# Machine Learning - Lecture 6 

## Linear Discriminants II

$$
09.05 .2016
$$

Bastian Leibe<br>RWTH Aachen<br>http://www.vision.rwth-aachen.de

leibe@vision.rwth-aachen.de

## Course Outline

- Fundamentals (2 weeks)
, Bayes Decision Theory
, Probability Density Estimation
- Discriminative Approaches (5 weeks)
, Linear Discriminant Functions
, Support Vector Machines

, Ensemble Methods \& Boosting
> Randomized Trees, Forests \& Ferns
- Generative Models (4 weeks)
, Bayesian Networks
, Markov Random Fields



## Recap: Linear Discriminant Functions

- Basic idea
, Directly encode decision boundary
, Minimize misclassification probability directly.
- Linear discriminant functions

$$
y(\mathbf{x})=\mathbf{w}^{\mathrm{T}} \mathbf{x}+w_{\text {weight vector }}^{w_{0}}
$$

> $\mathbf{w}, w_{0}$ define a hyperplane in $\mathbb{R}^{D}$.

, If a data set can be perfectly classified by a linear discriminant, then we call it linearly separable.

## Recap: Least-Squares Classification

- Simplest approach
, Directly try to minimize the sum-of-squares error

$$
\begin{aligned}
E(\mathbf{w}) & =\sum_{n=1}^{N}\left(y\left(\mathbf{x}_{n} ; \mathbf{w}\right)-\mathbf{t}_{n}\right)^{2} \\
E_{D}(\widetilde{\mathbf{W}}) & =\frac{1}{2} \operatorname{Tr}\left\{(\widetilde{\mathbf{X}} \widetilde{\mathbf{W}}-\mathbf{T})^{\mathrm{T}}(\widetilde{\mathbf{X}} \widetilde{\mathbf{W}}-\mathbf{T})\right\}
\end{aligned}
$$

, Setting the derivative to zero yields

$$
\widetilde{\mathbf{W}}=\left(\widetilde{\mathbf{X}}^{\mathrm{T}} \widetilde{\mathbf{X}}\right)^{-1} \widetilde{\mathbf{X}}^{\mathrm{T}} \mathbf{T}=\widetilde{\mathbf{X}}^{\dagger} \mathbf{T}=\left(\widetilde{\mathbf{X}}^{\mathrm{T}} \widetilde{\mathbf{X}}\right)^{-1} \widetilde{\mathbf{X}}^{\mathrm{T}} \mathbf{T}
$$

, We then obtain the discriminant function as

$$
\mathbf{y}(\mathbf{x})=\widetilde{\mathbf{W}}^{\mathrm{T}} \widetilde{\mathbf{x}}=\mathbf{T}^{\mathrm{T}}\left(\widetilde{\mathbf{X}}^{\dagger}\right)^{\mathrm{T}} \widetilde{\mathbf{x}}
$$

$\Rightarrow$ Exact, closed-form solution for the discriminant function parameters.

## Recap: Problems with Least Squares




- Least-squares is very sensitive to outliers!
, The error function penalizes predictions that are "too correct".


## Recap: Generalized Linear Models

- Generalized linear model

$$
y(\mathbf{x})=g\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}+w_{0}\right)
$$

- $g(\cdot)$ is called an activation function and may be nonlinear.
, The decision surfaces correspond to

$$
y(\mathbf{x})=\text { const } . \quad \Leftrightarrow \quad \mathbf{w}^{\mathrm{T}} \mathbf{x}+w_{0}=\text { const } .
$$

> If $g$ is monotonous (which is typically the case), the resulting decision boundaries are still linear functions of x .

- Advantages of the non-linearity
, Can be used to bound the influence of outliers and "too correct" data points.
, When using a sigmoid for $g(\cdot)$, we can interpret the $y(\mathbf{x})$ as posterior probabilities.

$$
g(a) \equiv \frac{1}{1+\exp (-a)}
$$

## Recap: Linear Separability

- Up to now: restrictive assumption
, Only consider linear decision boundaries
- Classical counterexample: XOR


Slide credit: Bernt Schiele
B. Leibe

## Linear Separability

- Even if the data is not linearly separable, a linear decision boundary may still be "optimal".
, Generalization
, E.g. in the case of Normal distributed data (with equal covariance matrices)

- Choice of the right discriminant function is important and should be based on
, Prior knowledge (of the general functional form)
, Empirical comparison of alternative models
, Linear discriminants are often used as benchmark.


## Generalized Linear Discriminants

- Generalization
, Transform vector $\mathbf{x}$ with $M$ nonlinear basis functions $\phi_{j}(\mathbf{x})$ :

$$
y_{k}(\mathbf{x})=\sum_{j=1}^{M} w_{k j} \phi_{j}(\mathbf{x})+w_{k 0}
$$

, Purpose of $\phi_{j}(\mathbf{x})$ : basis functions
, Allow non-linear decision boundaries.
, By choosing the right $\phi_{j}$, every continuous function can (in principle) be approximated with arbitrary accuracy.

- Notation

$$
y_{k}(\mathbf{x})=\sum_{j=0}^{M} w_{k j} \phi_{j}(\mathbf{x}) \quad \text { with } \phi_{0}(\mathbf{x})=1
$$

[^0]
## Generalized Linear Discriminants

- Model

$$
y_{k}(\mathbf{x})=\sum_{j=0}^{M} w_{k j} \phi_{j}(\mathbf{x})=y_{k}(\mathbf{x} ; \mathbf{w})
$$

, $K$ functions (outputs) $y_{k}(\mathbf{x} ; \mathbf{w})$

- Learning in Neural Networks
, Single-layer networks: $\phi_{j}$ are fixed, only weights $\mathbf{w}$ are learned.
, Multi-layer networks: both the $\mathbf{w}$ and the $\phi_{j}$ are learned.
, In the following, we will not go into details about neural networks in particular, but consider generalized linear discriminants in general...


## Gradient Descent

- Learning the weights w :
, $N$ training data points:
> $K$ outputs of decision functions:
, Target vector for each data point:

$$
\begin{aligned}
& \mathbf{X}=\left\{\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right\} \\
& y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right) \\
& \mathbf{T}=\left\{\mathbf{t}_{1}, \ldots, \mathbf{t}_{N}\right\}
\end{aligned}
$$

, Error function (least-squares error) of linear model

$$
\begin{aligned}
E(\mathbf{w}) & =\frac{1}{2} \sum_{n=1}^{N} \sum_{k=1}^{K}\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right)^{2} \\
& =\frac{1}{2} \sum_{n=1}^{N} \sum_{k=1}^{K}\left(\sum_{j=1}^{M} w_{k j} \phi_{j}\left(\mathbf{x}_{n}\right)-t_{k n}\right)^{2}
\end{aligned}
$$

## Gradient Descent

- Problem
, The error function can in general no longer be minimized in closed form.
- Idea (Gradient Descent)
, Iterative minimization
, Start with an initial guess for the parameter values $w_{k j}^{(0)}$.
> Move towards a (local) minimum by following the gradient.

$$
w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\left.\eta \frac{\partial E(\mathbf{w})}{\partial w_{k j}}\right|_{\mathbf{w}^{(\tau)}}
$$

$\eta$ : Learning rate
 (There are more complex procedures available).

## R

## Gradient Descent - Basic Strategies

- "Batch learning"

$$
w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\left.\eta \frac{\partial E(\mathbf{w})}{\partial w_{k j}}\right|_{\mathbf{w}^{(\tau)}}
$$

$\eta$ : Learning rate
, Compute the gradient based on all training data:

$$
\frac{\partial E(\mathbf{w})}{\partial w_{k j}}
$$

## r

## Gradient Descent - Basic Strategies

- "Sequential updating"

$$
\begin{aligned}
& E(\mathbf{w})=\sum_{n=1}^{N} E_{n}(\mathbf{w}) \\
& w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\left.\eta \frac{\partial E_{n}(\mathbf{w})}{\partial w_{k j}}\right|_{\mathbf{w}^{(\tau)}} \\
& \eta: \text { Learning rate }
\end{aligned}
$$

## Gradient Descent

- Error function

$$
\begin{aligned}
E(\mathbf{w})=\sum_{n=1}^{N} E_{n}(\mathbf{w}) & =\frac{1}{2} \sum_{n=1}^{N} \sum_{k=1}^{K}\left(\sum_{j=1}^{M} w_{k j} \phi_{j}\left(\mathbf{x}_{n}\right)-t_{k n}\right)^{2} \\
E_{n}(\mathbf{w}) & =\frac{1}{2} \sum_{k=1}^{K}\left(\sum_{j=1}^{M} w_{k j} \phi_{j}\left(\mathbf{x}_{n}\right)-t_{k n}\right)^{2} \\
\frac{\partial E_{n}(\mathbf{w})}{\partial w_{k j}} & =\left(\sum_{\tilde{j}=1}^{M} w_{k \tilde{j}} \phi_{\tilde{j}}\left(\mathbf{x}_{n}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right) \\
& =\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right)
\end{aligned}
$$

## Gradient Descent

- Delta rule (=LMS rule)

$$
\begin{aligned}
w_{k j}^{(\tau+1)} & =w_{k j}^{(\tau)}-\eta\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right) \\
& =w_{k j}^{(\tau)}-\eta \delta_{k n} \phi_{j}\left(\mathbf{x}_{n}\right)
\end{aligned}
$$

, where

$$
\delta_{k n}=y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}
$$

$\Rightarrow$ Simply feed back the input data point, weighted by the classification error.

## Gradient Descent

- Cases with differentiable, non-linear activation function

$$
y_{k}(\mathbf{x})=g\left(a_{k}\right)=g\left(\sum_{j=0}^{M} w_{k i} \phi_{j}\left(\mathbf{x}_{n}\right)\right)
$$

- Gradient descent

$$
\begin{aligned}
\frac{\partial E_{n}(\mathbf{w})}{\partial w_{k j}} & =\frac{\partial g\left(a_{k}\right)}{\partial w_{k j}}\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right) \\
w_{k j}^{(\tau+1)} & =w_{k j}^{(\tau)}-\eta \delta_{k n} \phi_{j}\left(\mathbf{x}_{n}\right) \\
\delta_{k n} & =\frac{\partial g\left(a_{k}\right)}{\partial w_{k j}}\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right)
\end{aligned}
$$

## Summary: Generalized Linear Discriminants

- Properties
, General class of decision functions.
> Nonlinearity $g(\cdot)$ and basis functions $\phi_{j}$ allow us to address linearly non-separable problems.
, Shown simple sequential learning approach for parameter estimation using gradient descent.
, Better $2^{\text {nd }}$ order gradient descent approaches available (e.g. Newton-Raphson).
- Limitations / Caveats
, Flexibility of model is limited by curse of dimensionality
- $g(\cdot)$ and $\phi_{j}$ often introduce additional parameters.
- Models are either limited to lower-dimensional input space or need to share parameters.
, Linearly separable case often leads to overfitting.
- Several possible parameter choices minimize training error.


## Topics of This Lecture

- Fisher's linear discriminant (FLD)
, Classification as dimensionality reduction
, Linear discriminant analysis
, Multiple discriminant analysis
, Applications
- Logistic Regression
, Probabilistic discriminative models
, Logistic sigmoid (logit function)
, Cross-entropy error
, Gradient descent
, Iteratively Reweighted Least Squares
- Note on Error Functions


## Classification as Dimensionality Reduction

- Classification as dimensionality reduction
, We can interpret the linear classification model as a projection onto a lower-dimensional space.
> E.g., take the $D$-dimensional input vector x and project it down to one dimension by applying the function

$$
y=\mathbf{w}^{\mathrm{T}} \mathbf{x}
$$

, If we now place a threshold at $y \geq-w_{0}$, we obtain our standard two-class linear classifier.
, The classifier will have a lower error the better this projection separates the two classes.
$\Rightarrow$ New interpretation of the learning problem
, Try to find the projection vector $\mathbf{w}$ that maximizes the class separation.

## Classification as Dimensionality Reduction

bad separation

good separation


- Two questions
, How to measure class separation?
, How to find the best projection (with maximal class separation)?


## Classification as Dimensionality Reduction

- Measuring class separation
, We could simply measure the separation of the class means.
$\Rightarrow$ Choose w so as to maximize

$$
\left(m_{2}-m_{1}\right)=\mathbf{w}^{\mathrm{T}}\left(\mathbf{m}_{2}-\mathbf{m}_{1}\right)
$$



- Problems with this approach

1. This expression can be made arbitrarily large by increasing $\|\mathbf{w}\|$.
$\Rightarrow$ Need to enforce additional constraint $\|\mathbf{w}\|=1$.
$\Rightarrow$ This constrained minimization results in $\mathbf{w} \propto\left(\mathbf{m}_{2}-\mathbf{m}_{1}\right)$
2. The criterion may result in bad separation if the clusters have elongated shapes.

## Fisher's Linear Discriminant Analysis (FLD)

- Better idea:
, Find a projection that maximizes the ratio of the between-class variance to the within-class variance:

$$
J(\mathbf{w})=\frac{\left(m_{2}-m_{1}\right)^{2}}{s_{1}^{2}+s_{2}^{2}} \quad \text { with } \quad s_{k}^{2}=\sum_{n \in \mathcal{C}_{k}}\left(y_{n}-m_{k}\right)^{2}
$$

, Usually, this is written as

$$
J(\mathbf{w})=\frac{\mathbf{w}^{\mathrm{T}} \mathbf{S}_{B} \mathbf{w}}{\mathbf{w}^{\mathrm{T}} \mathbf{S}_{W} \mathbf{w}}
$$

, where

$$
\begin{array}{ll}
\mathbf{S}_{B}=\left(\mathbf{m}_{2}-\mathbf{m}_{1}\right)\left(\mathbf{m}_{2}-\mathbf{m}_{1}\right)^{\mathrm{T}} & \begin{array}{c}
\text { between-class } \\
\text { scatter matrix }
\end{array} \\
\mathbf{S}_{W}=\sum_{k=1}^{2} \sum_{n \in \mathcal{C}_{k}}\left(\mathbf{x}_{n}-\mathbf{m}_{k}\right)\left(\mathbf{x}_{n}-\mathbf{m}_{k}\right)^{\mathrm{T}} & \begin{array}{c}
\text { within-class } \\
\text { scatter matrix }
\end{array}
\end{array}
$$

## Fisher's Linear Discriminant Analysis (FLD)



- Maximize distance between classes
- Minimize distance within a class
- Criterion: $J(\mathbf{w})=\frac{\mathbf{w}^{\mathrm{T}} \mathbf{S}_{B} \mathbf{w}}{\mathbf{w}^{\mathrm{T}} \mathbf{S}_{W} \mathbf{w}}$
$\mathbf{S}_{B} \ldots$ between-class scatter matrix $\mathbf{S}_{W} \ldots$ within-class scatter matrix
- The optimal solution for $\mathbf{w}$ can be obtained as:

$$
\mathbf{w} \propto \mathbf{S}_{W}^{-1}\left(\mathbf{m}_{2}-\mathbf{m}_{1}\right)
$$

- Classification function:

$$
y(\mathbf{x})=\mathbf{w}^{T} \mathbf{x}+w_{0} \stackrel{\text { Class } 1}{\gtrless} 0
$$

## Multiple Discriminant Analysis

- Generalization to $K$ classes

$$
J(\mathbf{W})=\frac{\left|\mathbf{W}^{\mathrm{T}} \mathbf{S}_{B} \mathbf{W}\right|}{\left|\mathbf{W}^{\mathrm{T}} \mathbf{S}_{W} \mathbf{W}\right|}
$$

, where

$$
\begin{aligned}
& \mathbf{W}=\left[\mathbf{w}_{1}, \ldots, \mathbf{w}_{K}\right] \quad \mathbf{m}=\frac{1}{N} \sum_{n=1}^{N} \mathbf{x}_{n}=\frac{1}{N} \sum_{k=1}^{K} N_{k} \mathbf{m}_{k} \\
& \mathbf{S}_{B}=\sum_{k=1}^{K} N_{k}\left(\mathbf{m}_{k}-\mathbf{m}\right)\left(\mathbf{m}_{k}-\mathbf{m}\right)^{\mathrm{T}} \\
& \mathbf{S}_{W}=\sum_{k=1}^{K} \sum_{n \in \mathcal{C}_{k}}\left(\mathbf{x}_{n}-\mathbf{m}_{k}\right)\left(\mathbf{x}_{n}-\mathbf{m}_{k}\right)^{\mathrm{T}}
\end{aligned}
$$

## Maximizing J(W)

- "Rayleigh quotient" $\Rightarrow$ Generalized eigenvalue problem

$$
J(\mathbf{W})=\frac{\left|\mathbf{W}^{\mathrm{T}} \mathbf{S}_{B} \mathbf{W}\right|}{\left|\mathbf{W}^{\mathrm{T}} \mathbf{S}_{W} \mathbf{W}\right|}
$$

, The columns of the optimal $\mathbf{W}$ are the eigenvectors corresponding to the largest eigenvalues of

$$
\mathbf{S}_{B} \mathbf{w}_{i}=\lambda_{i} \mathbf{S}_{W} \mathbf{w}_{i}
$$

, Defining $\mathbf{v}=\mathbf{S}_{B}^{\frac{1}{2}} \mathbf{w}$, we get

$$
\mathbf{S}_{B}^{\frac{1}{2}} \mathbf{S}_{W}^{-1} \mathbf{S}_{B}^{\frac{1}{2}} \mathbf{v}=\lambda \mathbf{v}
$$

which is a regular eigenvalue problem.
$\Rightarrow$ Solve to get eigenvectors of $v$, then from that of $w$.

- For the K-class case we obtain (at most) $K-1$ projections.
> (i.e. eigenvectors corresponding to non-zero eigenvalues.)


## What Does It Mean?

- What does it mean to apply a linear classifier?

- Classifier interpretation
> The weight vector has the same dimensionality as x .
> Positive contributions where $\operatorname{sign}\left(x_{i}\right)=\operatorname{sign}\left(w_{i}\right)$.
$\Rightarrow$ The weight vector identifies which input dimensions are important for positive or negative classification (large $\left|w_{i}\right|$ ) and which ones are irrelevant (near-zero $w_{i}$ ).
$\Rightarrow$ If the inputs x are normalized, we can interpret w as a "template" vector that the classifier tries to match.

$$
\mathbf{w}^{\mathrm{T}} \mathbf{x}=\|\mathbf{w}\|\|\mathbf{x}\| \cos \theta
$$

## Example Application: Fisherfaces

- Visual discrimination task
, Training data:
$C_{1}$ : Subjects with glasses

, Test:

$C_{2}$ : Subjects without glasses



Take each image as a vector of pixel values and apply FLD...

## Fisherfaces: Interpretability

- Resulting weight vector for "Glasses/NoGlasses"



## Summary: Fisher's Linear Discriminant

- Properties
- Simple method for dimensionality reduction, preserves class discriminability.
, Can use parametric methods in reduced-dim. space that might not be feasible in original higher-dim. space.
, Widely used in practical applications.
- Restrictions / Caveats
, Not possible to get more than $K-1$ projections.
, FLD reduces the computation to class means and covariances.
$\Rightarrow$ Implicit assumption that class distributions are unimodal and well-approximated by a Gaussian/hyperellipsoid.


## Topics of This Lecture

- Fisher's linear discriminant (FLD)
, Classification as dimensionality reduction
, Linear discriminant analysis
, Multiple discriminant analysis
- Applications
- Logistic Regression
, Probabilistic discriminative models
, Logistic sigmoid (logit function)
, Cross-entropy error
, Gradient descent
, Iteratively Reweighted Least Squares
- Note on Error Functions


## R

## Probabilistic Discriminative Models

- We have seen that we can write

$$
p\left(\mathcal{C}_{1} \mid \mathbf{x}\right)=\sigma(a)
$$

logistic sigmoid function

- We can obtain the familiar probabilistic model by setting

$$
a=\ln \frac{p\left(\mathbf{x} \mid \mathcal{C}_{1}\right) p\left(\mathcal{C}_{1}\right)}{p\left(\mathbf{x} \mid \mathcal{C}_{2}\right) p\left(\mathcal{C}_{2}\right)}
$$

- Or we can use generalized linear discriminant models

$$
\begin{aligned}
& a=\mathbf{w}^{T} \mathbf{x} \\
& \text { or } \quad a=\mathbf{w}^{T} \boldsymbol{\phi}(\mathbf{x})
\end{aligned}
$$

## R

## Probabilistic Discriminative Models

- In the following, we will consider models of the form

$$
\begin{aligned}
p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right) & =y(\boldsymbol{\phi})=\sigma\left(\mathrm{w}^{T} \boldsymbol{\phi}\right) \\
\text { with } & p\left(\mathcal{C}_{2} \mid \boldsymbol{\phi}\right)
\end{aligned}=1-p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right) \mathrm{l}
$$

- This model is called logistic regression.
- Why should we do this? What advantage does such a model have compared to modeling the probabilities?

$$
p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right)=\frac{p\left(\boldsymbol{\phi} \mid \mathcal{C}_{1}\right) p\left(\mathcal{C}_{1}\right)}{p\left(\boldsymbol{\phi} \mid \mathcal{C}_{1}\right) p\left(\mathcal{C}_{1}\right)+p\left(\boldsymbol{\phi} \mid \mathcal{C}_{2}\right) p\left(\mathcal{C}_{2}\right)}
$$

- Any ideas?


## Comparison

- Let's look at the number of parameters...
, Assume we have an $M$-dimensional feature space $\phi$.
, And assume we represent $p\left(\phi \mid \mathcal{C}_{k}\right)$ and $p\left(\mathcal{C}_{k}\right)$ by Gaussians.
, How many parameters do we need?
- For the means: $2 M$
- For the covariances: $\quad M(M+1) / 2$
- Together with the class priors, this gives $M(M+5) / 2+1$ parameters!
, How many parameters do we need for logistic regression?

$$
p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right)=y(\boldsymbol{\phi})=\sigma\left(\mathrm{w}^{T} \boldsymbol{\phi}\right)
$$

- Just the values of $\mathbf{w} \Rightarrow M$ parameters.
$\Rightarrow$ For large $M$, logistic regression has clear advantages!


## Logistic Sigmoid

- Properties
, Definition: $\sigma(a)=\frac{1}{1+\exp (-a)}$
, Inverse:

$$
a=\ln \left(\frac{\sigma}{1-\sigma}\right)
$$

"logit" function

$$
\sigma(a)=\frac{1}{1+\exp (-a)}
$$

, Symmetry property:

$$
\sigma(-a)=1-\sigma(a)
$$

, Derivative: $\frac{d \sigma}{d a}=\sigma(1-\sigma)$

## Logistic Regression

- Let's consider a data set $\left\{\phi_{n}, t_{n}\right\}$ with $n=1, \ldots, N$, where $\phi_{n}=\phi\left(\mathbf{x}_{n}\right)$ and $t_{n} \in\{0,1\}, \mathbf{t}=\left(t_{1}, \ldots, t_{N}\right)^{T}$.
- With $y_{n}=p\left(\mathcal{C}_{1} \mid \phi_{n}\right)$, we can write the likelihood as

$$
p(\mathbf{t} \mid \mathbf{w})=\prod_{n=1}^{N} y_{n}^{t_{n}}\left\{1-y_{n}\right\}^{1-t_{n}}
$$

- Define the error function as the negative log-likelihood

$$
\begin{aligned}
E(\mathbf{w}) & =-\ln p(\mathbf{t} \mid \mathbf{w}) \\
& =-\sum_{n=1}^{N}\left\{t_{n} \ln y_{n}+\left(1-t_{n}\right) \ln \left(1-y_{n}\right)\right\}
\end{aligned}
$$

, This is the so-called cross-entropy error function.

## Gradient of the Error Function

$$
y_{n}=\sigma\left(\mathbf{w}^{T} \boldsymbol{\phi}_{n}\right)
$$

- Error function

$$
E(\mathbf{w})=-\sum_{n=1}^{N}\left\{t_{n} \ln y_{n}+\left(1-t_{n}\right) \ln \left(1-y_{n}\right)\right\}
$$

- Gradient

$$
\begin{aligned}
\nabla E(\mathbf{w}) & =-\sum_{n=1}^{N}\left\{t_{n} \frac{\frac{d}{d \mathbf{w}} y_{n}}{y_{n}}+\left(1-t_{n}\right) \frac{\frac{d}{d \mathbf{w}}\left(1-y_{n}\right)}{\left(1-y_{n}\right)}\right\} \\
& =-\sum_{n=1}^{N}\left\{t_{n} \frac{y_{n}\left(1-y_{n}\right)}{\phi_{n}} \phi_{n}-\left(1-t_{n}\right) \frac{y_{n}\left(1-y_{n}\right)}{\left(1-y_{n}\right)} \boldsymbol{\phi}_{n}\right\} \\
& =-\sum_{n=1}^{N}\left\{\left(t_{n}-t_{n} y_{n}-y_{n}+t_{\pi} y_{n}\right) \phi_{n}\right\} \\
& =\sum_{n=1}^{N}\left(y_{n}-t_{n}\right) \phi_{n}
\end{aligned}
$$

## Gradient of the Error Function

- Gradient for logistic regression

$$
\nabla E(\mathbf{w})=\sum_{n=1}^{N}\left(y_{n}-t_{n}\right) \phi_{n}
$$

- Does this look familiar to you?
- This is the same result as for the Delta (=LMS) rule

$$
w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\eta\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right)
$$

- We can use this to derive a sequential estimation algorithm.
, However, this will be quite slow...


## A More Efficient Iterative Method...

- Second-order Newton-Raphson gradient descent scheme

$$
\mathbf{w}^{(\tau+1)}=\mathbf{w}^{(\tau)}-\mathbf{H}^{-1} \nabla E(\mathbf{w})
$$

where $\mathbf{H}=\nabla \nabla E(\mathbf{w})$ is the Hessian matrix, i.e. the matrix of second derivatives.

- Properties
, Local quadratic approximation to the log-likelihood.
, Faster convergence.


## Newton-Raphson for Least-Squares Estimation

- Let's first apply Newton-Raphson to the least-squares error function:

$$
\begin{aligned}
& E(\mathbf{w})=\frac{1}{2} \sum_{n=1}^{N}\left(\mathbf{w}^{T} \boldsymbol{\phi}_{n}-t_{n}\right)^{2} \\
& \nabla E(\mathbf{w})=\sum_{n=1}^{N}\left(\mathbf{w}^{T} \boldsymbol{\phi}_{n}-t_{n}\right) \boldsymbol{\phi}_{n}=\boldsymbol{\Phi}^{T} \boldsymbol{\Phi} \mathbf{w}-\boldsymbol{\Phi}^{T} \mathbf{t} \\
& \mathbf{H}=\nabla \nabla E(\mathbf{w})=\sum_{n=1}^{N} \boldsymbol{\phi}_{n} \boldsymbol{\phi}_{n}^{T}=\boldsymbol{\Phi}^{T} \boldsymbol{\Phi} \quad \text { where } \boldsymbol{\Phi}=\left[\begin{array}{c}
\boldsymbol{\phi}_{1}^{T} \\
\vdots \\
\boldsymbol{\phi}_{N}^{T}
\end{array}\right]
\end{aligned}
$$

- Resulting update scheme:

$$
\begin{aligned}
\mathbf{w}^{(\tau+1)} & =\mathbf{w}^{(\tau)}-\left(\boldsymbol{\Phi}^{T} \boldsymbol{\Phi}\right)^{-1}\left(\boldsymbol{\Phi}^{T} \boldsymbol{\Phi} \mathbf{w}^{(\tau)}-\boldsymbol{\Phi}^{T} \mathbf{t}\right) \\
& =\left(\boldsymbol{\Phi}^{T} \boldsymbol{\Phi}\right)^{-1} \boldsymbol{\Phi}^{T} \mathbf{t} \quad \text { Closed-form solution! }
\end{aligned}
$$

## Newton-Raphson for Logistic Regression

- Now, let's try Newton-Raphson on the cross-entropy error function:

$$
\begin{aligned}
& E(\mathbf{w})=-\sum_{n=1}^{N}\left\{t_{n} \ln y_{n}+\left(1-t_{n}\right) \ln \left(1-y_{n}\right)\right\} \\
& \nabla E(\mathbf{w})=\sum_{n=1}^{N}\left(y_{n}-t_{n}\right) \boldsymbol{\phi}_{n}=\boldsymbol{\Phi}^{T}(\mathbf{y}-\mathbf{t}) \\
& \mathbf{H}=\nabla \nabla E(\mathbf{w})=\sum_{n=1}^{N} y_{n}\left(1-y_{n}\right) \boldsymbol{\phi}_{n} \boldsymbol{\phi}_{n}^{T}=\boldsymbol{\Phi}^{T}\left(1-y_{n}\right) \boldsymbol{\phi}_{n} \\
& \mathbf{R} \boldsymbol{\Phi}
\end{aligned}
$$

$$
\text { where } \mathbf{R} \text { is an } N \times N \text { diagonal matrix with } R_{n n}=y_{n}\left(1-y_{n}\right) \text {. }
$$

$\Rightarrow$ The Hessian is no longer constant, but depends on $w$ through the weighting matrix $\mathbf{R}$.

## Iteratively Reweighted Least Squares

- Update equations

$$
\begin{aligned}
\mathbf{w}^{(\tau+1)}= & \mathbf{w}^{(\tau)}-\left(\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi}\right)^{-1} \boldsymbol{\Phi}^{T}(\mathbf{y}-\mathbf{t}) \\
= & \left(\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi}\right)^{-1}\left\{\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi} \mathbf{w}^{(\tau)}-\boldsymbol{\Phi}^{T}(\mathbf{y}-\mathbf{t})\right\} \\
= & \left(\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi}\right)^{-1} \boldsymbol{\Phi}^{T} \mathbf{R} \mathbf{z} \\
& \quad \text { with } \mathbf{z}=\boldsymbol{\Phi} \mathbf{w}^{(\tau)}-\mathbf{R}^{-1}(\mathbf{y}-\mathbf{t})
\end{aligned}
$$

- Again very similar form (normal equations)
, But now with non-constant weighing matrix $\mathbf{R}$ (depends on $\mathbf{w}$ ).
, Need to apply normal equations iteratively.
$\Rightarrow$ Iteratively Reweighted Least-Squares (IRLS)


## Summary: Logistic Regression

- Properties
, Directly represent posterior distribution $p\left(\phi \mid \mathcal{C}_{k}\right)$
> Requires fewer parameters than modeling the likelihood + prior.
, Very often used in statistics.
> It can be shown that the cross-entropy error function is concave
- Optimization leads to unique minimum
- But no closed-form solution exists
- Iterative optimization (IRLS)
> Both online and batch optimizations exist
, There is a multi-class version described in (Bishop Ch.4.3.4).
- Caveat
, Logistic regression tends to systematically overestimate odds ratios when the sample size is less than $\sim 500$.


## Topics of This Lecture

- Fisher's linear discriminant (FLD)
, Classification as dimensionality reduction
- Linear discriminant analysis
, Multiple discriminant analysis
- Applications
- Logistic Regression
, Probabilistic discriminative models
, Logistic sigmoid (logit function)
, Cross-entropy error
, Gradient descent
> Iteratively Reweighted Least Squares
- Note on Error Functions


## Note on Error Functions

$$
t_{n} \in\{-1,1\}
$$



Not differentiable!


- Ideal misclassification error function (black)
, This is what we want to approximate,
, Unfortunately, it is not differentiable.
, The gradient is zero for misclassified points.
$\Rightarrow$ We cannot minimize it by gradient descent.


## Note on Error Functions

$$
t_{n} \in\{-1,1\}
$$



- Squared error used in Least-Squares Classification
, Very popular, leads to closed-form solutions.
, However, sensitive to outliers due to squared penalty.
, Penalizes "too correct" data points
$\Rightarrow$ Generally does not lead to good classifiers.


## Comparing Error Functions (Loss Functions)



- Cross-Entropy Error
, Minimizer of this error is given by posterior class probabilities.
, Concave error function, unique minimum exists.
, Robust to outliers, error increases only roughly linearly
. But no closed-form solution, requires iterative estimation.


## Overview: Error Functions

- Ideal Misclassification Error
, This is what we would like to optimize.
, But cannot compute gradients here.
- Quadratic Error

, Easy to optimize, closed-form solutions exist.
> But not robust to outliers.
- Cross-Entropy Error
, Minimizer of this error is given by posterior class probabilities.
- Concave error function, unique minimum exists.
, But no closed-form solution, requires iterative estimation.
$\Rightarrow$ Analysis tool to compare classification approaches


## References and Further Reading

- More information on Linear Discriminant Functions can be found in Chapter 4 of Bishop's book (in particular Chapter 4.1-4.3).

Christopher M. Bishop
Pattern Recognition and Machine Learning Springer, 2006



[^0]:    Slide credit: Bernt Schiele

