# Machine Learning for Signal Processing 

 Bayes Classification and RegressionInstructor: Bhiksha Raj

## Recap: KNN

- A very effective and simple way of performing classification
- Simple model: For any instance, select the class from the instances close to it in feature space


## Multi-class Image Classification



## k-Nearest Neighbors

Given a query item: Find k closest matches in a labeled dataset $\downarrow$


## k-Nearest Neighbors

Given a query item:
Find k closest matches


Return the most
Frequent label


## k-Nearest Neighbors

$$
k=3 \text { votes for "cat" }
$$



## k-Nearest Neighbors

2 votes for cat, 1 each for Buffalo, Deer, Lion


Cat wins...


## Nearest neighbor method

- Weighted majority vote within the k nearest neighbors



## But what happens if..

- You have many training instances at exactly that value of $x$ ?
- Majority vote on nearest neighbors:



## But what happens if..

- You have many training instances at exactly that value of $x$ ?
- Majority vote on nearest neighbors:


$$
\hat{y}=\underset{y}{\operatorname{argmax}} \frac{N_{y}(x)}{N(x)}
$$

## But what happens if..

- You have many training instances at exactly that value of $x$ ?
- Majority vote on nearest neighbors:



## But what happens if..

- You have many training instances at exactly that value of $x$ ?
- Majority vote on nearest neighbors:

- Bayes Classification Rule


## Bayes Classification Rule

- For any observed feature $X$, select the class value $Y$ that is most frequent
- Also applies to continuous valued predicted variables
- I.e. regression
- Select $Y$ to maximize the a posteriori probability $P(Y \mid X)$
- Bayes classification is an instance of maximum a posteriori estimation


## Bayes classification

- What happens if there are no exact neighbors
- No training instances with exactly the same $X$ value?



## Bayes Classification Rule

- Given
- a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$
- Conditional probability distributions $P(C \mid X)$
- Classification performed as

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C \mid X)
$$

- Problem: How do you characterize $P(C \mid X)$
- Require a function that, given an $X$, computes $P(C \mid X)$ for every class $C$


## Modelling $\mathrm{P}(\mathrm{C} \mid \mathrm{X})$

Each pixel is a combination of red green and blue weighted by the a posteriori probability of the classes


Blue: Class 1
Red: Class 2
Green: Class 3

- Assumption: there's a continuous function that, at every $X$, produces a vector of outputs $P(C \mid X)$ for every class $C$
- The "decision boundary" for any class is the boundary within which its own posterior has the highest value
- This function accurately represents the actual a posteriori probabilities for the classes
- Objective: Estimate this function


## Modelling the posterior

- To model the posterior, we need a functional form for $P(C \mid X)$ which can be learned
- Typically this functional form is expressed in terms of distance from a decision boundary
- The simplest decision boundary is the linear boundary


## Bayesian Linear Classification:

## Two class case



- First: Two-class classification
- Assumption: the decision boundary between the classes is a simple hyperplane
- As you go away from the hyperplane, the fraction of data from one class increases, while that from the other decreases
- Will also hold for any sample of data


## 1-D binary class example



- One-dimensional example for visualization
- Only two classes (represented by $y=0$ and $y=1$ )
- All (red) dots at $Y=1$ represent instances of class $Y=1$
- All (blue) dots at $\mathrm{Y}=0$ are from class $\mathrm{Y}=0$
- The data are not linearly separable
- In this 1-D example, a linear separator is a threshold
- No threshold will cleanly separate red and blue dots


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of $\mathrm{Y}=1$ at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The probability of $\mathrm{y}=1$



- Consider this differently: at each point look at a small window around that point
- Plot the average value within the window
- This is an approximation of the probability of 1 at that point


## The logistic regression model



- Class 1 becomes increasingly probable going left to right
- Very typical in many problems
- The logistic is a function of the distance from the $P(y \mid x)=0.5$ boundary


## For two-dimensional input




- The decision boundary for $P(Y \mid X)=0.5$ is a hyperplane
- It is a linear model


## The logistic regression model



- Note how it varies with $w_{0}+w_{1} \mathbf{x}$


## Estimating the model



- Given the training data (many $(x, y)$ pairs represented by the dots), estimate $w_{0}$ and $w_{1}$ for the curve


## Estimating the model



- Easier to represent using a $y=+1 /-1$ notation

$$
\begin{gathered}
P(y=1 \mid x)=\frac{1}{1+e^{-\left(w_{0}+w_{1} x\right)}} \quad P(y=-1 \mid x)=\frac{1}{1+e^{\left(w_{0}+w_{1} x\right)}} \\
P(y \mid x)=\frac{1}{1+e^{-y\left(w_{0}+w_{1} x\right)}}
\end{gathered}
$$

## Estimating the model

- Given: Training data

$$
\left(X_{1}, y_{1}\right),\left(X_{2}, y_{2}\right), \ldots,\left(X_{N}, y_{N}\right)
$$

- Xs are vectors, $y$ s are binary (0/1) class values
- Total probability of data

$$
\begin{aligned}
& P\left(\left(X_{1}, y_{1}\right),\left(X_{2}, y_{2}\right), \ldots,\left(X_{N}, y_{N}\right)\right)=\prod_{i} P\left(X_{i}, y_{i}\right) \\
& =\prod_{i} P\left(y_{i} \mid X_{i}\right) P\left(X_{i}\right)=\prod_{i} \frac{1}{1+e^{-y_{i}\left(w_{0}+w^{T} X_{i}\right)}} P\left(X_{i}\right)
\end{aligned}
$$

## Estimating the model

- Likelihood

$$
P(\text { Training data })=\prod_{i} \frac{1}{1+e^{-y_{i}\left(w_{0}+w^{T} X_{i}\right)}} P\left(X_{i}\right)
$$

- Log likelihood

$$
\log P(\text { Training data })=
$$

$$
\sum_{i} \log P\left(X_{i}\right)-\sum_{i} \log \left(1+e^{-y_{i}\left(w_{0}+w^{T} X_{i}\right)}\right)
$$

## Maximum Likelihood Estimate

$$
\widehat{w}_{0}, \widehat{w}_{1}=\underset{w_{0}, w_{1}}{\operatorname{argmax}} \log P(\text { Training data })
$$

- Equals (note argmin rather than argmax)

$$
\widehat{w}_{0}, \widehat{w}_{1}=\underset{w_{0}, w}{\operatorname{argmin}} \sum_{i} \log \left(1+e^{-y_{i}\left(w_{0}+w^{T} X_{i}\right)}\right)
$$

- Minimizing the KL divergence between the desired output $y$ and actual output $\frac{1}{1+e^{-\left(w_{0}+w^{T} X_{i}\right)}}$
- Cannot be solved directly, needs gradient descent


## Model learned by logistic

 regressionPure Red: 0
Pure Green: 1


- The figure shows the class probability over a twodimensional feature space
- Any decision threshold $P(C \mid X)=$ Const is a hyperplane
- Diagonal line in this case


## Multi-class logistic regression

- The simple logistic regression model can be extended to multiple classes:

$$
P(C \mid X)=\frac{\exp \left(W_{C}^{T} X\right)}{\sum_{C^{\prime}} \exp \left(W_{C^{\prime}}^{T} X\right)}
$$

- $W_{C}^{T} X$ is, in fact, the discriminant function of the classes
- We've encountered discriminant functions earlier
- Also called a softmax
- Each class $C_{i}$ has a probability that is exponentially related to the closeness of the vector $X$ to a "representative" vector $w_{i}$ for the class
- This too can be learned from training data via maximum likelihood estimation
- Just like the two-class case


## Multi-class logistic regression



- The boundary between adjacent classes is a hyperplane (line)
- The decision boundary for any class is convex polytope with hyperplane segments
- I.e. still a linear classifier


## Multi-class Bayes




- In many classification problems, linear boundaries are not sufficient
- We need to be able to model more complex boundaries
- This too can be supported by the logistic regression classifier


## Logistic regression with non-linea ${ }^{\text {MLSP }}$ boundaries

- The logistic regression can be modified to have non-linear discriminants:

$$
P(C \mid X)=\frac{\exp \left(f\left(X ; \theta_{C}\right)\right)}{\sum_{C^{\prime}} \exp \left(f\left(X ; \theta_{C^{\prime}}\right)\right)}
$$

- $f\left(X ; \theta_{C}\right)$ is the discriminant for class C , and has parameter $\theta_{C}$
- The discriminants determine the shape of the decision boundary
- Non-linear discriminants result in non-linear decision boundaries
- The parameters $\theta_{C}$ for all classes can be learned by maximum likelihood (or MAP) estimation as before


## Quadratic discriminant

$$
P(C \mid X)=\frac{\exp \left(f\left(X ; \theta_{C}\right)\right)}{\sum_{C^{\prime}} \exp \left(f\left(X ; \theta_{C^{\prime}}\right)\right)}
$$



- With quadratic discriminants:

$$
f\left(X ; \theta_{C}\right)=\left(X-\alpha_{C}\right) \beta_{C}\left(X-\alpha_{C}\right)^{T}
$$

- Note that decision boundaries are quadratic
- The probability of a class increases (or decreases) as we go away from a boundary


## Logistic regression with non-linear ${ }^{\text {mLSP }}$ boundaries

$$
P(C \mid X)=\frac{\exp \left(f\left(X ; \theta_{C}\right)\right)}{\sum_{C^{\prime}} \exp \left(f\left(X ; \theta_{C^{\prime}}\right)\right)}
$$



- For complex decision boundaries, the function $f\left(X ; \theta_{C}\right)$ must be correspondingly complex
- Currently the most successful approach in these cases is to model $f\left(X ; \theta_{C^{\prime}}\right)$ by a neural network
- In fact neural networks with soft-max decision layers may be seen as an instance of a logistic regression with a non-linear discriminant
- Topic for another class


## Logistic regression with non-linear ${ }^{\text {mLSP }}$ boundaries

- The logistic regression can be modified to have nonlinear discriminants:

$$
P(C \mid X)=\frac{\exp \left(f\left(X ; \theta_{C}\right)\right)}{\sum_{C^{\prime}} \exp \left(f\left(X ; \theta_{C^{\prime}}\right)\right)}
$$

- Note: This can also be viewed as non-linearly transforming the data $X$ into a space where a simple linear logistic regression models posteriors well
- $Z(X)=\left[f\left(X ; \theta_{1}\right) f\left(X ; \theta_{2}\right) \ldots f\left(X ; \theta_{K}\right)\right]^{T}$
- I.e. into a space where the data are most linearly separable
- We will discuss this in a later lecture on neural networks


## Problem with modelling $\mathrm{P}(\mathrm{C} \mid \mathrm{X})$

- We have considered modelling the a posteriori probability of the classes directly
- This implicitly assumes that
- The characteristics of the data for any class remain the same between train and test
- The relative proportions of the classes too remain the same
- Often the second assumption will not hold
- The data characteristics remain, but the relative proportions change
- E.g. the shapes of the differently colored coins don't change, but the relative proportions of the colors changes between train and test
- We must then modify our approach to Bayes classification to a generative framework


## The Bayesian Classifier..

- $\hat{C}=\operatorname{argmax} P(C \mid X)$ $C \in \mathcal{C}$
- Choose the class that is most frequent for the given $X$

$$
P(C \mid X)=\frac{P(C) P(X \mid C)}{P(X)}
$$

- $\operatorname{argmax} P(C \mid X)=\operatorname{argmax} P(C) P(X \mid C)$ $C \in \mathcal{C}$
$C \in \mathcal{C}$
- Choose the class that is most likely to have produced $X$
- While accounting for the relative frequency of $C$


## Bayes Classification Rule

- Given a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)
$$

$\mathrm{P}\left(\mathrm{X} \mid \mathrm{C}_{\mathrm{i}}\right)$ measures the probability that a random instance of class $\mathrm{C}_{\mathrm{i}}$ will take the value X


## Bayes Classification Rule

- Given a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)
$$

$\mathrm{P}\left(\mathrm{X} \mid \mathrm{C}_{\mathrm{i}}\right)$ measures the probability that a random instance of class $\mathrm{C}_{\mathrm{i}}$ will take the value X


## Bayes Classification Rule

- Given a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)
$$

$\mathrm{P}\left(\mathrm{C}_{\mathrm{i}}\right)$ scales them up to match the expected relative proportions of the classes


## Bayes Classification Rule

- Given a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)
$$

$\mathrm{P}\left(\mathrm{C}_{\mathrm{i}}\right)$ scales them up to match the expected relative proportions of the classes


## Bayes Classification Rule

- Given a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)
$$

$\mathrm{P}\left(\mathrm{C}_{\mathrm{i}}\right)$ scales them up to match the expected relative proportions of the classes


## Bayes Classification Rule

- Given a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)
$$

$\mathrm{P}\left(\mathrm{C}_{\mathrm{i}}\right)$ scales them up to match the expected relative proportions of the classes


## Bayes Classification Rule

- Given a set of classes $\mathcal{C}=\left\{C_{1}, C_{2}, \ldots, C_{N}\right\}$

$$
\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)
$$

$\mathrm{P}\left(\mathrm{C}_{\mathrm{i}}\right)$ scales them up to match the expected relative proportions of the classes


## Bayes Classification Rule

- The Bayes classification rule is the statistically optimal classification rule
- Moving the boundary in either direction will always increase the classification error



## The Bayesian Classifier..

- $\hat{C}=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C \mid X)=\underset{C \in \mathcal{C}}{\operatorname{argmax}} P(C) P(X \mid C)$

$$
C \in \mathcal{C} \quad C \in \mathcal{C}
$$

- We can now directly learn the class-specific statistical characteristics $P(X \mid C)$ from the training data
- The relative frequency of $C, P(C)$, can be independently adjusted to our expectations of these frequencies in the test data
- These need not match the training data


## Modeling $P(X \mid C)$

- Challenge: How to learn $P(X \mid C)$
- This will not be known beforehand and must be learned from examples of $X$ that belong to class $C$
- Will generally have unknown and unknowable shape
- We only observe samples of $X$
- Must make some assumptions about the form of $P(X \mid C)$


## The problem of dependent variables

- $P(X \mid C)=P\left(X_{1}, X_{2}, \ldots, X_{D} \mid C\right)$ must be defined for every combination of $X_{1}, X_{2}, \ldots, X_{D}$
- Too many parameters to describe explicitly
- Most combinations unseen in training data
- $P(X \mid C)$ may have an arbitrary scatter/shape
- Hard to characterize mathematically
- Typically do so by assigning a functional form to it


## The problem of dependent variables

- $P(X \mid C)=P\left(X_{1}, X_{2}, \ldots, X_{D} \mid C\right)$ must be defined for every combination of $X_{1}, X_{2}, \ldots, X_{D}$
- Too many parameters to describe explicitly
- Most combinations unseen in training data
- $P(X \mid C)$ may have an arbitrary scatter/shape
- Hard to characterize mathematically
- Typically do so by assigning a functional form to it


## The Naïve Bayes assumption



- Assume all the components are independent of one another
- The joint probability is the product of the marginal

$$
P(X \mid C)=P\left(X_{1}, X_{2}, \ldots, X_{D} \mid C\right)=\prod_{i} P\left(X_{i} \mid C\right)
$$

- Sufficient to learn marginal distributions $P\left(X_{i} \mid C\right)$
- The problem of having to observe all combinations of $X_{1}, X_{2}, \ldots, X_{D}$ never arises


## Naïve Bayes - estimating $P\left(X_{i} \mid C\right)$

- $P\left(X_{i} \mid C\right)$ may be estimated using conventional maximum likelihood estimation
- Given a number of training instances belonging to class $C$
- Select the i-th component of all instances
- Estimate $P\left(X_{i} \mid C\right)$
- For discrete-valued $X_{i}$ this will be a multinomial distribution
- For continuous valued $X_{i}$ a form must be assumed
- E.g Gaussian, Laplacian etc


## Naïve Bayes - Binary Case

$$
P(X \mid C)=P\left(X_{1}, X 2, \ldots, X_{D} \mid C\right)
$$

if $X_{i} \in\{0,1\} \quad 2^{D}-1$ parameters for each $C$

$$
P(X \mid C)=P\left(X_{1} \mid C\right) \cdot \ldots \cdot P\left(X_{D} \mid C\right)
$$

$D$ parameters for each $C$

## The problem of dependent variables

- $P(X \mid C)=P\left(X_{1}, X_{2}, \ldots, X_{D} \mid C\right)$ must be defined for every combination of $X_{1}, X_{2}, \ldots, X_{D}$
- Too many parameters
- Most combinations unseen in training data
- $P(X \mid C)$ may have an arbitrary scatter/shape
- Hard to characterize mathematically
- Typically do so by assigning a functional form to it


## Assigning a functional form to $\mathrm{P}(\mathrm{X} \mid \mathrm{C})$

- Assign a functional form to $\mathrm{P}(\mathrm{X} \mid \mathrm{C})$
- Common assumptions:
- Unimodal forms: Gaussian, Laplacian
- Multimodal forms: Gaussian mixtures
- Time series: Hidden Markov models
- Multi-dimensional structures: Markov random fields


## Assigning a functional form to $\mathrm{P}(\mathrm{X} \mid \mathrm{C})$

- Assign a functional form to $\mathrm{P}(\mathrm{X} \mid \mathrm{C})$
- Common assumptions:
- Unimodal forms: Gaussian, Laplacian
- Most common of all
- Multimodal forms: Gaussian mixtures
- Time series: Hidden Markov models
- Multi-dimensional structures: Markov random fields


## Gaussian Distribution

$p(\mathbf{x}, \mu, \Sigma)=\frac{1}{(2 \pi)^{n / 2}|\Sigma|^{1 / 2}} \exp \left(-\frac{1}{2}(\mathbf{x}-\mu)^{\top} \mathbf{\Sigma}^{-\mathbf{1}}(\mathbf{x}-\mu)\right)$
$\mathbf{x} \in \mathbb{R}^{n}$
$\mu$ Mean Vector


$\Sigma$

Covariance Matrix

$$
n \times n
$$

- Symmetric
- Positive Definite


## Gaussian Distribution

$$
\mu=\left[\begin{array}{l}
0 \\
0
\end{array}\right] \Sigma=\left[\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right]
$$

$$
\mu=\left[\begin{array}{l}
0 \\
0
\end{array}\right] \Sigma=\left[\begin{array}{cc}
1 & 0.5 \\
0.5 & 1
\end{array}\right]
$$

$$
\mu=\left[\begin{array}{l}
0 \\
0
\end{array}\right] \Sigma=\left[\begin{array}{cc}
1 & 0.8 \\
0.8 & 1
\end{array}\right]
$$







## Parameter Estimation

$$
p(\mathbf{x}, \mu, \Sigma)=\frac{1}{(2 \pi)^{n / 2}|\Sigma|^{1 / 2}} \exp \left(-\frac{1}{2}(\mathbf{x}-\mu)^{\top} \mathbf{\Sigma}^{-\mathbf{1}}(\mathbf{x}-\mu)\right)
$$



Maximum Likelihood Estimators

$$
\begin{gathered}
\hat{\mu}=\frac{1}{n} \sum_{i=1}^{n} \mathbf{x}_{i} \\
\hat{\Sigma}=\frac{1}{n} \sum_{i=1}^{n}\left(\mathbf{x}_{i}-\hat{\mu}\right)\left(\mathbf{x}_{i}-\hat{\mu}\right)^{\top}
\end{gathered}
$$

## Gaussian classifier

$$
p(\mathbf{x}, \mu, \Sigma)=\frac{1}{(2 \pi)^{n / 2}|\Sigma|^{1 / 2}} \exp \left(-\frac{1}{2}(\mathbf{x}-\mu)^{\top} \mathbf{\Sigma}^{-\mathbf{1}}(\mathbf{x}-\mu)\right)
$$

## Different Classes, different Gaussians

$$
\begin{gathered}
p\left(\mathbf{x} \mid C_{1}\right)=p\left(\mathbf{x}, \mu_{1}, \Sigma_{1}\right)=\frac{1}{(2 \pi)^{n / 2}\left|\Sigma_{1}\right|^{1 / 2}} \exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{\top} \mathbf{\Sigma}_{\mathbf{1}}^{-\mathbf{1}}\left(\mathbf{x}-\mu_{1}\right)\right) \\
p\left(\mathbf{x} \mid C_{2}\right)=p\left(\mathbf{x}, \mu_{2}, \Sigma_{2}\right)=\frac{1}{(2 \pi)^{n / 2}\left|\Sigma_{2}\right|^{1 / 2}} \exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{\mathbf{2}}\right)^{\top} \mathbf{\Sigma}_{\mathbf{2}}^{-\mathbf{1}}\left(\mathbf{x}-\mu_{\mathbf{2}}\right)\right) \\
\vdots \\
p\left(\mathbf{x} \mid C_{k}\right)=p\left(\mathbf{x}, \mu_{k}, \Sigma_{k}\right)=\frac{1}{(2 \pi)^{n / 2}\left|\Sigma_{k}\right|^{1 / 2}} \exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{k}\right)^{\top} \Sigma_{k}^{-1}\left(\mathbf{x}-\mu_{k}\right)\right)
\end{gathered}
$$

## Gaussian Classifier

- For each class we need:
- Mean Vector
- Covariance Matrix
- Training
- "Fit" a Gaussian to each class
- Find the best Gaussian to explain the distribution for the class

- Classification:

$$
\arg \max _{i} P\left(C_{i}\right) p\left(\mathbf{x}, \mu_{i}, \Sigma_{i}\right)
$$

- Problem:
- Many parameters to train!
- Dominated by covariance: for $D$-dimensional data the covariance matrices requires $D^{2}$ parameters each
- For $N_{c}$ classes, a total of $N_{c} D^{2}$ parameters


## Homo-skedasticity assumption

- Assume all distributions have the same covariance
$-\Sigma_{i}=\Sigma \forall i$
- Assumption, may not be true
- But still works in many cases

- Fewer parameters to train
- One common covariance matrix for all classes
- Only $D^{2}$ total parameters
- As opposed to $N_{c} D^{2}$ if each class has its own covariance matrix


## Homo-skedastic Gaussians

$$
\Sigma_{1}=\Sigma_{2}=\cdots=\Sigma_{K}=\Sigma
$$

- For the binary classification case ( $K=2$ )


## Decision boundary:

$$
\begin{aligned}
& P\left(X \mid C_{2}\right) P\left(C_{2}\right)=P\left(X \mid C_{2}\right) P\left(C_{2}\right) \\
& \frac{1}{(2 \pi)^{n / 2}|\Sigma|^{1 / 2}} \exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{1}\right)\right) \cdot p\left(C_{1}\right)=\frac{1}{(2 \pi)^{n / 2}|\Sigma|^{1 / 2}} \exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{2}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{2}\right)\right) \cdot p\left(C_{2}\right) \\
& \exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{1}\right)\right) \cdot p\left(C_{1}\right)=\exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{2}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{2}\right)\right) \cdot p\left(C_{2}\right)
\end{aligned}
$$

## Homo-skedastic Gaussians

$$
\exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{1}\right)\right) \cdot p\left(C_{1}\right)=\exp \left(-\frac{1}{2}\left(\mathbf{x}-\mu_{2}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{2}\right)\right) \cdot p\left(C_{2}\right)
$$

taking log

$$
\begin{gathered}
\left(-\frac{1}{2}\left(\mathbf{x}-\mu_{1}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{1}\right)\right)+\log p\left(C_{1}\right)=\left(-\frac{1}{2}\left(\mathbf{x}-\mu_{2}\right)^{\top} \Sigma^{-1}\left(\mathbf{x}-\mu_{2}\right)\right)+\log p\left(C_{2}\right) \\
\mu_{1}^{\top} \Sigma^{-1} \mathbf{x}-\frac{1}{2} \mu_{1}^{\top} \Sigma^{-1} \mu_{1}+\log \cdot p\left(C_{1}\right)=\mu_{2}^{\top} \Sigma^{-1} \mathbf{x}-\frac{1}{2} \mu_{2}^{\top} \Sigma^{-1} \mu_{2}+\log p\left(C_{2}\right) \\
\left(\mu_{1}-\mu_{2}\right)^{\top} \Sigma^{-1} \mathbf{x}-\frac{1}{2} \mu_{1}^{\top} \Sigma^{-1} \mu_{1}+\frac{1}{2} \mu_{2}^{\top} \Sigma^{-1} \mu_{2}+\log p\left(C_{1}\right)-\log p\left(C_{2}\right)=0 \\
\mathbf{W} \mathbf{X}+\mathbf{C}=\mathbf{Q} \text { Linear } \\
\text { Boundaryl!!! }
\end{gathered}
$$

## Homo-skedastic Gaussians





## Homo-skedastic Gaussians, K > 2

- Case K > 2 (more than two classes)
- Classification performed as:

$$
\underset{i}{\operatorname{argmax}} P\left(C_{i}\right) p\left(\mathbf{x}, \mu_{i}, \Sigma_{i}\right)
$$

- Taking logs and ignoring the common constant

$$
\underset{i}{\operatorname{argmax}}-\frac{1}{2}\left(\mathbf{x}-\mu_{i}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{i}\right)+\log P\left(C_{i}\right)
$$

- Expanding out and ignoring common terms

$$
\underset{i}{\operatorname{argmax}}-\frac{1}{2} \mathbf{x}^{T} \Sigma^{-1} \mu_{i}-\frac{1}{2} \mu_{i}^{T} \Sigma^{-1} \mu_{i}+\log P\left(C_{i}\right)
$$

- This is just a linear classifier


## Homo-skedastic Gaussians, K > 2



- Decision boundaries for

$$
\underset{i}{\operatorname{argmax}} P\left(C_{i}\right) p\left(\mathbf{x}, \mu_{i}, \Sigma_{i}\right)
$$

- Linear classifier: Decision boundaries are hyperplanes


## Homo-skedastic Gaussians, K > 2

- Case K > 2 (more than two classes)
- Classification performed as:

$$
\underset{i}{\operatorname{argmax}} P\left(C_{i}\right) p\left(\mathbf{x}, \mu_{i}, \Sigma_{i}\right)
$$

- Taking logs and ignoring the common constant

$$
\underset{i}{\operatorname{argmax}}-\frac{1}{2}\left(\mathbf{x}-\mu_{i}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{i}\right)+\log P\left(C_{i}\right)
$$

- Changing the sign and rewriting it as argmin

$$
\underset{i}{\operatorname{argmin}}\left(\mathbf{x}-\mu_{i}\right)^{T} \Sigma^{-1}\left(\mathbf{x}-\mu_{i}\right)-2 \log P\left(C_{i}\right)
$$

## Homo-skedastic Gaussians

## Mahalanobis Distance

$$
D_{M}(\mathbf{x}, \mathbf{y})=\sqrt{(\mathbf{x}-\mathbf{y})^{\top} \Sigma^{-1}(\mathbf{x}-\mathbf{y})}
$$



- A Gaussian Classifier with common Covariance Matrix is similar to a Nearest Neighbor Classifier
- Classification corresponds to the nearest mean vector


## How to estimate the Covariance

## Matrix?

- Maximum likelihood estimate of covariances of individual classes:
- Estimate of common covariance for all classes

$$
\Sigma=\frac{1}{\sum_{C \prime=1}^{K} N_{C^{\prime}}} \sum_{C=1}^{K} N_{C} \Sigma_{C}
$$

## Hetero skedastic Gaussians..

- Homoskedastic Gaussians do not capture nonlinear decision boundaries
- Also, the assumption that all Gaussians have the same covariance is questionable
- Permitting each Gaussian to have its own covariance results in non-linear decision boundaries
- "Hetero skedastic" Gaussians


## Hetero-skedastic Gaussians

## Different Covariance Matrices

1D case. $\mathrm{K}=2$


Decision Boundary $\quad p\left(\mathbf{x} \mid C_{1}\right) p\left(C_{1}\right)=p\left(\mathbf{x} \mid C_{2}\right) p\left(C_{2}\right)$

$$
\begin{array}{ll}
\log \frac{\sigma_{2}}{\sigma_{1}}+\frac{1}{2}\left(\frac{x-\mu_{2}}{\sigma_{2}}\right)^{2}-\frac{1}{2}\left(\frac{x-\mu_{1}}{\sigma_{1}}\right)^{2}-\log \frac{P\left(C_{2}\right)}{P\left(C_{1}\right)}=0 \\
\left(x-x_{1}\right)\left(x-x_{2}\right)=0 & \text { Quadratic }
\end{array}
$$

## Hetero-skedastic Gaussians

$$
\mathrm{x}_{1} \sim \mathcal{N}\left(\left[\begin{array}{l}
0 \\
0
\end{array}\right],\left[\begin{array}{ll}
1 & 1 \\
1 & 8
\end{array}\right]\right), \mathrm{x}_{2} \sim \mathcal{N}\left(\left[\begin{array}{l}
4 \\
0
\end{array}\right],\left[\begin{array}{ll}
2 & 0 \\
0 & 2
\end{array}\right]\right)
$$



## Digit recognition



## Gaussian Classifier for Digit recognition



$$
p(\mathbf{x}, \mu, \Sigma)=\frac{1}{(2 \pi)^{n / 2}|\Sigma|^{1 / 2}} \exp \left(-\frac{1}{2}(\mathbf{x}-\mu)^{\top} \boldsymbol{\Sigma}^{-\mathbf{1}}(\mathbf{x}-\mu)\right)
$$

## Showing the average of each digit

- Average digit



## Bayes classifcation

$$
\arg \max _{i} p\left(C_{i}\right) p\left(\mathbf{x}, \mu_{i}, \Sigma\right)
$$

- Normalize the Posterior




## Inadequacy of Gaussian classifiers




- Gaussian classifiers can only capture simple linear or quadratic decision boundaries
- Often, the decision boundaries required are more complex
- In this case we must employ a Gaussian Mixture classifier

$$
\underset{i}{\operatorname{argmax}} P\left(C_{i}\right) p\left(\mathbf{x} \mid C_{i}\right)
$$

- $p\left(\mathbf{x} \mid C_{i}\right)$ is modelled by a Gaussian mixture


## GMM classifier



- For each class, train a GMM (with EM)
$p\left(\mathbf{x} \mid C_{i}\right)=\sum_{j=1}^{K} \pi_{j}^{(i)} p\left(\mathbf{x} \mid \mu_{j}^{(i)}, \Sigma_{j}^{(i)}\right)$
- Classify according to $\arg \max _{i} p\left(\mathbf{x} \mid C_{i}\right) \cdot p\left(C_{i}\right)$


## Bayesian Classification with Gaussian Mixtures



- Plotting $P\left(C_{i}\right) p\left(\mathbf{x} \mid C_{i}\right)$ for all classes
- Left: Two-class classification, Right: Three-class classification
- Each class modelled by a mixture of three Gaussians
- Note the complex decision boundary


## Estimating $\mathrm{P}(\mathrm{C})$

 $\operatorname{argmax} P\left(C_{i}\right) p\left(\mathbf{x}, \mu_{i}, \Sigma_{i}\right)$$i$

- Have not explained where the class prior $P\left(C_{i}\right)$ comes from
- This can be dependent on the test data
- Typical solutions:
- Estimate from training data
- Optimize on development or held-out test data
- Heuristic guess
- Conservative estimates
- Set the prior of classes that have high cost if incorrectly detected to be low
- Set prior of classes that have high cost if incorrectly missed to be low
- Etc..


## Topics not covered

- Maximum a posteriori estimation
- When we make assumptions about the parameters (means, covariances) themselves
- MAP regression with Gaussians
- Predicting continuous-valued RVs assuming Gaussian distributions
- MAP regression with Gaussian Mixtures
- Predicting continuous-valued RVs with Gaussian mixture distributions
- Time-series and other structured data
- Partially covered in a later lecture

