Module 5: Classification

Linear Methods: Logistic Regression

STAT/BIOSTAT 527, University of Washington
Emily Fox
May 22nd, 2014

Very convenient!

$$p(y=0 \mid x,\beta) = \frac{1}{1+\exp(\beta_0 + \sum_j \beta_j x_j)}$$

implies

$$p(y=1 \mid x,\beta) = \frac{\exp(\beta_0 + \sum_j \beta_j x_j)}{1+\exp(\beta_0 + \sum_j \beta_j x_j)}$$

$1 - P(y=0 \mid x)$

Examine ratio:

$$\frac{p(y=1 \mid x,\beta)}{p(y=0 \mid x,\beta)} = \exp(\beta_0 + \sum_j \beta_j x_j) > 1$$

$\Rightarrow$ class 1 wins, else class 0 (under 0-1 loss)

linear classification rule!

implies

log odds

$$\log \frac{p(y=1 \mid x,\beta)}{p(y=0 \mid x,\beta)} = \beta_0 + \sum_j \beta_j x_j$$

linear

$> 0 \Rightarrow$ class 1 wins, as before

Maximizing Conditional Log Likelihood

$$p(y=0 \mid x,\beta) = \frac{1}{1+\exp(\beta_0 + \sum_j \beta_j x_j)}$$

$$p(y=1 \mid x,\beta) = \frac{\exp(\beta_0 + \sum_j \beta_j x_j)}{1+\exp(\beta_0 + \sum_j \beta_j x_j)}$$

$$l(\beta) = \sum_i \log p(y_i \mid x_i, \beta)$$

$$= \sum_i y_i(\beta_0 + \sum_j \beta_j x_{ij}) - \log(1 + \exp(\beta_0 + \sum_j \beta_j x_{ij}))$$

$x \in \mathbb{R}^d$

fixed in training data

Good news: $l(\boldsymbol{\beta})$ is concave function of $\boldsymbol{\beta}$, no local optima problems

Bad news: no closed-form solution to maximize $l(\boldsymbol{\beta})$

Good news: concave functions easy to optimize

Optimizing Concave Function – Gradient Ascent

- Conditional likelihood for logistic regression is concave
- Find optimum with gradient ascent

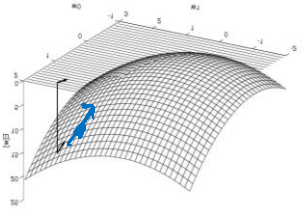

Gradient: $\nabla_\beta l(\beta) = \left[\frac{\partial l(\beta)}{\partial \beta_0}, \ldots, \frac{\partial l(\beta)}{\partial \beta_d}\right]'$

Step size, $\eta > 0$

Update rule: $\Delta\beta = \eta \nabla_\beta l(\beta)$

$$\beta_j^{(t+1)} \leftarrow \beta_j^{(t)} + \eta \frac{\partial l(\beta)}{\partial \beta_j}$$

new est. prev. est. stepsize

- Gradient ascent is simplest of optimization approaches
  - e.g., Conjugate gradient ascent can be much better

Often, esp. proofs, η gets smaller w/ iterations

e.g. $\eta_t = \frac{\alpha}{t}$ ← const.

Gradient Ascent for LR

revisit soon

start w/ $\beta^{(0)}$ (e.g. 0)

Gradient ascent algorithm: iterate until change < ε

$$\beta_0^{(t+1)} \leftarrow \beta_0^{(t)} + \eta \sum_i \left(y_i - \hat{p}(y = 1 \mid x_i, \beta^{(t)}) \right)$$

For j=1,…,d,

$$\beta_j^{(t+1)} \leftarrow \beta_j^{(t)} + \eta \sum_i x_{ij} \left(y_i - \hat{p}(y = 1 \mid x_i, \beta^{(t)}) \right)$$

repeat

Regularization in Linear Regression

- Overfitting usually leads to very large parameter choices, e.g.:

-2.2 + 3.1 X – 0.30 X^2

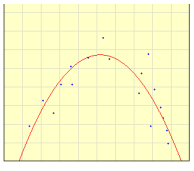

-1.1 + 4,700,910.7 X – 8,585,638.4 X^2 + …

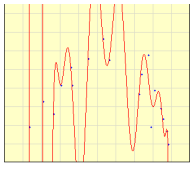

even for n≈p, p large

- ***Regularized*** or ***penalized*** regression aims to impose a "complexity" penalty by penalizing large weights
  - "Shrinkage" method

$$\hat{\beta} = \arg\min_{\beta} \sum_{i=1}^{n} (y_i - (\beta_0 + \beta^T x_i))^2 + \lambda ||\beta||$$

Linear Separability

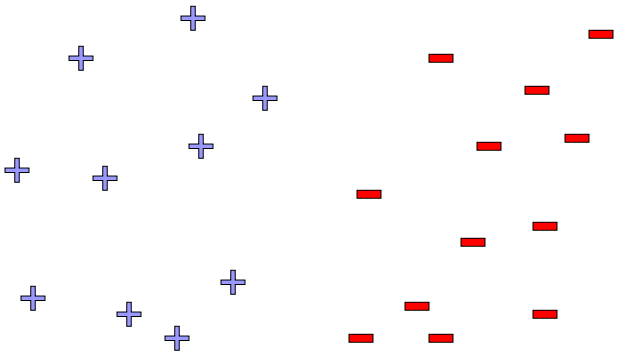

Large Parameters → Overfitting

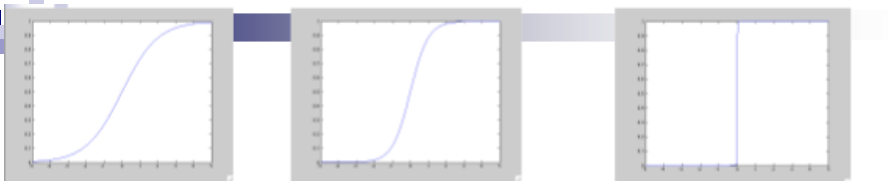

$$\frac{1}{1+e^{-x}} \qquad \frac{1}{1+e^{-2x}} \qquad \frac{1}{1+e^{-100x}}$$

- If data is linearly separable, weights go to infinity
  - In general, leads to overfitting:
- Penalizing high weights can prevent overfitting…

Regularized Conditional Log Likelihood

- Add regularization penalty, e.g., L_2:

$$l(\beta) = \log \prod_{i=1}^{n} p(y_i \mid x_i, \beta) - \frac{\lambda}{2}||\beta||_2^2$$

- Practical note about β_0:

- Gradient of regularized likelihood:

Standard v. Regularized Updates

- Maximum conditional likelihood estimate

$$\hat{\beta} = \arg\max_{\beta} \log \prod_{i=1}^{n} p(y_i \mid x_i, \beta)$$

$$\beta_j^{(t+1)} \leftarrow \beta_j^{(t)} + \eta \sum_i x_{ij} \left(y_i - \hat{p}(y = 1 \mid x_i, \beta^{(t)}) \right)$$

- Regularized maximum conditional likelihood estimate

$$\hat{\beta} = \arg\max_{\beta} \log \prod_{i=1}^{n} p(y_i \mid x_i, \beta) - \frac{\lambda}{2} \sum_{j=1}^{d} \beta_j^2$$

$$\beta_j^{(t+1)} \leftarrow \beta_j^{(t)} + \eta \left\{ -\lambda \beta_j^{(t)} + \sum_i x_{ij} \left(y_i - \hat{p}(y = 1 \mid x_i, \beta^{(t)}) \right) \right\}$$

Stopping Criterion

$$l(\beta) = \log \prod_{i=1}^{n} p(y_i \mid x_i, \beta) - \frac{\lambda}{2}||\beta||_2^2$$

- When do we stop doing gradient ascent?
- Because *l*(**w**) is strongly concave:
  - i.e., because of some technical condition

$$l(\beta^*) - l(\beta) \leq \frac{1}{2\lambda}||\nabla l(\beta)||_2^2$$

- Thus, stop when:

Digression: Logistic Regression for *K* > 2

- Logistic regression in more general case (*K* classes), where *Y* in {1,…,*K*}

Digression: Logistic Regression for $K > 2$

- Logistic regression in more general case, where Y *in* $\{1,\ldots,K\}$

  for $k<K$

$$p(y=k|\mathbf{x},\beta) = \frac{\exp(\beta_{k0} + \sum_{j=1}^{d} \beta_{kj} x_j)}{1 + \sum_{k'=1}^{K-1} \exp(\beta_{k'0} + \sum_{j=1}^{d} \beta_{k'j} x_j)}$$

  for $k=K$ (normalization, so no weights for this class)

$$p(y=K|\mathbf{x},\beta) = \frac{1}{1 + \sum_{k'=1}^{K-1} \exp(\beta_{k'0} + \sum_{j=1}^{d} \beta_{k'j} x_j)}$$

Estimation procedure is basically the same as what we derived!

The Cost, The Cost!!! Think about the cost…

- What's the cost of a gradient update step for LR???

$$\beta_j^{(t+1)} \leftarrow \beta_j^{(t)} + \eta \left\{ -\lambda \beta_j^{(t)} + \sum_i x_{ij} \left(y_i - \hat{p}(y=1 \mid x_i, \beta^{(t)}) \right) \right\}$$

Gradient ascent in Terms of Expectations

- "True" objective function:

$$l(\beta) = E_x[l(\beta, x)] = \int p(x) l(\beta, x) dx$$

- Taking the gradient:

- "True" gradient ascent rule:

- How do we estimate expected gradient?

SGD: Stochastic Gradient Ascent (or Descent)

- "True" gradient: $\nabla l(\beta) = E_x[\nabla l(\beta, x)]$

- Sample based approximation:

- What if we estimate gradient with just one sample???
  - Unbiased estimate of gradient
  - Very noisy!
  - Called stochastic gradient ascent (or descent)
    - Among many other names
  - VERY useful in practice!!!

Stochastic Gradient Ascent for Logistic Regression

- Logistic loss as a stochastic function:

$$E_x[l(\beta, x)] = E_x\left[\log p(y \mid x, \beta) - \frac{\lambda}{2}||\beta||_2^2\right]$$

- Batch gradient ascent updates:

$$\beta_j^{(t+1)} \leftarrow \beta_j^{(t)} + \eta\left\{-\lambda\beta_j^{(t)} + \frac{1}{n}\sum_{i=1}^{n} x_{ij}\left(y_i - \hat{p}(y=1 \mid x_i, \beta^{(t)})\right)\right\}$$

- Stochastic gradient ascent updates:
  - Online setting:

$$\beta_j^{(t+1)} \leftarrow \beta_j^{(t)} + \eta\left\{-\lambda\beta_j^{(t)} + x_{i(t),j}\left(y_{i(t)} - \hat{p}(y=1 \mid x_{i(t)}, \beta^{(t)})\right)\right\}$$

What you should know…

- Classification: predict discrete classes rather than real values
- Logistic regression model: Linear model
  - Logistic function maps real values to [0,1]
- Optimize conditional likelihood
- Gradient computation
- Overfitting
- Regularization
- Regularized optimization
- Cost of gradient step is high, use stochastic gradient descent

Module 5: Classification

Linear Methods: LDA and QDA

STAT/BIOSTAT 527, University of Washington

Emily Fox

May 22nd, 2014

Discriminative vs. Generative

- So far, we have considered modeling/fitting

$$p(Y \mid X)$$

- There are also a large set of ***generative*** methods
- Model:
  - Class-conditional densities $f_k(X) =$
  - Class prior probabilities π_k
- Via Bayes' rule:

Generative Classifiers

$$p(Y = k \mid X = x) = \frac{\pi_k f_k(x)}{\sum_\ell \pi_\ell f_\ell(x)}$$

- Examples include:
  - Linear and quadratic discriminative analysis (LDA and QDA)
  - Mixture of Gaussians (saw in BNP module)
  - Nonparametric density estimation for $f_k(x)$
  - Naïve Bayes

Linear Discriminative Analysis

- Assume Gaussian class-conditional densities

$$f_k(X) =$$

- Furthermore, consider equal covariances
- Log odds

$$\log \frac{p(Y = k \mid X = x)}{p(Y = \ell \mid X = x)} =$$

Linear Discriminative Analysis

$$\log \frac{p(Y=k \mid X=x)}{p(Y=\ell \mid X=x)} = \log \frac{\pi_k}{\pi_\ell} - \frac{1}{2}(\mu_k + \mu_\ell)^T \Sigma^{-1}(\mu_k - \mu_\ell) + x^T \Sigma^{-1}(\mu_k - \mu_\ell)$$

- Equivalently,

$$\log \frac{p(Y=k \mid X=x)}{p(Y=\ell \mid X=x)} = \delta_k(x) - \delta_\ell(x)$$

  where

$$\delta_k(x) =$$

- Decision rule:

- Linear decision boundaries

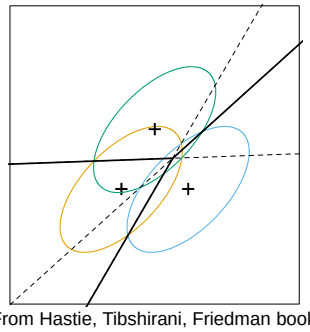

From Hastie, Tibshirani, Friedman book

LDA Parameter Estimation

$$\log \frac{p(Y=k \mid X=x)}{p(Y=\ell \mid X=x)} = \log \frac{\pi_k}{\pi_\ell} - \frac{1}{2}(\mu_k + \mu_\ell)^T \Sigma^{-1}(\mu_k - \mu_\ell) + x^T \Sigma^{-1}(\mu_k - \mu_\ell)$$

- Based on the training class labels, estimate parameters:

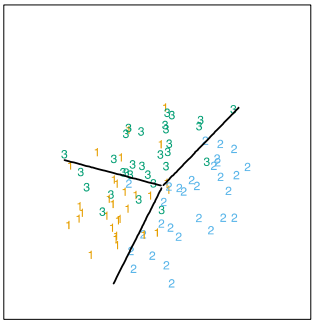

From Hastie, Tibshirani, Friedman book

Quadratic Discriminative Analysis

- Same setup as LDA, but allow class-specific covariances
- Quadratic discriminant functions:

$\delta_k(x) =$

- Quadratic decision boundaries

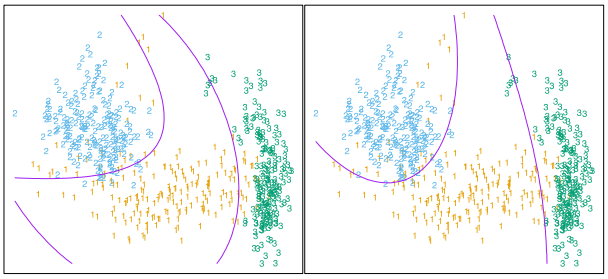

From Hastie, Tibshirani, Friedman book

QDA Parameter Estimation

- Based on the training class labels, estimate parameters:
- Number of parameters:
- Can also consider shrinkage estimators

$\hat{\Sigma}_k(\alpha) = \alpha\hat{\Sigma}_k + (1-\alpha)\hat{\Sigma}$ $\quad$ $\hat{\Sigma}(\gamma) = \gamma\hat{\Sigma} + (1-\gamma)\sigma^2 I$

Notes on QDA and LDA

- LDA + QDA tend to perform very well in practice
- It is not true that data are Gaussian or, furthermore, that covariances are equal (LDA)
- Performance is likely attributed to the fact that the data can only support simple decision boundaries
  - Also, estimates for Gaussian models are stable

LDA vs. Logistic Regression

- Both have linear log odds:

$$\log \frac{p(Y=k \mid X=x)}{p(Y=K \mid X=x)} = \alpha_{k0} + \alpha_k^T x$$

$$\log \frac{p(Y=k \mid X=x)}{p(Y=K \mid X=x)} = \beta_{k0} + \beta_k^T x$$

- Difference is in how the coefficients are estimated

$$p(X, Y=k) =$$

LDA vs. Logistic Regression

$$p(X, Y = k) = p(X)p(Y = k \mid X)$$

- Marginal likelihood term
  - Logistic regression:
  - LDA:

LDA vs. Logistic Regression

- In LDA, the data inform the parameters more
  - If data are indeed Gaussian, then asymptotically maximizing just conditional likelihood requires 30% more data to perform as well
- Data far from boundary affect $\sum$ in LDA, but are ignored by logistic regression
- Observations without class labels can be used in mixture model case, but not in logistic regression
- Marginal likelihood p(*X*) acts as a regularizer
- Logistic regression tends to be more robust than LDA and can handle qualitative *X* variables, but performance is often similar.

Module 5: Classification

Nonparametric Methods: KDE and Naïve Bayes

STAT/BIOSTAT 527, University of Washington
Emily Fox
May 22nd, 2014

KDE for Classification

$$p(Y = k \mid X = x) = \frac{\pi_k f_k(x)}{\sum_\ell \pi_\ell f_\ell(x)}$$

- Use KDE to estimate class-conditional densities
- Recall commonly used Gaussian KDE in 1D

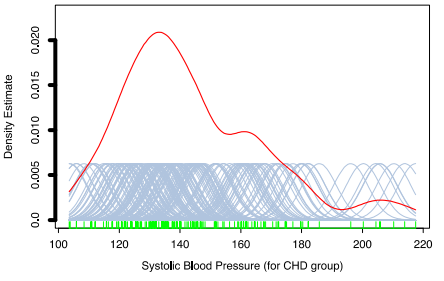

From Hastie, Tibshirani, Friedman book

Example: Heart Disease Data

- Binary response = CHD (coronary heart disease)
- Predictor = systolic blood pressure

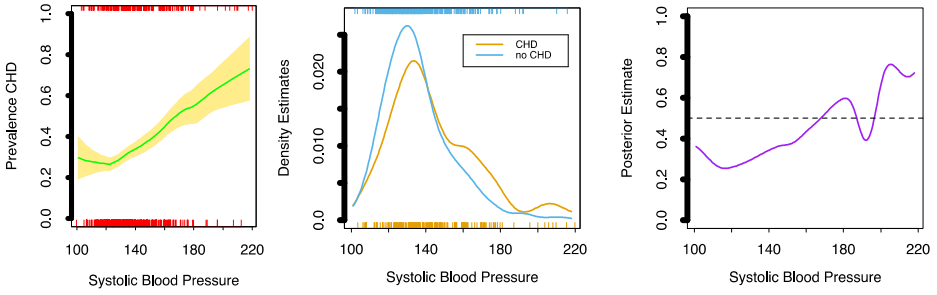

From Hastie, Tibshirani, Friedman book

Example: Heart Disease Data

From Hastie, Tibshirani, Friedman book

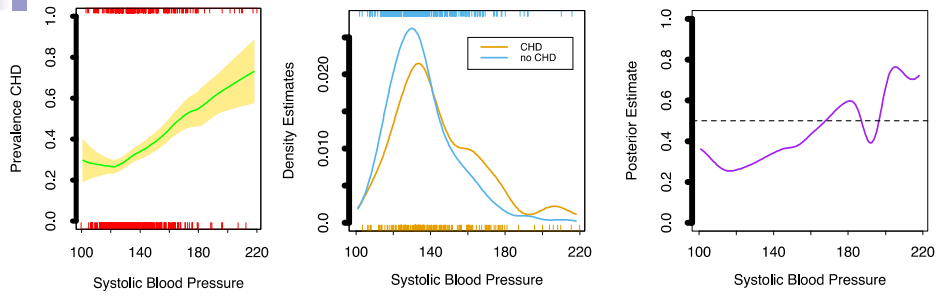

- KDE estimates are poor in regions with little data
- Local linear model uses variable bandwidth based on k-NN → smooths out over these regions
- For classification tasks, do not need to estimate each class-conditional density well. Just need good estimates of the posterior near the decision boundary

Class-Conditionals vs. Posterior

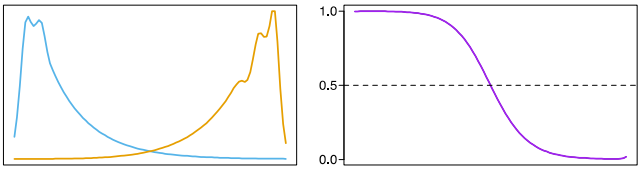

- Example:
  - Both densities are multimodal
  - Might opt for rougher, high-variance estimator to capture features
  - However, posterior is quite smooth
  - Fine-scale features are irrelevant for classification here

Multivariate KDE

- In 1d $$\hat{p}(x_0) = \frac{1}{n\lambda} \sum_{i=1}^{n} K_\lambda(x_0, x_i)$$
- In R^d, assuming a product kernel, $x \in \mathbb{R}^d$

$$\hat{p}(x_0) = \frac{1}{n\lambda_1 \cdots \lambda_d} \sum_{i=1}^{n} \left\{ \prod_{j=1}^{d} K_{\lambda_j}(x_{0j}, x_{ij}) \right\}$$

lots of params to choose

- Typical choice = Gaussian RBF → Gaussian KDE

$$e^{-\frac{\|x_0 - x\|}{\lambda}}$$

Naïve Bayes Classifier

$$p(Y = k \mid X = x) = \frac{\pi_k f_k(x)}{\sum_\ell \pi_\ell f_\ell(x)}$$

- Useful in high-dimensional settings (*d* large)
- Assumes factored form for class-conditional densities

$$f_k(X) =$$

- Benefits:
  - Estimate $f_{kj}(X_j)$ separately for each *j* using only 1D KDE
  - If X_j of *X* is discrete, then can combine using a histogram estimate

Naïve Bayes Classifier

$$p(Y = k \mid X = x) = \frac{\pi_k \prod_j f_{kj}(x_j)}{\sum_\ell \pi_\ell \prod_j f_{\ell j}(x_j)}$$

- Log odds

$$\log \frac{p(Y = k \mid X = x)}{p(Y = \ell \mid X = x)} =$$

- Has form of GAM, but fit very differently
  - Analogous to difference between LDA and logistic regression

Module 5: Classification

Mixture Models for Classification

STAT/BIOSTAT 527, University of Washington
Emily Fox
May 22nd, 2014

Density as Mixture of Gaussians

- Approximate density with a mixture of Gaussians

Mixture of 3 Gaussians

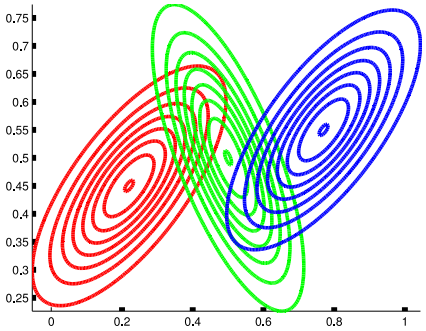

K = 3

$$p(x_i \mid \pi, \mu, \Sigma) = \sum_{k=1}^{K} \pi_k N(x_i \mid \mu_k, \Sigma_k)$$

Gauss. kernel, just like in KDE, but not centered at obs.

K: # of mix comp.

π_k: mix. weights

μ_k, Σ_k: shape params

In 1D:

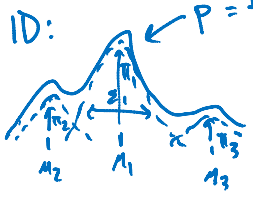

$\Sigma \pi_k = 1$

Clustering our Observations

- Imagine we have an assignment of each x_i to a Gaussian

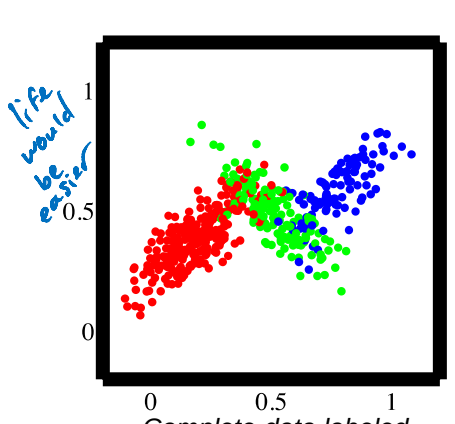

Complete data labeled by true cluster assignments

Our actual observations

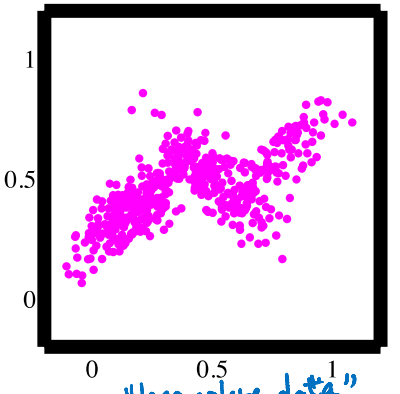

C. Bishop, Pattern Recognition & Machine Learning

Clustering our Observations

- Imagine we have an assignment of each x_i to a Gaussian

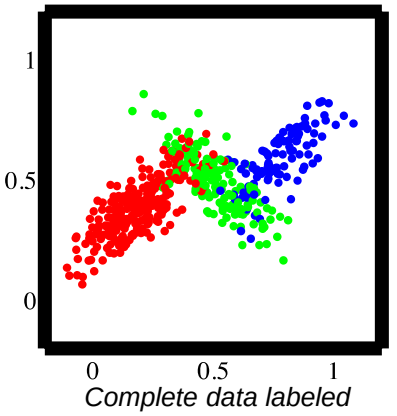

Complete data labeled by true cluster assignments

- Introduce latent cluster indicator variable z_i

$z_i \in \{1, \ldots, K\}$

$\Pr(z_i = k) = \pi_k$

- Then we have

$$p(x_i \mid z_i^k, \pi, \mu, \Sigma) = N(x_i \mid \mu_{z_i^k}, \Sigma_{z_i^k})$$

param. est. is easy if we have $\{z_i\}$ ⇒ decoupled into K Gauss. est.

C. Bishop, Pattern Recognition & Machine Learning

Clustering our Observations

- We must infer the cluster assignments from the observations

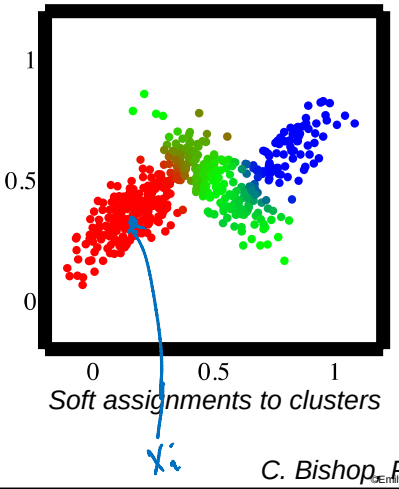

Soft assignments to clusters

- Posterior probabilities of assignments to each cluster *given* model parameters:

$$r_{ik} = p(z_i = k \mid x_i, \pi, \theta) =$$

$$= \frac{\pi_k N(x_i \mid \mu_k, \Sigma_k)}{\sum_j \pi_j N(x_i \mid \mu_j, \Sigma_j)}$$

motivates an iterative alg.

C. Bishop, Pattern Recognition & Machine Learning

Mixture Models for Classification

- Can use mixture models as a generative classifier in the unsupervised setting

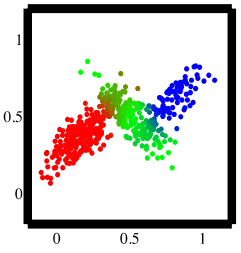

- EM algorithm = iteratively:
  - Estimate responsibilities given parameter estimates

$$\hat{r}_{ik} = \frac{\hat{\pi}_k N(x_i, \hat{\mu}_k, \hat{\Sigma}_k)}{\sum_\ell \hat{\pi}_\ell N(x_i, \hat{\mu}_\ell, \hat{\Sigma}_\ell)}$$

  - Maximize parameters given responsibilities

- For classification, threshold the estimated responsibilities
  - E.g., $\hat{g}(x_i) = \arg\max_k \hat{r}_{ik}$

- Note: allows non-linear boundaries as in QDA

Example: Heart Disease Data

- Binary response = CHD (coronary heart disease)
- Predictor = systolic blood pressure

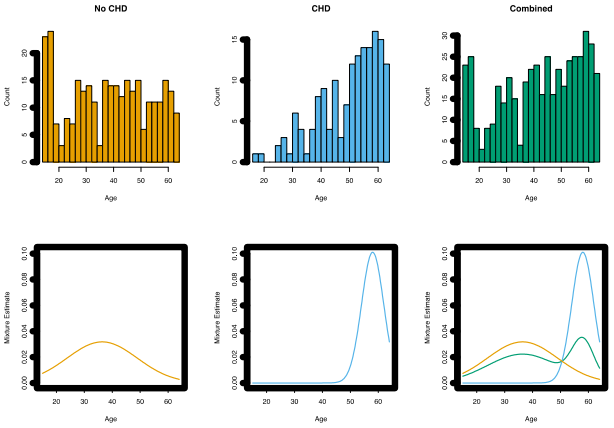

From Hastie, Tibshirani, Friedman book

What you need to know

- Discriminative vs. Generative classifiers
- LDA and QDA assume Gaussian class-conditional densities
  - Results in linear and quadratic decision boundaries, respectively
- KDE for classification
  - Challenging in areas with little data or in high dimensions
  - Estimating class-conditionals is not optimizing classification objective
- Naïve Bayes assumes factored form
  - Results in log odds that have GAM form
- Mixture models allow for unsupervised generative approach

Readings

- Hastie, Tibshirani, Friedman – 4.3, 4.4.5, 6.6.2-6.6.3, 6.8