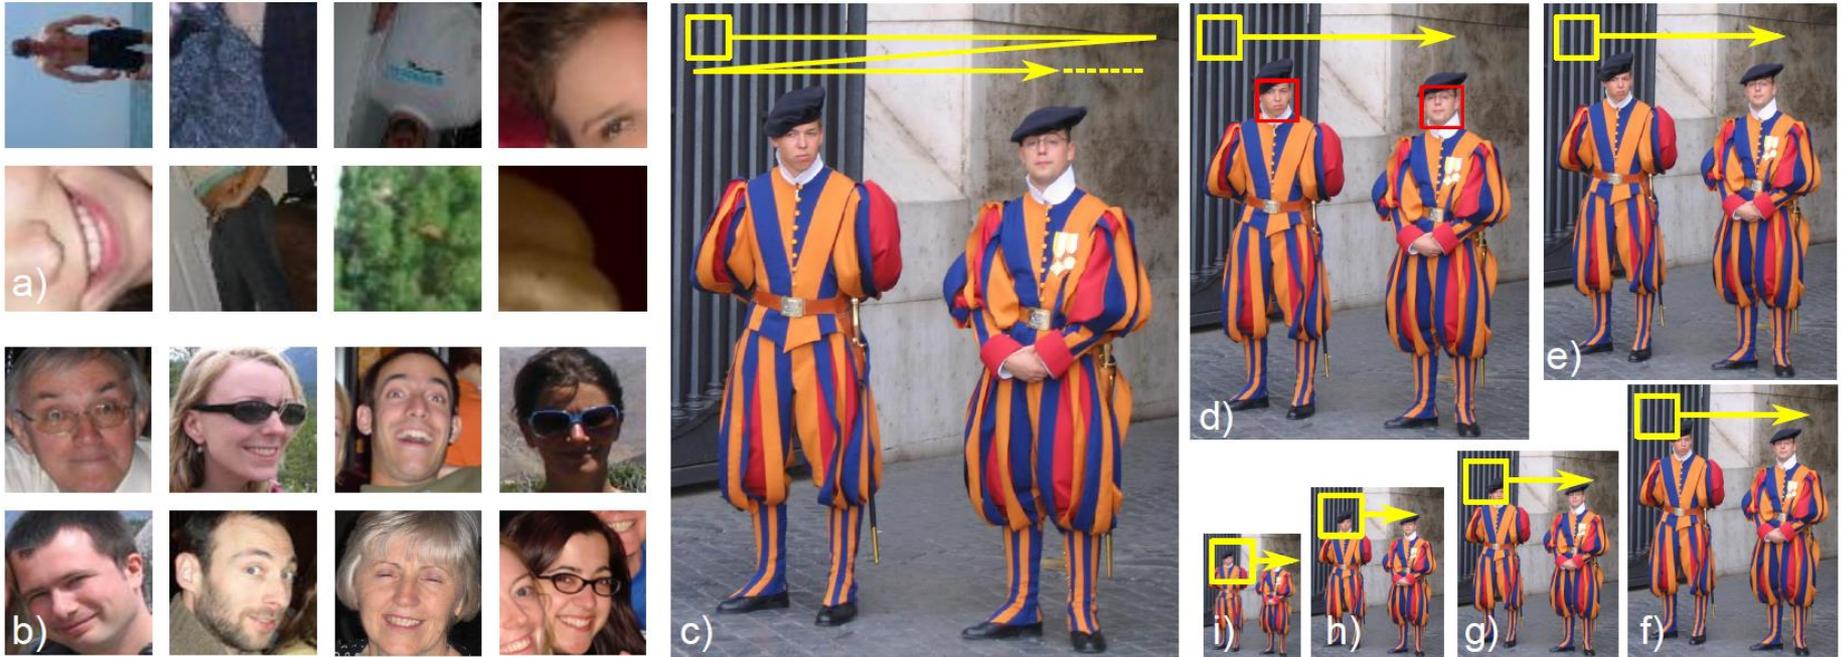
- Densities for classification
- Models with hidden variables
 - Mixture of Gaussians
 - t-distributions
 - Factor analysis
- EM algorithm in detail
- Applications

Models for machine vision

	Model $Pr(w x)$	Model $Pr(x w)$
Regression $x \in [-\infty, \infty], w \in [-\infty, \infty]$	Linear regression	Linear regression
Classification $x \in [-\infty, \infty], w \in \{0, 1\}$	Logistic regression	Probability density function

Table 5.1: Example models in this chapter. These can be categorized into those that are based on modelling probability density functions, those that are based on linear regression and those that are based on logistic regression.

Face Detection



Type 3: $\Pr(\mathbf{x} | \mathbf{w})$ - Generative

How to model $\Pr(\mathbf{x} | \mathbf{w})$?

- Choose an appropriate form for $\Pr(\mathbf{x})$
- Make parameters a function of \mathbf{w}
- Function takes parameters θ that define its shape

Learning algorithm: learn parameters θ from training data \mathbf{x}, \mathbf{w}

Inference algorithm: Define prior $\Pr(\mathbf{w})$ and then compute $\Pr(\mathbf{w} | \mathbf{x})$ using Bayes' rule

$$\Pr(w = 1 | \mathbf{x}) = \frac{\Pr(\mathbf{x} | w = 1) \Pr(w = 1)}{\sum_{k=0}^1 \Pr(\mathbf{x} | w = k) \Pr(w = k)}$$

Classification Model

$$Pr(\mathbf{x}|w) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w]$$

Or writing in terms of class conditional density functions

$$Pr(\mathbf{x}|w = 0) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0]$$

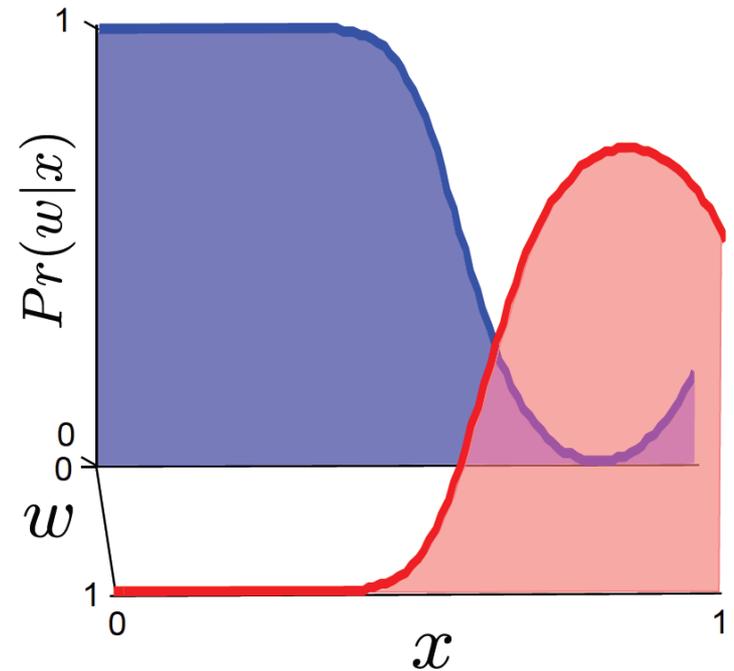
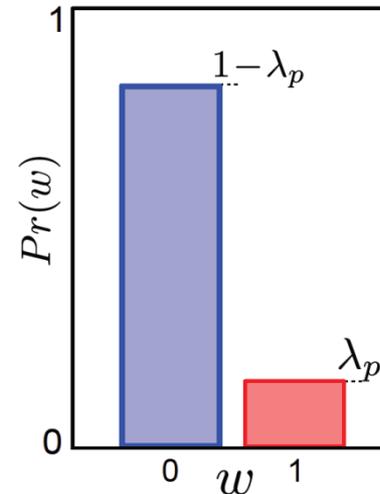
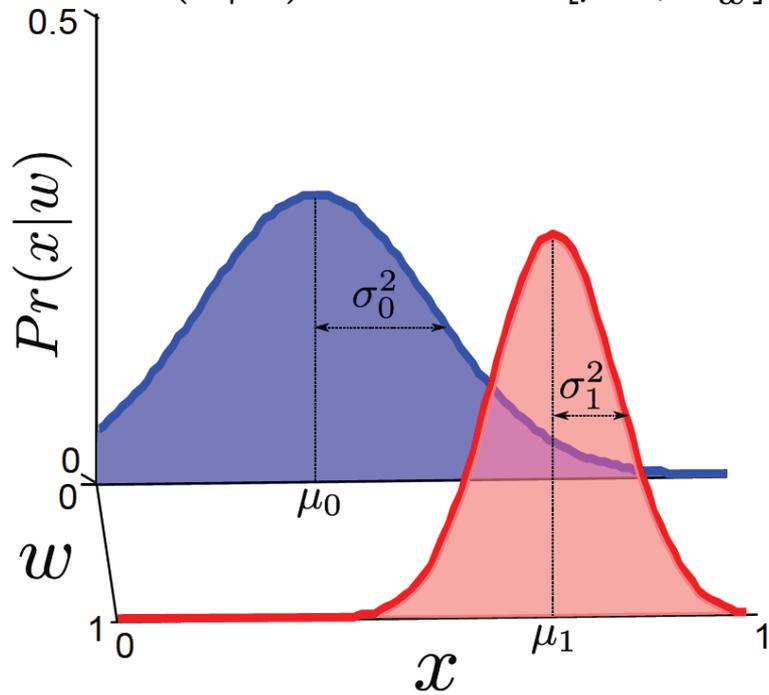
$$Pr(\mathbf{x}|w = 1) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1]$$

Parameters $\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0$ learnt just from data \mathcal{S}_0 where $w=0$

$$\begin{aligned}\hat{\boldsymbol{\mu}}_0, \hat{\boldsymbol{\Sigma}}_0 &= \underset{\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0}{\text{argmax}} \left[\prod_{i \in \mathcal{S}_0} Pr(\mathbf{x}_i | \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) \right] \\ &= \underset{\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0}{\text{argmax}} \left[\prod_{i \in \mathcal{S}_0} \text{Norm}_{\mathbf{x}_i}[\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0] \right]\end{aligned}$$

Similarly, parameters $\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1$ learnt just from data \mathcal{S}_1 where $w=1$

$$Pr(x|w) = \text{Norm}_x[\mu_w, \sigma_w^2]$$



Inference algorithm: Define prior $Pr(\mathbf{w})$ and then compute $Pr(\mathbf{w}|\mathbf{x})$ using Bayes' rule

$$Pr(\mathbf{w}|\mathbf{x}) = \frac{Pr(\mathbf{x}|\mathbf{w})Pr(\mathbf{w})}{\int Pr(\mathbf{x}|\mathbf{w})Pr(\mathbf{w})d\mathbf{w}}$$

Experiment

1000 non-faces

1000 faces

60x60x3 Images = 10800 x1 vectors

Equal priors $\Pr(y=1)=\Pr(y=0) = 0.5$

75% performance on test set. Not very good!



Results (diagonal covariance)

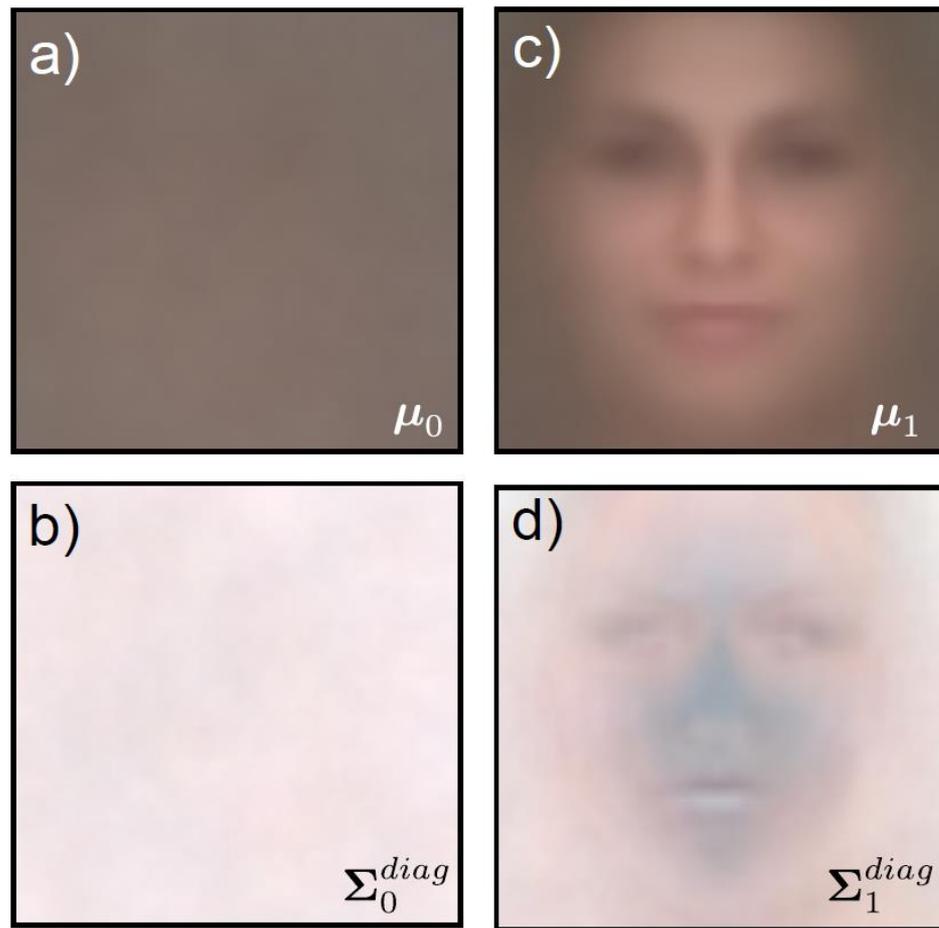
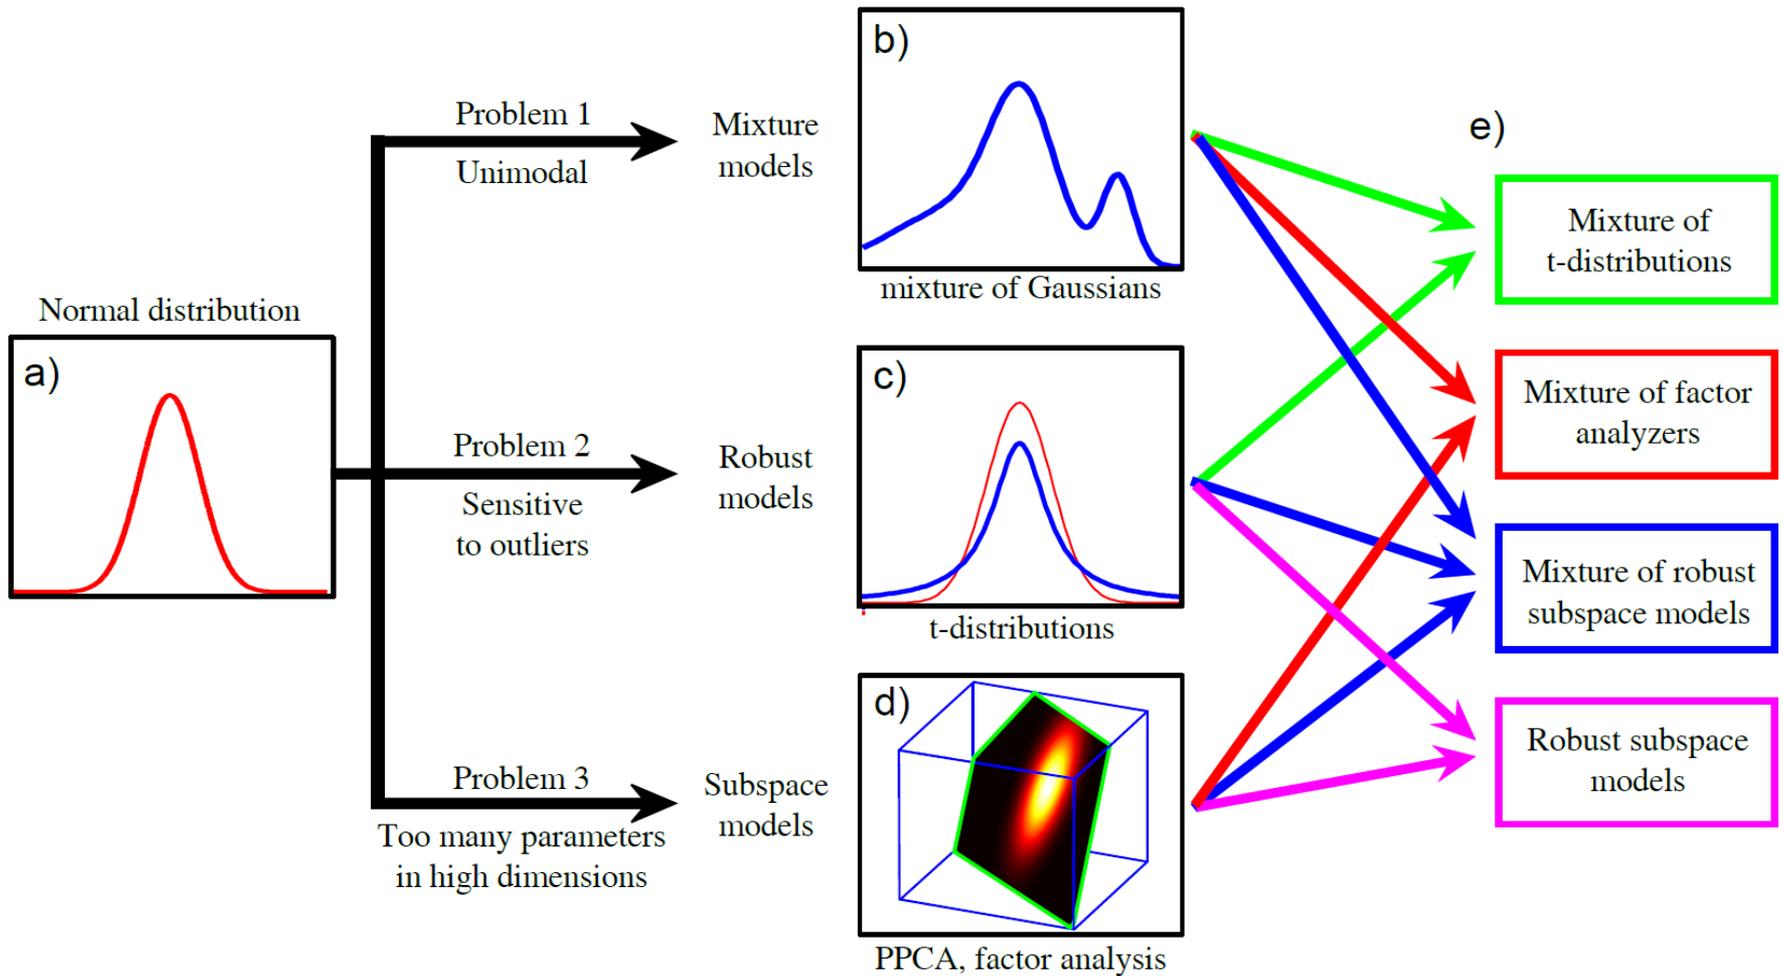


Figure 7.2 Class conditional density functions for normal model with diagonal covariance. Maximum likelihood fits based on 1000 training examples per class. a) Mean for background data μ_0 (reshaped from 10800×1 vector to 60×60 RGB image). b) Reshaped square root of diagonal covariance for background data Σ_0 . c) Mean for face data μ_1 d) Covariance for face data Σ_1 . The background model has little structure: the mean is uniform and the variance is high everywhere. The mean of the face model clearly captures class-specific information. The covariance of the face is larger at the edges of the image which usually contain hair or background.

The plan...



Structure

- Densities for classification
- **Models with hidden variables**
 - Mixture of Gaussians
 - t-distributions
 - Factor analysis
- EM algorithm in detail
- Applications

Hidden (or latent) Variables

Key idea: represent density $\Pr(\mathbf{x})$ as marginalization of joint density with another variable \mathbf{h} that we do not see

$$\Pr(\mathbf{x}) = \int \Pr(\mathbf{x}, \mathbf{h}) d\mathbf{h}$$

Will also depend on some parameters:

$$\Pr(\mathbf{x}|\boldsymbol{\theta}) = \int \Pr(\mathbf{x}, \mathbf{h}|\boldsymbol{\theta}) d\mathbf{h}.$$

Hidden (or latent) Variables

$$Pr(\mathbf{x}|\boldsymbol{\theta}) = \int Pr(\mathbf{x}, \mathbf{h}|\boldsymbol{\theta}) d\mathbf{h}.$$

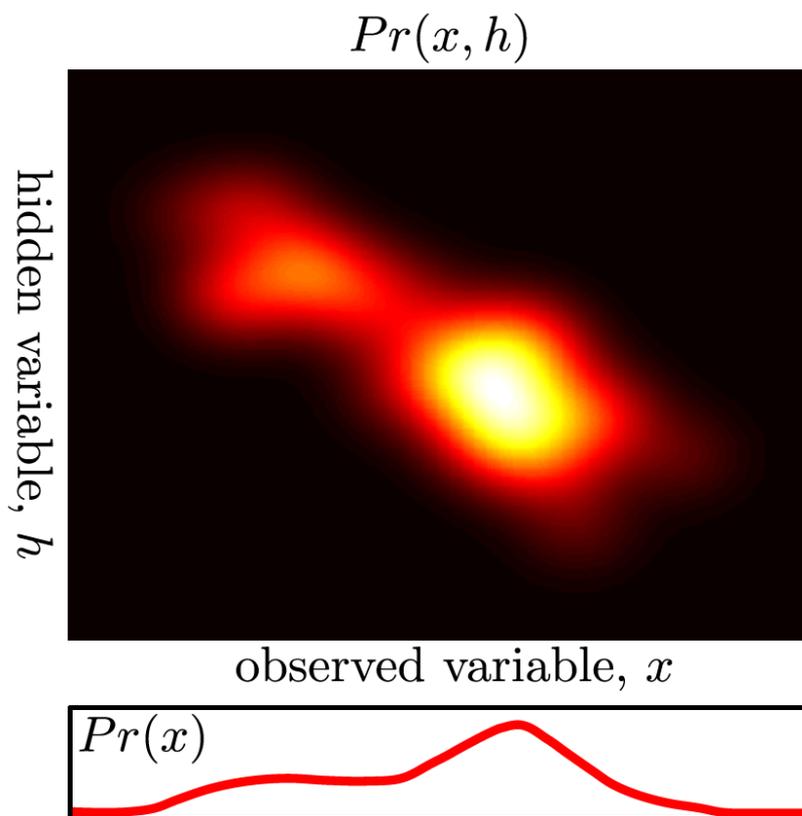


Figure 7.4 Using hidden variables to help model complex densities. One way to model the density $Pr(x)$ is to consider the joint probability distribution $Pr(x, h)$ between the observed data x and a hidden variable h . The density $Pr(x)$ can be considered as the marginalization of (integral over) this distribution with respect to the hidden variable h . As we manipulate the parameters $\boldsymbol{\theta}$ of this joint distribution, the marginal changes and the agreements with the observed data $\{x_i\}_{i=1}^I$ increases or decreases. Sometimes it is easier to fit the distribution in this indirect way than to directly manipulate $Pr(x)$.

Expectation Maximization

An algorithm specialized to fitting pdfs which are the marginalization of a joint distribution

$$\hat{\boldsymbol{\theta}} = \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \log \left[\int Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta}) d\mathbf{h}_i \right] \right]$$

Defines a lower bound on log likelihood and increases bound iteratively

$$\begin{aligned} \mathcal{B}[\{q_i(\mathbf{h}_i)\}, \boldsymbol{\theta}] &= \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_{1\dots I} \\ &\leq \sum_{i=1}^I \log \left[\int Pr(\mathbf{x}, \mathbf{h} | \boldsymbol{\theta}) d\mathbf{h}_{1\dots I} \right]. \end{aligned}$$

Lower bound

$$\mathcal{B} [\{q_i(\mathbf{h}_i)\}, \boldsymbol{\theta}] = \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_{1\dots I}$$

Lower bound is a *function* of parameters $\boldsymbol{\theta}$ and a set of probability distributions $q_i(\mathbf{h}_i)$

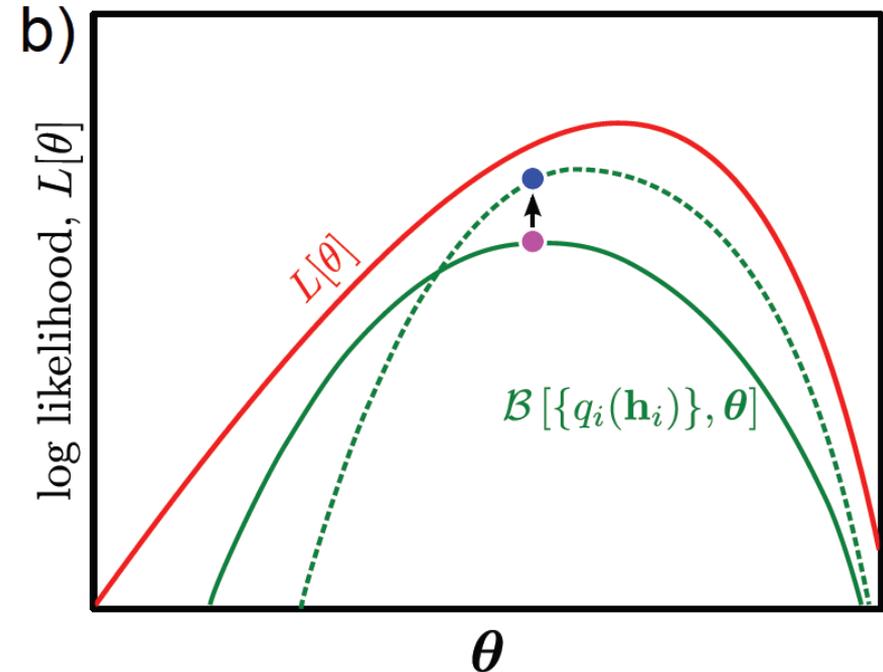
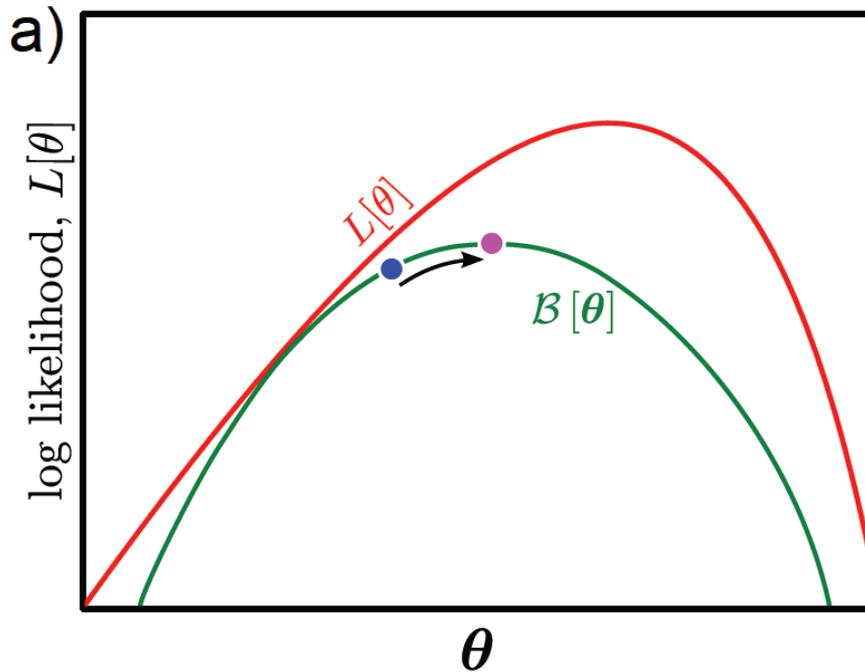
Expectation Maximization (EM) algorithm alternates E-steps and M-Steps

E-Step – Maximize bound w.r.t. distributions $q(\mathbf{h}_i)$

M-Step – Maximize bound w.r.t. parameters $\boldsymbol{\theta}$

Lower bound

$$\mathcal{B} [\{q_i(\mathbf{h}_i)\}, \boldsymbol{\theta}] = \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_{1\dots I}$$



E-Step & M-Step

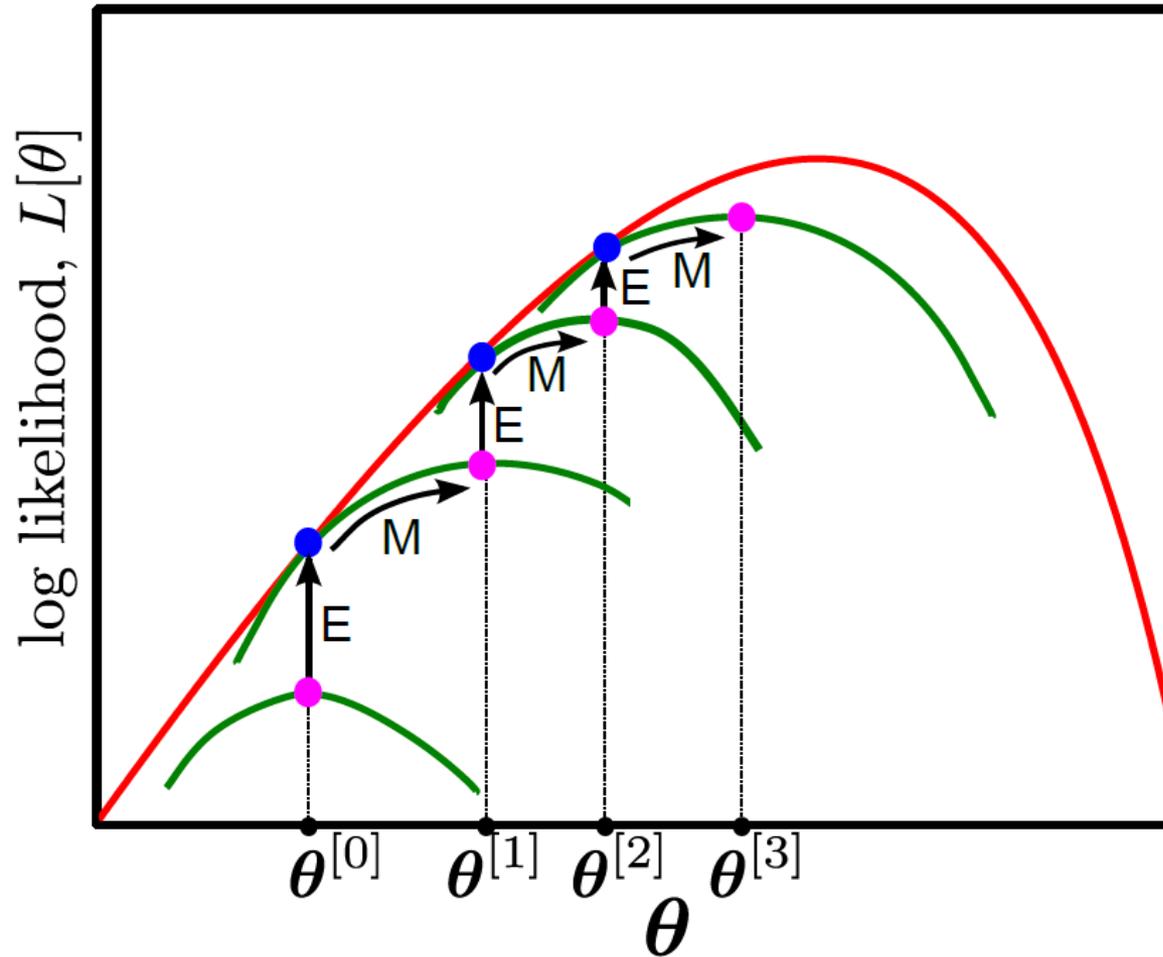
E-Step – Maximize bound w.r.t. distributions $q_i(\mathbf{h}_i)$

$$\hat{q}_i(\mathbf{h}_i) = Pr(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta}^{[t]}) = \frac{Pr(\mathbf{x}_i | \mathbf{h}_i, \boldsymbol{\theta}^{[t]}) Pr(\mathbf{h}_i | \boldsymbol{\theta}^{[t]})}{Pr(\mathbf{x}_i)}$$

M-Step – Maximize bound w.r.t. parameters $\boldsymbol{\theta}$

$$\hat{\boldsymbol{\theta}}^{[t+1]} = \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int \hat{q}_i(\mathbf{h}_i) \log [Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta})] d\mathbf{h}_i \right]$$

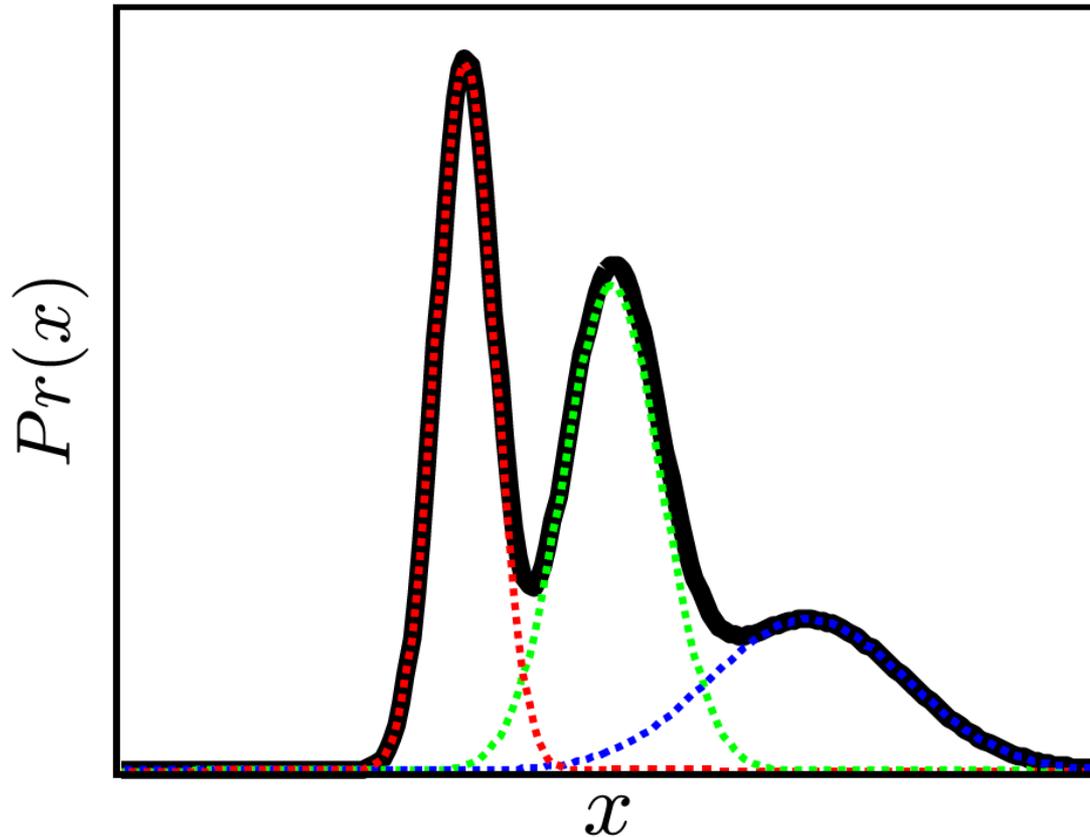
E-Step & M-Step



Structure

- Densities for classification
- Models with hidden variables
 - Mixture of Gaussians
 - t-distributions
 - Factor analysis
- EM algorithm in detail
- Applications

Mixture of Gaussians (MoG)



$$Pr(\mathbf{x}|\boldsymbol{\theta}) = \sum_{k=1}^K \lambda_k \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k]$$

ML Learning

Q. Can we learn this model with maximum likelihood?

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \log [Pr(\mathbf{x}_i | \boldsymbol{\theta})] \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \log \left[\sum_{k=1}^K \lambda_k \operatorname{Norm}_{\mathbf{x}_i} [\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k] \right] \right]\end{aligned}$$

- A. Yes, but using brute force approach is tricky
- If you take derivative and set to zero, can't solve
 - the log of the sum causes problems
 - Have to enforce constraints on parameters
 - covariances must be positive definite
 - weights must sum to one

MoG as a marginalization

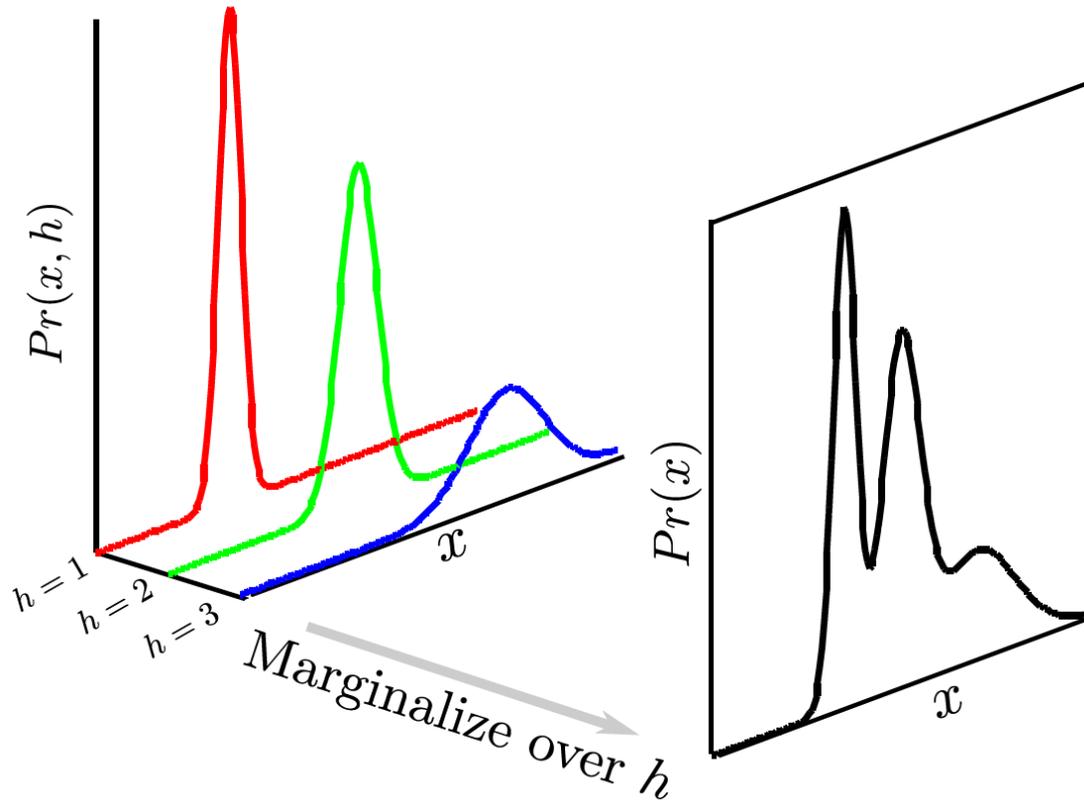
Define a variable $h \in \{1 \dots K\}$ and then write

$$\begin{aligned}Pr(\mathbf{x}|h, \boldsymbol{\theta}) &= \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h] \\Pr(h|\boldsymbol{\theta}) &= \text{Cat}_h[\boldsymbol{\lambda}]\end{aligned}$$

Then we can recover the density by marginalizing $\Pr(\mathbf{x}, h)$

$$\begin{aligned}Pr(\mathbf{x}|\boldsymbol{\theta}) &= \sum_{k=1}^K Pr(\mathbf{x}, h = k|\boldsymbol{\theta}) \\&= \sum_{k=1}^K Pr(\mathbf{x}|h = k, \boldsymbol{\theta})Pr(h = k|\boldsymbol{\theta}) \\&= \sum_{k=1}^K \lambda_k \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k].\end{aligned}$$

MoG as a marginalization



$$Pr(\mathbf{x}|\boldsymbol{\theta}) = \sum_{k=1}^K Pr(\mathbf{x}, h = k|\boldsymbol{\theta})$$

MoG as a marginalization

Define a variable $h \in \{1 \dots K\}$ and then write

$$Pr(\mathbf{x}|h, \boldsymbol{\theta}) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h]$$

$$Pr(h|\boldsymbol{\theta}) = \text{Cat}_h[\boldsymbol{\lambda}]$$

Note :

- This gives us a method to generate data from MoG
 - First sample $Pr(h)$, then sample $Pr(\mathbf{x}|h)$
- The hidden variable h has a clear interpretation – it tells you which Gaussian created data point \mathbf{x}

Expectation Maximization for MoG

GOAL: to learn parameters $\theta = \{\lambda_{1...K}, \mu_{1...K}, \Sigma_{1...K}\}$ from training data $\mathbf{x}_{1...I}$

E-Step – Maximize bound w.r.t. distributions $q(\mathbf{h}_i)$

$$\hat{q}_i(\mathbf{h}_i) = Pr(\mathbf{h}_i | \mathbf{x}_i, \theta^{[t]}) = \frac{Pr(\mathbf{x}_i | \mathbf{h}_i, \theta^{[t]}) Pr(\mathbf{h}_i | \theta^{[t]})}{Pr(\mathbf{x}_i)}$$

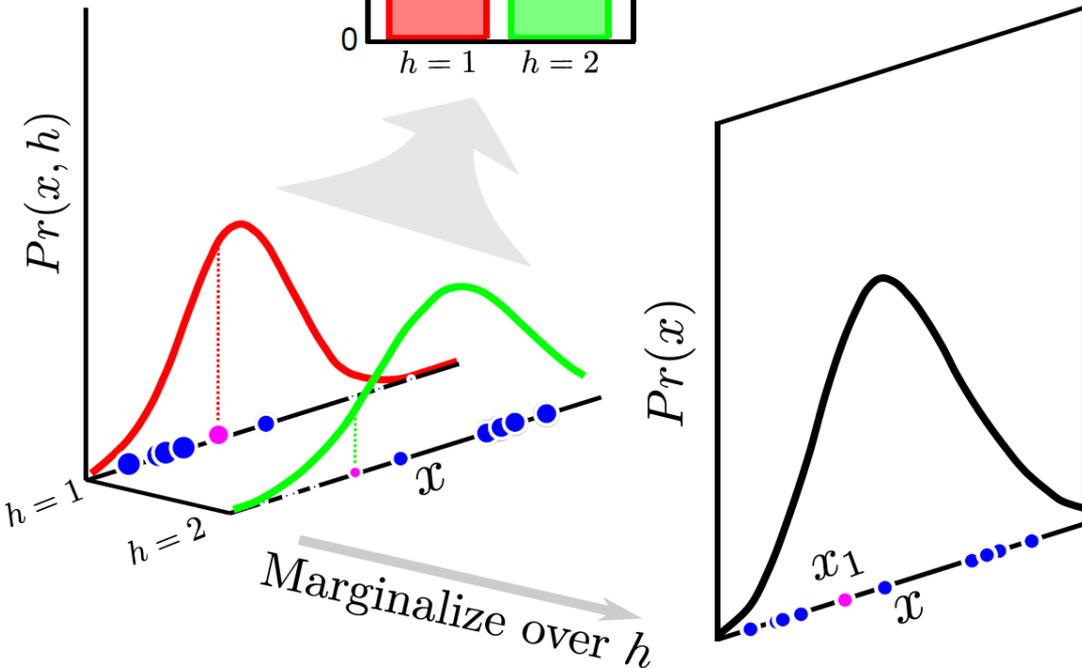
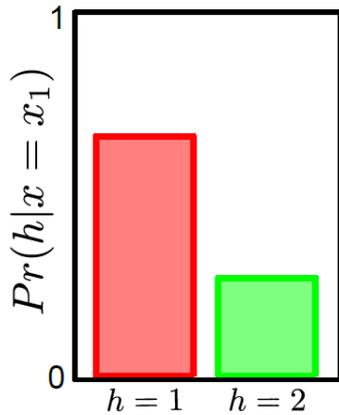
M-Step – Maximize bound w.r.t. parameters θ

$$\hat{\theta}^{[t+1]} = \operatorname{argmax}_{\theta} \left[\sum_{i=1}^I \sum_{k=1}^K \hat{q}_i(\mathbf{h}_i = k) \log [Pr(\mathbf{x}_i, \mathbf{h}_i = k | \theta)] \right]$$

E-Step

$$Pr(h_i = k | \mathbf{x}_i, \boldsymbol{\theta}^{[t]})$$

$$\begin{aligned}
 &= \frac{Pr(\mathbf{x}_i | h_i = k, \boldsymbol{\theta}^{[t]}) Pr(h_i = k, \boldsymbol{\theta}^{[t]})}{\sum_{j=1}^K Pr(\mathbf{x}_i | h_i = j, \boldsymbol{\theta}^{[t]}) Pr(h_i = j, \boldsymbol{\theta}^{[t]})} \\
 &= \frac{\lambda_k \text{Norm}_{\mathbf{x}_i} [\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k]}{\sum_{j=1}^K \lambda_j \text{Norm}_{\mathbf{x}_i} [\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j]} \\
 &= r_{ik}
 \end{aligned}$$



We'll call this the responsibility of the k^{th} Gaussian for the i^{th} data point

Repeat this procedure for every datapoint!

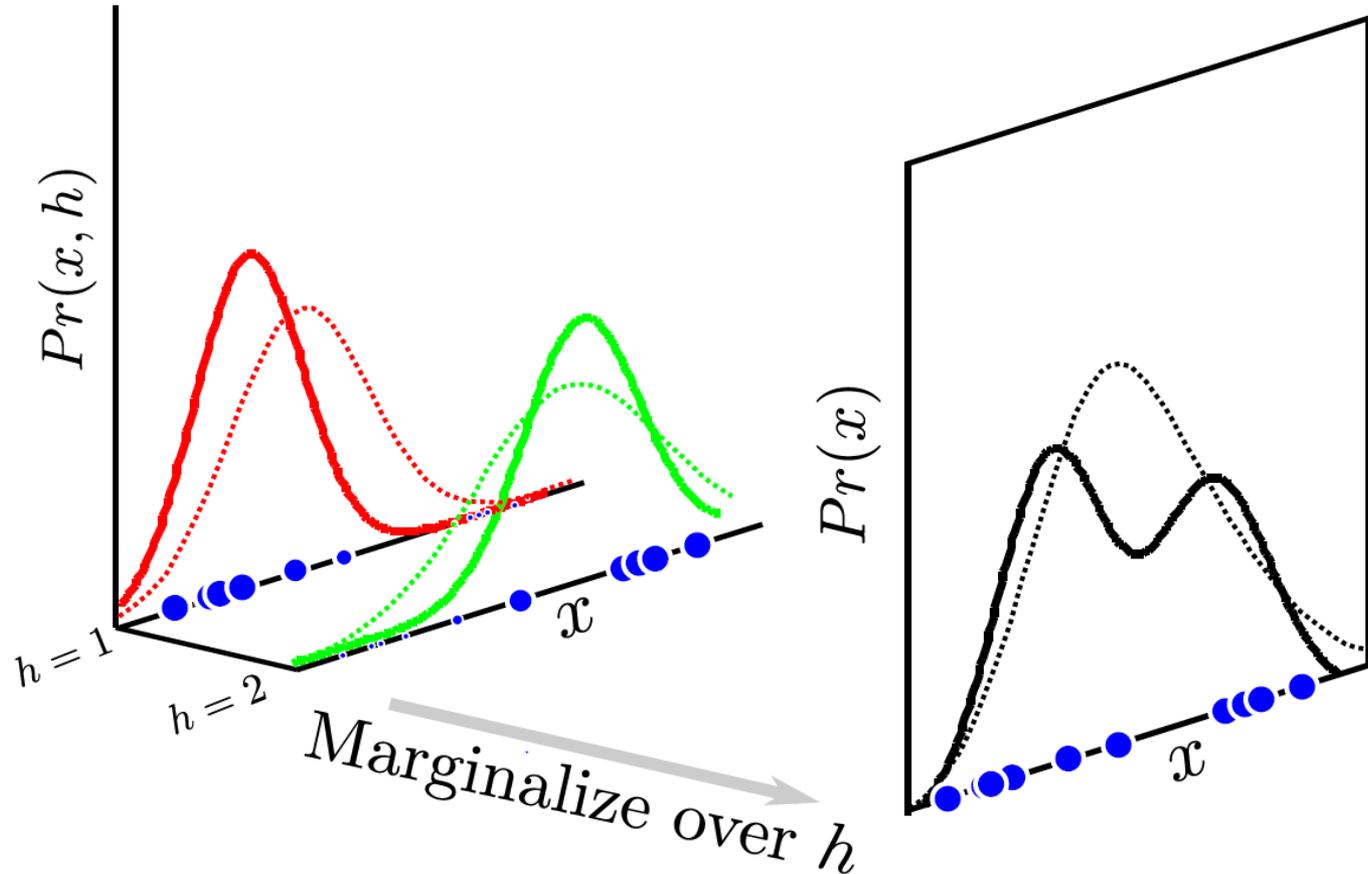
M-Step

$$\begin{aligned}\hat{\boldsymbol{\theta}}^{[t+1]} &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \sum_{k=1}^K \hat{q}_i(h_i = k) \log [Pr(\mathbf{x}_i, h_i = k | \boldsymbol{\theta})] \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \sum_{k=1}^K r_{ik} \log [\lambda_k \operatorname{Norm}_{\mathbf{x}_i} [\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k]] \right].\end{aligned}$$

Take derivative, equate to zero and solve (Lagrange multipliers for λ)

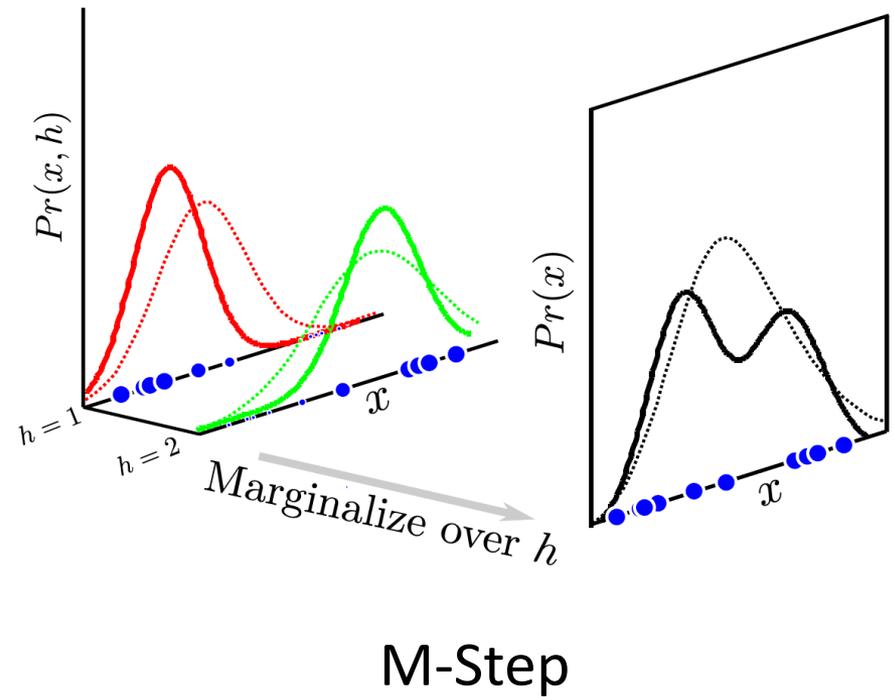
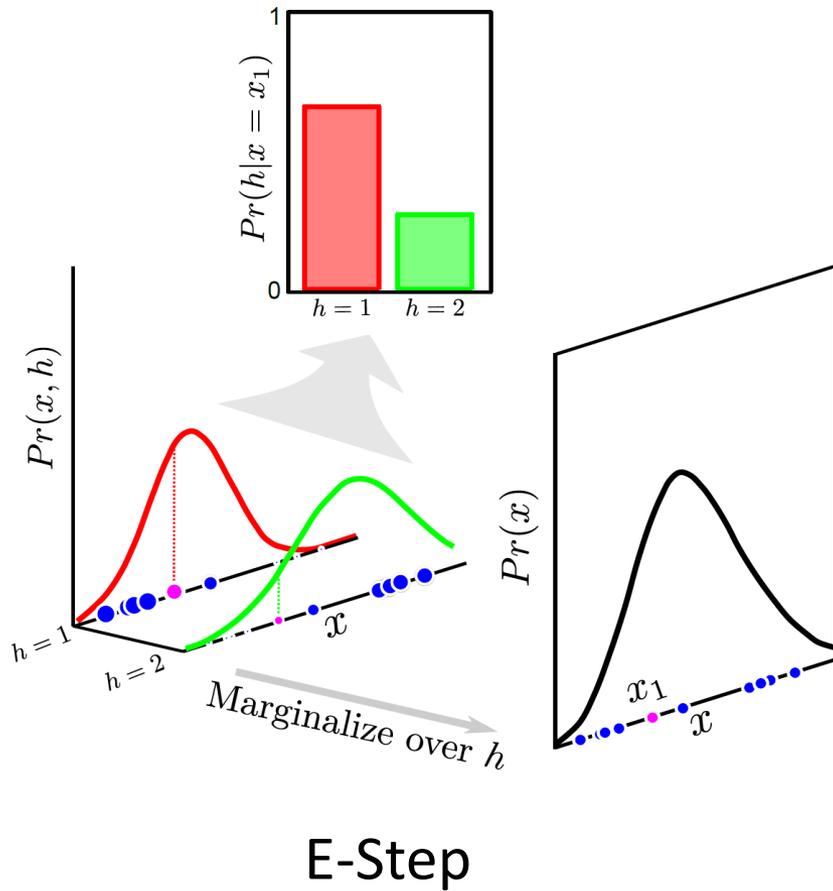
$$\begin{aligned}\lambda_k^{[t+1]} &= \frac{\sum_{i=1}^I r_{ik}}{\sum_{j=1}^K \sum_{i=1}^I r_{ij}} \\ \boldsymbol{\mu}_k^{[t+1]} &= \frac{\sum_{i=1}^I r_{ik} \mathbf{x}_i}{\sum_{i=1}^I r_{ik}} \\ \boldsymbol{\Sigma}_k^{[t+1]} &= \frac{\sum_{i=1}^I r_{ik} (\mathbf{x}_i - \boldsymbol{\mu}_k^{[t+1]})(\mathbf{x}_i - \boldsymbol{\mu}_k^{[t+1]})^T}{\sum_{i=1}^I r_{ik}}\end{aligned}$$

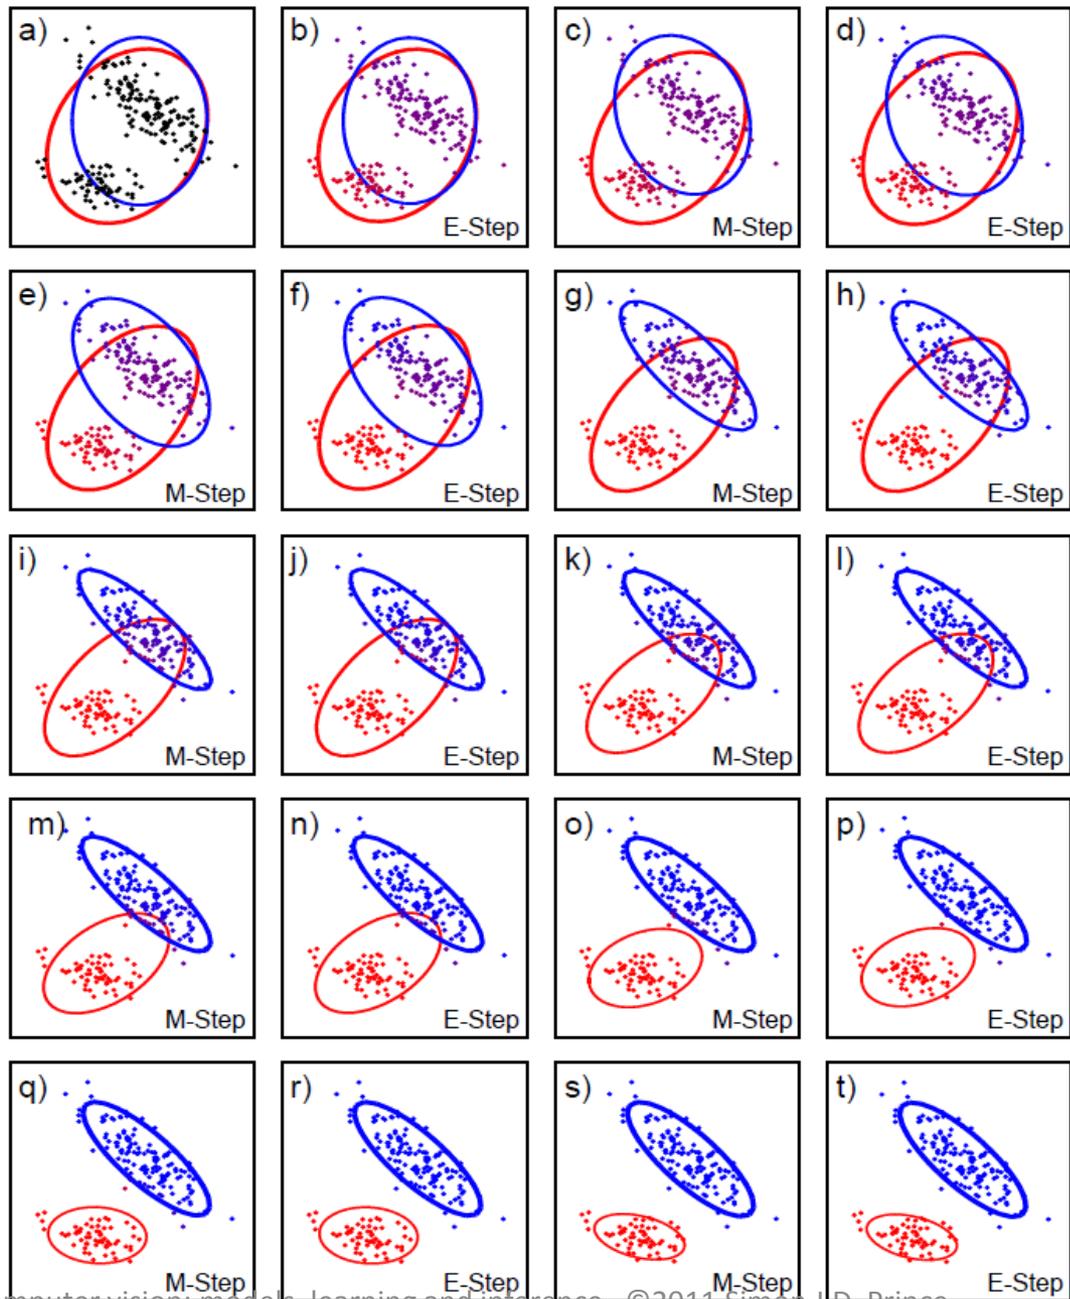
M-Step



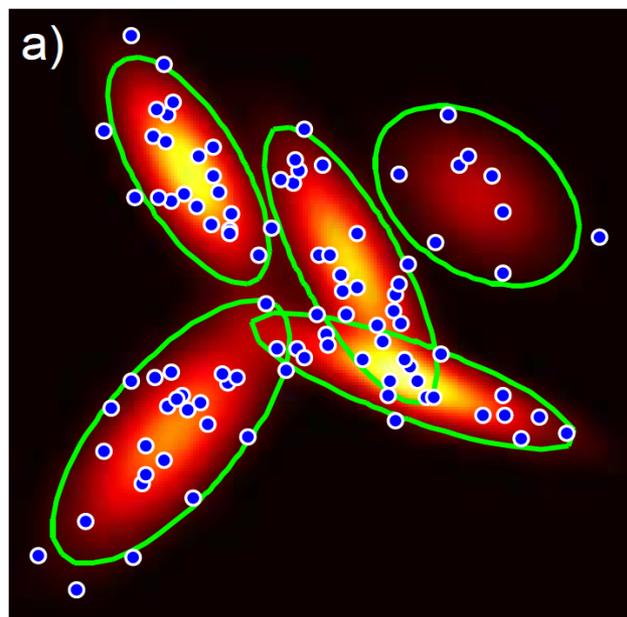
Update means , covariances and weights according to responsibilities of datapoints

Iterate until no further improvement

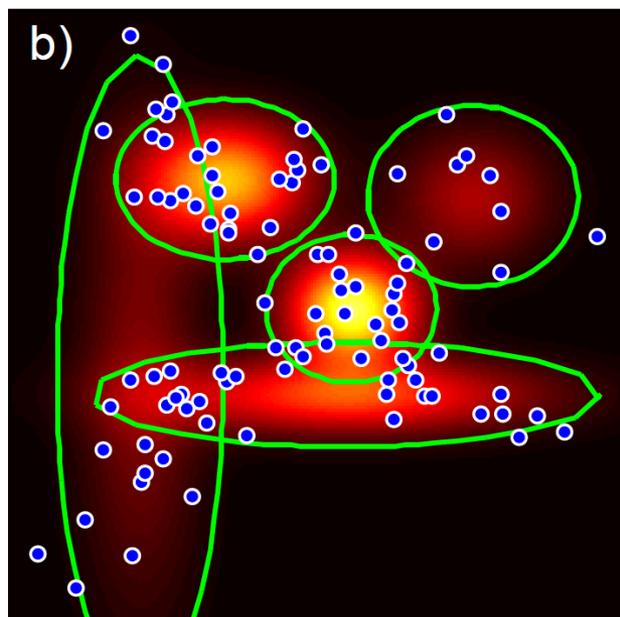




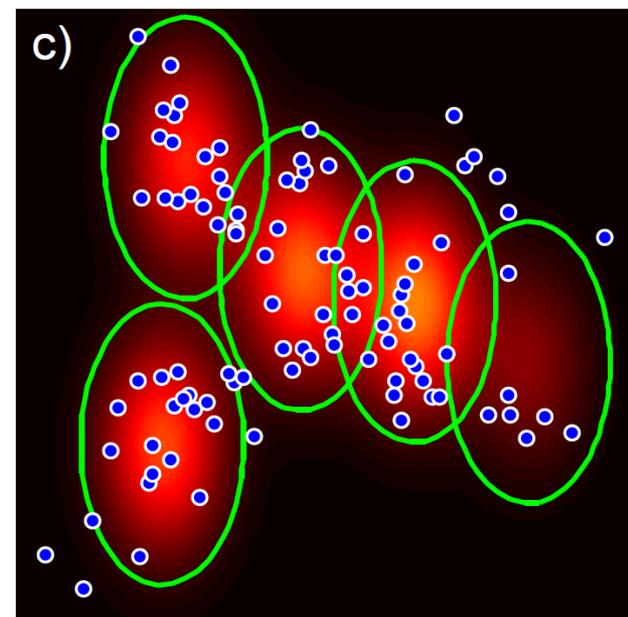
Different flavours...



Full
covariance



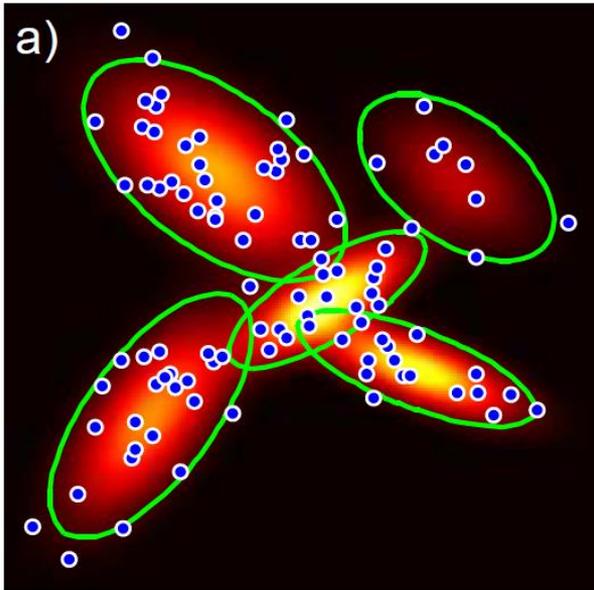
Diagonal
covariance



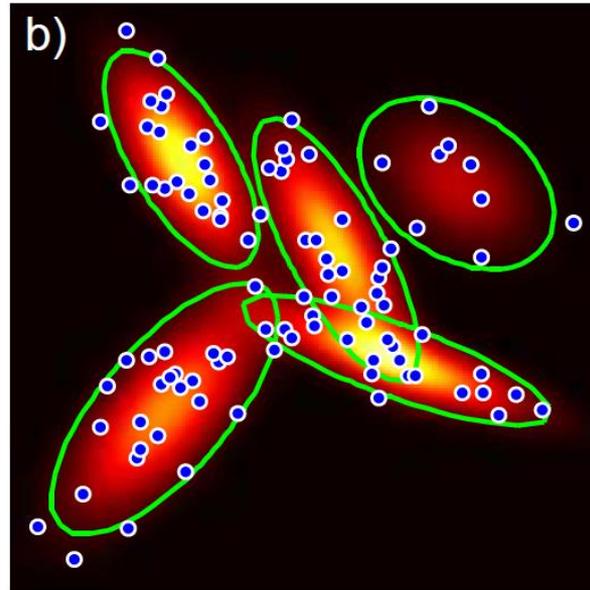
Same
covariance

Local Minima

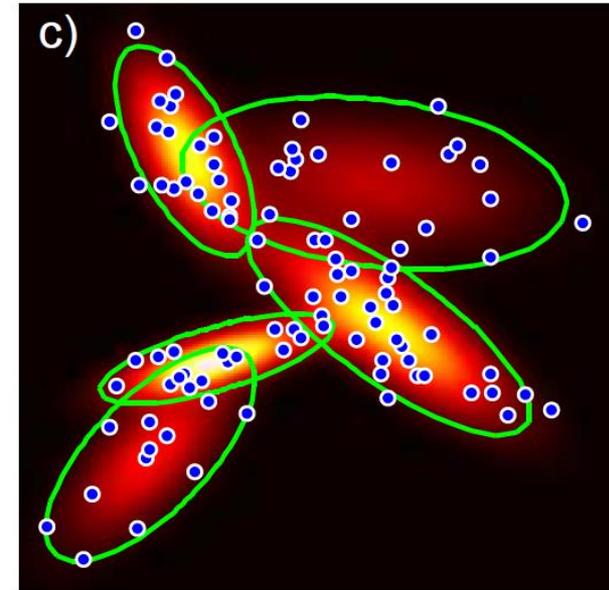
Start from three random positions



$L = 98.76$

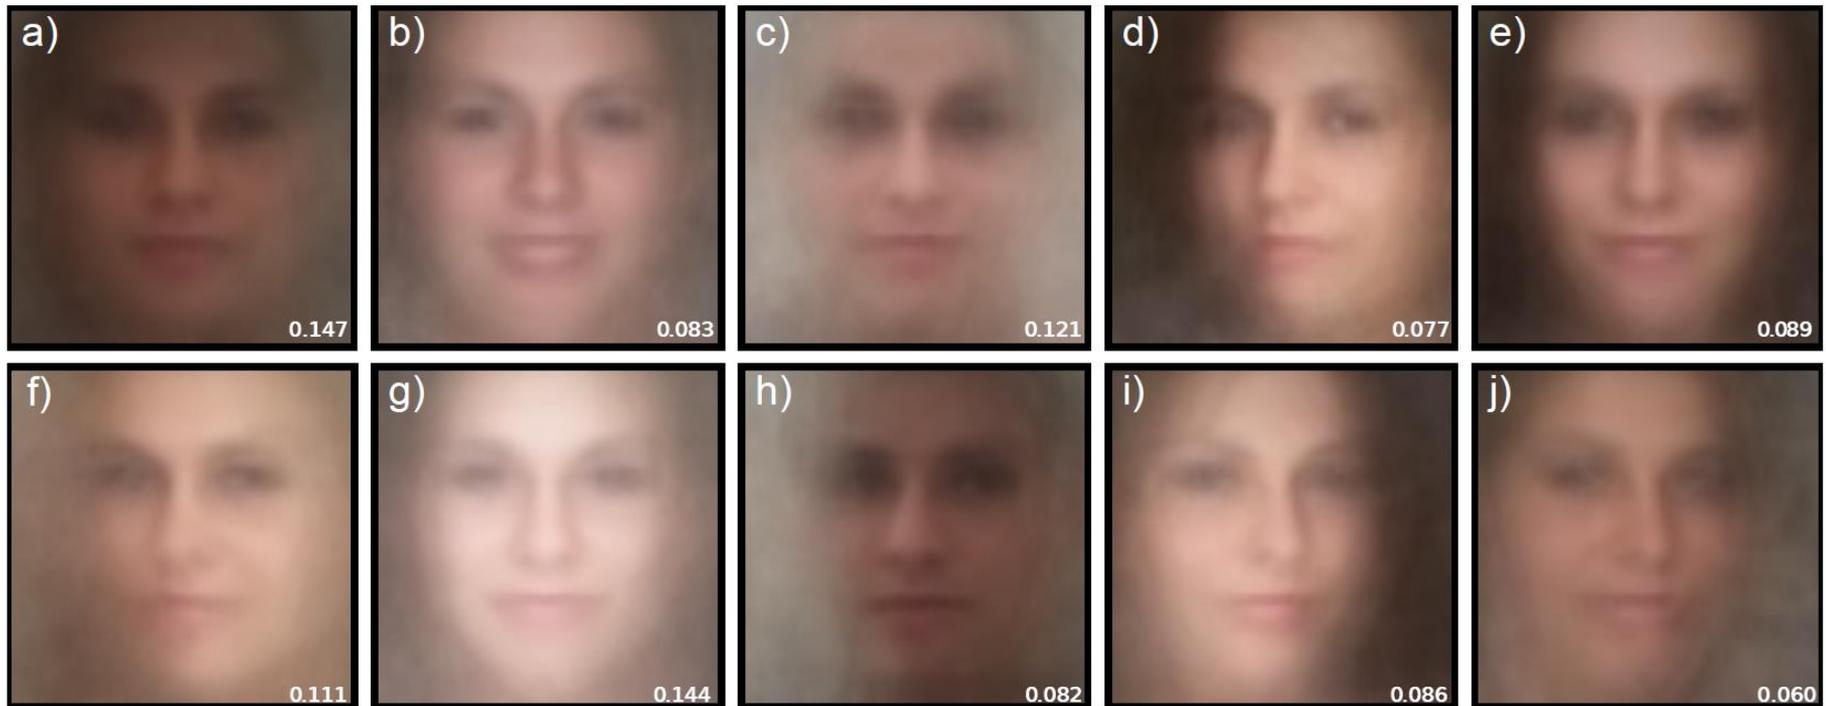


$L = 96.97$



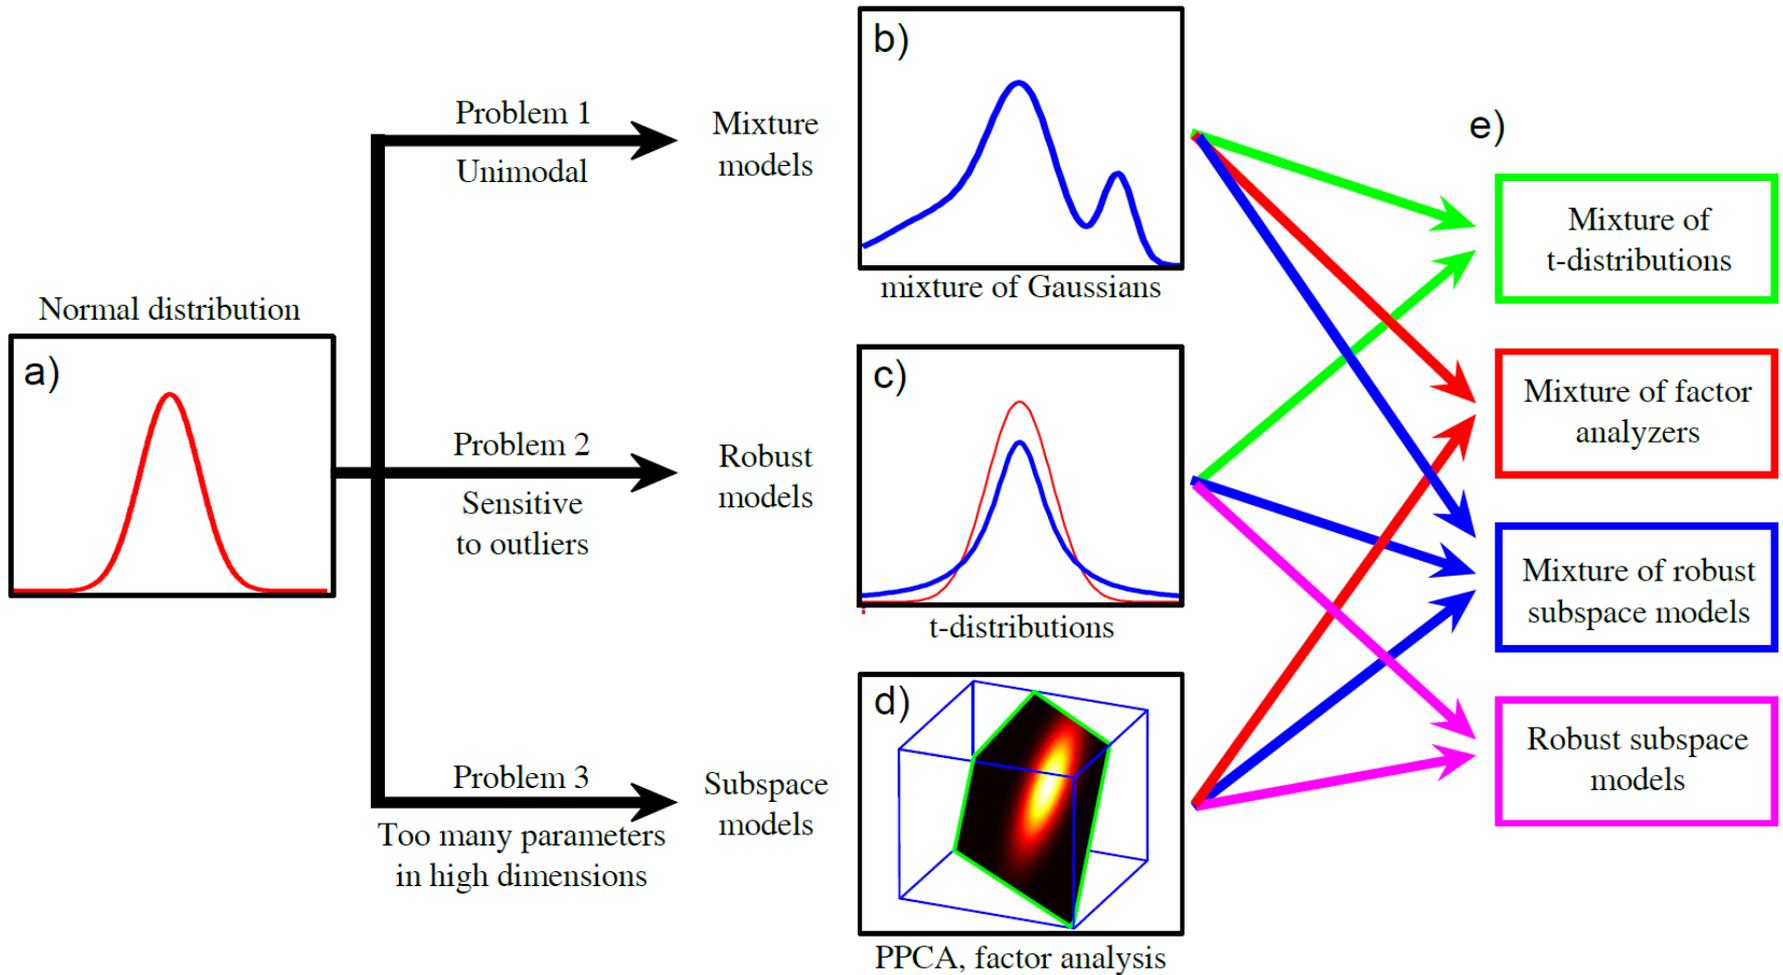
$L = 94.35$

Means of face/non-face model



Classification → 84% (9% improvement!)

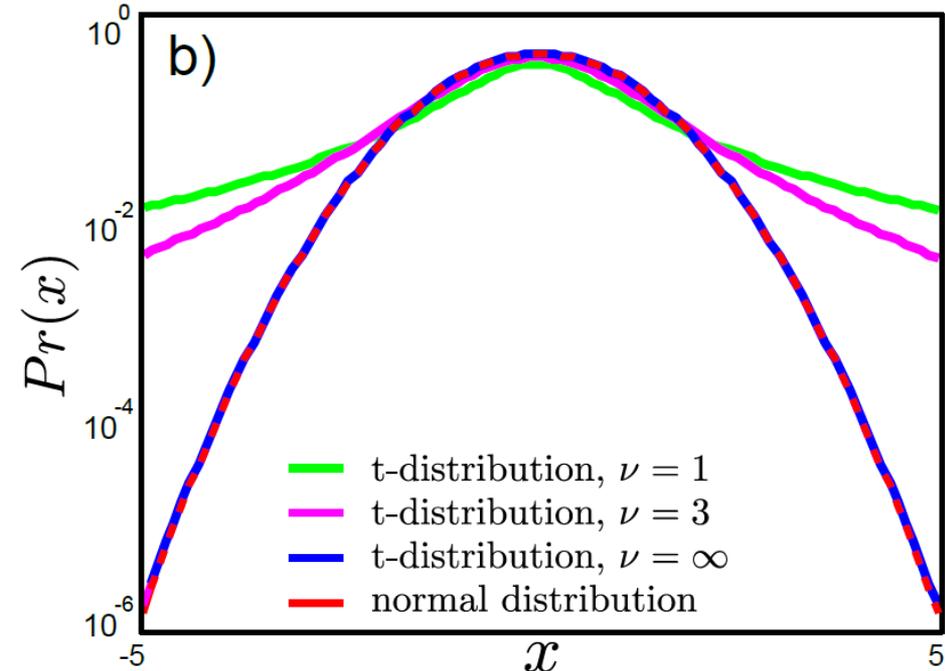
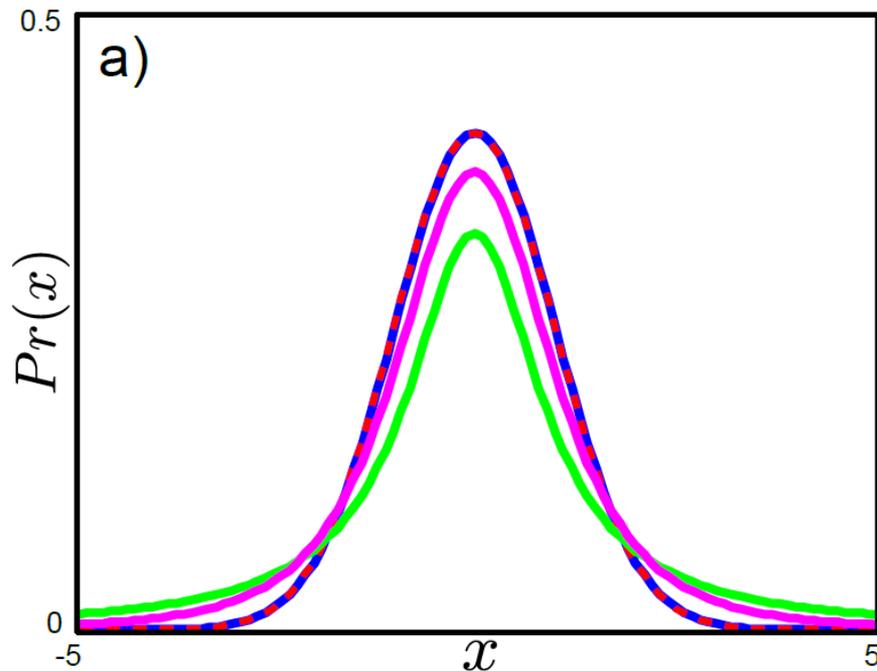
The plan...



Structure

- Densities for classification
- Models with hidden variables
 - Mixture of Gaussians
 - t-distributions
 - Factor analysis
- EM algorithm in detail
- Applications

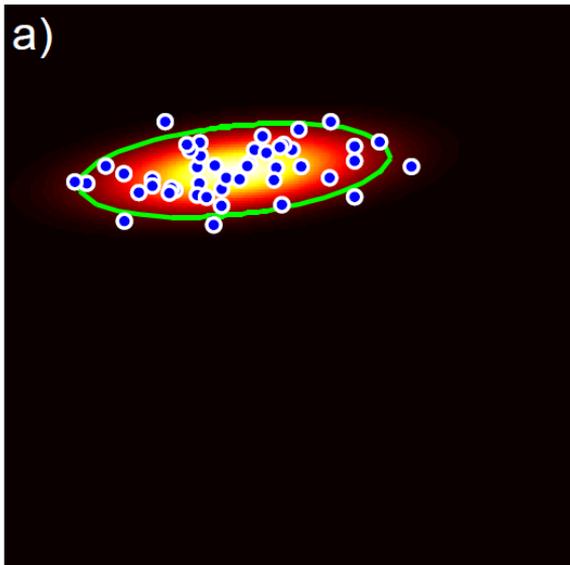
Student t-distributions



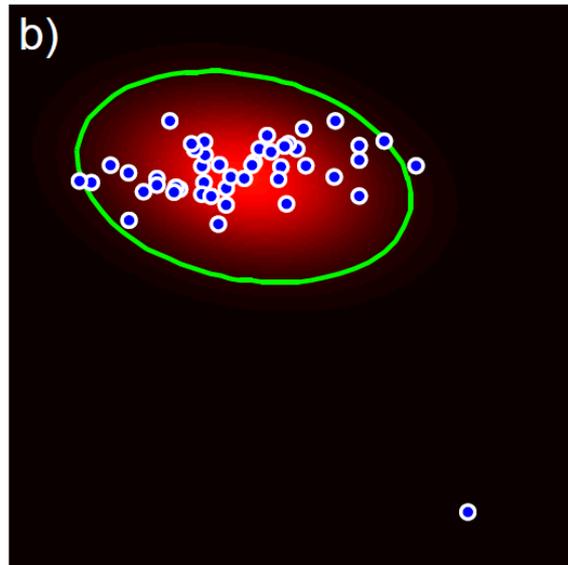
$$\begin{aligned} Pr(x) &= \text{Stud}_{\mathbf{x}} [\mu, \sigma^2, \nu] \\ &= \frac{\Gamma \left[\frac{\nu+1}{2} \right]}{\sqrt{\nu\pi\sigma^2} \Gamma \left[\frac{\nu}{2} \right]} \left(1 + \frac{(x - \mu)^2}{\nu\sigma^2} \right)^{-\frac{\nu+1}{2}} \end{aligned}$$

Student t-distributions motivation

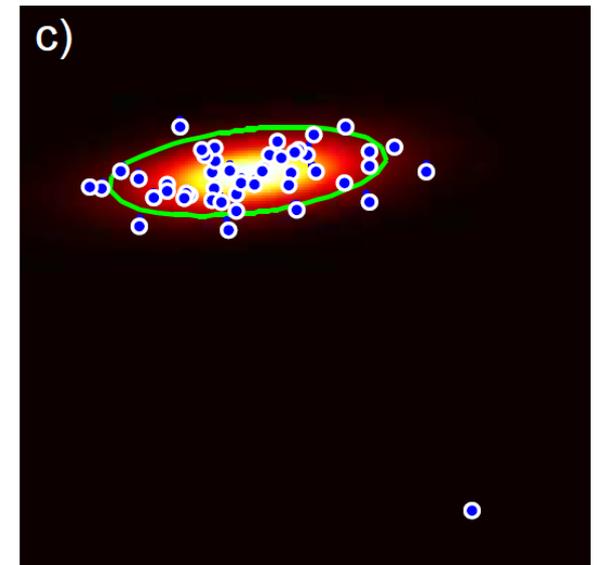
The normal distribution is not very robust – a single outlier can completely throw it off because the tails fall off so fast...



Normal distribution



Normal distribution
w/ one extra datapoint!



t-distribution

Student t-distributions

Univariate student t-distribution

$$\begin{aligned} Pr(x) &= \text{Stud}_{\mathbf{x}} [\mu, \sigma^2, \nu] \\ &= \frac{\Gamma \left[\frac{\nu+1}{2} \right]}{\sqrt{\nu\pi\sigma^2} \Gamma \left[\frac{\nu}{2} \right]} \left(1 + \frac{(x - \mu)^2}{\nu\sigma^2} \right)^{-\frac{\nu+1}{2}} \end{aligned}$$

Multivariate student t-distribution

$$\begin{aligned} Pr(\mathbf{x}) &= \text{Stud}_{\mathbf{x}} [\boldsymbol{\mu}, \boldsymbol{\Sigma}, \nu] \\ &= \frac{\Gamma \left[\frac{\nu+D}{2} \right]}{(\nu\pi)^{D/2} |\boldsymbol{\Sigma}|^{1/2} \Gamma \left[\frac{\nu}{2} \right]} \left(1 + \frac{(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})}{\nu} \right)^{-\frac{\nu+D}{2}} \end{aligned}$$

t-distribution as a marginalization

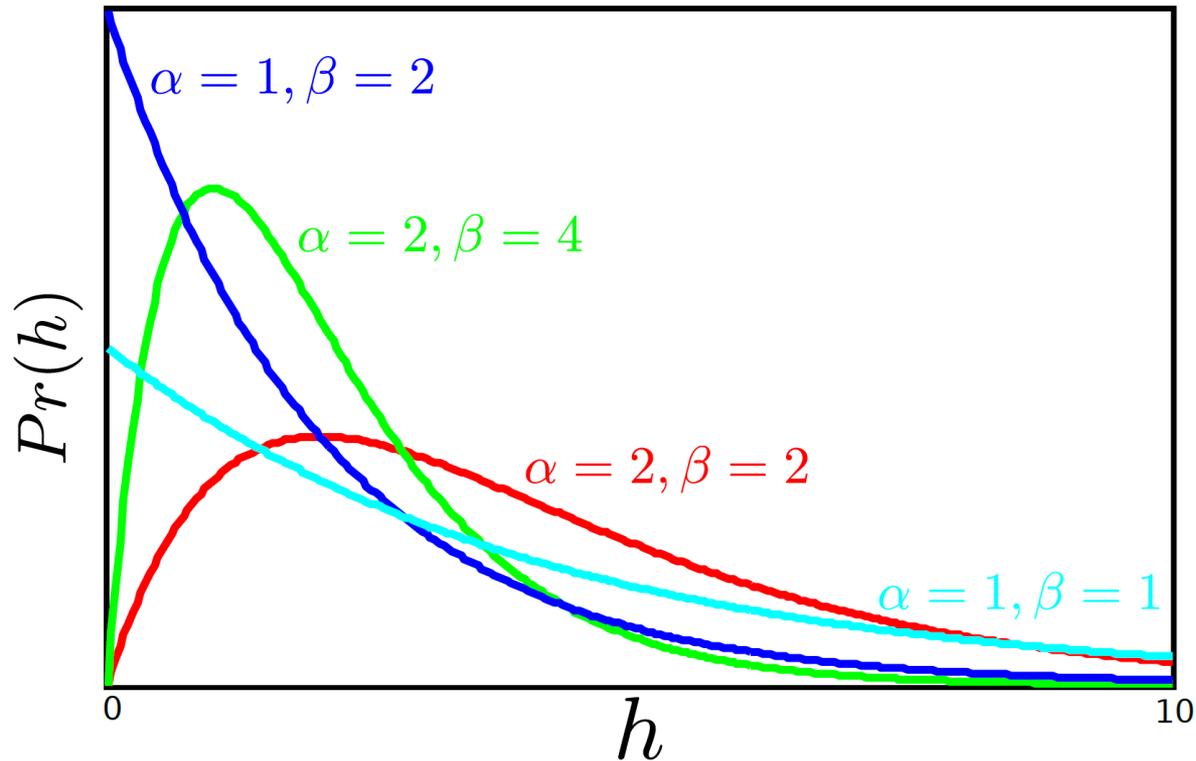
Define hidden variable h

$$\begin{aligned}Pr(\mathbf{x}|h) &= \text{Norm}_x[\boldsymbol{\mu}, \boldsymbol{\Sigma}/h] \\Pr(h) &= \text{Gam}_h[\nu/2, \nu/2]\end{aligned}$$

Can be expressed as a marginalization

$$\begin{aligned}Pr(\mathbf{x}) &= \int Pr(\mathbf{x}, h)dh = \int Pr(\mathbf{x}|h)Pr(h)dh \\&= \int \text{Norm}_x[\boldsymbol{\mu}, \boldsymbol{\Sigma}/h]\text{Gam}_h[\nu/2, \nu/2]dh \\&= \text{Stud}_{\mathbf{x}}[\boldsymbol{\mu}, \boldsymbol{\Sigma}, \nu].\end{aligned}$$

Gamma distribution



$$\text{Gam}_h[\alpha, \beta] = \frac{\beta^\alpha}{\Gamma[\alpha]} \exp[-\beta h] h^{\alpha-1}$$

t-distribution as a marginalization

Define hidden variable h

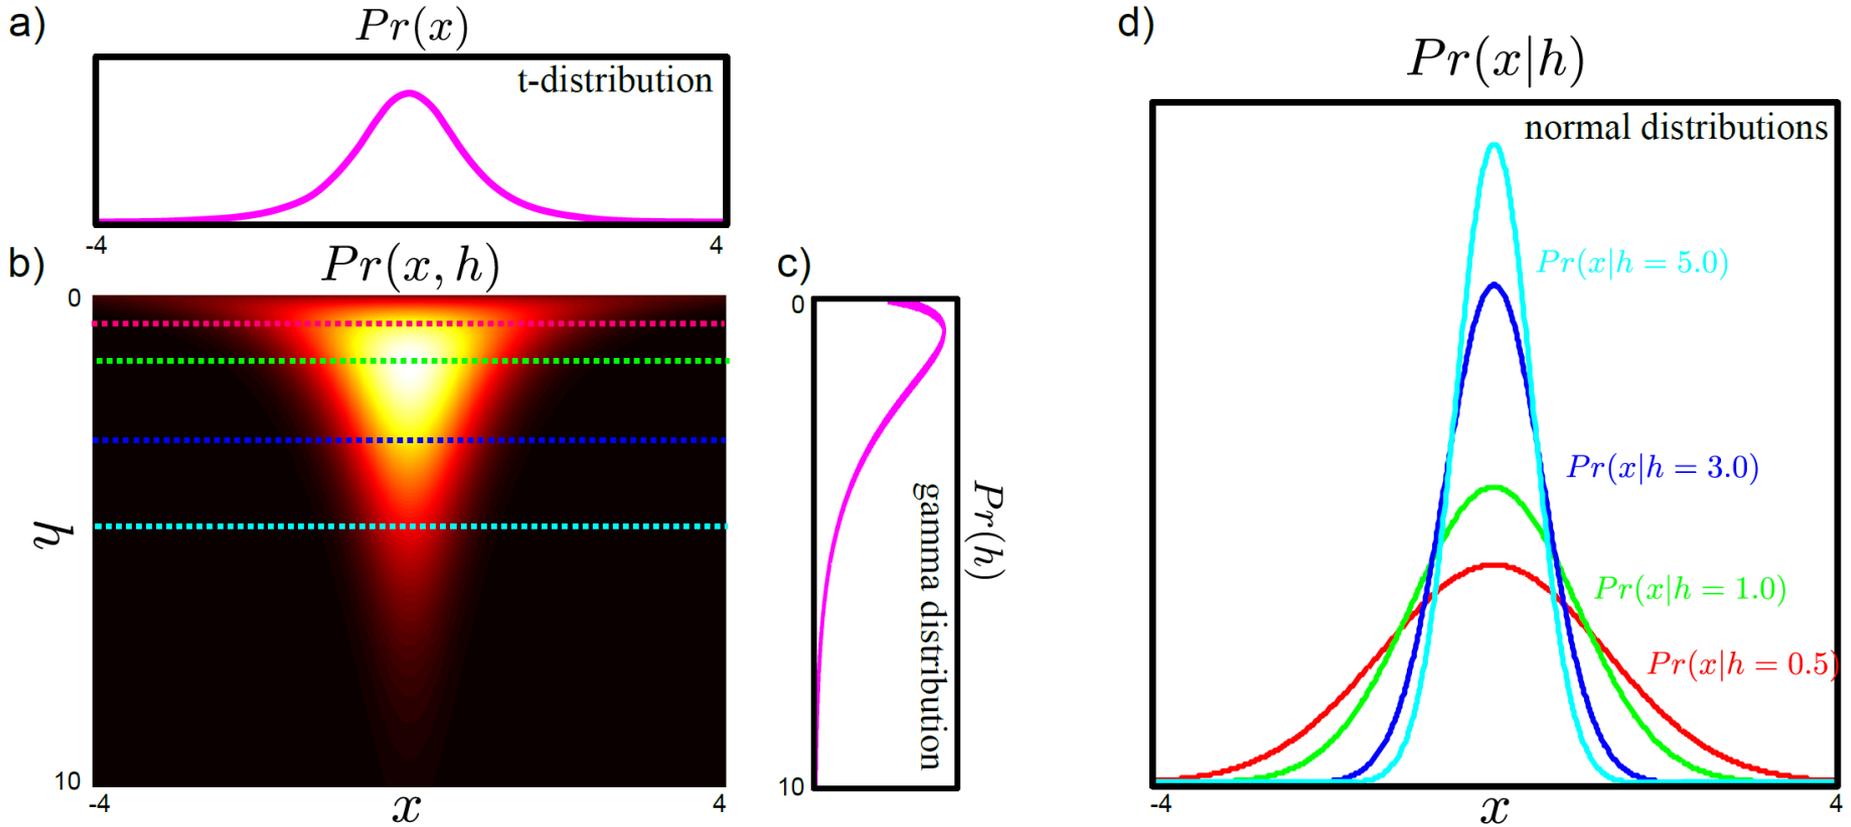
$$Pr(\mathbf{x}|h) = \text{Norm}_x[\boldsymbol{\mu}, \boldsymbol{\Sigma}/h]$$

$$Pr(h) = \text{Gam}_h[\nu/2, \nu/2]$$

Things to note:

- Again this provides a method to sample from the t-distribution
- Variable h has a clear interpretation:
 - Each datum drawn from a Gaussian, mean $\boldsymbol{\mu}$
 - Covariance depends inversely on h
- Can think of this as an infinite mixture (sum becomes integral) of Gaussians w/ same mean, but different variances

t-distribution



EM for t-distributions

GOAL: to learn parameters $\theta = \{\mu, \sigma^2, \nu\}$ from training data $\mathbf{x}_{1\dots I}$

E-Step – Maximize bound w.r.t. distributions $q(h_i)$

$$q_i(h_i) = Pr(h_i|\mathbf{x}_i, \boldsymbol{\theta}^{[t]}) = \frac{Pr(\mathbf{x}_i|h_i, \boldsymbol{\theta}^{[t]})Pr(h_i)}{Pr(\mathbf{x}_i|\boldsymbol{\theta}^{[t]})}$$

M-Step – Maximize bound w.r.t. parameters $\boldsymbol{\theta}$

$$\hat{\boldsymbol{\theta}}^{[t+1]} = \arg \max_{\boldsymbol{\theta}} \sum_{i=1}^I \int \hat{q}_i(h_i) \log [Pr(\mathbf{x}, h_i, \boldsymbol{\theta})] dh_i$$

E-Step

$$\begin{aligned} q_i(h_i) = Pr(h_i | \mathbf{x}_i, \boldsymbol{\theta}^{[t]}) &= \frac{Pr(\mathbf{x}_i | h_i, \boldsymbol{\theta}^{[t]}) Pr(h_i)}{Pr(\mathbf{x}_i | \boldsymbol{\theta}^{[t]})} \\ &= \frac{\text{Norm}_{\mathbf{x}_i}[\boldsymbol{\mu}, \boldsymbol{\Sigma} / h_i] \text{Gam}_{h_i}[\nu/2, \nu/2]}{Pr(\mathbf{x}_i)} \\ &= \text{Gam}_{h_i} \left[\frac{\nu + D}{2}, \frac{(\mathbf{x}_i - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu})}{2} + \frac{\nu}{2} \right] \end{aligned}$$

Extract expectations

$$\begin{aligned} E[h_i] &= \frac{(\nu + D)}{\nu + (\mathbf{x}_i - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu})} \\ E[\log[h_i]] &= \Psi \left[\frac{\nu + D}{2} \right] - \log \left[\frac{\nu + (\mathbf{x}_i - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu})}{2} \right] \end{aligned}$$

M-Step

$$\begin{aligned}\hat{\boldsymbol{\theta}}^{[t+1]} &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int \hat{q}_i(h_i) \log [Pr(\mathbf{x}_i, h_i | \boldsymbol{\theta})] dh_i \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int \hat{q}_i(h_i) (\log [Pr(\mathbf{x}_i | h_i, \boldsymbol{\theta})] + \log [Pr(h_i)]) dh_i \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int Pr(h_i | \mathbf{x}_i, \boldsymbol{\theta}^{[t]}) (\log [Pr(\mathbf{x}_i | h_i, \boldsymbol{\theta})] + \log [Pr(h_i)]) dh_i \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I E [\log [Pr(\mathbf{x}_i | h_i, \boldsymbol{\theta})]] + E [\log [Pr(h_i)]] \right]\end{aligned}$$

Where...

$$E [\log [Pr(\mathbf{x}_i | h_i, \boldsymbol{\theta})]] = \frac{DE[\log h_i] - D \log 2\pi - \log |\boldsymbol{\Sigma}| - (\mathbf{x}_i - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu}) E[h_i]}{2}$$
$$E [\log [Pr(h_i)]] = \frac{\nu}{2} \log \left[\frac{\nu}{2} \right] - \log \Gamma \left[\frac{\nu}{2} \right] + \left(\frac{\nu}{2} - 1 \right) E[\log h_i] - \frac{\nu}{2} E[h_i].$$

Updates

$$\begin{aligned}\boldsymbol{\mu}^{[t+1]} &= \frac{\sum_{i=1}^I \mathbb{E}[h_i] \mathbf{x}_i}{\sum_{i=1}^I \mathbb{E}[h_i]} \\ \boldsymbol{\Sigma}^{[t+1]} &= \frac{\sum_{i=1}^I \mathbb{E}[h_i] (\mathbf{x}_i - \boldsymbol{\mu}^{[t+1]})(\mathbf{x}_i - \boldsymbol{\mu}^{[t+1]})^T}{\sum_{i=1}^I \mathbb{E}[h_i]}\end{aligned}$$

No closed form solution for v . Must optimize bound – since it is only one number we can optimize by just evaluating with a fine grid of values and picking the best

EM algorithm for t-distributions

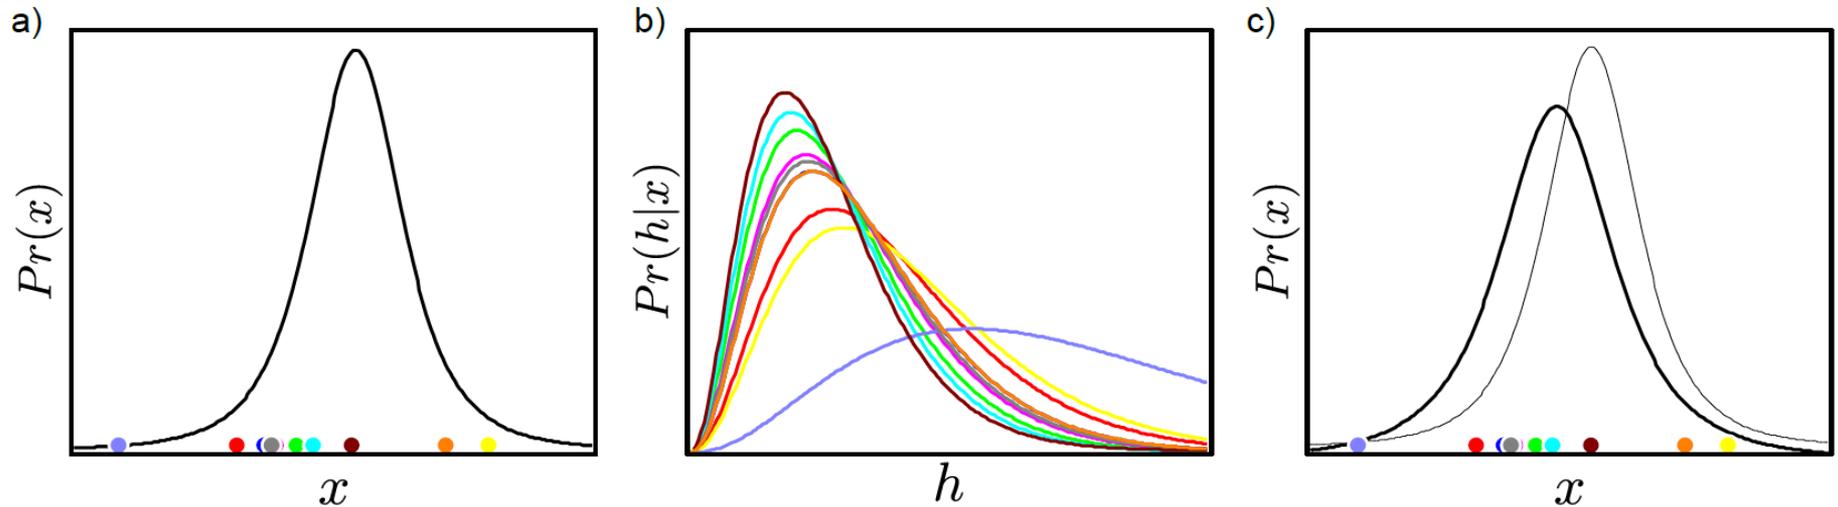
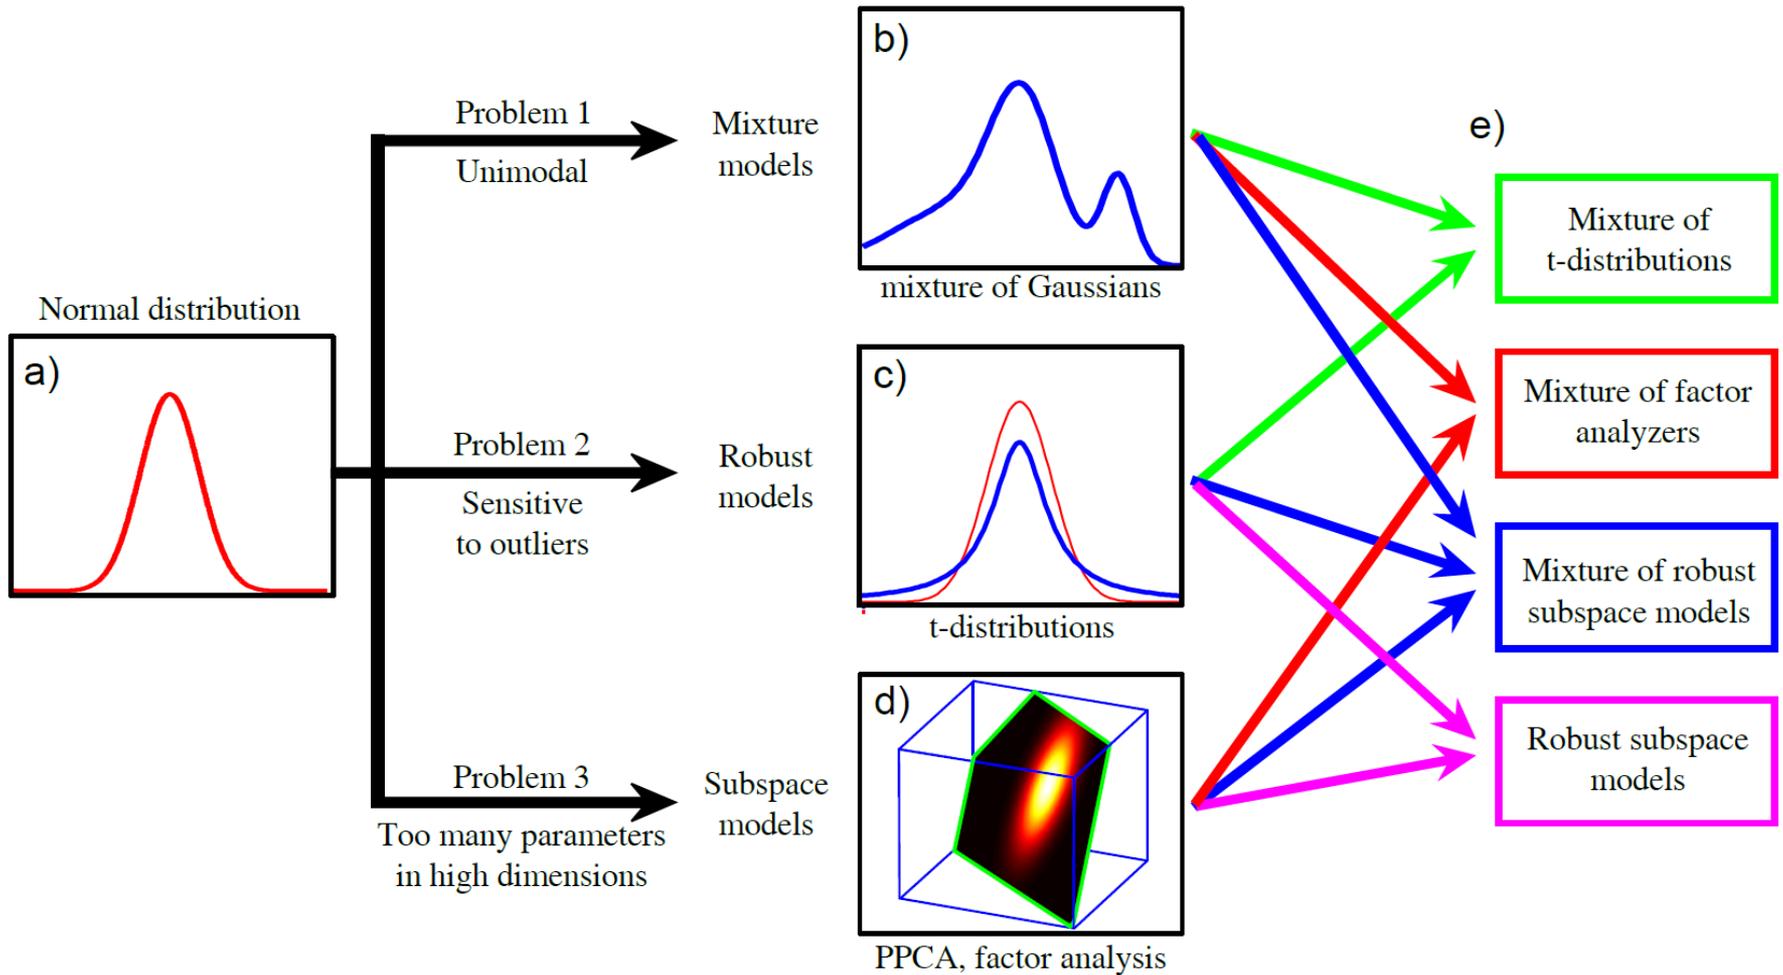


Figure 7.18 Expectation maximization for fitting t-distributions. a) Estimate of distribution before update. b) In the E-Step we calculate the posterior distribution $Pr(h_i|x_i)$ over the hidden variable h_i for each data point x_i . The color of each curve corresponds to that of the original data point in (a). c) In the M-Step we use these distributions over h to update the estimate of the parameters $\theta = \{\mu, \sigma^2, \nu\}$.

The plan...



Structure

- Densities for classification
- Models with hidden variables
 - Mixture of Gaussians
 - t-distributions
 - Factor analysis
- EM algorithm in detail
- Applications

Factor Analysis

Compromise between

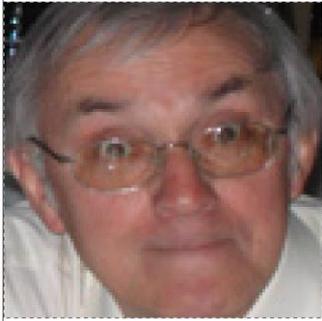
- Full covariance matrix (desirable but many parameters)
- Diagonal covariance matrix (no modelling of covariance)

Models full covariance in a subspace Φ

Mops everything else up with diagonal component Σ

$$Pr(\mathbf{x}) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}, \Phi\Phi^T + \Sigma]$$

Data Density

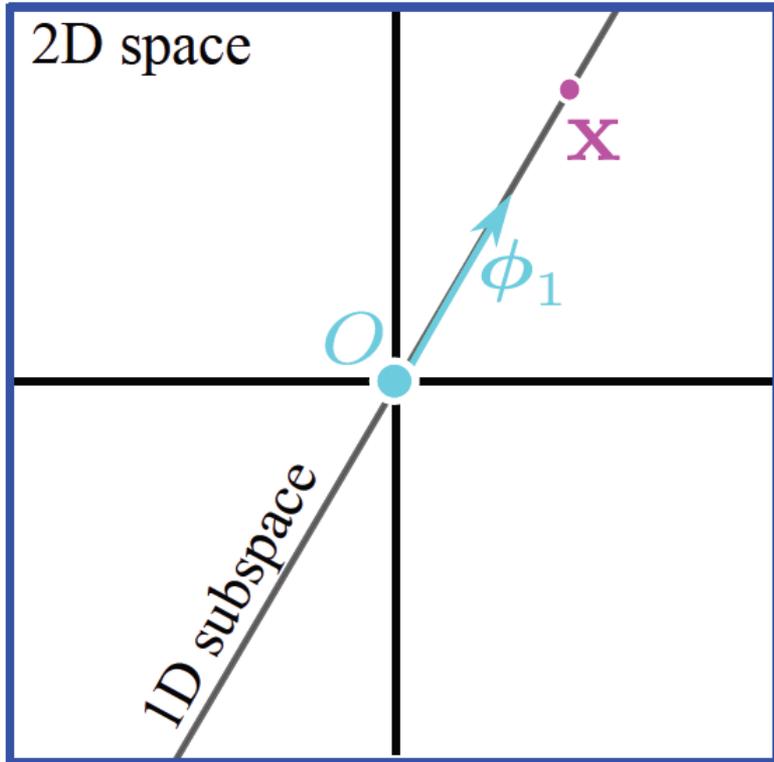


Consider modelling distribution of 100x100x3 RGB Images

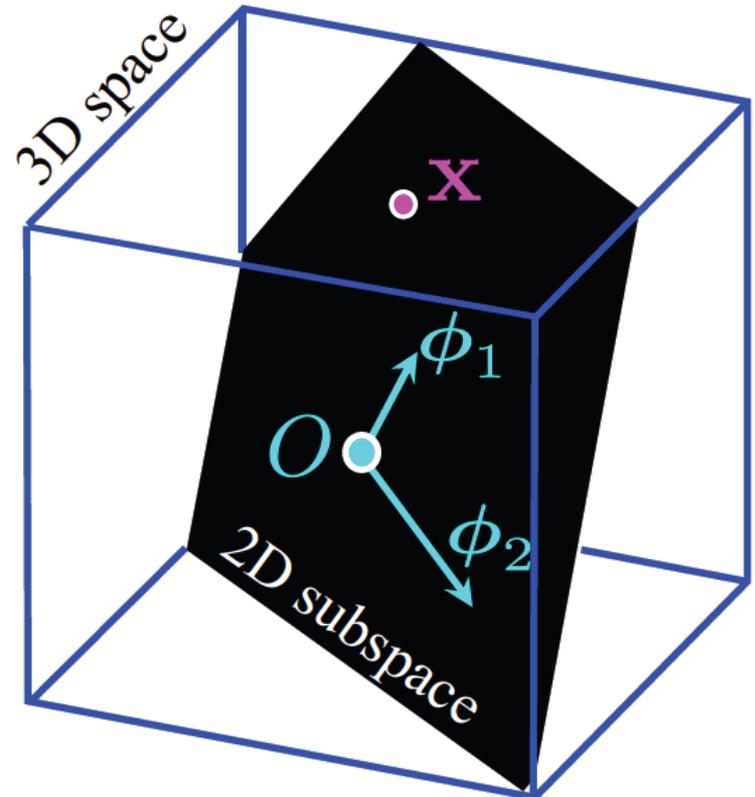
Gaussian w/ spherical covariance:	1	covariance matrix parameters
Gaussian w/ diagonal covariance:	D_x	covariance matrix parameters
Full Gaussian:	$\sim D_x^2$	covariance matrix parameters
PPCA:	$\sim D_x D_h$	covariance matrix parameters
Factor Analysis:	$\sim D_x (D_h + 1)$	covariance matrix parameters

Subspaces

a)



b)



Factor Analysis

Compromise between

- Full covariance matrix (desirable but many parameters)
- Diagonal covariance matrix (no modelling of covariance)

Models full covariance in a subspace Φ

Mops everything else up with diagonal component Σ

$$Pr(\mathbf{x}) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}, \boldsymbol{\Phi}\boldsymbol{\Phi}^T + \boldsymbol{\Sigma}]$$

Factor Analysis as a Marginalization

Let's define

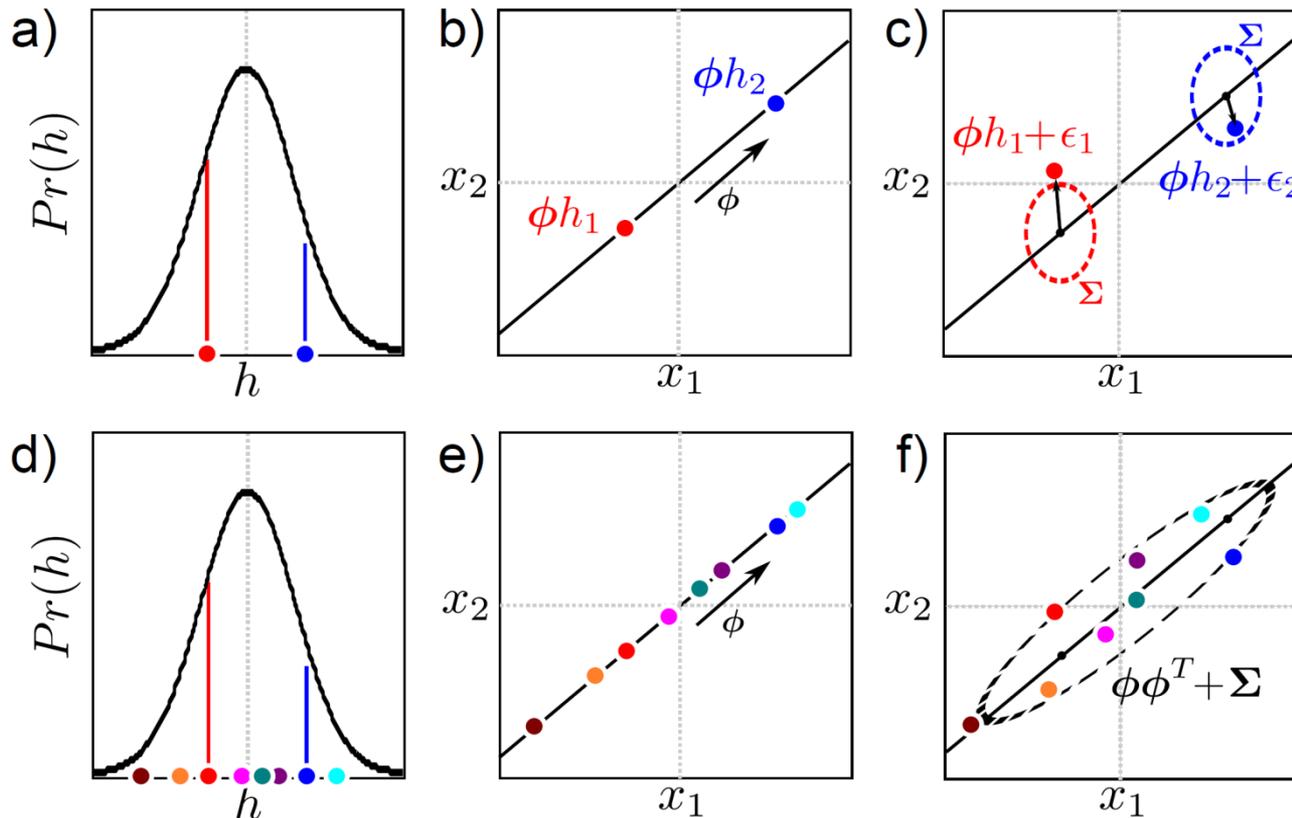
$$Pr(\mathbf{x}|\mathbf{h}) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu} + \boldsymbol{\Phi}\mathbf{h}, \boldsymbol{\Sigma}]$$

$$Pr(\mathbf{h}) = \text{Norm}_{\mathbf{h}}[\mathbf{0}, \mathbf{I}]$$

Then it can be shown (not obvious) that

$$\begin{aligned} Pr(\mathbf{x}) &= \int Pr(\mathbf{x}, \mathbf{h})d\mathbf{h} = \int Pr(\mathbf{x}|\mathbf{h})Pr(\mathbf{h})d\mathbf{h} \\ &= \int \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu} + \boldsymbol{\Phi}\mathbf{h}, \boldsymbol{\Sigma}]\text{Norm}_{\mathbf{h}}[\mathbf{0}, \mathbf{I}]d\mathbf{h} \\ &= \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu}, \boldsymbol{\Phi}\boldsymbol{\Phi}^T + \boldsymbol{\Sigma}] \end{aligned}$$

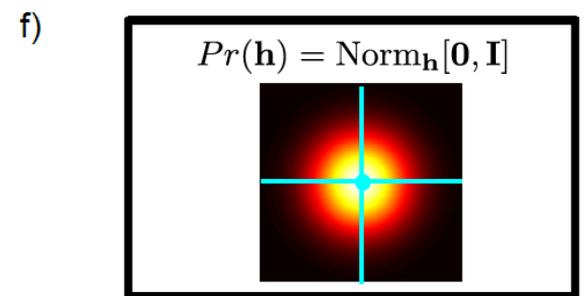
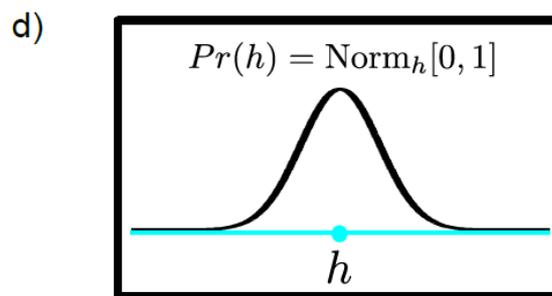
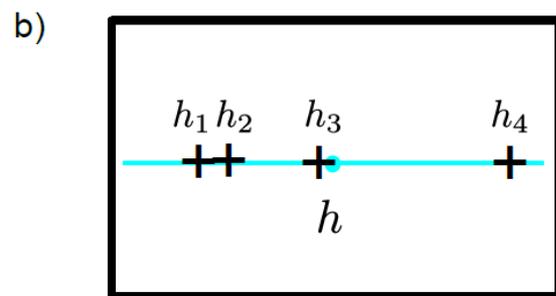
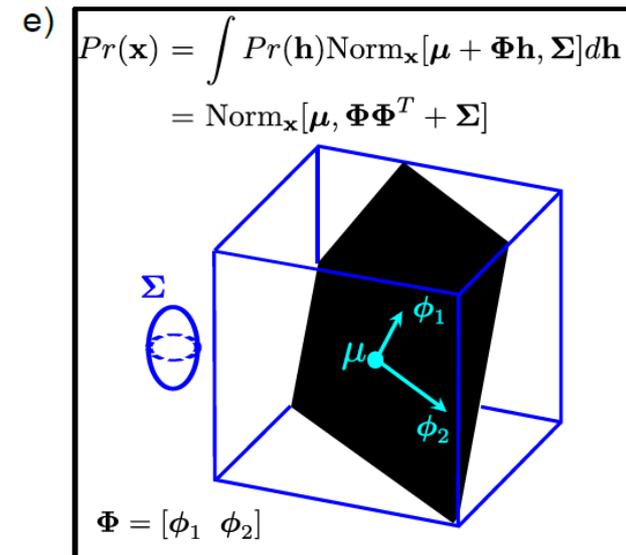
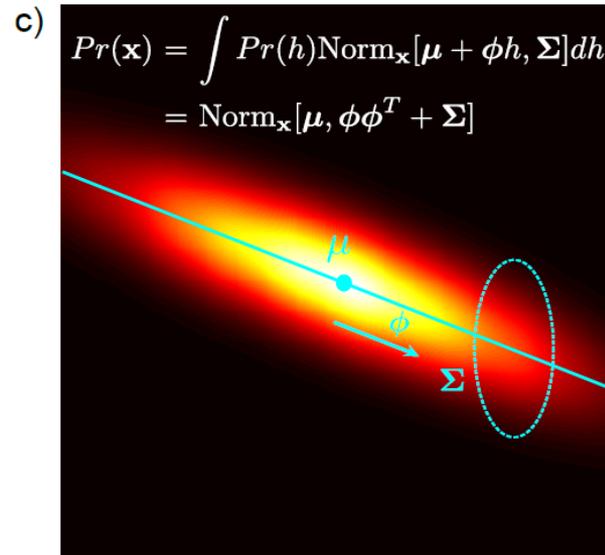
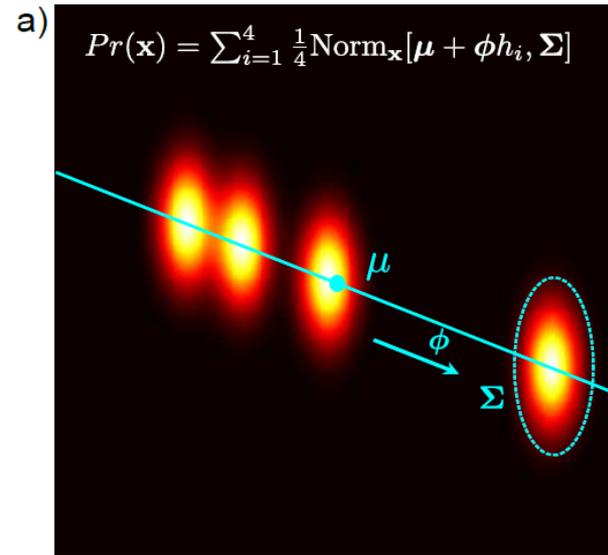
Sampling from factor analyzer



$$Pr(\mathbf{x}|\mathbf{h}) = \text{Norm}_{\mathbf{x}}[\boldsymbol{\mu} + \boldsymbol{\Phi}\mathbf{h}, \boldsymbol{\Sigma}]$$

$$Pr(\mathbf{h}) = \text{Norm}_{\mathbf{h}}[\mathbf{0}, \mathbf{I}]$$

Factor analysis vs. MoG



E-Step

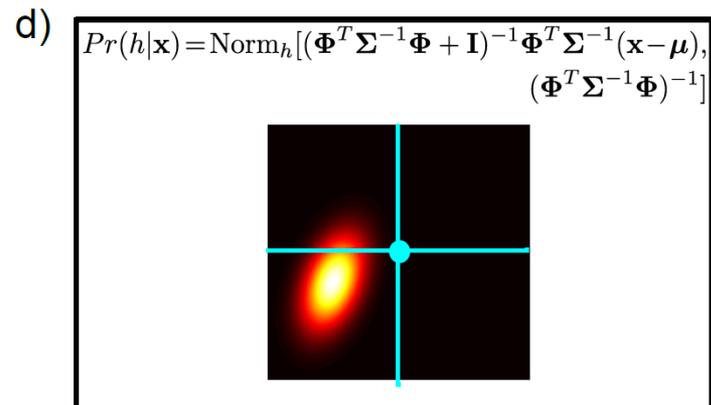
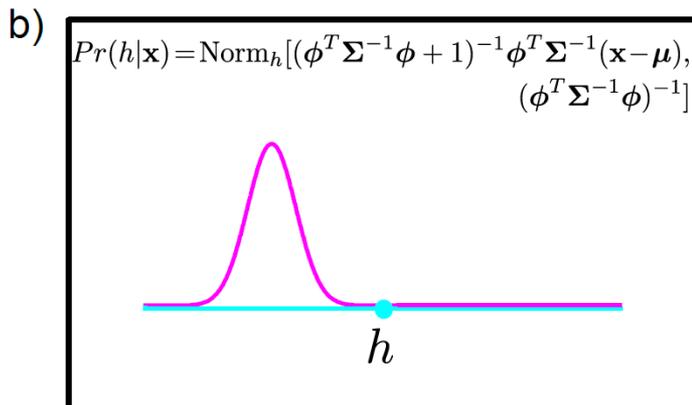
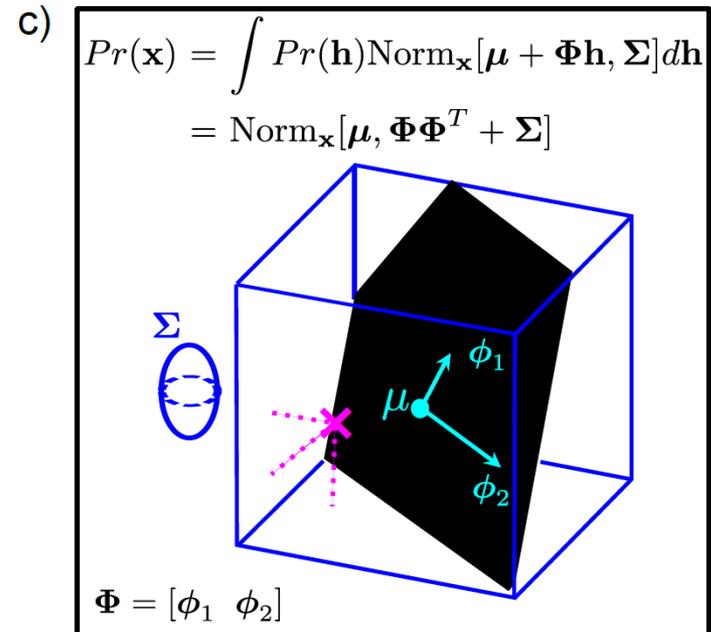
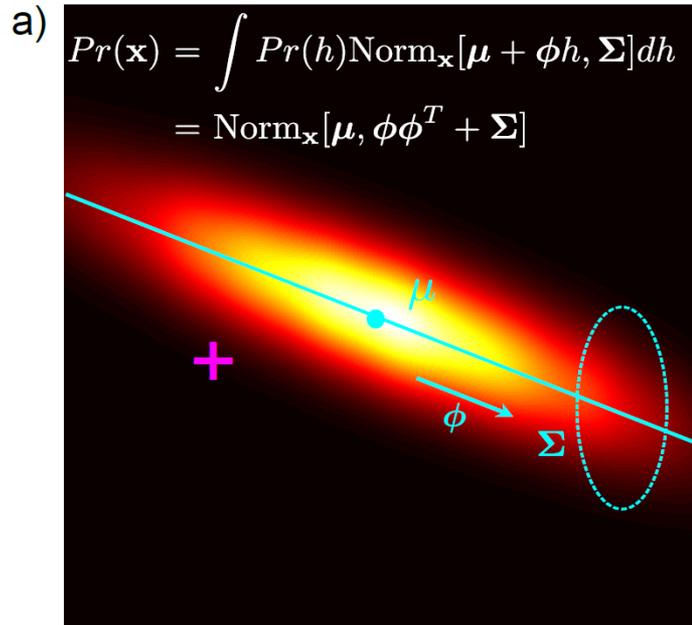
- Compute posterior over hidden variables

$$\begin{aligned}\hat{q}(\mathbf{h}_i) &= Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta}^{[t]}) \\ &= \frac{Pr(\mathbf{x}_i|\mathbf{h}_i, \boldsymbol{\theta}^{[t]})Pr(\mathbf{h}_i)}{Pr(\mathbf{x}_i|\boldsymbol{\theta}^{[t]})} \\ &= \frac{\text{Norm}_{\mathbf{x}_i}[\boldsymbol{\mu} + \boldsymbol{\Phi}\mathbf{h}_i, \boldsymbol{\Sigma}]\text{Norm}_{\mathbf{h}_i}[\mathbf{0}, \mathbf{I}]}{Pr(\mathbf{x}_i|\boldsymbol{\theta}^{[t]})} \\ &= \text{Norm}_{\mathbf{h}_i}[(\boldsymbol{\Phi}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\Phi} + \mathbf{I})^{-1} \boldsymbol{\Phi}^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu}), (\boldsymbol{\Phi}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\Phi} + \mathbf{I})]\end{aligned}$$

- Compute expectations needed for M-Step

$$\begin{aligned}\mathbf{E}[\mathbf{h}_i] &= (\boldsymbol{\Phi}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\Phi} + \mathbf{I})^{-1} \boldsymbol{\Phi}^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu}) \\ \mathbf{E}[\mathbf{h}_i \mathbf{h}_i^T] &= E [(\mathbf{h}_i - \mathbf{E}[\mathbf{h}_i])(\mathbf{h}_i - \mathbf{E}[\mathbf{h}_i])^T] + \mathbf{E}[\mathbf{h}_i] \mathbf{E}[\mathbf{h}_i]^T \\ &= (\boldsymbol{\Phi}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\Phi} + \mathbf{I})^{-1} + \mathbf{E}[\mathbf{h}_i] \mathbf{E}[\mathbf{h}_i]^T.\end{aligned}$$

E-Step



M-Step

- Optimize parameters w.r.t. EM bound

$$\begin{aligned}\hat{\boldsymbol{\theta}}^{[t+1]} &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int \hat{q}_i(\mathbf{h}_i) \log [Pr(\mathbf{x}, \mathbf{h}_i, \boldsymbol{\theta})] d\mathbf{h}_i \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int \hat{q}_i(\mathbf{h}_i) [\log [Pr(\mathbf{x}|\mathbf{h}_i, \boldsymbol{\theta})] + \log [Pr(\mathbf{h}_i)]] d\mathbf{h}_i \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int \hat{q}_i(\mathbf{h}_i) \log [Pr(\mathbf{x}|\mathbf{h}_i, \boldsymbol{\theta})] d\mathbf{h}_i \right] \\ &= \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I E [\log Pr(\mathbf{x}|\mathbf{h}_i, \boldsymbol{\theta})] \right],\end{aligned}$$

where....

$$\log Pr(\mathbf{x}_i|\mathbf{h}_i) = -\frac{D \log(2\pi) + \log |\boldsymbol{\Sigma}| + (\mathbf{x}_i - \boldsymbol{\mu} - \boldsymbol{\Phi}\mathbf{h}_i)^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu} - \boldsymbol{\Phi}\mathbf{h}_i)}{2}$$

M-Step

- Optimize parameters w.r.t. EM bound

$$\hat{\boldsymbol{\theta}}^{[t+1]} = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} \left[\sum_{i=1}^I E [\log Pr(\mathbf{x}|\mathbf{h}_i, \boldsymbol{\theta})] \right]$$

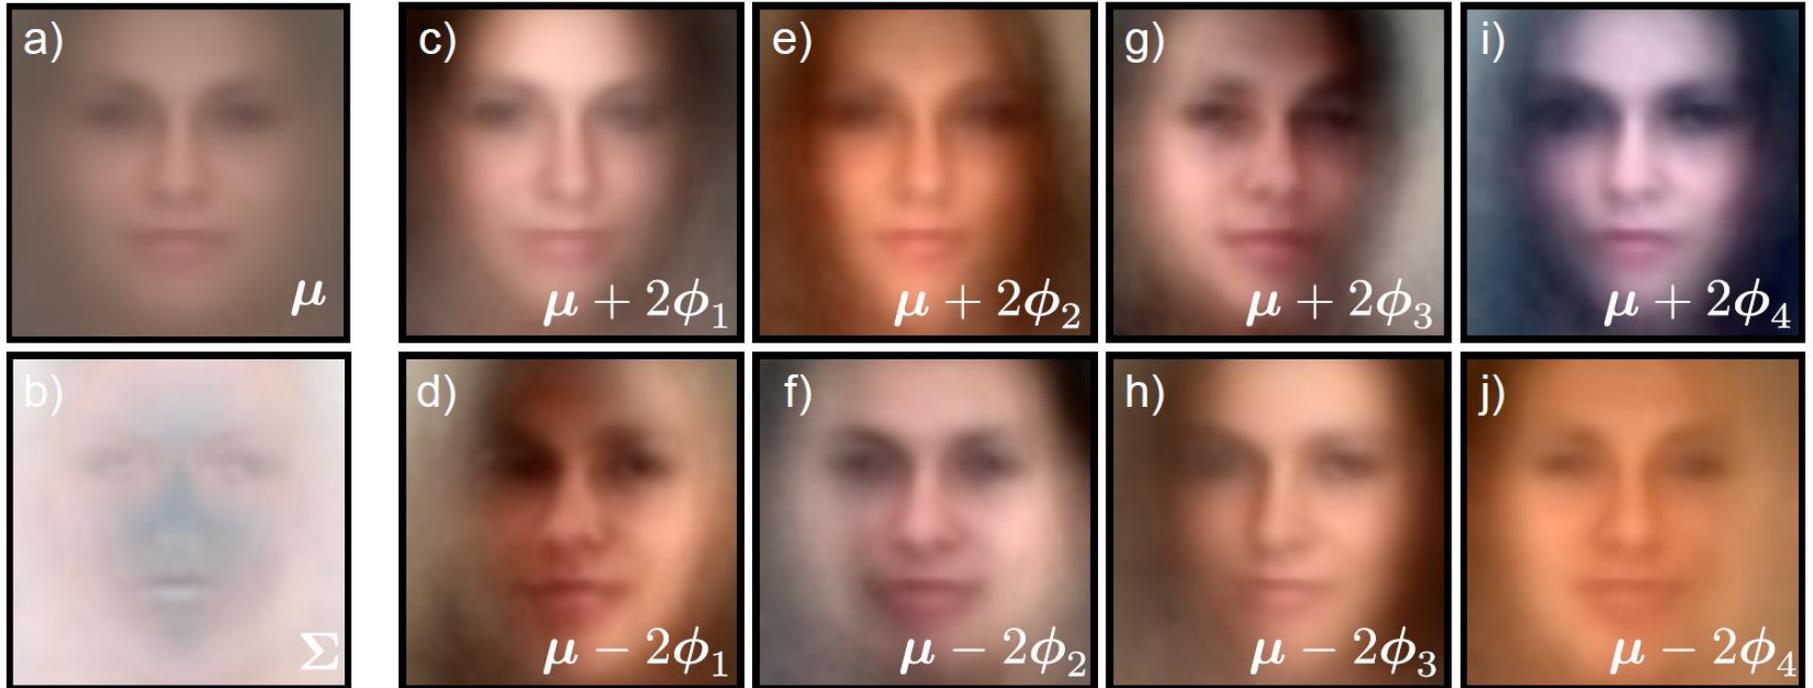
where...

$$\log Pr(\mathbf{x}_i|\mathbf{h}_i) = -\frac{D \log(2\pi) + \log |\boldsymbol{\Sigma}| + (\mathbf{x}_i - \boldsymbol{\mu} - \boldsymbol{\Phi}\mathbf{h}_i)^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu} - \boldsymbol{\Phi}\mathbf{h}_i)}{2}$$

giving expressions:

$$\begin{aligned} \hat{\boldsymbol{\mu}} &= \frac{\sum_{i=1}^I \mathbf{x}_i}{I} \\ \hat{\boldsymbol{\Phi}} &= \left(\sum_{i=1}^I (\mathbf{x}_i - \hat{\boldsymbol{\mu}}) \mathbf{E}[\mathbf{h}_i]^T \right) \left(\sum_{i=1}^I \mathbf{E}[\mathbf{h}_i \mathbf{h}_i^T] \right)^{-1} \\ \hat{\boldsymbol{\Sigma}} &= \frac{1}{I} \sum_{i=1}^I \operatorname{diag} \left[(\mathbf{x}_i - \hat{\boldsymbol{\mu}})^T (\mathbf{x}_i - \hat{\boldsymbol{\mu}}) - \hat{\boldsymbol{\Phi}} \mathbf{E}[\mathbf{h}_i] \mathbf{x}_i^T \right] \end{aligned}$$

Face model



Sampling from 10 parameter model

To generate:

- Choose factor loadings, \mathbf{h}_i from standard normal distribution
- Multiply by factors, Φ
- Add mean, μ
- (should add random noise component ε_i w/ diagonal cov Σ)



Combining Models

Can combine three models in various combinations

- Mixture of t-distributions = mixture of Gaussians + t-distributions
- Robust subspace models = t-distributions + factor analysis
- Mixture of factor analyzers = mixture of Gaussians + factor analysis

Or combine all models to **create mixture of robust subspace models**

$$Pr(\mathbf{x}) = \sum_{k=1}^K \lambda_k \text{Stud}_{\mathbf{x}} \left[\boldsymbol{\mu}_k, \boldsymbol{\Phi}_k \boldsymbol{\Phi}_k^T + \boldsymbol{\Sigma}_k, \nu_k \right]$$

Structure

- Densities for classification
- Models with hidden variables
 - Mixture of Gaussians
 - t-distributions
 - Factor analysis
- EM algorithm in detail
- Applications

Expectation Maximization

Problem: Optimize cost functions of the form

$$\hat{\theta} = \arg \max_{\theta} \sum_{i=1}^I \log \left[\sum_h Pr(\mathbf{x}_i, h_i) \right] \quad \leftarrow \text{Discrete case}$$

$$\hat{\theta} = \arg \max_{\theta} \sum_{i=1}^I \log \left[\int Pr(\mathbf{x}_i, \mathbf{h}_i) d\mathbf{h}_i \right] \quad \leftarrow \text{Continuous case}$$

Solution: Expectation Maximization (EM) algorithm

(Dempster, Laird and Rubin 1977)

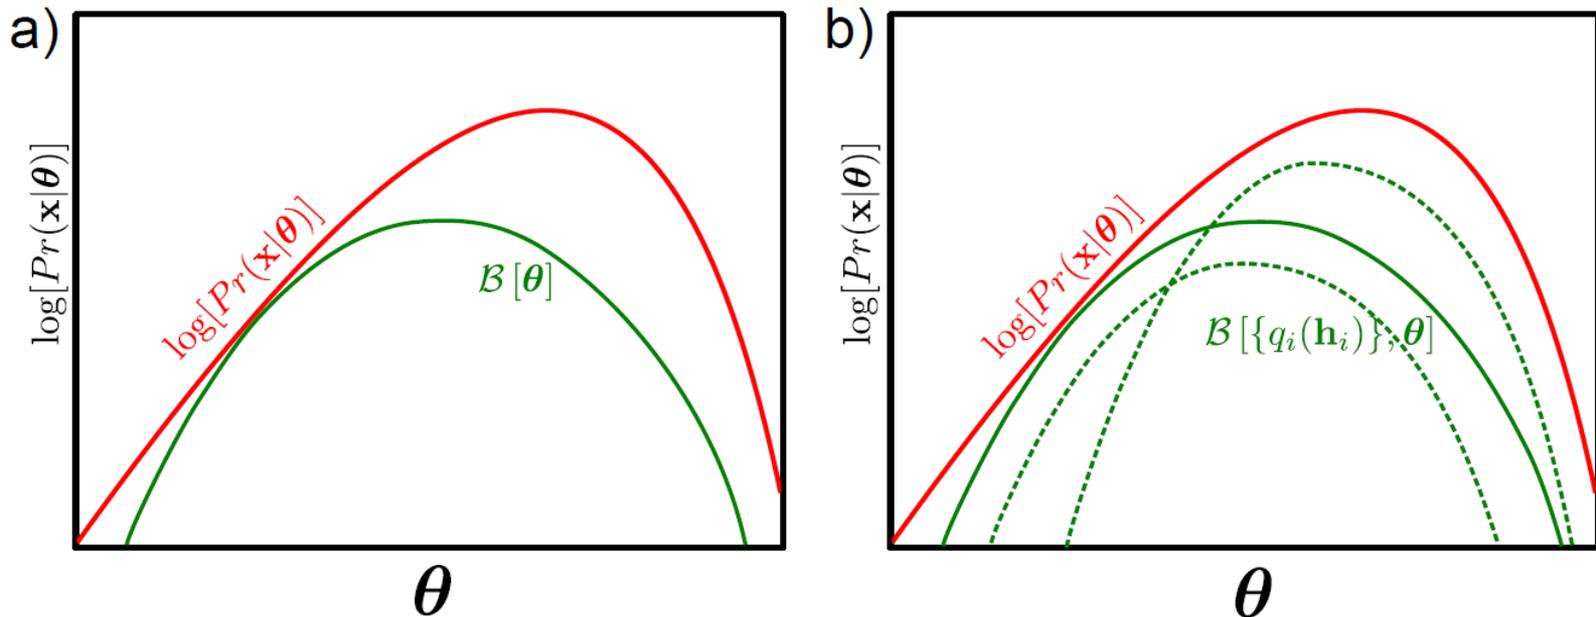
Key idea: Define lower bound on log-likelihood and increase at each iteration

Expectation Maximization

Defines a lower bound on log likelihood and increases bound iteratively

$$\begin{aligned}\mathcal{B}[\{q_i(\mathbf{h}_i)\}, \boldsymbol{\theta}] &= \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \\ &\leq \sum_{i=1}^I \log \left[\int q_i(\mathbf{h}_i) \frac{Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} d\mathbf{h}_i \right] \\ &= \sum_{i=1}^I \log \left[\int Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta}) d\mathbf{h}_i \right],\end{aligned}$$

Lower bound



$$B[\{q_i(\mathbf{h}_i)\}, \theta] = \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}_i, \mathbf{h}_i | \theta)}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i$$

Lower bound is a *function* of parameters θ and a *set* of probability distributions $\{q_i(\mathbf{h}_i)\}$

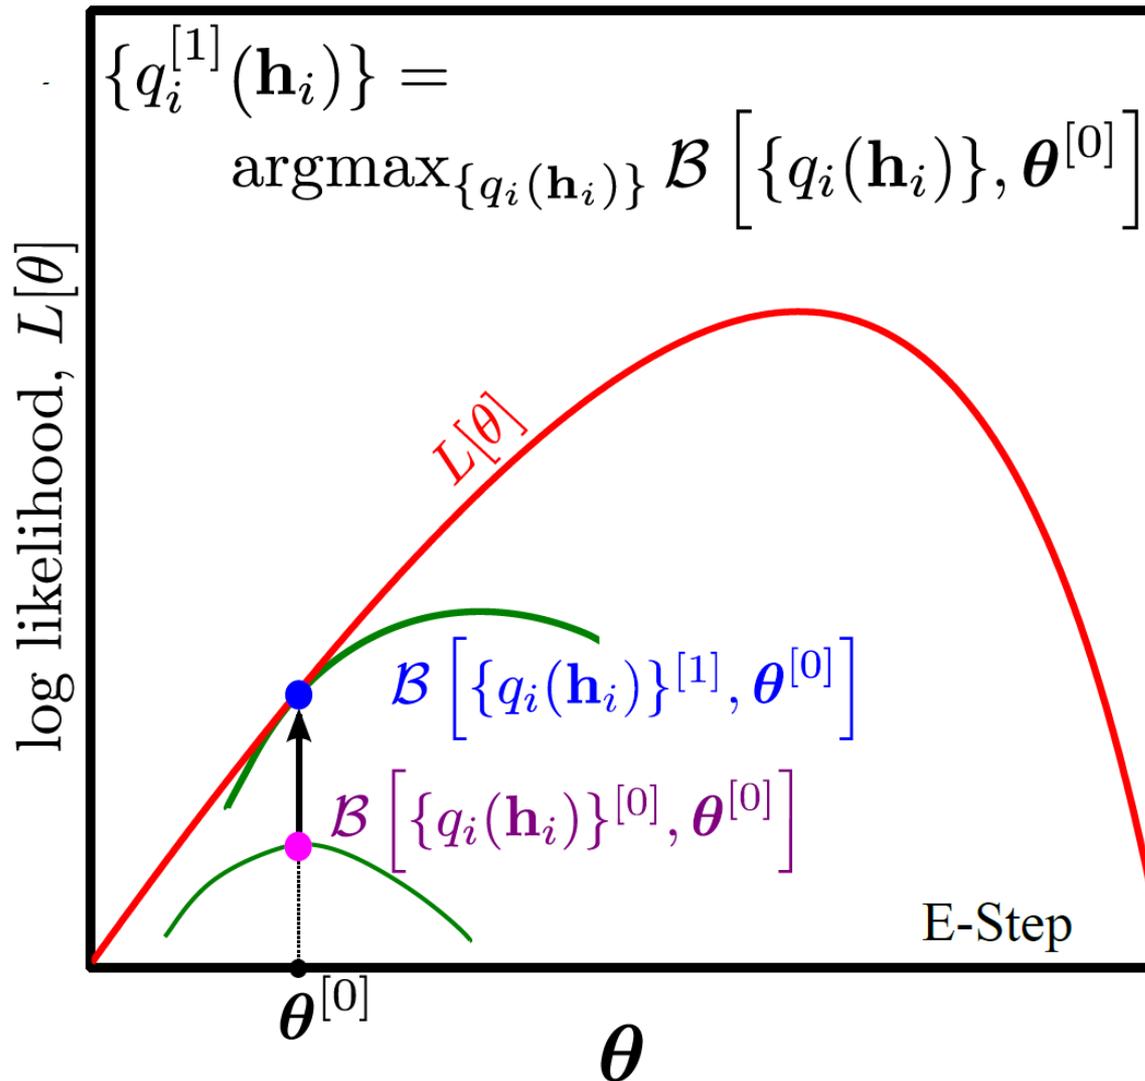
E-Step & M-Step

E-Step – Maximize bound w.r.t. distributions $\{q_i(\mathbf{h}_i)\}$

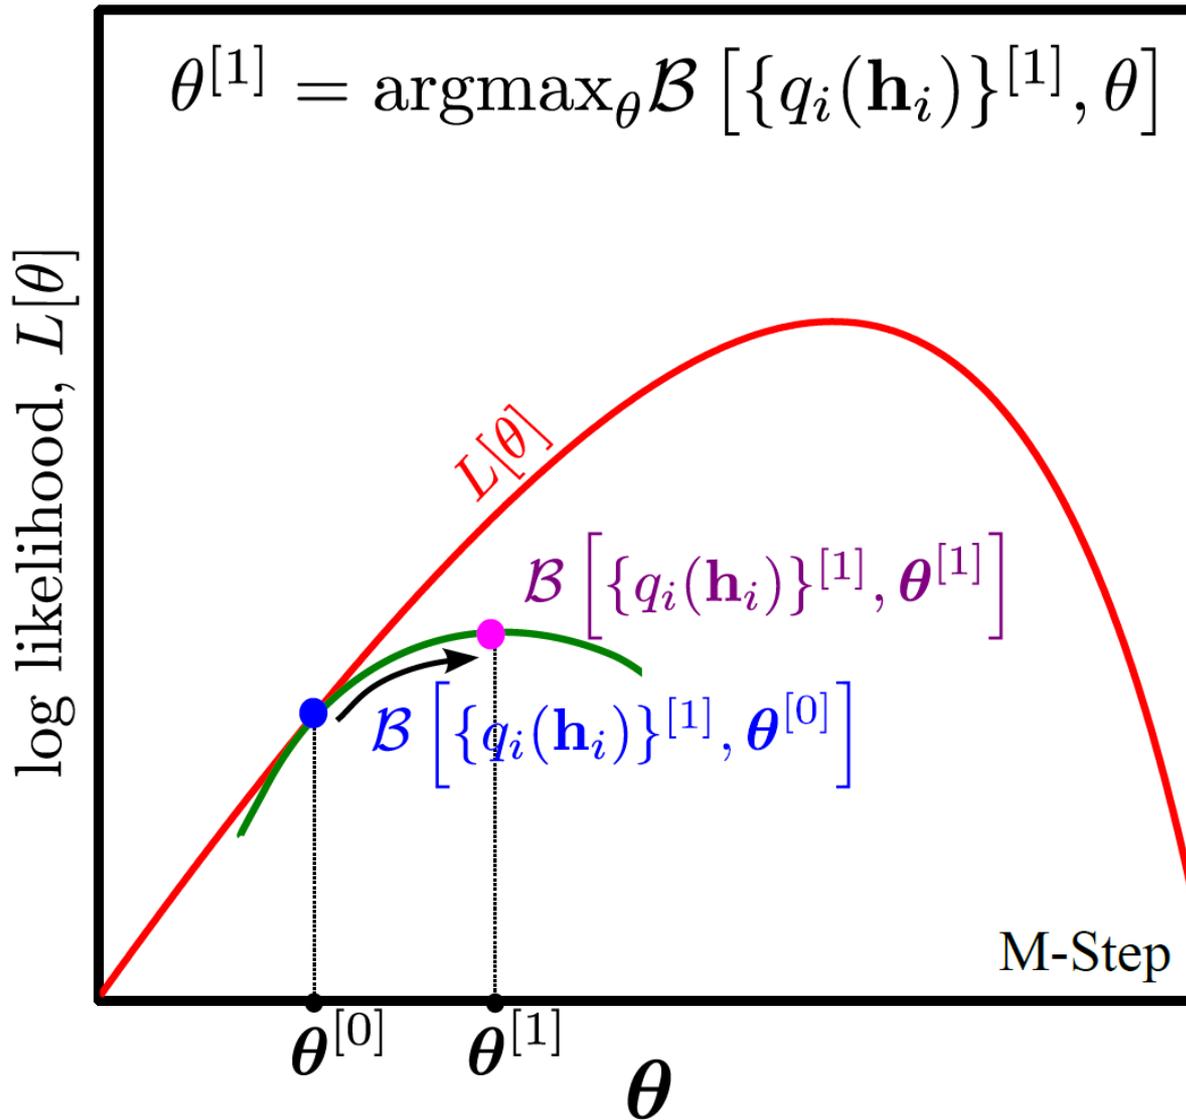
$$q_i^{[t]}[\mathbf{h}_i] = \operatorname{argmax}_{q_i[\mathbf{h}_i]} \left[\mathcal{B}[\{q_i(\mathbf{h}_i)\}, \theta^{[t-1]}] \right]$$

M-Step – Maximize bound w.r.t. parameters θ

$$\theta^{[t]} = \operatorname{argmax}_{\theta} \left[\mathcal{B}[\{q_i^{[t]}(\mathbf{h}_i)\}, \theta] \right]$$

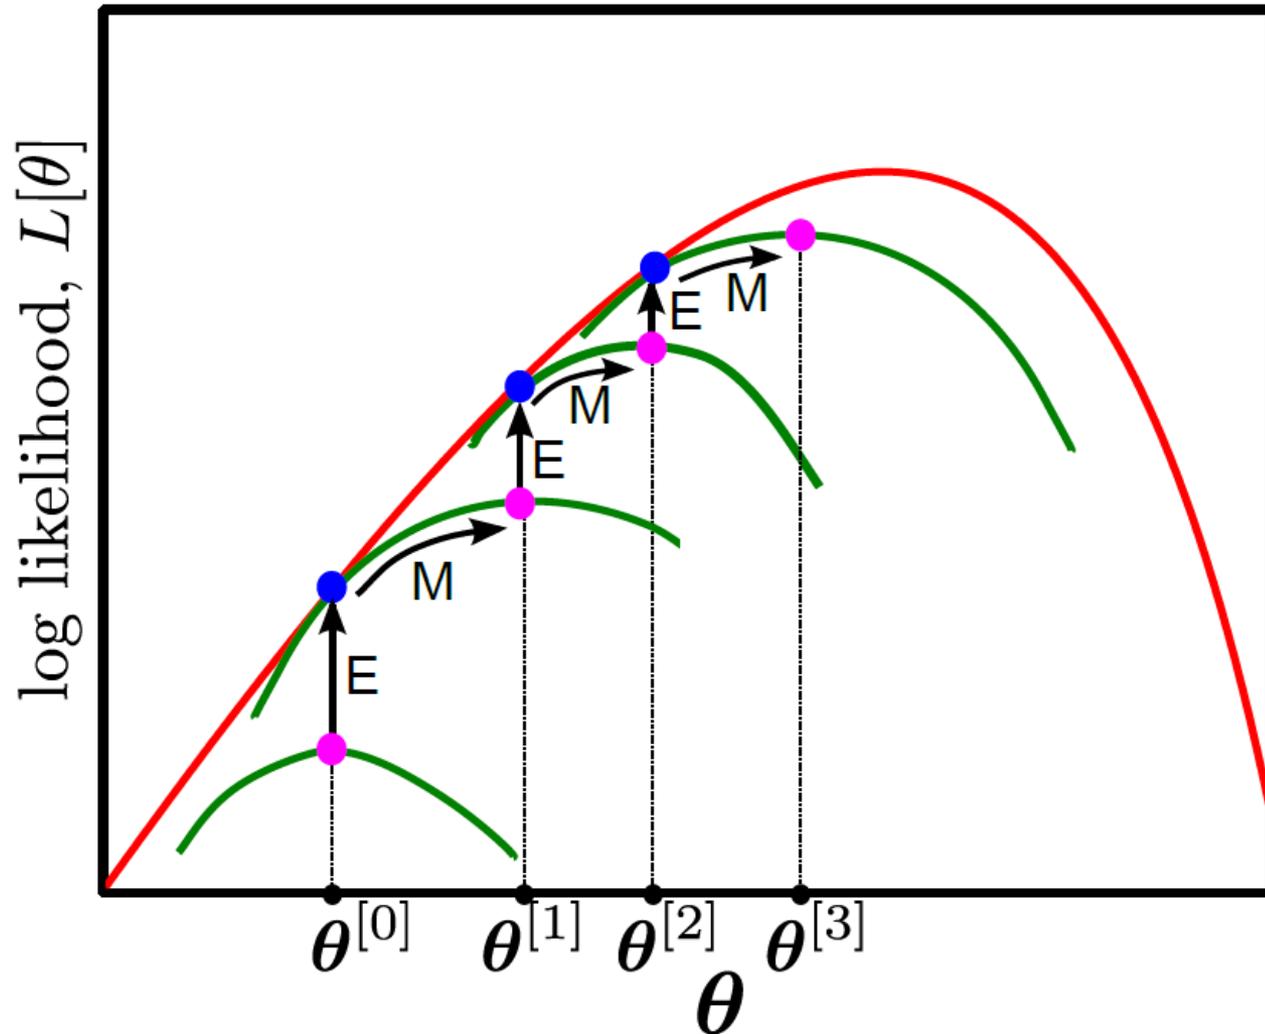


E-Step: Update $\{q_i[\mathbf{h}_i]\}$ so that bound equals log likelihood for this $\boldsymbol{\theta}$



M-Step: Update θ to maximum

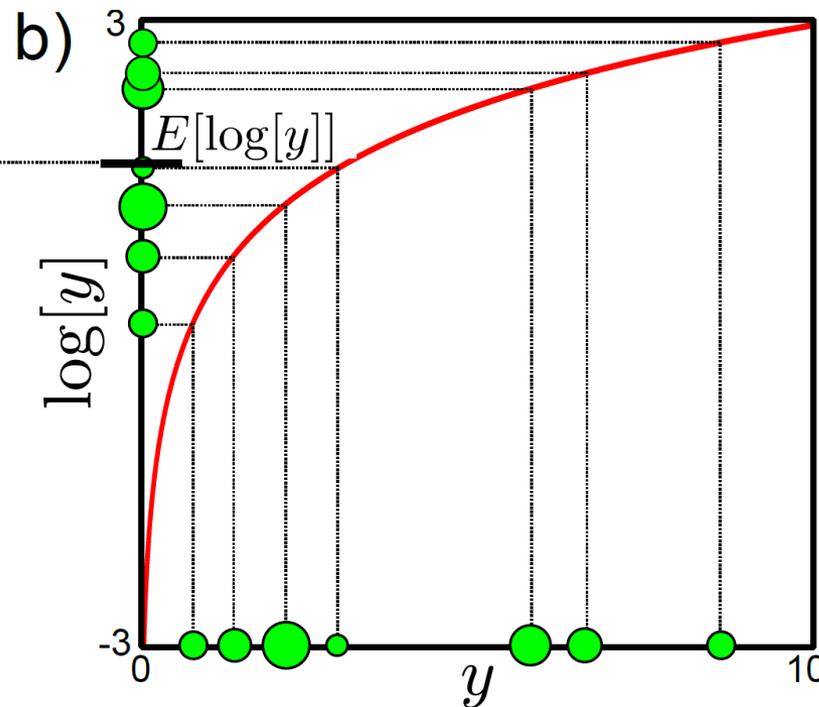
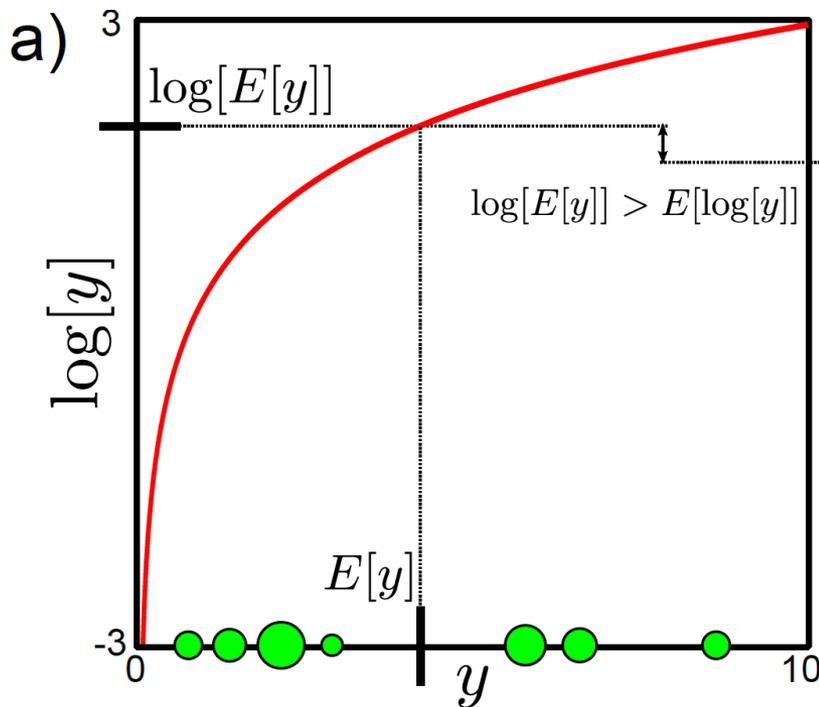
Missing Parts of Argument



Missing Parts of Argument

1. Show that this is a lower bound
2. Show that the E-Step update is correct
3. Show that the M-Step update is correct

Jensen's Inequality



$$E[\log[y]] \leq \log[E[y]] \quad \text{or...}$$

$$\sum_y q(y) \log[y] \leq \log \left[\sum_y q(y) y \right]$$

similarly

$$\int q(y) \log[y] dy \leq \int \log[q(y)] dy$$

Lower Bound for EM Algorithm

$$\begin{aligned}\mathcal{B}[\{q_i(\mathbf{h}_i)\}, \boldsymbol{\theta}] &= \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \\ &\leq \sum_{i=1}^I \log \left[\int q_i(\mathbf{h}_i) \frac{Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} d\mathbf{h}_i \right] \\ &= \sum_{i=1}^I \log \left[\int Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta}) d\mathbf{h}_i \right],\end{aligned}$$

Where we have used Jensen's inequality

$$\int q(y) \log[y] dy \leq \int \log[q(y)] dy$$

E-Step – Optimize bound w.r.t $\{q_i(\mathbf{h}_i)\}$

$$\begin{aligned}\mathcal{B}[\{q_i(\mathbf{h}_i)\}, \boldsymbol{\theta}] &= \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \\ &= \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta}) Pr(\mathbf{x}_i | \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \\ &= \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log [Pr(\mathbf{x}_i | \boldsymbol{\theta})] d\mathbf{h}_i - \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{q_i(\mathbf{h}_i)}{Pr(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta})} \right] d\mathbf{h}_i \\ &= \sum_{i=1}^I \log [Pr(\mathbf{x}_i | \boldsymbol{\theta})] - \sum_{i=1}^I \int q_i(\mathbf{h}_i) \log \left[\frac{q_i(\mathbf{h}_i)}{Pr(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta})} \right] d\mathbf{h}_i \quad (7.46)\end{aligned}$$

Constant w.r.t. $q(\mathbf{h})$

Only this term matters

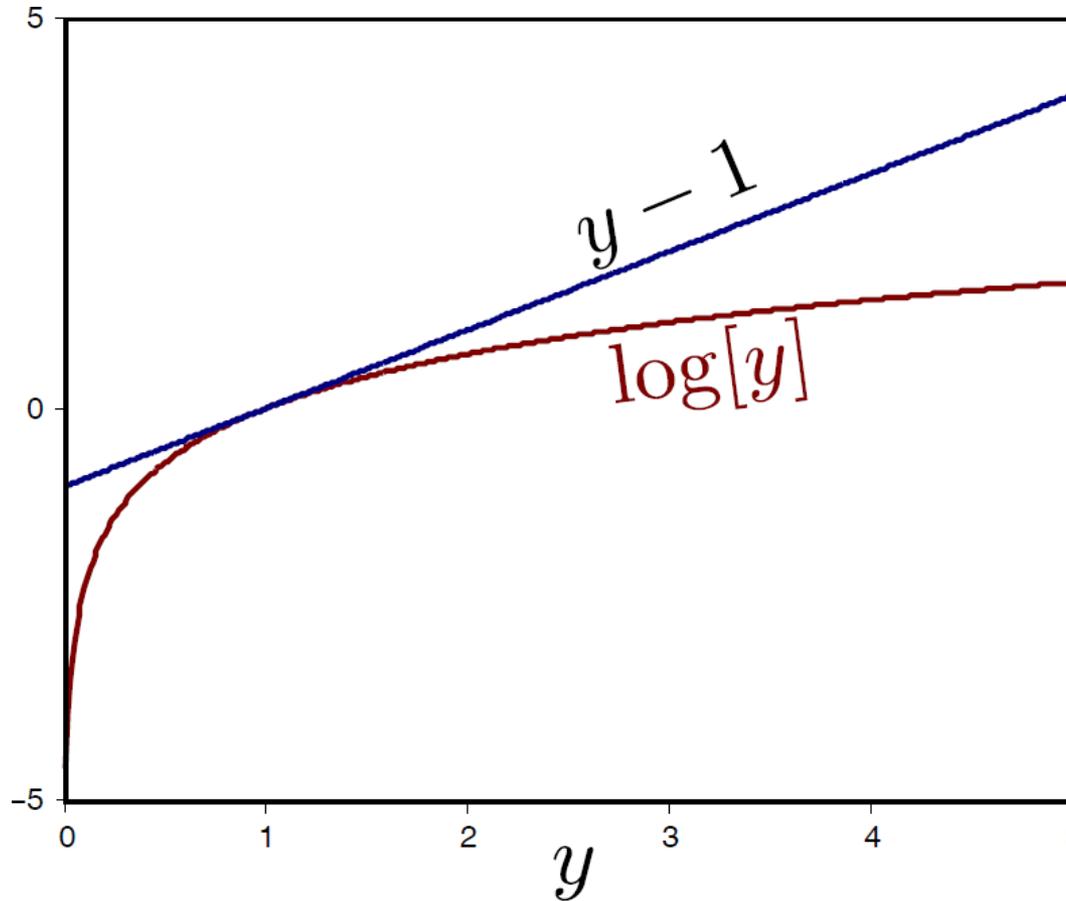
E-Step

$$\begin{aligned}\hat{q}_i(\mathbf{h}_i) &= \operatorname{argmax}_{q_i(\mathbf{h}_i)} \left[- \int q_i(\mathbf{h}_i) \log \left[\frac{q_i(\mathbf{h}_i)}{Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta})} \right] d\mathbf{h}_i \right] \\ &= \operatorname{argmax}_{q_i(\mathbf{h}_i)} \left[\int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \right] \\ &= \operatorname{argmin}_{q_i(\mathbf{h}_i)} \left[- \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \right]\end{aligned}$$



Kullback Leibler divergence – distance between probability distributions . We are maximizing the negative distance (i.e. Minimizing distance)

Use this relation



$$\log[y] \leq y - 1$$

Kullback Leibler Divergence

$$\begin{aligned} \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i &\leq \int q_i(\mathbf{h}_i) \left(\frac{Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} - 1 \right) d\mathbf{h}_i \\ &= \int Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta}) - q_i(\mathbf{h}_i) d\mathbf{h}_i \\ &= 1 - 1 = 0, \end{aligned}$$

So the cost function must be positive

$$\hat{q}_i(\mathbf{h}_i) = \operatorname{argmin}_{q_i(\mathbf{h}_i)} \left[- \int q_i(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{h}_i|\mathbf{x}_i, \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \right]$$

In other words, the best we can do is choose $q_i(\mathbf{h}_i)$ so that this is zero

E-Step

So the cost function must be positive

$$\hat{q}_i(\mathbf{h}_i) := \operatorname{argmin}_{q_i(\mathbf{h}_i)} \left[- \int q_i(\mathbf{h}_i) \log \left[\frac{\operatorname{Pr}(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i \right]$$

The best we can do is choose $q_i(\mathbf{h}_i)$ so that this is zero.

How can we do this? Easy – choose posterior $\operatorname{Pr}(\mathbf{h} | \mathbf{x})$

$$\begin{aligned} \int q_i(\mathbf{h}_i) \log \left[\frac{\operatorname{Pr}(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta})}{q_i(\mathbf{h}_i)} \right] d\mathbf{h}_i &= \int \operatorname{Pr}(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta}) \log \left[\frac{\operatorname{Pr}(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta})}{\operatorname{Pr}(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta})} \right] d\mathbf{h}_i \\ &= \int \operatorname{Pr}(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta}) \log [1] d\mathbf{h}_i = 0. \end{aligned}$$

M-Step – Optimize bound w.r.t. θ

$$\begin{aligned}\theta^{[t]} &= \operatorname{argmax}_{\theta} \left[\mathcal{B}[\{q_i^{[t]}(\mathbf{h}_i)\}, \theta] \right] \\ &= \operatorname{argmax}_{\theta} \left[\sum_{i=1}^I \int q_i^{[t]}(\mathbf{h}_i) \log \left[\frac{Pr(\mathbf{x}_i, \mathbf{h}_i | \theta)}{q_i^{[t]}(\mathbf{h}_i)} \right] d\mathbf{h}_i \right] \\ &= \operatorname{argmax}_{\theta} \left[\sum_{i=1}^I \int q_i^{[t]}(\mathbf{h}_i) \log [Pr(\mathbf{x}_i, \mathbf{h}_i | \theta)] - q_i^{[t]}(\mathbf{h}_i) \log [q_i^{[t]}(\mathbf{h}_i)] d\mathbf{h}_i \right] \\ &= \operatorname{argmax}_{\theta} \left[\sum_{i=1}^I \int q_i^{[t]}(\mathbf{h}_i) \log [Pr(\mathbf{x}_i, \mathbf{h}_i | \theta)] d\mathbf{h}_i \right]\end{aligned}$$

In the M-Step we optimize expected joint log likelihood with respect to parameters θ (Expectation w.r.t distribution from E-Step)

E-Step & M-Step

E-Step – Maximize bound w.r.t. distributions $q_i(\mathbf{h}_i)$

$$\hat{q}_i(\mathbf{h}_i) = Pr(\mathbf{h}_i | \mathbf{x}_i, \boldsymbol{\theta}^{[t]}) = \frac{Pr(\mathbf{x}_i | \mathbf{h}_i, \boldsymbol{\theta}^{[t]}) Pr(\mathbf{h}_i)}{Pr(\mathbf{x}_i)}$$

M-Step – Maximize bound w.r.t. parameters $\boldsymbol{\theta}$

$$\hat{\boldsymbol{\theta}}^{[t+1]} = \operatorname{argmax}_{\boldsymbol{\theta}} \left[\sum_{i=1}^I \int \hat{q}_i(\mathbf{h}_i) \log [Pr(\mathbf{x}_i, \mathbf{h}_i | \boldsymbol{\theta})] d\mathbf{h}_i \right]$$

Structure

- Densities for classification
- Models with hidden variables
 - Mixture of Gaussians
 - t-distributions
 - Factor analysis
- EM algorithm in detail
- Applications

Face Detection

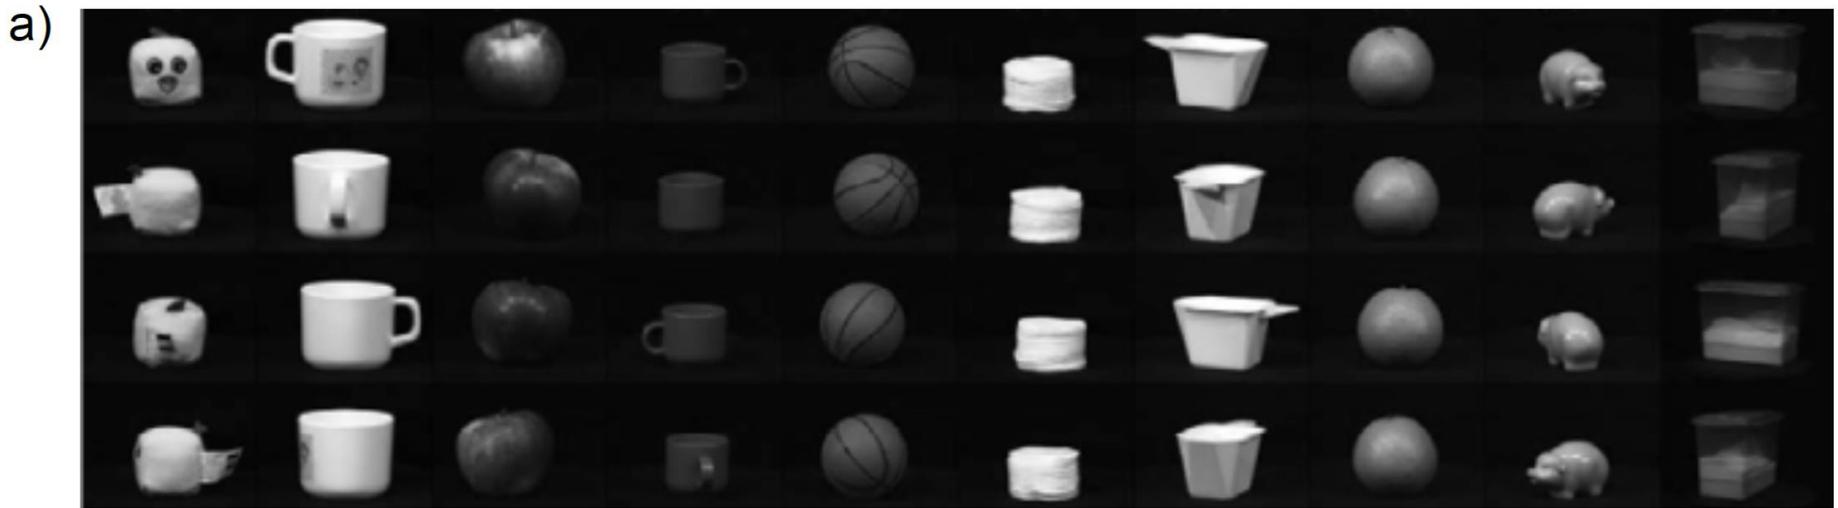


	Color	Grayscale	Equalized
Single Gaussian	76%	79%	80%
Mixture of 10 Gaussians	81%	85%	89%

(not a very realistic model – face detection really done using discriminative methods)

Object recognition

Aeschliman *et al.* (2010)



b)

74	66	66	90	61	97	47	62	43	71
74	25	44	91	43	80	63	24	08	55

t-distributions

normal
distributions

Model as J independent patches

$$Pr(\mathbf{x}_i | w = m) = \prod_{j=1}^J \text{Stud}_{\mathbf{x}_{ij}} [\boldsymbol{\mu}_{jm}, \boldsymbol{\Sigma}_{jm}, \nu_{jm}]$$

Segmentation

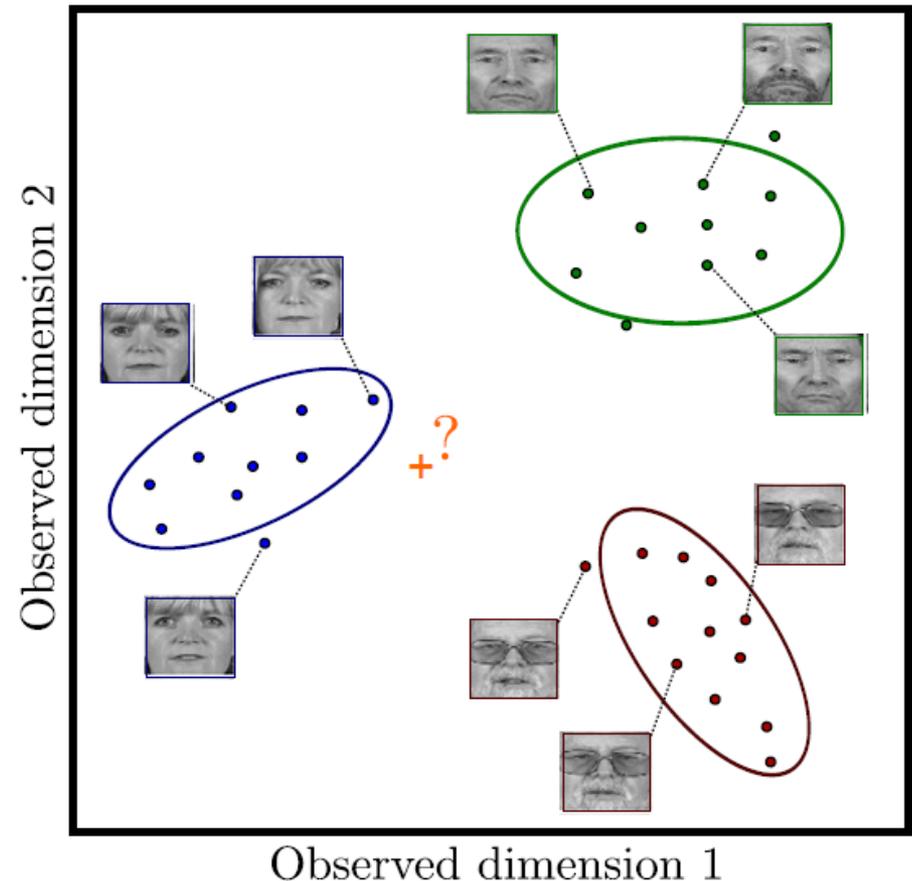


$$Pr(w_n) = \text{Cat}_{w_n}[\boldsymbol{\lambda}] \quad \text{Sfikas et al. (2007) } \text{\textcircled{C}}\text{IEEE 2007}$$

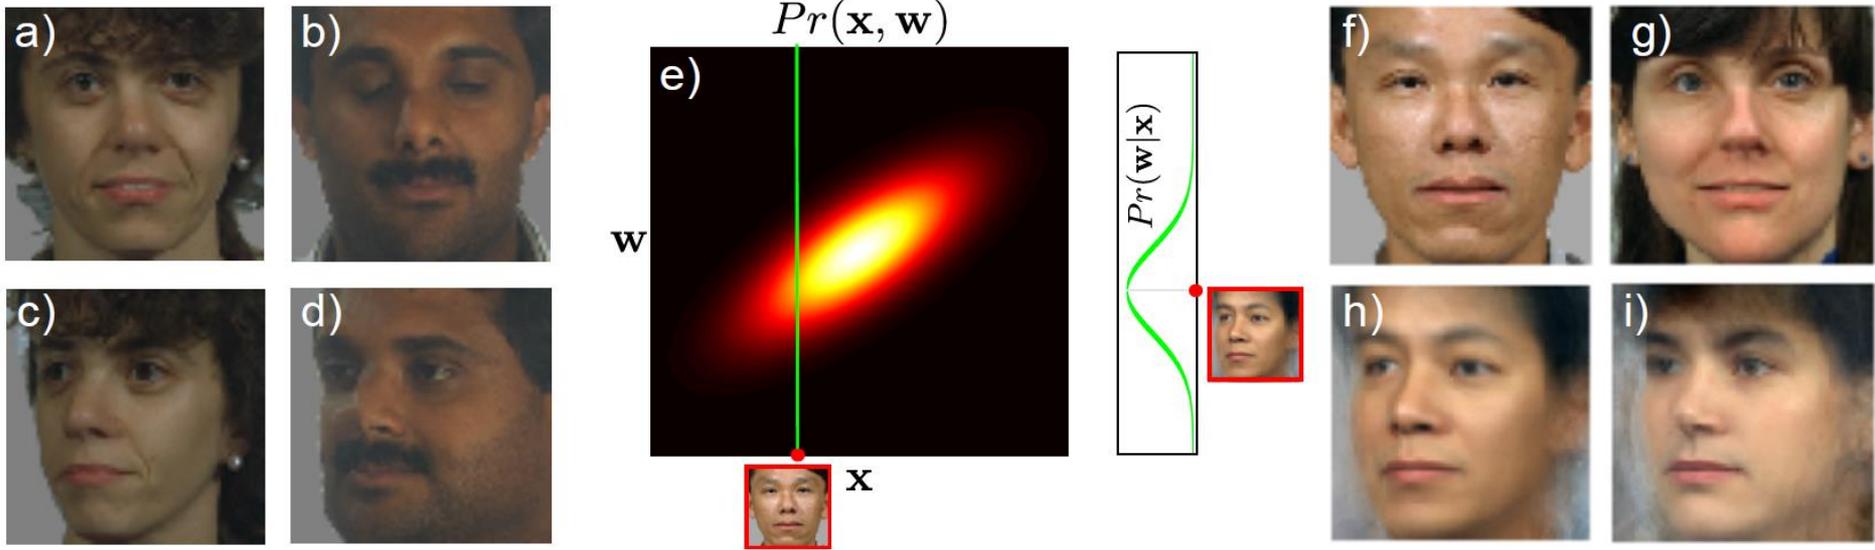
$$Pr(\mathbf{x}_i | w_i = k) = \text{Norm}_{\mathbf{x}_i}[\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k].$$

Face Recognition

Figure 6.28 Face recognition. Our goal is to take the RGB values of a facial image \mathbf{x} and assign a label $y \in \{1 \dots K\}$. Since the data is high dimensional, we model the class conditional density function $Pr(\mathbf{x}|y = k)$ for each individual in the database as a factor analyzer. To classify a new face, we apply Bayes' rule with suitable priors $Pr(y)$ to compute the posterior distribution $Pr(y|\mathbf{x})$. Finally, we choose the label $\hat{y} = \arg \max_k Pr(y = k|\mathbf{x})$ that maximizes the posterior. This is not suitable if there are not plentiful training examples of each individual or where there are significant pose or lighting changes.



Pose Regression



Training data

Learning

Inference

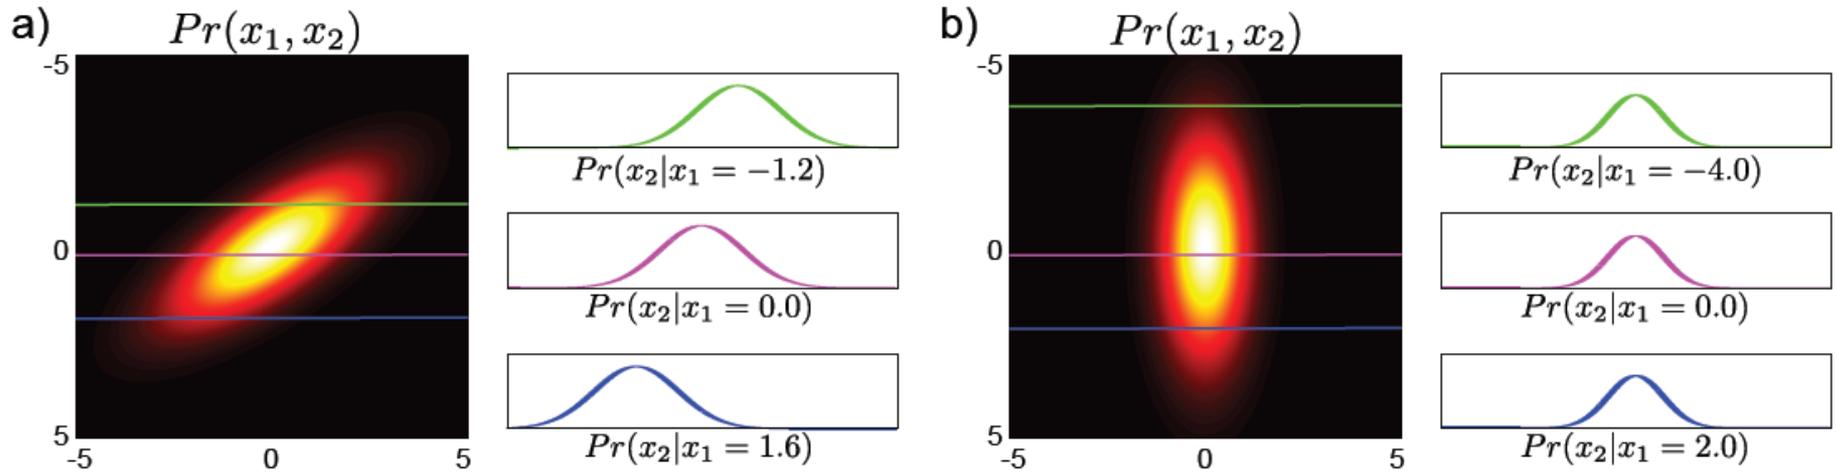
Frontal faces x

Fit factor analysis model to joint distribution of x, w

Given new x , infer w (cond. dist of normal)

Non-frontal faces w

Conditional Distributions



If

$$Pr(\mathbf{x}) = Pr\left(\begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}\right) = \text{Norm}_{\mathbf{x}}\left(\begin{bmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12}^T \\ \boldsymbol{\Sigma}_{12} & \boldsymbol{\Sigma}_{22} \end{bmatrix}\right)$$

then

$$Pr(\mathbf{x}_1|\mathbf{x}_2) = \text{Norm}_{\mathbf{x}_1}\left(\boldsymbol{\mu}_1 + \boldsymbol{\Sigma}_{12}^T \boldsymbol{\Sigma}_{22}^{-1}(\mathbf{x}_2 - \boldsymbol{\mu}_2), \boldsymbol{\Sigma}_{11} - \boldsymbol{\Sigma}_{12}^T \boldsymbol{\Sigma}_{22}^{-1} \boldsymbol{\Sigma}_{12}\right)$$

$$Pr(\mathbf{x}_2|\mathbf{x}_1) = \text{Norm}_{\mathbf{x}_2}\left(\boldsymbol{\mu}_2 + \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{11}^{-1}(\mathbf{x}_1 - \boldsymbol{\mu}_1), \boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{11}^{-1} \boldsymbol{\Sigma}_{12}^T\right)$$

Modeling transformations with hidden variables

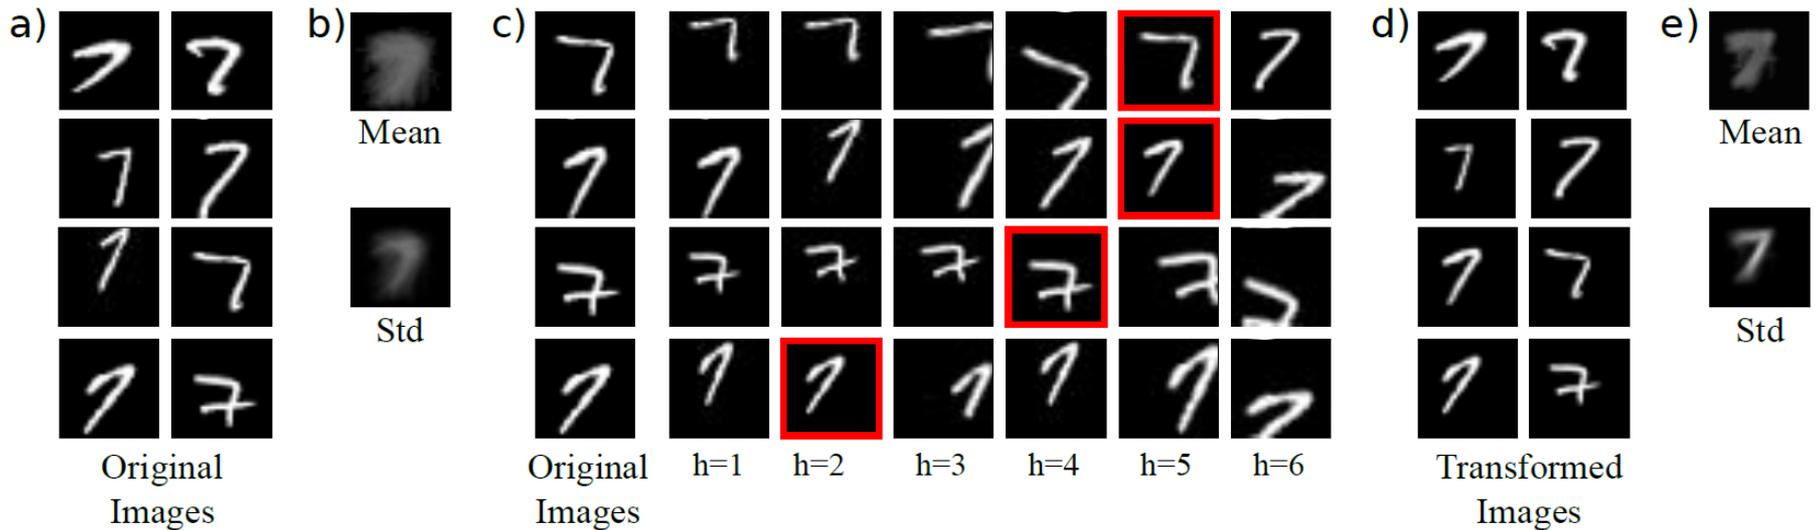


Figure 7.30 Modeling transformations with hidden variables. a) Original set of digit images are only weakly aligned. b) Mean and standard deviation images are consequently blurred out. The probability density model does not fit well. c) Each possible value of a discrete hidden variable represents a different transformation (here inverse transformations are shown). Red square highlights most likely choice of hidden variable after 10 iterations. d) Transformed digits (based on most likely hidden variable). e) New mean and standard deviation images are more focussed: the probability density function fits better.