Mixture of Gaussians

CS229: Machine Learning

Carlos Guestrin
Stanford University

©2022 Carlos Guestrin

Failure modes of k-means

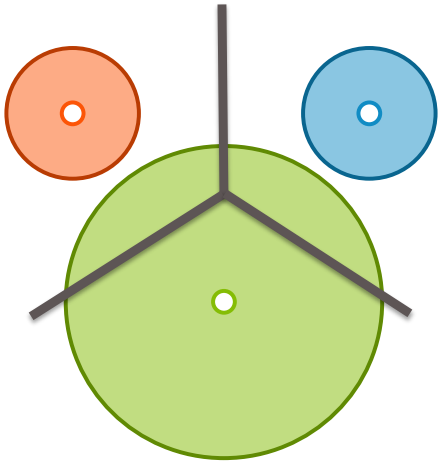

disparate cluster sizes

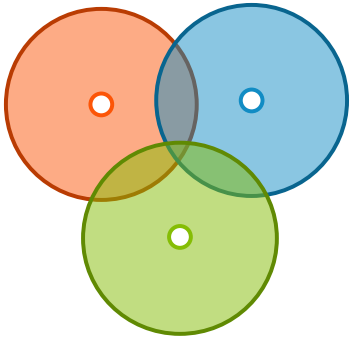

overlapping clusters

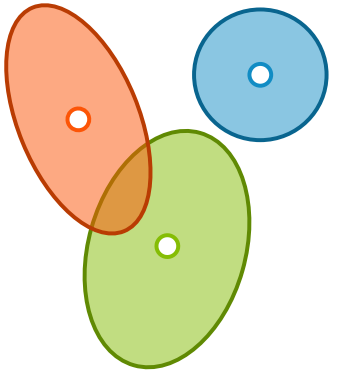

different shaped/oriented clusters

©2022 Carlos Guestrin

(One) bad case for k-means

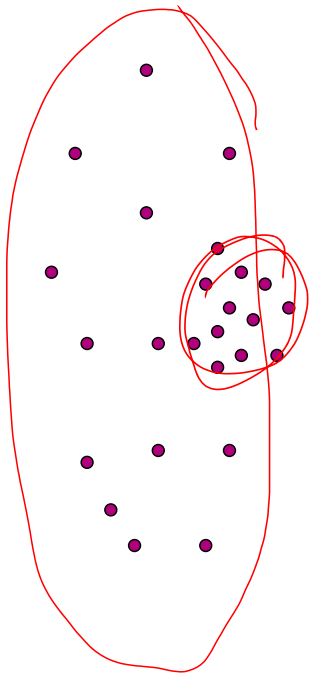

- Clusters may overlap
- Some clusters may be "wider" than others

©2022 Carlos Guestrin

Gaussians in *m* Dimensions

$$P(\mathbf{x}) = \frac{1}{(2\pi)^{m/2}\,|\Sigma|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{x}-\mu)^T \Sigma^{-1}(\mathbf{x}-\mu)\right] = N(\mu, \Sigma)$$

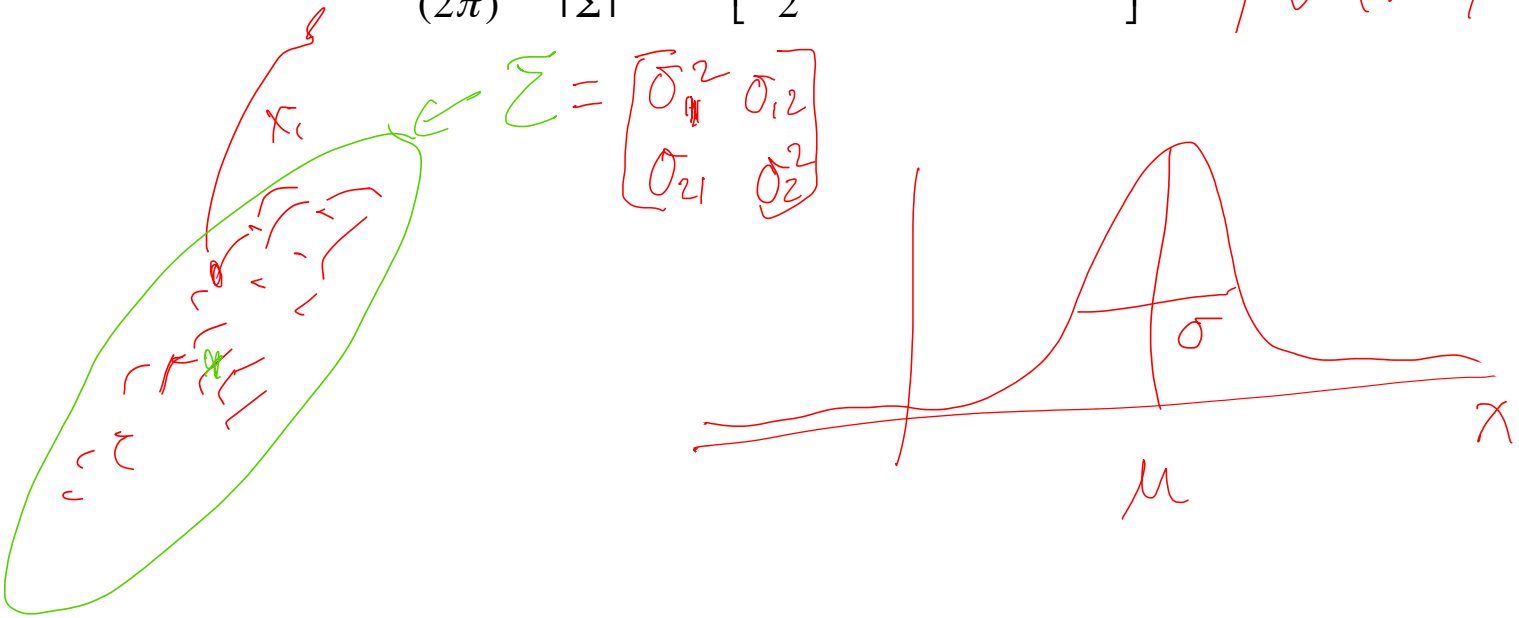

©2022 Carlos Guestrin

Suppose You Have a Gaussian For Each Class

$$p(x|y=i) = \frac{1}{(2\pi)^{m/2}\,|\Sigma_i|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{x}-\mu_i)^T \Sigma_i^{-1}(\mathbf{x}-\mu_i)\right]$$

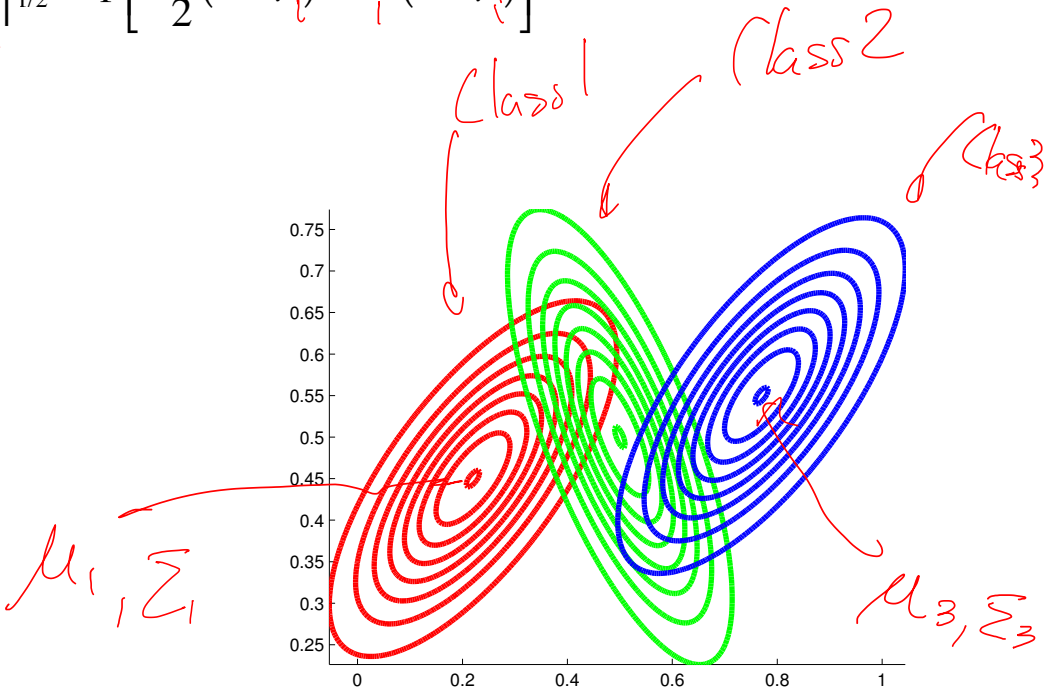

©2022 Carlos Guestrin

Gaussian Bayes Classifier

- You have a Gaussian over **x** for each class y=i:

$$P(\mathbf{x} \mid y=i) = \frac{1}{(2\pi)^{m/2}\,|\Sigma_i|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{x}-\mu_i)^T \Sigma_i^{-1}(\mathbf{x}-\mu_i)\right] = N(\mathbf{x} \mid \mu_i, \Sigma_i)$$

- But you need probability of class y=i given **x**:

$$\text{classification}: \quad \hat{y}_i = \operatorname{argmax}_i P(y=i \mid \mathbf{x})$$

- Thank you Bayes Rule!!

$$P(y=i \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid y=i)P(y=i)}{p(\mathbf{x})}$$

$N(\mu_i, \Sigma_i)$ ← prior prob. per class

normalize

©2022 Carlos Guestrin

Learning modelyear , mpg ---> maker

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} & \cdots & \sigma_{1m} \\ \sigma_{12} & \sigma_2^2 & \cdots & \sigma_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{1m} & \sigma_{2m} & \cdots & \sigma_m^2 \end{pmatrix}$$

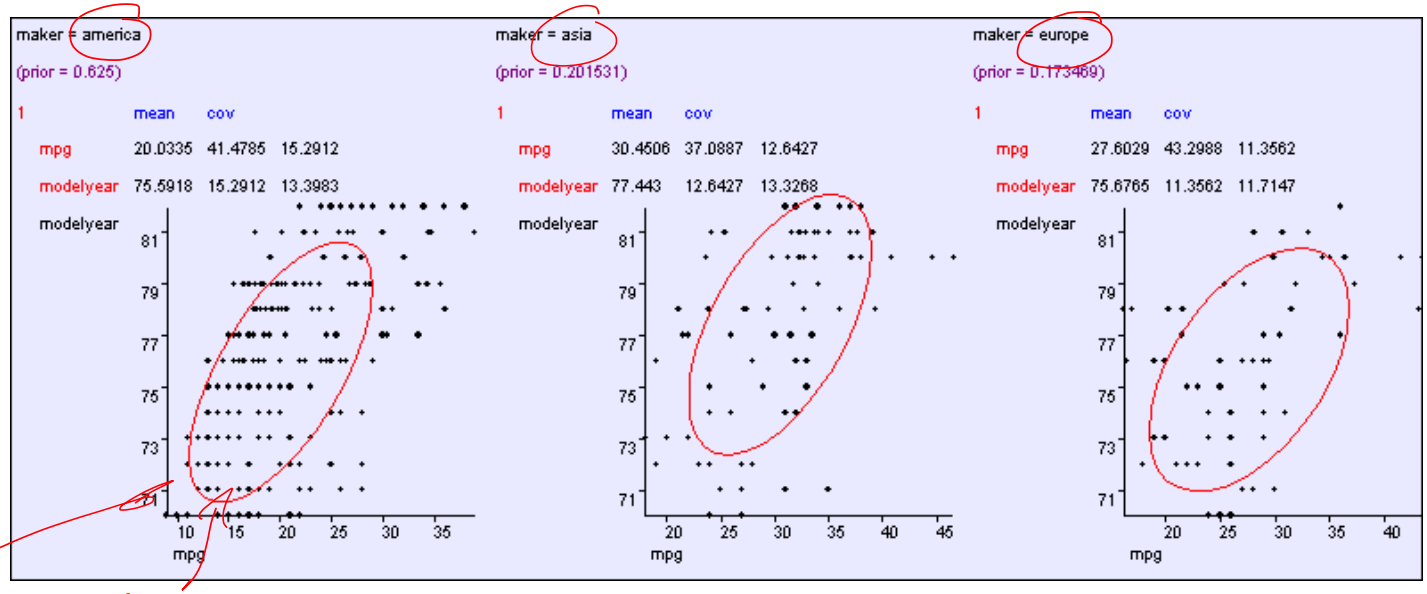

$P(x \mid y = \text{america}) \; N(\mu_{\text{America}}, \Sigma_{\text{America}})$

General: $O(m^2)$ parameters

many parameters, need a lot of data to learn

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} & \cdots & \sigma_{1m} \\ \sigma_{12} & \sigma_2^2 & \cdots & \sigma_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{1m} & \sigma_{2m} & \cdots & \sigma_m^2 \end{pmatrix}$$

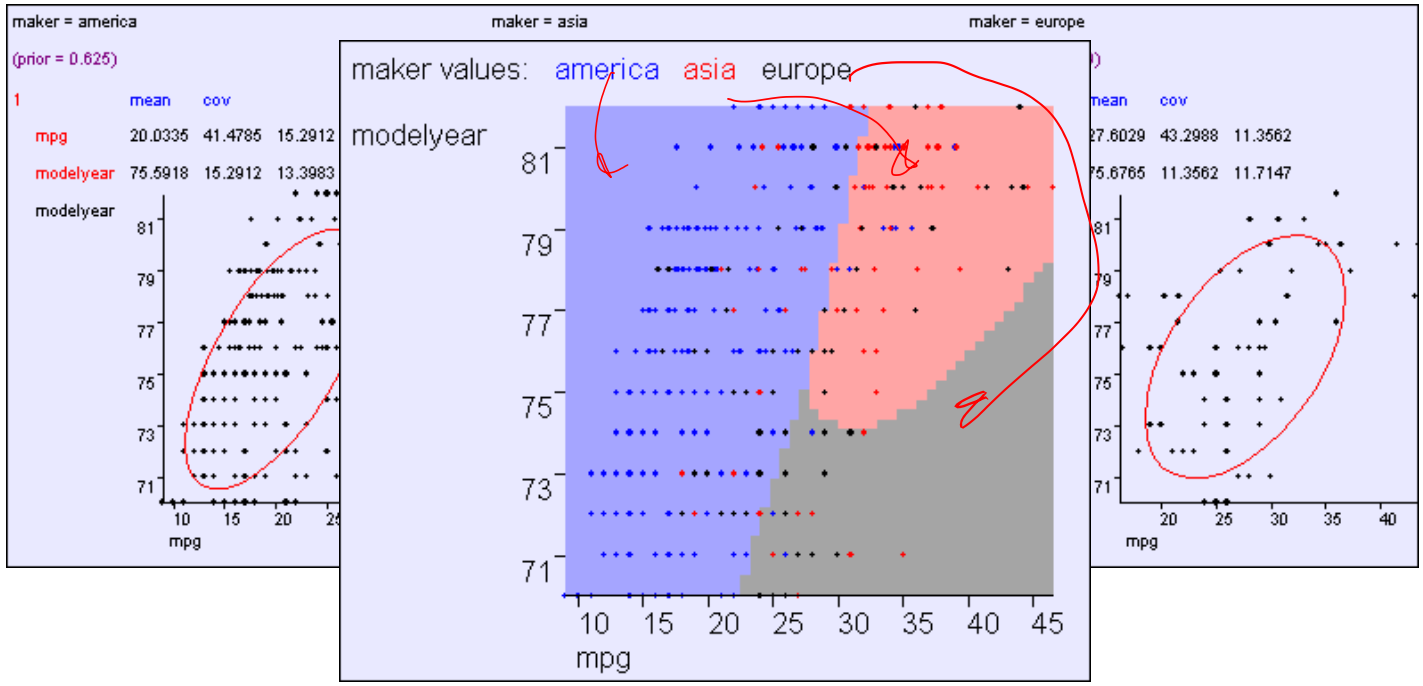

Aligned: *O(m)* parameters

$$\Sigma = \begin{pmatrix} \sigma_1^2 & 0 & 0 & \cdots & 0 & 0 \\ 0 & \sigma_2^2 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \sigma_3^2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \sigma_{m-1}^2 & 0 \\ 0 & 0 & 0 & \cdots & 0 & \sigma_m^2 \end{pmatrix}$$

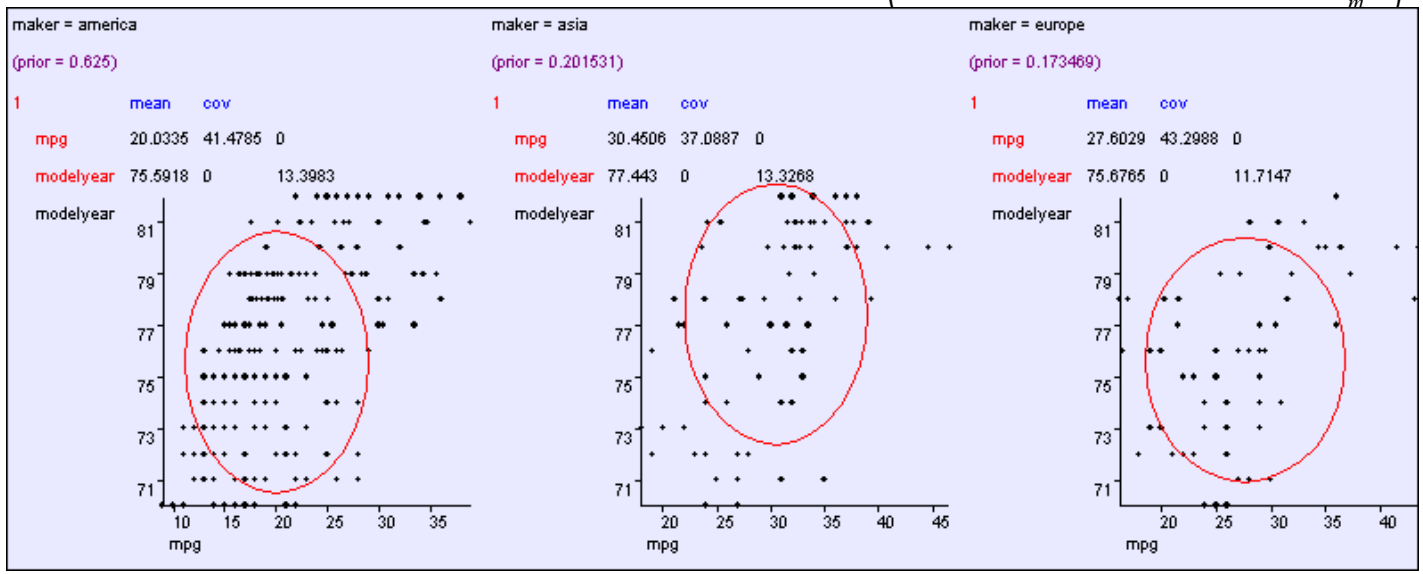

Aligned: *O(m)* parameters

$$\Sigma = \begin{pmatrix} \sigma_1^2 & 0 & 0 & \cdots & 0 & 0 \\ 0 & \sigma_2^2 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \sigma_3^2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \sigma_{m-1}^2 & 0 \\ 0 & 0 & 0 & \cdots & 0 & \sigma_m^2 \end{pmatrix}$$

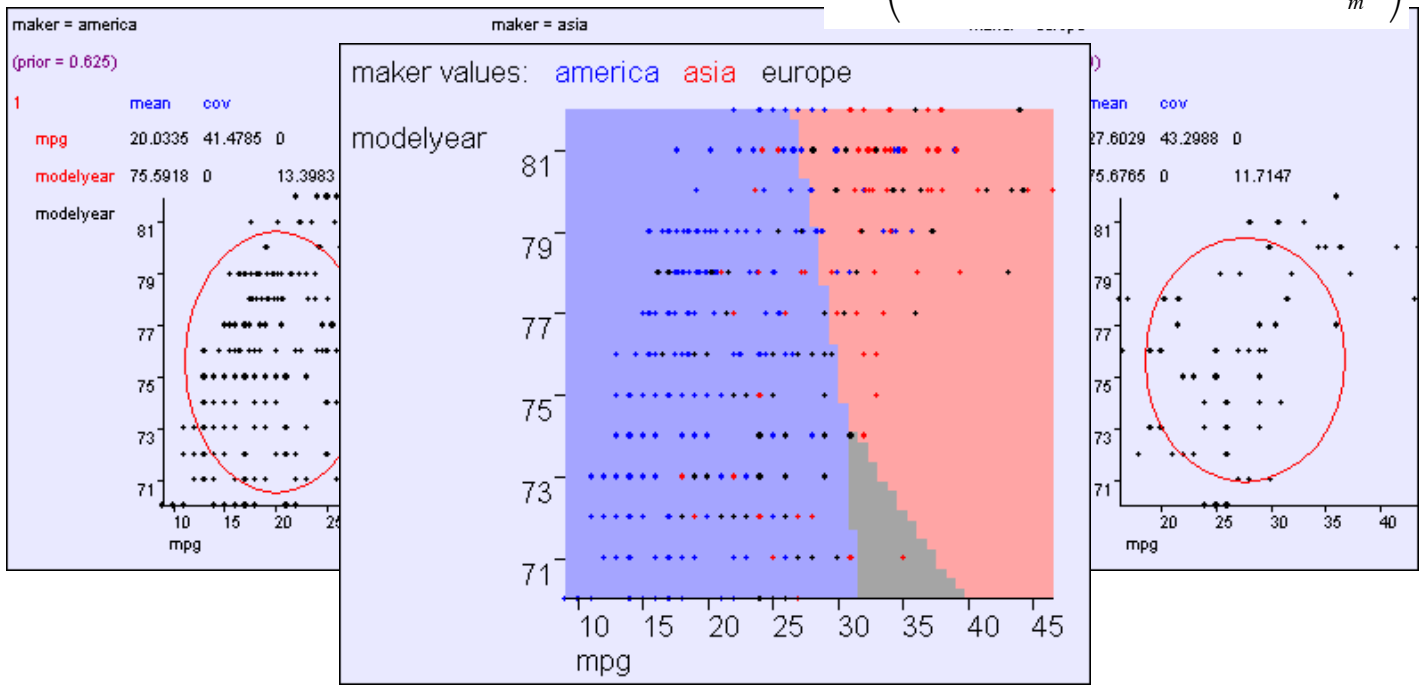

Spherical: *O(1)* cov parameters

$$\Sigma = \begin{pmatrix} \sigma^2 & 0 & 0 & \cdots & 0 & 0 \\ 0 & \sigma^2 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \sigma^2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \sigma^2 & 0 \\ 0 & 0 & 0 & \cdots & 0 & \sigma^2 \end{pmatrix}$$

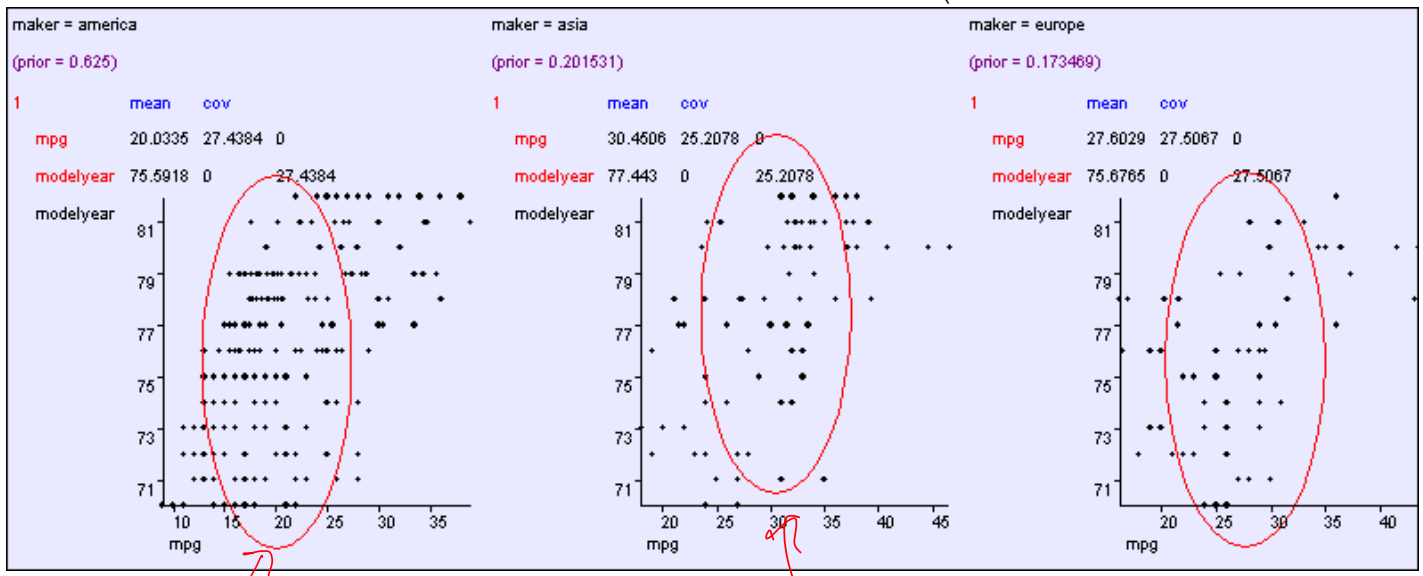

$N(\mu_1, \sigma^2)$ $N(\mu_2, \sigma^2)$

Spherical: *O(1)* cov parameters

$$\Sigma = \begin{pmatrix} \sigma^2 & 0 & 0 & \cdots & 0 & 0 \\ 0 & \sigma^2 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \sigma^2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \sigma^2 & 0 \\ 0 & 0 & 0 & \cdots & 0 & \sigma^2 \end{pmatrix}$$

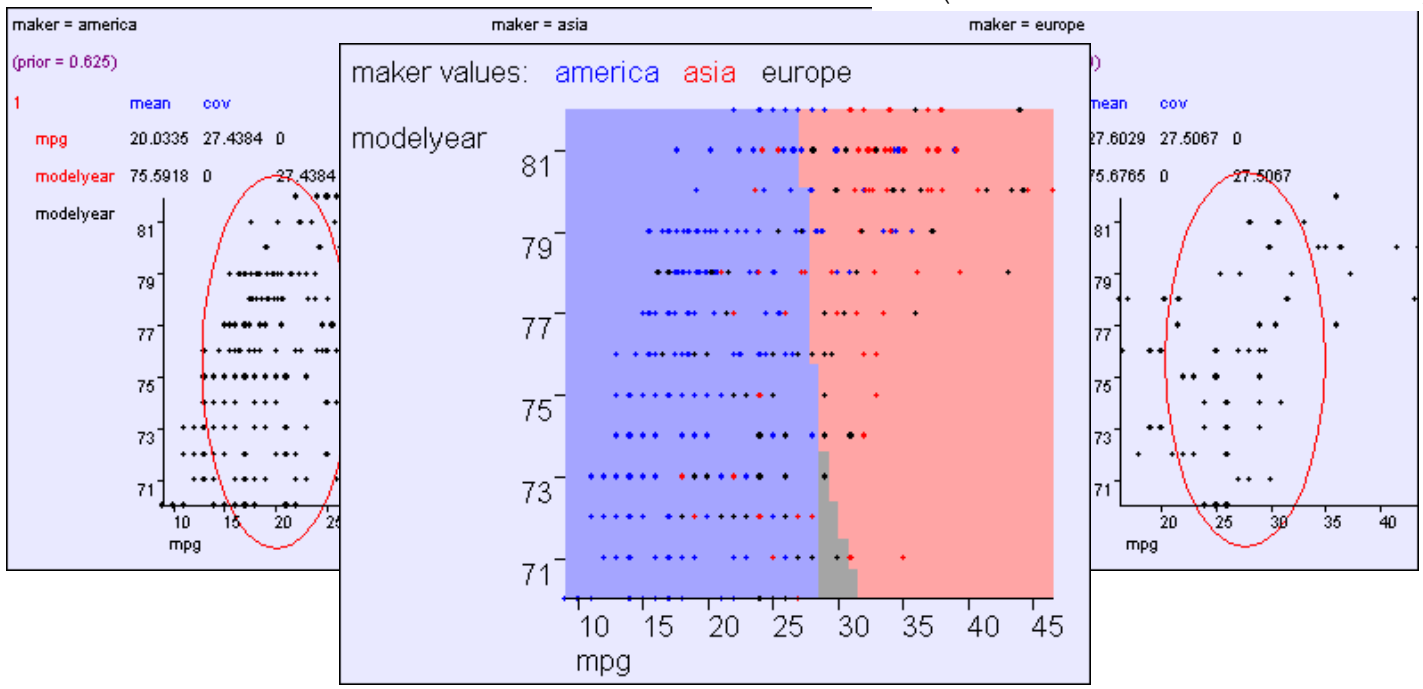

Clustering our Observations

- Imagine we have an assignment of each x^i to a Gaussian

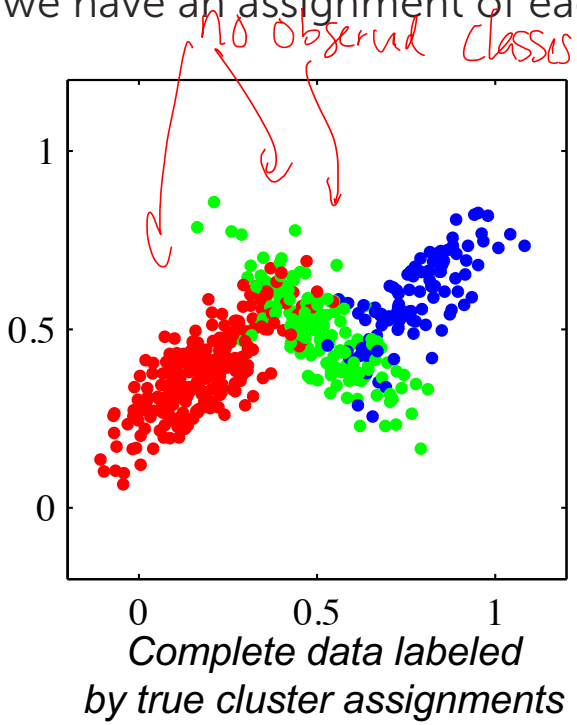

Complete data labeled by true cluster assignments

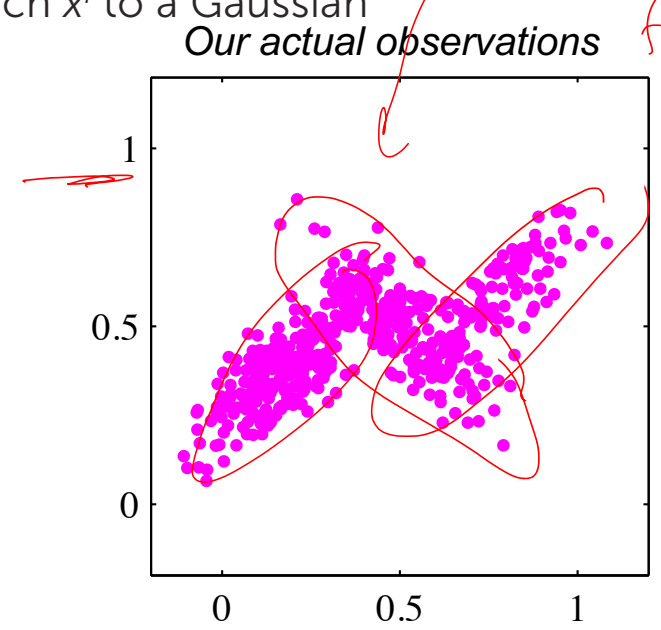

Our actual observations

©2022 Carlos Guestrin

Density as Mixture of Gaussians

- Approximate with density with a mixture of Gaussians

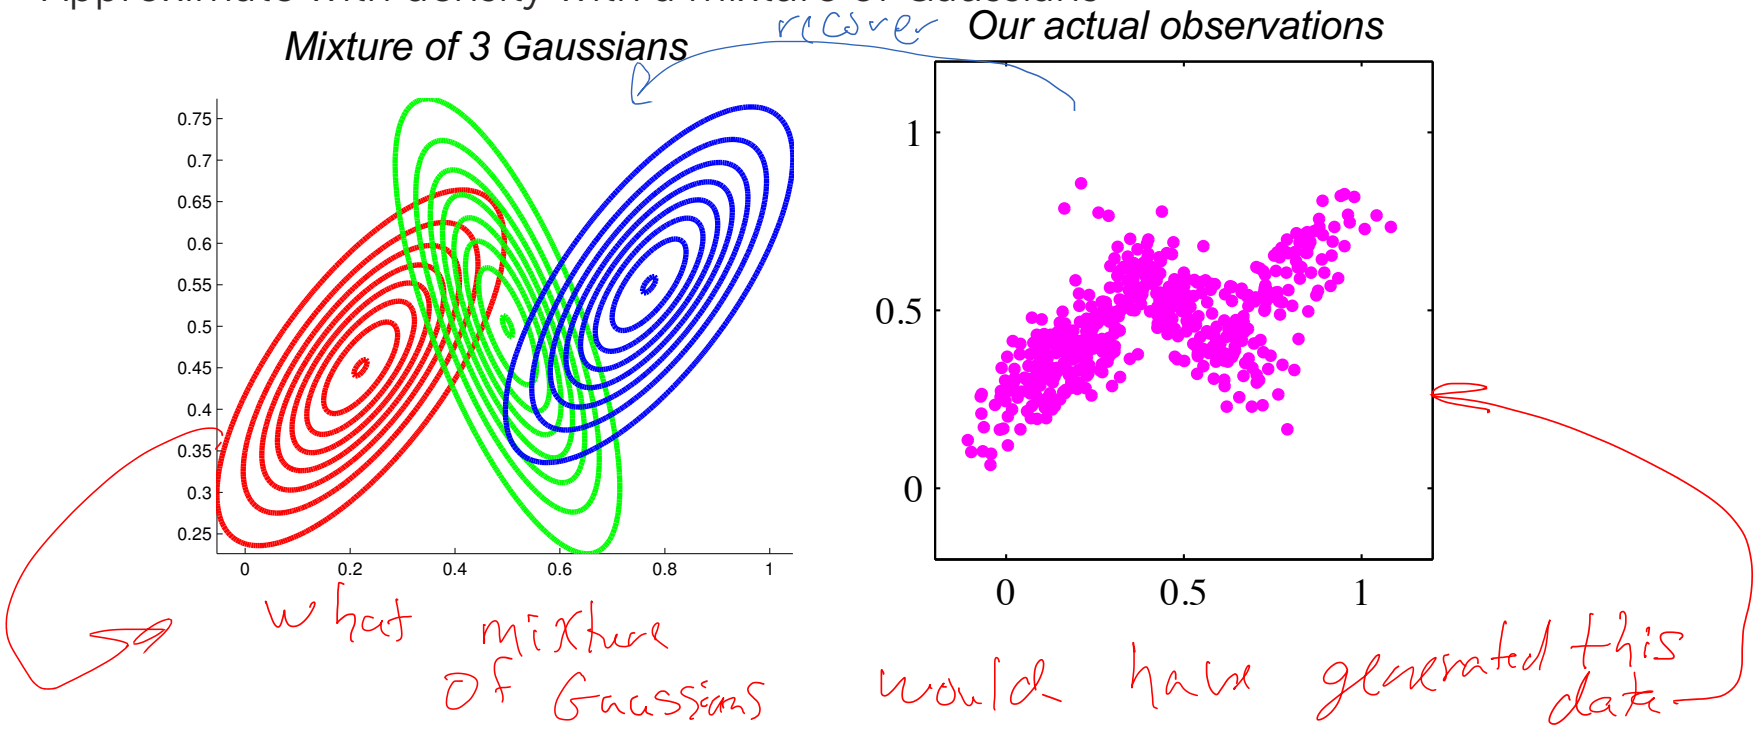

©2022 Carlos Guestrin

The General GMM assumption

- There are k components
- Component *i* has an associated mean vector μ_i
- Each component generates data from a Gaussian with mean μ_i and covariance matrix Σ_i

Each data point is generated according to the following recipe:

1. Pick a component at random: Choose component i with probability *P(z=i)*
2. Datapoint ~ N(μ_i, Σ_i)

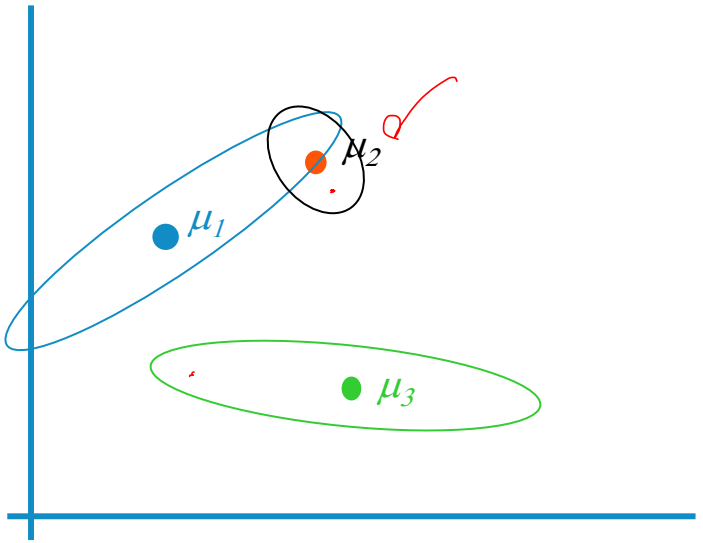

Density Estimation

- Estimate a density based on $x^1,\ldots,x^N$

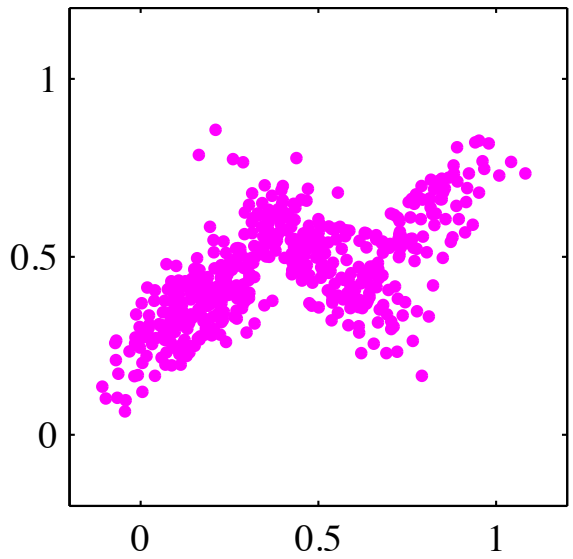

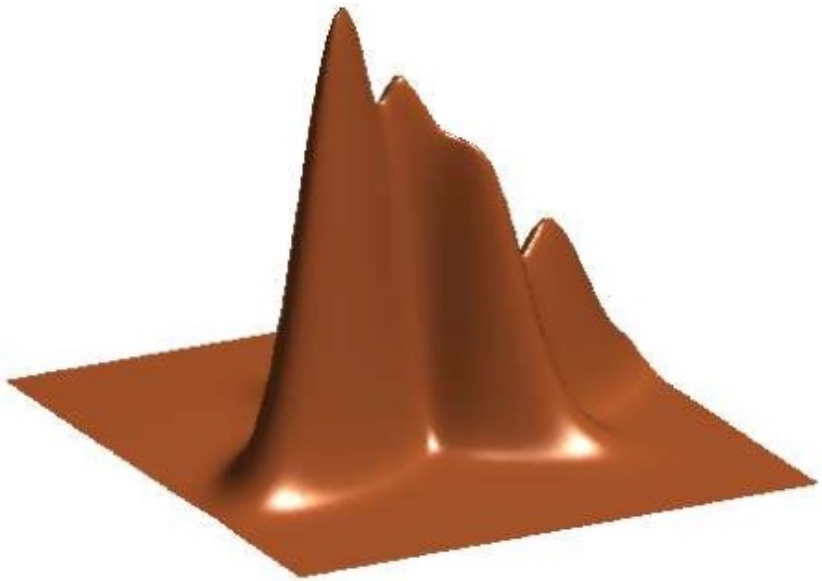

©2022 Carlos Guestrin

Density as Mixture of Gaussians

- Approximate density with a mixture of Gaussians

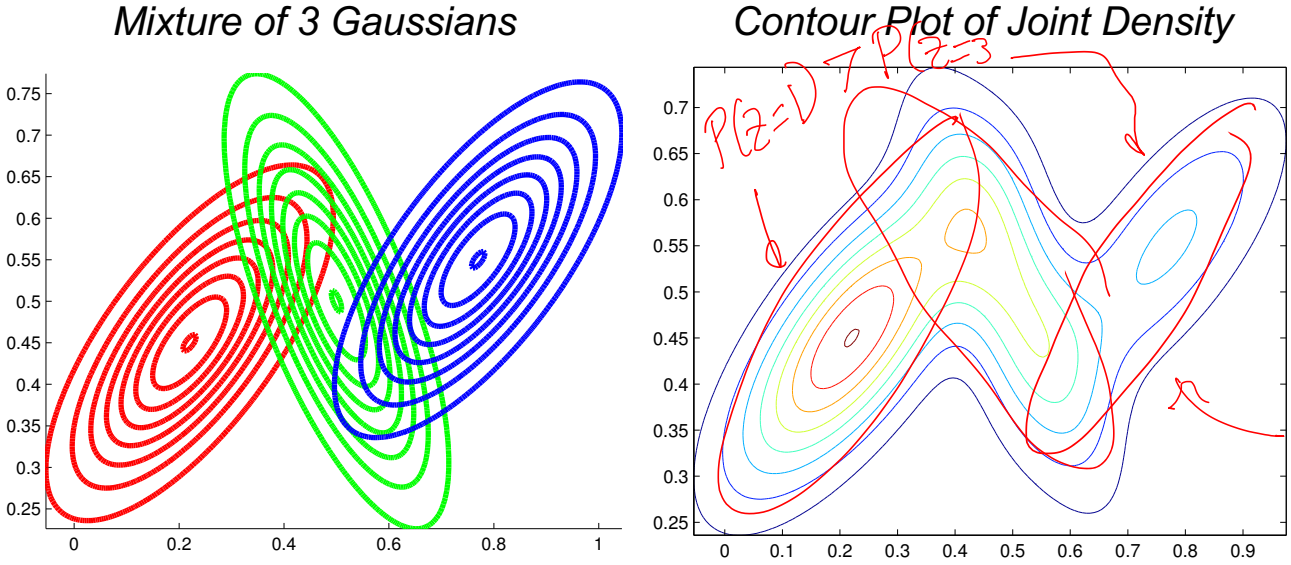

©2022 Carlos Guestrin

Summary of GMM Components

- Observations $x^i \in \mathbb{R}^d, \quad i = 1, 2, \ldots, N$
- Hidden cluster labels $z_i \in \{1, 2, \ldots, K\}, \quad i = 1, 2, \ldots, N$
- Hidden mixture means $\mu_k \in \mathbb{R}^d, \quad k = 1, 2, \ldots, K$
- Hidden mixture covariances $\Sigma_k \in \mathbb{R}^{d \times d}, \quad k = 1, 2, \ldots, K$
- Hidden mixture probabilities $\pi_k, \quad \sum_{k=1}^{K} \pi_k = 1$

Gaussian mixture marginal and conditional likelihood **:**

$$p(x^i|\pi, \mu, \Sigma) = \sum_{z^i=1}^{K} \pi_{z^i} \; p(x^i|z^i, \mu, \Sigma)$$

$$p(x^i|z^i, \mu, \Sigma) = \mathcal{N}(x^i|\mu_{z^i}, \Sigma_{z^i})$$

©2022 Carlos Guestrin