## 数字图像处理

第A章
机器学习模式识别基础

# 机器学习／模式识别基础 

- 预测与识别
- 预测算法
- 特征提取
- 预测：根据当前的观测，预测未观测事件
- 识别：根据当前的观测，判断是否是预定模式
- 如何定义＂事件＂或＂模式＂？
- 机器学习方法：基于数据的定义

1 预测与识别


## 1 预测与识别

## color＝\｛0，1，2，3\} weight=\{0,1,2,3,4\}

| color | weight | sweet？ |
| :---: | :---: | :---: |
| 3 | 4 | yes |
| 2 | 3 | yes |
| 0 | 3 | no |
| 3 | 2 | no |
| 1 | 4 | no |

# 机器学习／模式识别基础 

- 预测与识别
- 预测算法
- 特征提取


## Classification

Features: color, weight
Label: taste is sweet (positive/+) or not (negative/-)

(color, weight) $\rightarrow$ sweet ?

$$
\mathcal{X} \quad \rightarrow\{-1,+1\}
$$

ground-truth function $f$
examples/training data:

$$
\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{m}, y_{m}\right)\right\}
$$

$$
y_{i}=f\left(\boldsymbol{x}_{i}\right)
$$

## Regression

Features: color, weight Label: sweetness [0,1]

(color, weight) $\rightarrow$ sweetness

$$
\mathcal{X} \quad \rightarrow[-1,+1]
$$

ground-truth function $f$
examples/training data: $\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{m}, y_{m}\right)\right\}$

$$
y_{i}=f\left(\boldsymbol{x}_{i}\right)
$$

## Decision tree model

decision process with a tree structure


## Decision tree model


find a decision tree that matches the data?

## Split-criterion: classification

for every possible split of every feature:


## Training error:

#  prediction: - " prediction: + error: 1 error: 3 <br> total error: 4 <br> prediction: - prediction: + error: 3 error: 2 <br> total error: 5 



## Split-criterion: classification

## Information gain (ID3):

Entropy: $H(X)=-\sum p_{i} \ln \left(p_{i}\right)$
Entropy after split: $I(X ;$ split $)=\frac{\text { \#left }}{\# \text { all }} H($ left $)+\frac{\text { \#right }}{\# \text { all }} H$ (right)
Information gain: $H(X)-I(X$;split $)$

$$
\begin{aligned}
& H(\text { left })=-\frac{1}{8} \ln \frac{1}{8}-\frac{7}{8} \ln \frac{7}{8}=0.3768 \\
& H(\text { right })=-\frac{5}{8} \ln \frac{5}{8}-\frac{3}{8} \ln \frac{3}{8}=0.6616 \\
& \mathrm{IG}=H(X)-(0.5 \times 0.3768+0.5 \times 0.6616) \\
& \quad=H(X)-0.5192
\end{aligned}
$$

## Split-criterion: classification

## Gain ratio (C4.5):

$$
\begin{aligned}
& \text { Gain } \operatorname{ratio}(X)=\frac{H(X)-I(X ; \text { split })}{I V(\mathrm{split})} \\
& I V(\text { split })=H(\mathrm{split})
\end{aligned}
$$


e.g. student ID

$$
\mathrm{IG}=H(X)-0
$$

## Split-criterion: classification

Gini index (CART):
Gini: $\operatorname{Gini}(X)=1-\sum_{i} p_{i}^{2}$
Gini after split: $\frac{\text { \#left }}{\text { \#all }} \operatorname{Gini}($ left $)+\frac{\text { \#right }}{\text { \#all }}$ Gini(right)


## Split-criterion: regression

## Training error:



MSE: $8.75+22.83=31.583$
MSE: 43.5833

## Split-criterion: stop

Stop criterion: no feature to use

Classification: examples are pure of class
Regression: variance small enough

## Make-leaf



Classification: major class
Regression: mean value

## DT boundary visualization


decision stump

max depth=2

max depth=12

## Neural networks




NEURON
NEURON
(receivers

## Node of

Node of

## Ranvier

Myelin Sheath
(insulating fatty layer that speeds transmission)

## Neuron / perceptron

output a function of sum of input
linear function:

$$
f\left(\sum_{i} w_{i} x_{i}\right)=\sum_{i} w_{i} x_{i}
$$

threshold function:


$$
f\left(\sum_{i} w_{i} x_{i}\right)=I\left(\sum_{i} w_{i} x_{i}>0\right)
$$

sigmoid function:

$$
f\left(\sum_{i} w_{i} x_{i}\right)=\frac{1}{1+e^{-\Sigma}}
$$

## Limitation of single neuron


[Minsky and Papert, Perceptrons, 1969]

Marvin Minsky Turing Award 1969

AI Winter

## Multi-layer perceptrons

feed-forward network

sigmoid network with one hidden layer can approximate arbitrary function [Cybenko 1989]

## Back-propagation algorithm



$$
\hat{y}=F(\boldsymbol{x})
$$

## gradient descent

$$
\text { error: } E(\boldsymbol{w})=(F(\boldsymbol{x})-y)^{2}
$$

update one weight: $\Delta w_{i, j}=-\eta \frac{\partial E(\boldsymbol{w})}{\partial w_{i, j}}$ weight of the laster layer

$$
\frac{\partial E(\boldsymbol{w})}{\partial w_{i, j}}=\frac{\partial E(\boldsymbol{w})}{\partial F(\boldsymbol{x})} \frac{\partial F(\boldsymbol{x})}{\partial w_{i, j}}
$$

weight of the first layer

$$
\frac{\partial E(\boldsymbol{w})}{\partial w_{i, j}}=\frac{\partial E(\boldsymbol{w})}{\partial F(\boldsymbol{x})} \frac{\partial F(\boldsymbol{x})}{\partial \mathrm{HL} 2} \frac{\partial \mathrm{HL} 2}{\partial \mathrm{HL} 1} \frac{\partial \mathrm{HL} 1}{\partial w_{i, j}}
$$

[Rumelhart, Hinton, Williams, Nature 1986]

## Back-propagation algorithm

For each given training example ( $\mathbf{x}, \mathbf{y}$ ), do

1. Input the instance $\mathbf{x}$ to the NN and compute the output value $o_{u}$ of every output unit $u$ of the network
2. For each network output unit $k$, calculate its error term $\delta_{k}$

$$
\delta_{k} \leftarrow o_{k}\left(1-o_{k}\right)\left(y_{k}-o_{k}\right)
$$

3. For each hidden unit $k$, calculate its error term $\delta_{h}$

$$
\delta_{h} \leftarrow o_{k}\left(1-o_{k}\right) \sum_{k \in \text { outputs }} w_{k h} \delta_{k}
$$

4. Update each network weight $w_{j i}$ which is the weight associated with the $i$-th input value to the unit $j$


$$
w_{j i} \leftarrow w_{j i}+\eta \delta_{j} x_{j i}
$$

[Rumelhart, Hinton, Williams, Nature 1986]

## Advantage and disadvantages

## Smooth and nonlinear decision boundary



Slow convergence
Many local optima
Best network structure unknown


Hard to handle nominal features

## Deep network

## autoencoder:


[Hinton and Salakhutdinov, Science 2006]

## Bayes rule

classification using posterior probability
for binary classification

$$
f(x)= \begin{cases}+1, & P(y=+1 \mid \boldsymbol{x})>P(y=-1 \mid \boldsymbol{x}) \\ -1, & P(y=+1 \mid \boldsymbol{x})<P(y=-1 \mid \boldsymbol{x}) \\ \text { random, }, & \text { otherwise }\end{cases}
$$

in general

$$
\begin{aligned}
f(x) & =\underset{y}{\arg \max } P(y \mid \boldsymbol{x}) \\
& =\underset{y}{\arg \max } P(\boldsymbol{x} \mid y) P(y) / P(\boldsymbol{x}) \\
& =\underset{y}{\arg \max } P(\boldsymbol{x} \mid y) P(y)
\end{aligned}
$$

how the probabilities be estimated

## Naive Bayes

$$
f(x)=\underset{y}{\arg \max } P(\boldsymbol{x} \mid y) P(y)
$$

estimation the a priori by frequency:
$P(y) \leftarrow \tilde{P}(y)=\frac{1}{m} \sum_{i} I\left(y_{i}=y\right)$
assume features are conditional independence given the class (naive assumption):

$$
\begin{aligned}
P(\boldsymbol{x} \mid y) & =P\left(x_{1}, x_{2}, \ldots, x_{n} \mid y\right) \\
& =P\left(x_{1} \mid y\right) \cdot P\left(x_{2} \mid y\right) \cdot \ldots P\left(x_{n} \mid y\right)
\end{aligned}
$$

decision function:

$$
f(x)=\underset{y}{\arg \max } \tilde{P}(y) \prod_{i} \tilde{P}\left(x_{i} \mid y\right)
$$

## Naive Bayes

## color=\{0,1,2,3\} weight=\{0,1,2,3,4\}

| color | weight | sweet? | $P(y=y e s)=2 / 5$ |
| :---: | :---: | :---: | :--- |
| 3 | 4 | yes | $P(y=n o)=3 / 5$ |
| 2 | 3 | yes | $P($ color $=3$ |

$f(y \mid$ color $=3$, weight $=3) \rightarrow$
$P($ color $=3 \mid y=$ yes $) P($ weight $=3 \mid y=$ yes $) P(y=y e s)=0.5 \times 0.5 \times 0.4=0.1$
$P($ color $=3 \mid y=n o) P(w e i g h t=3 \mid y=n o) P(y=n o)=0.33 \times 0.33 \times 0.6=0.06$
$f(y \mid$ color $=0$, weight $=1) \rightarrow$

$$
\begin{aligned}
& P(\text { color }=0 \mid y=y e s) P(\text { weight }=1 \mid y=y e s) P(y=y e s)=0 \\
& P(\text { color }=0 \mid y=n o) P(\text { weight }=1 \mid y=n o) P(y=n o)=0
\end{aligned}
$$

## Naive Bayes

color $=\{0,1,2,3\}$ weight $=\{0,1,2,3,4\}$

| color | weight | sweet? |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 3 | 4 |  | color | sweet? |  |
| 2 | 3 | yes |  | 0 | yes |
| 0 | 3 | yes |  |  |  |
| 3 | 2 | no |  | 1 | yes |
| 1 | 4 | no |  | 2 | yes |
|  |  | no |  | 3 | yes |

smoothed (Laplacian correction) probabilities:
$P($ color $=0 \mid y=y e s)=(0+1) /(2+4)$
$P(y=y e s)=(2+1) /(5+2)$
for counting frequency, assume every event has happened once.
$f(y \mid$ color $=0$, weight $=1) \rightarrow$

$$
\begin{aligned}
& P(\text { color }=0 \mid y=\text { yes }) P(\text { weight }=1 \mid y=\text { yes }) P(y=y e s)=\frac{1}{6} \times \frac{1}{7} \times \frac{3}{7}=0.01 \\
& P(\text { color }=0 \mid y=n o) P(\text { weight }=1 \mid y=n o) P(y=n o)=\frac{2}{7} \times \frac{1}{8} \times \frac{4}{7}=0.02
\end{aligned}
$$

## Nearest neighbor classifier

1-nearest neighbor:

$k$-nearest neighbor:


- asymptotically less than 2 times of the optimal Bayes error
- naturally handle multi-class
- no training time
- nonlinear decision boundary
- slow testing speed for a large training data set
- have to store the training data
- sensitive to similarity function


## Linear model

model space: $\mathbb{R}^{n+1}$

$$
f(\boldsymbol{x})=\boldsymbol{w}^{\top} \boldsymbol{x}+b
$$

we sometimes omit the bias

$$
f(\boldsymbol{x})=\boldsymbol{w}^{\top} \boldsymbol{x}
$$



1. $w$ with a constant element
2. practically as good as with bias (centered data)

## Least square regression

Regression: $y \in \mathbb{R}$
Training data:

$$
\left\{\left(\boldsymbol{x}_{1}, y_{1}\right),\left(\boldsymbol{x}_{2}, y_{2}\right),\left(\boldsymbol{x}_{m}, y_{m}\right)\right\}
$$

Least square loss:

$$
\frac{1}{m} \sum_{i=1}^{m}\left(\boldsymbol{w}^{\top} \boldsymbol{x}_{i}+b-y_{i}\right)^{2}
$$

## Least square regression

$$
L(\boldsymbol{w}, b)=\frac{1}{m} \sum_{i=1}^{m}\left(\boldsymbol{w}^{\top} \boldsymbol{x}_{i}+b-y_{i}\right)^{2}
$$

$$
\frac{\partial L(\boldsymbol{w}, b)}{\partial b}=\frac{1}{m} \sum_{i=1}^{m} 2\left(\boldsymbol{w}^{\top} \boldsymbol{x}_{i}+b-y_{i}\right)=0
$$

$$
\frac{\partial L(\boldsymbol{w}, b)}{\partial \boldsymbol{w}}=\frac{1}{m} \sum_{i=1}^{m} 2\left(\boldsymbol{w}^{\top} \boldsymbol{x}_{i}+b-y_{i}\right) \boldsymbol{x}_{i}=0
$$

$$
b=\frac{1}{m} \sum_{i=1}^{m}\left(y_{i}-\boldsymbol{w}^{\top} \boldsymbol{x}_{i}\right)=\bar{y}-\boldsymbol{w}^{\top} \overline{\boldsymbol{x}}
$$

$$
\boldsymbol{w}=\left(\frac{1}{m} \sum_{i=1}^{m} \boldsymbol{x}_{i} \boldsymbol{x}_{i}^{\top}-\overline{\boldsymbol{x}} \overline{\boldsymbol{x}}^{\top}\right)^{-1}\left(\frac{1}{m} \sum_{i=1}^{m}\left(y_{i} \boldsymbol{x}_{i}\right)-\bar{y} \overline{\boldsymbol{x}}\right)
$$

$$
=\operatorname{var}(\boldsymbol{x})^{-1} \operatorname{cov}(\boldsymbol{x}, y)=\left(X^{\top} X\right)^{-1} X^{\top} Y
$$

## I.I.D. assumption

all training examples and future (test) examples are drawn independently from an identical distribution


## Hypothesis class



## box hypothesis class $\mathcal{H}$ contains all boxes

$h \in \mathcal{H}$ is a hypothesis
$h(\boldsymbol{x})=\left\{\begin{array}{l}+1, \text { if } x \text { is inside the box } \\ -1, \text { if } x \text { is outside the box }\end{array}\right.$

## Training and generalization errors


find a hypothesis minimizes the generalization error

## Generalization error

assume i.i.d. examples, and the ground-truth hypothesis is a box

the error of picking a consistent hypothesis:
with probability at least $1-\delta$

$$
\epsilon_{g}<\frac{1}{m} \cdot\left(\ln |\mathcal{H}|+\ln \frac{1}{\delta}\right)
$$

\author{

- more examples <br> - smaller hypothesis space
}


## Generalization error

for one $h$
What is the probability of $h$ is consistent
What is the probability of

$$
\epsilon_{g}(h) \geq \epsilon
$$

assume $h$ is bad: $\epsilon_{g}(h) \geq \epsilon$
$h$ is consistent with 1 example:

$$
P \leq 1-\epsilon
$$

$h$ is consistent with $\boldsymbol{m}$ example:

$$
P \leq(1-\epsilon)^{m}
$$

## Generalization error

$h$ is consistent with $\boldsymbol{m}$ example:

$$
P \leq(1-\epsilon)^{m}
$$

There are $\boldsymbol{k}$ consistent hypotheses

Probability of choosing a bad one: $h_{1}$ is chosen and $h_{1}$ is bad $P \leq(1-\epsilon)^{m}$
 $h_{2}$ is chosen and $h_{2}$ is bad $P \leq(1-\epsilon)^{m}$
$h_{k}$ is chosen and $h_{k}$ is bad $P \leq(1-\epsilon)^{m}$
overall:
$\exists h$ : $h$ can be chosen (consistent) but is bad

## Generalization error

$h_{1}$ is chosen and $h_{1}$ is bad $P \leq(1-\epsilon)^{m}$ $h_{2}$ is chosen and $h_{2}$ is bad $P \leq(1-\epsilon)^{m}$
$h_{k}$ is chosen and $h_{k}$ is bad $P \leq(1-\epsilon)^{m}$
overall:
$\exists h: h$ can be chosen (consistent) but is bad
Union bound: $P(A \cup B) \leq P(A)+P(B)$
$P(\exists h$ is consistent but bad $) \leq k \cdot(1-\epsilon)^{m} \leq|\mathcal{H}| \cdot(1-\epsilon)^{m}$

## Generalization error

$P(\exists h$ is consistent but bad $) \leq k \cdot(1-\epsilon)^{m} \leq|\mathcal{H}| \cdot(1-\epsilon)^{m}$

$$
P\left(\epsilon_{g} \geq \epsilon\right) \leq \frac{|\mathcal{H}| \cdot(1-\epsilon)^{m}}{\delta}
$$

with probability at least $1-\delta$

$$
\epsilon_{g}<\frac{1}{m} \cdot\left(\ln |\mathcal{H}|+\ln \frac{1}{\delta}\right)
$$

## Inconsistent hypothesis

What if the ground-truth hypothesis is NOT a box: non-zero training error

with probability at least $1-\delta$
$\epsilon_{g}<\epsilon_{t}+\sqrt{\frac{1}{m}\left(\ln |\mathcal{H}|+\ln \frac{1}{\delta}\right)}$
training error

- more examples
smaller generalization error: • smaller hypothesis space
- smaller training error


# 机器学习／模式识别基础 

- 预测与识别
- 预测算法
- 特征提取


## Feature extraction

disclosure the inner structure of the data to support a better mining performance
feature extraction construct new features
commonly followed by a feature selection
usually used for low-level features


## Linear methods

Principal components analysis (PCA)
rotate the data to align the directions of the variance



## Linear methods

Principal components analysis (PCA)
the first dimension $=$ the largest variance direction

$$
\begin{aligned}
& z=\boldsymbol{w}^{T} \boldsymbol{x} \\
& \operatorname{Var}\left(z_{1}\right)=\boldsymbol{w}_{1}^{T} \boldsymbol{\Sigma} \boldsymbol{w}_{1}
\end{aligned}
$$

find a unit $\boldsymbol{w}$ to maximize the variance

$$
\max _{\boldsymbol{w}_{1}} \boldsymbol{w}_{1}^{T} \boldsymbol{\Sigma} \boldsymbol{w}_{1}-\alpha\left(\boldsymbol{w}_{1}^{T} \boldsymbol{w}_{1}-1\right)
$$


$2 \boldsymbol{\Sigma} \boldsymbol{w}_{1}-2 \alpha \boldsymbol{w}_{1}=0$, and therefore $\boldsymbol{\Sigma} \boldsymbol{w}_{1}=\alpha \boldsymbol{w}_{1}$
$\boldsymbol{w}_{1}^{T} \boldsymbol{\Sigma} \boldsymbol{w}_{1}=\alpha \boldsymbol{w}_{1}^{T} \boldsymbol{w}_{1}=\alpha$
$w$ is the eigenvector with the largest eigenvalue

## Linear methods

Principal components analysis (PCA)
the second dimension = the largest variance direction orthogonal to the first dimension

$$
\begin{aligned}
& \max _{\boldsymbol{w}_{2}} \boldsymbol{w}_{2}^{T} \boldsymbol{\Sigma} \boldsymbol{w}_{2}-\alpha\left(\boldsymbol{w}_{2}^{T} \boldsymbol{w}_{2}-1\right)-\beta\left(\boldsymbol{w}_{2}^{T} \boldsymbol{w}_{1}-0\right) \\
& 2 \boldsymbol{\Sigma} \boldsymbol{w}_{2}-2 \alpha \boldsymbol{w}_{2}-\beta \boldsymbol{w}_{1}=0 \\
& \beta=0 \quad \boldsymbol{\Sigma} \boldsymbol{w}_{2}=\alpha \boldsymbol{w}_{2}
\end{aligned}
$$

## Linear methods

Optdigits after PCA


First Eigenvector
from [Intro. ML]

## Linear methods

(a) Scree graph for Optdigits

(b) Proportion of variance explained

from [Intro. ML]

## Linear methods

Multidimensional Scaling (MDS)
keep the distance into a lower dimensional space
for linear transformation, W is an n *k matrix
$\arg \min _{W} \sum_{i, j}\left(\left\|\boldsymbol{x}_{i}^{\top} W-\boldsymbol{x}_{j}^{\top} W\right\|-\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j}\right\|\right)^{2}$


## Linear methods


from [Intro. ML]

## Linear methods

Linear Discriminant Analysis (LDA)
find a direction such that the two classes are well separated

$$
z=\boldsymbol{w}^{T} \boldsymbol{x}
$$

$m$ be the mean of a class $s^{2}$ be the variance of a class

maximize the criterion

$$
J(\boldsymbol{w})=\frac{\left(m_{1}-m_{2}\right)^{2}}{s_{1}^{2}+s_{2}^{2}}
$$

## Linear methods

Linear Discriminant Analysis (LDA)

$$
\begin{aligned}
&\left(m_{1}-m_{2}\right)^{2}=\left(\boldsymbol{w}^{T} \boldsymbol{m}_{1}-\boldsymbol{w}^{T} \boldsymbol{m}_{2}\right)^{2} \\
&=\boldsymbol{w}^{T}\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)^{T} \boldsymbol{w} \\
&=\boldsymbol{w}^{T} \mathbf{S}_{B} \boldsymbol{w} \\
& s_{1}^{2}=\sum_{t}\left(\boldsymbol{w}^{T} \boldsymbol{x}^{t}-m_{1}\right)^{2} \boldsymbol{r}^{t} \\
&= \sum_{t} \boldsymbol{w}^{T}\left(\boldsymbol{x}^{t}-\boldsymbol{m}_{1}\right)\left(\boldsymbol{x}^{t}-\boldsymbol{m}_{1}\right)^{T} \boldsymbol{w} \boldsymbol{r}^{t} \\
&= \boldsymbol{w}^{T} \mathbf{S}_{1} \boldsymbol{w}
\end{aligned}
$$

$$
s_{1}^{2}+s_{2}^{2}=\boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w} \quad \mathbf{S}_{W}=\mathbf{S}_{1}+\mathbf{S}_{2}
$$

The objective becomes:

$$
J(\boldsymbol{w})=\frac{\left(m_{1}-m_{2}\right)^{2}}{s_{1}^{2}+s_{2}^{2}}=\frac{\boldsymbol{w}^{T} \mathbf{S}_{B} \boldsymbol{w}}{\boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w}}=\frac{\left|\boldsymbol{w}^{T}\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)\right|^{2}}{\boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w}}
$$

## Linear methods

Linear Discriminant Analysis (LDA)
The objective becomes:

$$
\begin{array}{r}
J(\boldsymbol{w})=\frac{\left(m_{1}-m_{2}\right)^{2}}{s_{1}^{2}+s_{2}^{2}}=\frac{\boldsymbol{w}^{T} \mathbf{S}_{B} \boldsymbol{w}}{\boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w}}=\frac{\left|\boldsymbol{w}^{T}\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)\right|^{2}}{\boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w}} \\
\frac{\boldsymbol{w}^{T}\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)}{\boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w}}\left(2\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)-\frac{\boldsymbol{w}^{T}\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)}{\boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w}} \mathbf{S}_{W} \boldsymbol{w}\right)=0
\end{array}
$$

Given that $\boldsymbol{w}^{T}\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right) / \boldsymbol{w}^{T} \mathbf{S}_{W} \boldsymbol{w}$ is a constant, we have

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

$$
\text { just take } c=1 \text { and find } w
$$

## Linear methods

Optdigits after LDA

from [Intro. ML]

## Manifold learning



## Manifold learning

A low intrinsic dimensional data embedded in a high dimensional space

cause a bad distance measure



## Manifold learning

## ISOMAP

1. construct a neighborhood graph (kNN and $\varepsilon$-NN)
2. calculate distance matrix as the shortest path on the graph
3. apply MDS on the distance matrix


## Manifold learning

Optdigits after Isomap (with neighborhood graph).


## Manifold learning

Local Linear Embedding (LLE):

1. find neighbors for each instance
2. calculate a linear reconstruction for an instance

$$
\sum_{r}\left\|\boldsymbol{X}^{r}-\sum_{s} \mathbf{W}_{r s} \boldsymbol{X}_{(r)}^{s}\right\|^{2}
$$

3. find low dimensional instances preserving the reconstruction

$$
\sum_{r}\left\|\boldsymbol{z}^{r}-\sum_{s} \mathbf{W}_{r s} \boldsymbol{z}^{s}\right\|^{2}
$$



## Manifold learning



## Manifold learning

## more manifold learning examples



## Manifold learning

more manifold learning examples


机器学习／模式识别基础
小结

- 预测与识别
- 用数据定义＂模式＂
- 预测算法
- 特征提取

