## Algorithmisches Lernen/Machine Learning

Part 1: Stefan Wermter

- Introduction
- Connectionist Learning (e.g. Neural Networks)
- Decision-Trees, Genetic Algorithms

Part 2: Norman Hendrich

- Support-Vector Machines
- Learning of Symbolic Structures
- Bayesian Learning
- Dimensionality Reduction

Part 3: Jianwei Zhang

- Function approximation
- Reinforcement Learning
- Applications in Robotics


## Bayesian Learning

- Bayesian Reasoning
- Bayes Optimal Classifier
- Naïve Bayes Classifier
- Cost-Sensitive Decisions
- Modelling with Probability Density Functions
- Parameter Estimation
- Bayesian Networks
- Markov Models
- Dynamic Bayesian Networks
- Conditional Random Fields


## Bayesian Learning

- Bayesian Reasoning
- Bayes Optimal Classifier
- Naïve Bayes Classifier
- Cost-Sensitive Decisions
- Modelling with Probability Density Functions
- Parameter Estimation
- Bayesian Networks
- Markov Models
- Dynamic Bayesian Networks
- Conditional Random Fields


## Bayesian Reasoning

- derive the probability of a hypothesis $h$ about some observation $\vec{x}$
- a priori probability: probability of the hypothesis prior to the observation $P(h)$
- a posteriori probability: probability of the hypothesis after observation $P(h \mid \vec{x})$
- observation can have discrete or continuous values
- continuous values: probability density functions $p(h \mid \vec{x})$ instead of probabilities
- error optimal decision: choose the hypothesis which maximizes the a posteriori probability (MAP-decision)


## Bayesian Reasoning

- a posteriori probability is difficult to estimate
- Bayes' rule provides the missing link

$$
P(h, \vec{x})=P(\vec{x}, h)=P(h) P(\vec{x} \mid h)=P(\vec{x}) P(h \mid \vec{x})
$$

$$
P(h \mid \vec{x})=\frac{P(h) P(\vec{x} \mid h)}{P(\vec{x})}
$$

## Bayesian Reasoning

- classification: using the posterior probability as a target function

$$
h_{M A P}=\arg \max _{h_{i} \in H} \frac{P\left(h_{i}\right) P\left(\vec{x} \mid h_{i}\right)}{P(\vec{x})}=\arg \max _{h_{i} \in H} P\left(h_{i}\right) P\left(\vec{x} \mid h_{i}\right)
$$

- simplified form: maximum likelihood decision (e.g. if the priors are uniform)

$$
h_{M A P}=\arg \max _{h_{i} \in H} P(\vec{x} \mid h)
$$

## Bayesian Reasoning

- allows
- to include domain knowledge (prior probabilities)
- to deal with inconsistent training data
- to provide probabilistic results (confidence)
- but: probability distributions have to be estimated
$\rightarrow$ usually many parameters


## Bayesian Reasoning

- derived results: Bayesian analysis of learning paradigms may uncover their hidden assumptions, even if they are not probabilistic:
- Every consistent learner outputs a MAP hypothesis under the assumption of uniform prior probabilities for all hypotheses and deterministic, noise-free training data
- If the real training data can be assumed to be produced out of ideal ones by adding a normal-distributed noise term, any learner that minimizes the mean-squared error yields a ML hypothesis
- If an observed Boolean value is a probabilistic function of the input value, minimizing cross entropy in neural networks yields a ML hypothesis


## Bayesian Reasoning

- derived results (cont.):
- If optimal encodings for the hypotheses and the training data given the hypothesis are chosen, selecting the hypothesis according to the principle of minimal description length gives a MAP hypothesis

$$
h_{M D L}=\arg \min _{h \in H} L_{C_{1}}(h)+L_{C_{2}}(D \mid h)
$$

## Bayesian Learning

- Bayesian Reasoning
- Bayes Optimal Classifier
- Naïve Bayes Classifier
- Modelling with Probability Density Functions
- Parameter Estimation
- Bayesian Networks
- Markov Models
- Dynamic Bayesian Networks
- Conditional Random Fields


## Bayes Optimal Classifier

- Bayes classifier does not always produce a true MAP decision
- e.g. for composite results

| hypothesis | $h_{1}$ | $h_{2}$ | $h_{3}$ |
| :--- | :---: | :---: | :---: |
| posterior probability | 0.3 | 0.4 | 0.3 |

- maximum of posteriors gives $h_{2}$
- but if a new observation is classified positive by $h_{2}$ but negative by $h_{1}$ and $h_{3}$ the MAP decision would be "negative"
- extension of the Bayes classifier to composite decisions separating hypotheses $h$ from decisions $v$

$$
v_{M A P}=\arg \max _{v_{j} \in V} \sum_{h_{i} \in H} P\left(v_{j} \mid h_{i}\right) P\left(h_{i} \mid \vec{x}\right)
$$

- simplification for $P(v \mid h) \in\{0,1\}$


## Bayesian Learning

- Bayesian Reasoning
- Bayes Optimal Classifier
- Naïve Bayes Classifier
- Cost-Sensitive Decisions
- Modelling with Probability Density Functions
- Parameter Estimation
- Bayesian Networks
- Markov Models
- Dynamic Bayesian Networks
- Conditional Random Fields


## Naïve Bayes Classifier

- Bayes Optimal Classifier is too expensive

$$
\begin{aligned}
v_{M A P}=\arg \max _{v_{j} \in V} P\left(v_{j} \mid \vec{x}\right) & =\arg \max _{v_{j} \in V} P\left(v_{j} \mid x_{1}, x_{2}, \ldots, x_{n}\right) \\
& =\arg \max _{v_{j} \in V} P\left(v_{j}\right) P\left(x_{1}, x_{2}, \ldots, x_{n} \mid v_{j}\right)
\end{aligned}
$$

- prohibitively many parameter to estimate
- independence assumption:
$P\left(x_{i} \mid v_{j}\right)$ is independent of $P\left(x_{k} \mid v_{j}\right)$ for $i \neq k$

$$
v_{N B}=\arg \max _{v_{j} \in V} P\left(v_{j}\right) \prod_{i} P\left(x_{i} \mid v_{j}\right)
$$

- simple training
- usually good results


## Bayesian Learning

- Bayesian Reasoning
- Bayes Optimal Classifier
- Naïve Bayes Classifier
- Cost-Sensitive Decisions
- Modelling with Probability Density Functions
- Parameter Estimation
- Bayesian Networks
- Markov Models
- Dynamic Bayesian Networks
- Conditional Random Fields


## Cost-Sensitive Decisions

- error optimal classification not always welcome: highly asymmetric distributions
- diseases, errors, failures, ...
- priors determine the decision
- including a cost function into the decision rule
- $c_{i j}$ cost of predicting $i$ when the true class is $j$
- cost matrix

$$
C=\left(\begin{array}{cccc}
c_{11} & c_{12} & \ldots & C_{1 n} \\
c_{21} & c_{22} & \ldots & C_{2 n} \\
\ldots & \ldots & \ldots & \ldots \\
c_{n 1} & c_{n 2} & \ldots & C_{n n}
\end{array}\right)
$$

## Cost-Sensitive Decisions

- Bayes classifier with cost function can help to reduce false positives/negatives

$$
h(\vec{x})=\arg \min _{h_{i} \in H} \sum_{j} c_{i j} p\left(h_{j} \mid \vec{x}\right)
$$

- alternative: biased sampling of training data
- not really effective


## Cost-Sensitive Decisions

- not every cost matrix is a reasonable one
$\rightarrow$ reasonableness conditions
- correct decisions should be less expensive than incorrect ones $c_{i i}<c_{i j} \quad i \neq j$
- a row in the cost matrix should not dominate another one
- row $m$ dominates row $n: \forall j . c_{m j} \geq c_{n j}$
- optimal policy: always decide for the dominated class
- e.g. asymmetric cost function for diseases:

|  | actually not ill | actually ill |
| ---: | :---: | :---: |
| predict not ill | 0 | 1 |
| predict ill | 9 | 0 |

## Cost-Sensitive Decisions

- any two-class cost matrix can be changed by
- adding a constant to every entry (shifting)
- multiplying every entry with a constant (scaling)
without affecting the optimal decision

$$
\left(\begin{array}{ll}
c_{00} & c_{01} \\
c_{10} & c_{11}
\end{array}\right) \quad \Longrightarrow \quad\left(\begin{array}{ll}
0 & c_{01}-c_{00} / c_{10}-c_{00} \\
1 & c_{11}-c_{00} / c_{10}-c_{00}
\end{array}\right)
$$

$\rightarrow$ actually only one degree of freedom!

## Cost-Sensitive Decisions

- optimal decision require the expected cost of the decision to be larger than the expected cost for the alternative decisions e.g. two-class case

$$
\begin{aligned}
& P(\oplus \mid x) c_{10}+P(\ominus \mid x) c_{11} \leq P(\oplus \mid x) c_{00}+P(\ominus \mid x) c_{01} \\
& \left(1-P(\ominus \mid x) c_{10}+P(\ominus \mid x) c_{11} \leq\left(1-P(\ominus \mid x) c_{00}+P(\ominus \mid x) c_{01}\right.\right.
\end{aligned}
$$

- threshold for making optimal cost-sensitive decisions

$$
\begin{aligned}
& \left(1-p^{*}\right) c_{10}+p^{*} c_{11}=\left(1-p^{*}\right) c_{00}+p^{*} c_{01} \\
& p^{*}=\frac{c_{10}-c_{00}}{c_{10}-c_{00}+c_{01}-c_{11}}
\end{aligned}
$$

can be used e.g. in decision tree learning

## Cost-Sensitive Decisions

- costs are a dangerous perspective for many applications
- e.g. rejecting a "good" bank loan application is a missed opportunity not an actual loss
$\rightarrow$ cost are easily measured against different baselines
- benefits provides a more natural (uniform) baseline: cash flow
- costs/benfits are usually not constant for every instance
- e.g. potential benefit/loss of a defaulted bank loan varies with the amount


## Bayesian Learning

- Bayesian Reasoning
- Bayes Optimal Classifier
- Naïve Bayes Classifier
- Cost-Sensitive Decisions
- Modelling with Probability Density Functions
- Parameter Estimation
- Bayesian Networks
- Markov Models
- Conditional Random Fields


## Modelling with Probability Density Functions

- probability density functions $p\left(\vec{x} \mid v_{j}\right)$ instead of $P\left(\vec{x} \mid v_{j}\right)$
- $P\left(\vec{x} \mid v_{j}\right)$ is always zero in a continuous domain
- choosing a distribution class, e.g. Gaussian or Laplace

$$
\begin{aligned}
& p(x \mid v)=\mathcal{N}[x, \mu, \sigma]=\frac{1}{\sqrt{2 \pi \sigma}} e^{-\frac{(x-\mu)^{2}}{2 \sigma^{2}}} \\
& p(x \mid v)=\mathcal{L}[x, \mu, \sigma]=\frac{1}{2 \sigma} e^{-\frac{|x-\mu|}{\sigma}}
\end{aligned}
$$

- parameters: mean $\mu$, variance $\sigma$


## Modelling with Probability Density Functions

- distributions for multidimensional observations
- e.g. multivariate normal distribution

$$
p(\vec{x} \mid v)=\mathcal{N}[\vec{x}, \vec{\mu}, \Sigma]
$$



- parameters
- vector of means $\vec{\mu}$
- co-variance matrix $\Sigma$


## Modelling with Probability Density Functions

 diagonal covariance matrix uniformly filled (rotation symmetry around the mean)

$$
\sigma_{i j}= \begin{cases}n & \text { for } i=j \\ 0 & \text { else }\end{cases}
$$

diagonal covariance matrix
filled with arbitrary values (reflection symmetry)

completely filled covariance matrix


## Modelling with Probability Density Functions

- diagonal covariance matrix: uncorrelated features relativly small number of parameters to be trained $\rightarrow$ naïve Bayes classifier
- completely filled covariance matrix: correlated features high number of parameters to be trained



## Modelling with Probability Density Functions

- decorrelation of the features: transformation of the feature space
- Principal Component Analysis
- Karhunen-Loève-Transformation


## Modelling with Probability Density Functions

- compromise: mixture densities
- superposition of several Gaussians with uncorrelated features

$$
p(\vec{x} \mid v)=\sum_{m=1}^{M} c_{m} \mathcal{N}\left[\vec{x}, \vec{\mu}_{m}, \Sigma_{m}\right]
$$



## Modelling with Probability Density Functions

- compromise: mixture densities
- superposition of several Gaussians with uncorrelated features

$$
p(\vec{x} \mid v)=\sum_{m=1}^{M} c_{m} \mathcal{N}\left[\vec{x}, \vec{\mu}_{m}, \Sigma_{m}\right]
$$



## Modelling with Probability Density Functions

- mixture density functions introduce a hidden variable: Which Gaussian produced the value?
- two step stochastic process:
- choosing a mixture randomly

$$
z_{i j}= \begin{cases}1 & \text { if } \vec{x}_{i} \text { was generated by } p_{j}(\vec{x} \mid v) \\ 0 & \text { otherwise }\end{cases}
$$

- choosing a value randomly

- direct estimation of distribution parameters is not possible


## Bayesian Learning

- Bayesian Reasoning
- Bayes Optimal Classifier
- Naïve Bayes Classifier
- Cost-Sensitive Decisions
- Modelling with Probability Density Functions
- Parameter Estimation
- Bayesian Networks
- Markov Models
- Dynamic Bayesian Networks
- Conditional Random Fields


## Parameter estimation

- complete data
- maximum likelihood estimation
- Bayesian estimation
- incomplete data
- expectation maximization
- (gradient descent techniques)


## Maximum Likelihood Estimation

- likelihood of the model $M$ given the (training) data $\mathcal{D}$

$$
L(M \mid \mathcal{D})=\prod_{d \in \mathcal{D}} P(d \mid M)
$$

- log-likelihood

$$
L L(M \mid \mathcal{D})=\prod_{d \in \mathcal{D}} \log _{2} P(d \mid M)
$$

- choose among several possible models for describing the data according to the principle of maximum likelihood

$$
\hat{\Theta}=\arg \max _{\Theta} L\left(M_{\Theta} \mid \mathcal{D}\right)=\arg \max _{\Theta} L L\left(M_{\Theta} \mid \mathcal{D}\right)
$$

- the models only differ in the set of parameters $\Theta$


## Maximum Likelihood Estimation

- complete data: estimating the parameters by counting

$$
\begin{aligned}
& P(A=a)=\frac{N(A=a)}{\sum_{v \in \operatorname{dom}(A)} N(A=v)} \\
& P(A=a \mid B=b, C=c)=\frac{N(A=a, B=b, C=c)}{N(B=b, C=c)}
\end{aligned}
$$

## Bayesian Estimation

- sparse data bases result in pessimistic estimations for unseen events
- if the count for an event in the data base is 0 , the event ios considered impossible by the model
- Bayesian estimation: using an estimate of the prior probability as starting point for the counting
- estimation of maximum a posteriori parameters
- no zero counts can occur
- if nothing else available use an even distribution as prior
- Bayesian estimate in the binary case with an even distribution

$$
P(\text { yes })=\frac{n+1}{n+m+2}
$$

$n$ : counts for yes, $m$ : counts for no

- effectively adding virtual counts to the estimate
- alternative: smoothing as a post processing step


## Incomplete Data

- missing at random:
- probability that a value is missing depends only on the observed value
- e.g. confirmation measurement: values are available only if the preceding measurement was positive/negative
- missing completely at random
- probability that a value is missing is also independent of the value
- e.g. stochastic failures of the measurement equipment
- e.g. hidden/latent variables (mixture coefficients of a Gaussian mixture distribution)
- nonignorable:
- neither MAR or MCAR
- probability depends on the unseen values, e.g. exit polls for extremist parties


## Expectation Maximization

Estimating the means of a Gaussian mixture distribution

- choose an initial hypothesis for $\Theta=\left(\mu_{1}, \ldots, \mu_{k}\right)$
- estimate the expected mean $E\left(z_{i j}\right)$ given $\Theta=\left(\mu_{1}, \ldots, \mu_{k}\right)$
- recalculate the maximum likelihood estimate of the means: $\Theta^{\prime}=\left(\mu_{1}^{\prime}, \ldots, \mu_{k}^{\prime}\right)$ assuming $z_{i j}$

$$
z_{i j}= \begin{cases}1 & \text { if } \vec{x}_{i} \text { was generated by } p_{j}(\vec{x} \mid v) \\ 0 & \text { otherwise }\end{cases}
$$

- replace $\mu_{j}$ by $\mu_{j}^{\prime}$ and repeat until convergence


## Expectation Maximization

- expectation:
- "complete" the data set using the current estimation $h=\Theta$ to calculate expectations for the missing values
- applies the model to be learned (Bayesian inference)
- maximization:
- use the "completed" data set to find a new maximum likelihood estimation $h^{\prime}=\Theta^{\prime}$


## Expectation Maximization

- generalizing the EM framework
- estimating the underlying distribution of not directly observable variables
- full data $\mathrm{n}+1$-tuples $\left\langle\vec{x}_{i}, z_{i 1}, \ldots, z_{i k}\right\rangle$ only $x_{i}$ can be observed
- training data: $X=\left\{\vec{x}_{1}, \ldots, \vec{x}_{m}\right\}$
- hidden information: $Z=\left\{z_{1}, \ldots, z_{m}\right\}$
- parameters of the distribution to be estimated: $\Theta$
- $Z$ can be treated as random variable with $p(Z)=f(\Theta, X)$
- full data: $Y=X \cup Z$
- hypothesis: $h$ of $\Theta$, needs to be revised into $h^{\prime}$


## Expectation Maximization

- goal of EM: $h^{\prime}=\arg \max E\left(\log _{2} p\left(Y \mid h^{\prime}\right)\right)$
- define a function $Q\left(h^{\prime} \mid h\right)=E\left(\log _{2} p\left(Y \mid h^{\prime}\right) \mid h, X\right)$


## Estimation (E) step

Calculate $Q\left(h^{\prime} \mid h\right)$ using the current hypothesis h and the observed data $X$ to estimate the probability distribution over $Y$

$$
Q\left(h^{\prime} \mid h\right) \leftarrow E\left(\log _{2} p\left(Y \mid h^{\prime}\right) \mid h, X\right)
$$

## Maximization (M) step

Replace hypothesis $h$ by $h^{\prime}$ that maximizes the function $Q$

$$
h \leftarrow \arg \max _{h^{\prime}} Q\left(h^{\prime} \mid h\right)
$$

## Expectation Maximization

- expectation step requires applying the model to be learned
- Bayesian inference
- gradient ascent search
- converges to the next local optimum
- global optimum is not guaranteed


## Expectation Maximization



- If Q is continuous, EM converges to the local maximum of the likelihood function $P\left(Y \mid h^{\prime}\right)$

