## CS340 Machine learning Gaussian classifiers

## Correlated features

- Height and weight are not independent



## Multivariate Gaussian

- Multivariate Normal (MVN)

$$
\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \Sigma) \stackrel{\text { def }}{=} \frac{1}{(2 \pi)^{p / 2}|\Sigma|^{1 / 2}} \exp \left[-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{T} \Sigma^{-1}(\mathbf{x}-\boldsymbol{\mu})\right]
$$

- Exponent is the Mahalanobis distance between $x$ and $\mu$

$$
\Delta=(\mathrm{x}-\mu)^{T} \Sigma^{-1}(\mathrm{x}-\mu)
$$

$\Sigma$ is the covariance matrix (positive definite)

$$
\mathbf{x}^{T} \Sigma \mathbf{x}>0 \forall \mathbf{x}
$$

## Bivariate Gaussian

- Covariance matrix is

$$
\Sigma=\left(\begin{array}{cc}
\sigma_{x}^{2} & \rho \sigma_{x} \sigma_{y} \\
\rho \sigma_{x} \sigma_{y} & \sigma_{y}^{2}
\end{array}\right)
$$

where the correlation coefficient is

$$
\rho=\frac{\operatorname{Cov}(X, Y)}{\sqrt{\operatorname{Var}(X) \operatorname{Var}(Y)}}
$$

and satisfies $-1 \leq \rho \leq 1$

- Density is

$$
p(x, y)=\frac{1}{2 \pi \sigma_{x} \sigma_{y} \sqrt{1-\rho^{2}}} \exp \left(-\frac{1}{2\left(1-\rho^{2}\right)}\left(\frac{x^{2}}{\sigma_{x}^{2}}+\frac{y^{2}}{\sigma_{y}^{2}}-\frac{2 \rho x y}{\left(\sigma_{x} \sigma_{y}\right)}\right)\right)
$$

## Spherical, diagonal, full covariance



$$
\Sigma=\left(\begin{array}{cc}
\sigma^{2} & 0 \\
0 & \sigma^{2}
\end{array}\right)
$$

$$
\Sigma=\left(\begin{array}{cc}
\sigma_{x}^{2} & 0 \\
0 & \sigma_{y}^{2}
\end{array}\right)
$$

$$
\Sigma=\left(\begin{array}{cc}
\sigma_{x}^{2} & \rho \sigma_{x} \sigma_{y} \\
\rho \sigma_{x} \sigma_{y} & \sigma_{y}^{2}
\end{array}\right)
$$

## Surface plots



## Generative classifier

- A generative classifier is one that defines a classconditional density $\mathrm{p}(\mathrm{x} \mid \mathrm{y}=\mathrm{c})$ and combines this with a class prior $\mathrm{p}(\mathrm{c})$ to compute the class posterior

$$
p(y=c \mid \mathbf{x})=\frac{p(\mathbf{x} \mid y=c) p(y=c)}{\sum_{c^{\prime}} p\left(\mathbf{x} \mid y=c^{\prime}\right) p\left(c^{\prime}\right)}
$$

- Examples:
- Naïve Bayes:

$$
p(\mathbf{x} \mid y=c)=\prod_{j=1}^{d} p\left(x_{j} \mid y=c\right)
$$

- Gaussian classifiers $p(\mathbf{x} \mid y=c)=\mathcal{N}\left(\mathbf{x} \mid \boldsymbol{\mu}_{c}, \boldsymbol{\Sigma}_{c}\right)$
- Alternative is a discriminative classifier, that estimates $p(y=c \mid x)$ directly.


## Naïve Bayes with Bernoulli features

- Consider this class-conditional density

$$
p(x \mid y=c)=\prod_{i=1}^{d} \theta_{i c}^{I\left(x_{i}=1\right)}\left(1-\theta_{i c}\right)^{I\left(x_{i}=0\right)}
$$

- The resulting class posterior (using plugin rule) has the form

$$
p(y=c \mid x)=\frac{p(y=c) p(x \mid y=c)}{p(x)}=\frac{\pi_{c} \prod_{i=1}^{d} \theta_{i c}^{I\left(x_{i}=1\right)}\left(1-\theta_{i c}\right)^{I\left(x_{i}=0\right)}}{p(x)}
$$

- This can be rewritten as

$$
\begin{aligned}
p(Y=c \mid x, \theta, \pi) & =\frac{p(x \mid y=c) p(y=c)}{\sum_{c^{\prime}} p\left(x \mid y=c^{\prime}\right) p\left(y=c^{\prime}\right)} \\
& =\frac{\exp [\log p(x \mid y=c)+\log p(y=c)]}{\sum_{c^{\prime}} \exp \left[\log p\left(x \mid y=c^{\prime}\right)+\log p\left(y=c^{\prime}\right)\right]} \\
& =\frac{\exp \left[\log \pi_{c}+\sum_{i} I\left(x_{i}=1\right) \log \theta_{i c}+I\left(x_{i}=0\right) \log \left(1-\theta_{i c}\right)\right]}{\sum_{c^{\prime}} \exp \left[\log \pi_{c^{\prime}}+\sum_{i} I\left(x_{i}=1\right) \log \theta_{i, c^{\prime}}+I\left(x_{i}=0\right) \log \left(1-\theta_{i c}\right)\right]}
\end{aligned}
$$

## Form of the class posterior

- From previous slide

$$
p(Y=c \mid x, \theta, \pi) \quad \propto \quad \exp \left[\log \pi_{c}+\sum_{i} I\left(x_{i}=1\right) \log \theta_{i c}+I\left(x_{i}=0\right) \log \left(1-\theta_{i c}\right)\right]
$$

- Define

$$
\begin{aligned}
x^{\prime} & =\left[1, I\left(x_{1}=1\right), I\left(x_{1}=0\right), \ldots, I\left(x_{d}=1\right), I\left(x_{d}=0\right)\right] \\
\beta_{c} & =\left[\log \pi_{c}, \log \theta_{1 c}, \log \left(1-\theta_{1 c}\right), \ldots, \log \theta_{d c}, \log \left(1-\theta_{d c}\right)\right]
\end{aligned}
$$

- Then the posterior is given by the softmax function

$$
p(Y=c \mid x, \beta)=\frac{\exp \left[\beta_{c}^{T} x^{\prime}\right]}{\sum_{c^{\prime}} \exp \left[\beta_{c^{\prime}}^{T} x^{\prime}\right]}
$$

- This is called softmax because it acts like the max function when $\left|\beta_{\mathrm{c}}\right| \rightarrow \infty$

$$
p(Y=c \mid \mathbf{x})= \begin{cases}1.0 & \text { if } c=\arg \max _{c^{\prime}} \beta_{c^{\prime}}^{T} \mathbf{x} \\ 0.0 & \text { otherwise }\end{cases}
$$

## Two-class case

- From previous slide

$$
p(Y=c \mid x, \beta)=\frac{\exp \left[\beta_{c}^{T} x^{\prime}\right]}{\sum_{c^{\prime}} \exp \left[\beta_{c^{\prime}}^{T} x^{\prime}\right]}
$$

- In the binary case, $\mathrm{Y} \in\{0,1\}$, the softmax becomes the logistic (sigmoid) function $\sigma(u)=1 /\left(1+e^{-u}\right)$

$$
\begin{aligned}
p(Y=1 \mid x, \theta) & =\frac{e^{\beta_{1}^{T} x^{\prime}}}{e^{\beta_{1}^{T} x^{\prime}}+e^{\beta_{0}^{T} x^{\prime}}} \\
& =\frac{1}{1+e^{\left(\beta_{0}-\beta_{1}\right)^{T} x^{\prime}}} \xrightarrow{\sim} \quad \cdots(u) \\
& =\frac{1}{1+e^{w^{T} x^{\prime}}} \\
& =\sigma\left(w^{T} x^{\prime}\right)
\end{aligned}
$$

## Sigmoid function

- $\sigma(a x+b)$, a controls steepness, $b$ is threshold.
- For small $a$ and $x \approx-b / 2$, roughly linear



## Sigmoid function in 2D

$\sigma\left(w_{1} x_{1}+w_{2} x_{2}\right)=\sigma\left(w^{\top} x\right): w$ is perpendicular to the decision boundary


Mackay 39.3

## Logit function

- Let $p=p(y=1)$ and $\eta$ be the log odds

$$
\eta=\log \frac{p}{1-p}
$$

- Then $p=\sigma(\eta)$ and $\eta=\operatorname{logit}(p)$

$$
\begin{aligned}
\sigma(\eta) & =\frac{1}{1+e^{-\eta}}=\frac{e^{\eta}}{e^{\eta}+1} \\
& =\frac{\frac{p}{(1-p)}}{\frac{p}{1-p}+1}=\frac{\frac{p}{(1-p)}}{\frac{p+1-p}{1-p}}=p
\end{aligned}
$$

$\eta$ is the natural parameter of the Bernoulli distribution, and $p=E[y]$ is the moment parameter

- If $\eta=w^{\top} x$, then $w_{i}$ is how much the log-odds increases by if we increase $x_{i}$


## Gaussian classifiers

- Class posterior (using plug-in rule)

$$
\begin{aligned}
p(Y=c \mid \mathbf{x}) & =\frac{p(\mathbf{x} \mid Y=c) p(Y=c)}{\sum_{c^{\prime}=1}^{C} p\left(\mathbf{x} \mid Y=c^{\prime}\right) p\left(Y=c^{\prime}\right)} \\
& =\frac{\pi_{c}\left|2 \pi \Sigma_{c}\right|^{-\frac{1}{2}} \exp \left[-\frac{1}{2}\left(\mathbf{x}-\mu_{c} T^{T} \Sigma_{c}^{-1}\left(\mathbf{x}-\mu_{c}\right)\right]\right.}{\sum_{c^{\prime}} \pi_{c^{\prime}}\left|2 \pi \Sigma_{c^{\prime}}\right|^{-\frac{1}{2}} \exp \left[-\frac{1}{2}\left(\mathbf{x}-\mu_{c^{\prime}}\right)^{T} \Sigma_{c^{\prime}}^{-1}\left(\mathbf{x}-\mu_{c^{\prime}}\right)\right]}
\end{aligned}
$$

- We will consider the form of this equation for various special cases:
- $\Sigma_{1}=\Sigma_{0}$,
- $\Sigma_{c}$ tied, many classes
- General case


## $\Sigma_{1}=\Sigma_{0}$

- Class posterior simplifies to

$$
\begin{aligned}
& p(Y=1 \mid \mathbf{x})=\frac{p(\mathbf{x} \mid Y=1) p(Y=1)}{p(\mathbf{x} \mid Y=1) p(Y=1)+p(\mathbf{x} \mid Y=0) p(Y=0)} \\
&=\frac{\pi_{1} \exp \left[-\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{1}\right)\right]}{\pi_{1} \exp \left[-\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{1}\right)\right]+\pi_{0} \exp \left[-\frac{1}{2}\left(\mathbf{x}-\mu_{0}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{0}\right)\right]} \\
&=\frac{\pi_{1} e^{a_{1}}}{\pi_{1} e^{a_{1}}+\pi_{0} e^{a_{0}}}=\frac{1}{1+\frac{\pi_{0}}{\pi_{1}} e^{a_{0}-a_{1}}} \\
& a_{c} \stackrel{\stackrel{\text { def }}{=}}{ }-\frac{1}{2}\left(\mathbf{x}-\mu_{c}\right)^{T} \Sigma\left(\mathbf{x}-\mu_{c}\right)
\end{aligned}
$$

## $\Sigma_{1}=\Sigma_{0}$

## - Class posterior simplifies to

$$
\begin{aligned}
p(Y=1 \mid \mathbf{x}) & =\frac{1}{1+\exp \left[-\log \frac{\pi_{1}}{\pi_{0}}+a_{0}-a_{1}\right]} \\
a_{0}-a_{1} & =-\frac{1}{2}\left(\mathbf{x}-\mu_{0}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{0}\right)+\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{1}\right) \\
& =-\left(\mu_{1}-\mu_{0}\right)^{T} \Sigma^{-1} \mathbf{x}+\frac{1}{2}\left(\mu_{1}-\mu_{0}\right)^{T} \Sigma^{-1}\left(\mu_{1}+\mu_{0}\right)
\end{aligned}
$$

SO
Linear function of $x$

$$
\begin{aligned}
p(Y=1 \mid \mathbf{x}) & =\frac{1}{1+\exp \left[-\beta^{T} \mathbf{x}-\gamma\right]}=\sigma\left(\beta^{T} \mathbf{x}+\gamma\right) \\
\beta & \stackrel{\text { def }}{=} \Sigma^{-1}\left(\mu_{1}-\mu_{0}\right) \\
\gamma & \stackrel{\text { def }}{=}-\frac{1}{2}\left(\mu_{1}-\mu_{0}\right)^{T} \Sigma^{-1}\left(\mu_{1}+\mu_{0}\right)+\log \frac{\pi_{1}}{\pi_{0}} \\
\sigma(\eta) & \stackrel{\text { def }}{=} \frac{1}{1+e^{-\eta}}=\frac{e^{\eta}}{e^{\eta}+1}
\end{aligned}
$$

## Decision boundary

- Rewrite class posterior as

$$
\begin{aligned}
p(Y=1 \mid \mathbf{x}) & =\sigma\left(\boldsymbol{\beta}^{T} \mathbf{x}+\gamma\right)=\sigma\left(\mathbf{w}^{T}\left(\mathbf{x}-\mathbf{x}_{0}\right)\right) \\
\mathbf{w} & =\beta=\Sigma^{-1}\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{0}\right) \\
\mathbf{x}_{0} & =-\frac{\gamma}{\boldsymbol{\beta}}=\frac{1}{2}\left(\boldsymbol{\mu}_{1}+\boldsymbol{\mu}_{0}\right)-\frac{\log \left(\pi_{1} / \pi_{0}\right)}{\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{0}\right)^{T} \Sigma^{-1}\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{0}\right)}\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{0}\right)
\end{aligned}
$$

- If $\Sigma=$ l, then $\mathrm{w}=\left(\mu_{1}-\mu_{0}\right)$ is in the direction of $\mu_{1}-\mu_{0}$, so the hyperplane is orthogonal to the line between the two means, and intersects it at $x_{0}$
- If $\pi_{1}=\pi_{0}$, then $x_{0}=0.5\left(\mu_{1}+\mu_{0}\right)$ is midway between the two means
- If $\pi_{1}$ increases, $x_{0}$ decreases, so the boundary shifts toward $\mu_{0}$ (so more space gets mapped to class 1)


## Decision boundary in 1d


$\mu_{1}=1.0_{1} \mu_{2}=-1.0, \pi_{1}=0.5_{2} \sigma_{1}=1.0, \sigma_{2}=1.0$
$\mu_{1}=10, \mu_{2}=-1.0, \pi_{1}=0.5, \sigma_{1}=3.0_{3} \sigma_{2}=1.0$



Discontinuous decision region

## Decision boundary in 2d



## Tied $\Sigma_{,}$many classes

- Similarly to before

$$
\begin{aligned}
& p(Y=c \mid \mathbf{x})=\frac{\pi_{c} \exp \left[-\frac{1}{2}\left(\mathbf{x}-\mu_{c}\right)^{T} \Sigma_{c}^{-1}\left(\mathbf{x}-\mu_{c}\right)\right]}{\sum_{c^{\prime}} \pi_{c^{\prime}} \exp \left[-\frac{1}{2}\left(\mathbf{x}-\mu_{c^{\prime}}\right)^{T} \Sigma_{c^{\prime}}^{-1}\left(\mathbf{x}-\mu_{c^{\prime}}\right)\right]} \\
&=\frac{\exp \left[\mu_{c}^{T} \Sigma^{-1} x-\frac{1}{2} \mu_{c}^{T} \Sigma^{-1} \mu_{c}+\log \pi_{c}\right]}{\sum_{c^{\prime}} \exp \left[\mu_{c^{\prime}}^{T} \Sigma^{-1} \mathbf{x}-\frac{1}{2} \mu_{c^{\prime}}^{T} \Sigma^{-1} \mu_{c^{\prime}}+\log \pi_{c^{\prime}}\right]} \\
& \theta_{c} \stackrel{\text { def }}{=}\binom{-\mu_{c}^{T} \Sigma^{-1} \mu_{c}+\log \pi_{c}}{\Sigma^{-1} \mu_{c}}=\binom{\gamma_{c}}{\beta_{c}} \\
& p(Y=c \mid \mathbf{x})=\frac{e^{\theta_{c}^{T} \mathbf{x}}}{\sum_{c^{\prime}} \theta^{\theta_{c^{\prime}}^{T} \mathbf{x}}=\frac{e^{\beta_{c}^{T} \mathbf{x}+\gamma_{c}}}{\sum_{c^{\prime}} e^{\beta_{c^{\prime}}^{T} \mathbf{x}+\gamma_{c^{\prime}}}}}
\end{aligned}
$$

- This is the multinomial logit or softmax function


## Tied $\Sigma_{,}$many classes

- Discriminant function

$$
\begin{aligned}
g_{c}(\mathbf{x}) & =-\frac{1}{2}\left(\mathbf{x}-\mu_{c}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{c}\right)+\log p(Y=c)=\beta_{c}^{T} \mathbf{x}+\beta_{c 0} \\
\beta_{c} & =\Sigma^{-1} \mu_{c} \\
\beta_{c 0} & =-\frac{1}{2} \mu_{c}^{T} \Sigma^{-1} \mu_{c}+\log \pi_{c}
\end{aligned}
$$

- Decision boundary is again linear, since $x^{\top} \Sigma x$ terms cancel
- If $\Sigma=I$, then the decision boundaries are orthogonal to $\mu_{\mathrm{i}}-\mu_{\mathrm{j}}$, otherwise skewed



## Decision boundaries

$$
g_{1}(x)-\max \left(g_{2}(x), g_{3}(x)\right)=0
$$



$[\mathrm{x}, \mathrm{y}]=\operatorname{meshgrid}(\operatorname{linspace}(-10,10,100)$, linspace $(-10,10,100))$;
g1 = reshape (mvnpdf(X, mu1(:)', S1), [m n]); ...
contour ( $\mathrm{x}, \mathrm{y}, \mathrm{g} 2 * \mathrm{p} 2-\max (\mathrm{g} 1 * \mathrm{p} 1, \mathrm{~g} 3 * \mathrm{p} 3),\left[\begin{array}{ll}0 & 0\end{array}\right],-\mathrm{k}$ ') ;

## $\Sigma_{0}, \Sigma_{1}$ arbitrary

- If the $\Sigma$ are unconstrained, we end up with cross product terms, leading to quadratic decision boundaries



General case

$$
\mu_{1}=(0,0), \mu_{2}=(0,5), \mu_{3}=(5,5), \pi=(1 / 3,1 / 3,1 / 3)
$$

All boundaries are linear


$$
\begin{aligned}
& \Sigma_{1}=\left(\begin{array}{ll}
1 & 0 \\
0 & 3
\end{array}\right) \\
& \Sigma_{2}=\left(\begin{array}{ll}
4 & 0 \\
0 & 1
\end{array}\right) \\
& \Sigma_{5}=\left(\begin{array}{ll}
1 & 0 \\
0 & 2
\end{array}\right)
\end{aligned}
$$

All boundaries are quadratic


Some linear, some quadratic


$$
\Sigma_{1}=\left(\begin{array}{ll}
4 & 0 \\
0 & 1
\end{array}\right), \Sigma_{2}=\Sigma_{3}=T
$$

$$
\Sigma_{1}=\left(\begin{array}{ll}
1 & 0 \\
0 & 3
\end{array}\right)
$$

$$
E_{2}=\left(\begin{array}{ll}
4 & 0 \\
0 & 1
\end{array}\right)
$$

$$
\Sigma_{j}=\left(\begin{array}{ll}
1 & 0 \\
0 & 2
\end{array}\right)
$$

